Induced Automation Innovation: Evidence from Firm-level Patent Data

Antoine Dechezleprêtre David Hémous Morten Olsen
Carlo Zanella*

June 2023 (First draft: March 2019)

Abstract

Do higher wages induce more automation innovation? We identify automation patents in machinery. We show that a higher automation intensity predicts a decline in routine tasks across US sectors. Then, we estimate how innovating firms respond to changes in their downstream firms' low- and high-skill wages. We compute these wages by combining macroeconomic data on 41 countries with innovating firms' global market exposure. Higher low-skill wages increase automation innovation (but not other machinery innovation) with an elasticity of 2-5. Finally, we show that the German Hartz labor market reforms reduced automation innovations by foreign firms more exposed to Germany.

JEL: O31, O33, J20

KEYWORDS: Automation, Innovation, Patents, Income Inequality

*Antoine Dechezleprêtre, OECD, David Hémous, University of Zurich and CEPR, Morten Olsen, University of Copenhagen, and Carlo Zanella, Swiss Life Asset Managers. David Hémous gratefully acknowledges the financial support of the European Commission under the ERC Starting Grant 805007 AUTOMATION. We thank Daron Acemoglu, Kirill Borusyak, Lorenzo Casaburi, Patrick Gaule, Michael McMahon, Pascual Restrepo, Joachim Voth and Fabrizio Zilibotti among others for helpful comments and suggestions. We also thank many seminar and conference participants. We thank Amedeo Andriollo, Jan Fasnacht, Selina Schön and Shi Suo for fantastic research assistance.

1 Introduction

Do higher wages lead to more labor-saving innovation? And if so, by how much? Economic theory posits that firms innovate more in automation technology when labor costs rise. This mechanism is key for the long-term consequences of policies such as the minimum wage and is at the heart of a growing theoretical literature on the role of automation in economic growth (e.g. Acemoglu and Restrepo, 2018, Hémous and Olsen, 2022). However, empirical evidence for induced automation innovation is lacking. Current research faces two challenges: identifying automation innovations and finding exogenous variation in labor costs.

Accordingly, our paper makes two contributions: First, we develop a new classification of automation innovations using patent data. Second, we isolate exogenous variation in labor costs from the perspective of the innovating firms and provide the first evidence on induced automation innovation at the firm-level. We find that a 1% increase in low-skill wages induces between 2 and 5% more automation innovations. Two policy shocks, changes in the minimum wage and Germany's Hartz reforms, display comparable effects.

Our classification aims to identify automation innovations that allow for the replacement of workers with equipment in some tasks. To classify patents, we follow a two-step procedure: First, we use the text of European patents to identify technology categories in machinery (IPC and CPC codes) associated with automation. Second, we use these technology categories to classify the universe of machinery patents. This procedure leverages that the combined wording of many patents improves the signal of automation characteristics and allows us to classify patents for which we do not have the text. The resulting classification is transparent, uses highly disaggregated technology categories, and covers a wide range of automation technologies. As a validation exercise, we map our patents to sectors of use and reproduce the cross-sectoral analysis of Autor, Levy and Murnane (2003). We find that in the United States, sectors that use more automation-intensive equipment experience a relative decline in routine tasks and the labor share, and a relative increase in the skill ratio.

We then proceed to our main empirical analysis which studies how automation innovations respond to changes in wages. We exploit plausibly exogenous variation in labor costs from the perspective of the innovating firms using a shift-share design. Automation innovators are often equipment manufacturers that sell their machines to downstream firms in different countries. Thus, the incentives of automation producers to innovate depend on the labor costs paid by their downstream firms. To proxy for these labor costs, we compute weighted averages of low- and high-skill labor costs using data on innovating firms' international exposure and country-level labor costs.

We implement this empirical strategy as follows. We rely on the PATSTAT database, which contains close to the universe of patents. We link patents to firms and apply our classification of automation and non-automation patents in machinery. To proxy for firms' international exposure, we use the geographic distribution of their machinery patents pre-sample. We combine these exposure weights with macroeconomic data from 41 countries. Given our focus on international innovation, we restrict attention to patents applied for in at least two countries. Our sample covers the period 1997-2011 and includes 3,255 firms, accounting for 53.4% of global automation innovations. We run Poisson regressions with a 2-year lag between patent applications and labor costs, and we include firm, industry-year, and country-year fixed effects.

We find a substantial effect of wages on automation innovations. Increases in low-skill labor costs (referred to as wages for simplicity) lead to more automation innovations with an elasticity between 2 and 5. In line with the capital-skill complementarity hypothesis (Krusell, Ohanian, Rios-Rull and Violante, 2000), increases in high-skill wages reduce automation innovations by a similar amount. We draw on the recent shift-share literature and interpret our results through the lens of Borusyak, Hull and Jaravel (2022). In our context, identification can be obtained from conditionally randomly assigned wage shocks. We leverage firm-level variation by including country-year fixed effects, which enables us to control for domestic shocks to wages and innovation. Since we control for high-skill wages, our regression coefficients must reflect the effect of foreign demand shocks for automation equipment producers with asymmetric effect on wages. We then argue that these foreign demand shocks are most likely regulatory shocks or labor supply shocks that allow us to identify a causal effect of wages on automation innovations. Importantly, we find that non-automation machinery innovations by the same firms, targeting the same sectors, do not respond to wage shocks.

We supplement this analysis by focusing on two cases where labor cost changes arise directly from policy interventions. First, we build a measure of minimum wages for a subset of countries. We find a positive effect of minimum wages on automation innovations. Second, we focus on a specific labor market shock, the Hartz reforms in

Germany. These were a series of labor market reforms implemented in 2003-2005. They are credited with increasing labor supply and reducing labor costs, notably for low-skill workers (Krause and Uhlig, 2012). Therefore, we predict that these reforms reduced automation innovation. In a difference-in-difference exercise, we find that foreign firms that are relatively more exposed to Germany innovated less in automation technologies after the Hartz reforms. Finally, in a triple-difference exercise, we find that the reforms also decreased automation innovations relative to non-automation innovations.

We are motivated by a theoretical literature on automation and endogenous growth and contribute to three literatures: on induced automation, on induced innovation more generally, and on the measurement of automation. The theoretical argument that higher wages should lead to more labor-saving technology adoption (e.g. Zeira, 1998) and innovation is well-understood. In Hémous and Olsen (2022) and Acemoglu and Restrepo (2018), wages affect the direction of innovation in the form of automation or the creation of new tasks. The quantitative behavior of these directed technical change models depends crucially on the elasticity of innovation with respect to labor costs (Acemoglu, 2023). We estimate such an elasticity and provide empirical support for this literature.[1]

Despite this theoretical literature, existing empirical evidence on the effect of wages on induced automation, adoption or innovation, remains limited.[2] A few papers show that labor market conditions affect adoption of labor-saving technologies in agriculture (Hornbeck and Naidu, 2014, Clemens, Lewis and Postel, 2018, and Voth, Caprettini and Trew, 2022), health care (Acemoglu and Finkelstein, 2008), and manufacturing (Lewis, 2011 and Acemoglu and Restrepo, 2022). Lordan and Neumark (2018) find that minimum wage increases displace workers in automatable jobs and Fan, Hu and Tang (2021) find that they induce Chinese firms to adopt industrial robots. None of these papers use firm-level variation in the manufacturing sector, as we do. More importantly, we focus on *innovation* rather than *adoption*. This is an important distinction: adoption generally refers to firms deciding whether to use an existing technology, while innovation is a continuous process of creating new technologies. While a shock to adoption must die out over time (when all firms have adopted a technology), a shock to innovation can

[1] Acemoglu (2023) uses estimates from our analysis to calibrate a directed technical change model.

[2] In contrast, there is an extensive empirical literature on the effects of technology on wages and employment: see e.g., Autor et al. (2003), Autor and Dorn (2013) or Gaggl and Wright (2017) for IT, Doms, Dunne and Troske (1997) for factory automation, Boustan, Choi and Clingingsmith (2022) for CNC, Graetz and Michaels (2017) or Acemoglu and Restrepo (2020) for robots, Mann and Püttmann (2021), Bessen, Goos, Salomons and van den Berge (2019) and Aghion, Antonin, Bunin and Jaravel (2022) for broader measures of automation.

build on itself. As a result, i) adoption and innovation matter at different time horizons, ii) knowledge spillovers play a different role for innovation than adoption and iii) the growth literature mentioned above focuses on innovation.

The literature on induced automation innovation is sparser. Acemoglu and Restrepo (2022) find a positive effect of aging on patenting in robotics and numerical control in cross-country regressions, though they focus mainly on adoption. Our paper differs in at least four ways: we build a broader measure of automation innovation in machinery; we are interested in the effect of all wage variations, not only variations arising from demographic trends; we consider policy-induced changes and, most importantly, we use firm-level panel regressions instead of a country-level cross-sectional analysis. In contemporaneous work, Danzer, Feuerbaum and Gaessler (2020) exploit an immigrant settlement policy in Germany to show that increases in labor supply discourage automation innovation at the level of local labor markets, and San (2023) shows that a negative shock to agricultural immigration from Mexico induced relatively more innovation related to labor-intensive crops.[3] Neither of these papers exploit firm-level variation, nor do they estimate the effect of labor cost changes on automation innovations.

In a broader context, an extensive literature shows that the direction of innovation is endogenous (e.g. Acemoglu and Linn, 2004, for pharmaceuticals and Popp, 2002, for energy-saving technologies). We build on Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen (2016), who show that an increase in gas prices leads firms in the auto industry to engage more in clean and less in dirty innovations. We use a similar shift-share design and also measure firms' international exposure with patent weights.[4]

Finally, a recent literature has emerged that classifies patents as automation or not (see, in particular, Mann and Püttmann, 2021, Webb, 2020, Kogan, Papanikolaou, Schmidt and Seegmiller, 2022, and Autor, Chin, Salomons and Seegmiller, 2022). We compare our approaches in detail. A key difference is that we classify readily available technology codes rather than patents directly, so that our classification can be easily

[3] Relatedly, Andersson, Karadja and Prawitz (2022) look at the effect of emigration to the US in the 19^{th}century in Sweden and find that more exposed municipalities experienced an increase in innovation (though they do not identify automation innovations). Bena and Simintzi (2019) show that firms with better access to the Chinese labor market decrease their share of process innovations after the 1999 U.S.-China trade agreement. Process innovations and automation innovations are not the same: some process innovations reduce costs other than labor (say, material cost) and many automation innovations are product innovations (a new industrial robot is a product innovation for its maker).

[4] Other papers have used their methodology, including Noailly and Smeets (2015) on innovation in electricity generation, Coelli, Moxnes and Ulltveit-Moe (2020) on the effect of trade policy on innovation and Aghion, Bénabou, Martin and Roulet (forthcoming) on the role of environmental preferences and competition in innovation in the auto industry.

applied to other patent data by other researchers.

Section 2 develops our classification of automation technologies. Section 3 describes the data and our empirical strategy. Section 4 presents the results of the main analysis on the effect of wages on automation innovations. Section 5 discusses policy shocks. Section 6 concludes. The Appendix provides an analytical model, additional robustness checks, and details on our methodology.

2 Classifying Automation Patents

In this section, we develop our classification of automation patents and use it to build a measure of automation at the industry level. We find that it predicts a decline in routine tasks, in the relative demand for low-skill workers, and in the labor share.

2.1 Our approach to classifying patents

Our goal is to identify automation innovations in machinery: that is, innovations embedded in equipment goods, such as machine tools or robots, which allow for the replacement of workers in some tasks. Non-automation innovations, in contrast, may improve energy efficiency, reduce the cost of producing certain machines or increase reliability.

We follow a well-established tradition in the empirical literature and use patent data as a measure of innovative activity.[5] We use two databases: the EP full-text database from 2018, which contains the full text of patent applications at the European Patent Office (EPO), and the World Patent Statistical Database (PATSTAT) from Autumn 2018, which contains the bibliographic information, but not the text, of close to the universe of patents. In these datasets, the technological characteristics of patents are recorded in technology codes (notably CPC and IPC codes, hereafter C/IPC codes, explained in footnote 9 below). Certain types of technologies, such as fossil fuel engines, can be readily identified with existing groupings of C/IPC codes. No such grouping exists for automation in machinery, and we use text analysis to create one.

We employ a dictionary method on patent data and proceed in four steps: i) We use the existing literature to identify automation-related keywords. ii) For each "technology category" (defined below based on C/IPC codes), we compute the share of patents at the EPO containing one of our automation keywords. iii) We use this measure to classify

[5] For our analysis, patent data present several advantages: they specify the countries where inventions are protected, contain highly disaggregated technology codes and can be matched to firms.

Table 1: Choice of automation keywords

Keywords	Comments	Source
Automat*	*Automation*, *automatization* or *automat** at least 5 times. Or automat* or autonomous with secondary words, *warehouse*, *operator*, arm, convey*, handling, inspect*, knitting, manipulat*, regulat*, sensor, storage, store, vehicle system, weaving, or welding) in the same sentence at least twice.	Own or Doms, Dunne and Troske (DDT) or Acemoglu Restrepo (AR).
Robot*	Not surgical or medical.	DDT and AR
Numerical Control	CNC or numeric* control* or NC in the same sentence as secondary words.	DDT and AR
Computer-aided design and manufacturing	Computer-aided/-assisted/-supported in the same patent as secondary words, also CAD or CAM and not "content addressable memory" in same sentence as secondary words.	DDT
Flexible manufacturing		DDT
Programmable logic control	"Programmable logic control" or (PLC and not (powerline or "power line")).	DDT
3D printer	"*3D print**" or "*additive manufacturing*" or "*additive layer manufacturing*".	Own
Labor	Including *laborious*.	Own
Secondary words	*Machine* or *manufacturing* or *equipment* or *apparatus* or machining.	

Notes: This table describes the keywords that we use to identify automation technologies. Keywords include i) natural adjacent words numerical control includes NC, numerically controlled and numeric control), ii) British/American spelling (i.e. labour/labor) and iii) hyphen adjectives (i.e. computer aided / computer-aided design). "In the same sentence as secondary words" refers to at least one secondary word added words in italics, the others come from AR or DDT. See Appendix A.2 for details.

technology categories as automation or not based on a cut-off. iv) We classify worldwide patents as automation if they belong to an automation technology category.

This strategy of first classifying technology categories and then patents has two advantages over classifying patents directly. First, it allows for the inclusion of patents without text from PATSTAT so that we – and other researchers – can use our classification on patents without text or future patents.[6] Second, the C/IPC codes themselves are informative about the characteristics of an innovation – including whether it relates to automation. Patents are written in different styles and applicants can often describe the same innovation with or without using our keywords. Conversely, if a patent uses one of our keywords but does not belong to any C/IPC code where this is common, the inclusion of that keyword is often uninformative about the nature of the innovation. That is, the wording of a given patent is a weak signal of whether that patent corresponds to automation, but the *combined* wording of many patents gives a strong signal of whether a technology code corresponds to automation.[7]

[6] To give an idea of the increase in the sample size, over the period 1997-2011 there are 3.19 million patent families with patent applications in at least two offices (a condition we will impose in our main analysis). Among these only around $740,000$ have an EPO patent with a description in English.

[7] Our strategy follows the World Intellectual Property Organization (WIPO), which offers a simple tool on its website based on a similar principle: a search engine allows one to identify up to 5 IPC codes most likely to correspond to a set of keywords in the text of the patents.

2.2 Choosing automation keywords

To tie our hands, we choose most of our keywords from the automation technologies identified in Doms, Dunne and Troske (DDT, 1997) and Acemoglu and Restrepo (AR, 2022), and supplement them with additional words as described below.[8] Most of our keywords correspond to the co-occurrence of several words in the same sentence or patent or the repetition of these words a sufficient number of times. Table 1 lists our keywords.

We have eight categories of keywords. Five of them are automation technologies in DDT or AR (robot*, numerical control, computer-aided design and manufacturing, flexible manufacturing, and programmable logic control). Directly using some of these keywords results in false positives. Therefore, we require that our keywords occur in the same patent or in the same sentence as secondary words, such as machinery or equipment, indicating that the text describes a machine. We also add "automation" and "automatization". The stem "automat*" gathers too many false positives such as "automatic transmission". We resolve this in two ways: we only count patents if the frequency of automat* is at least 5 or automat* is combined with a list of words in the same sentence at least twice. This list of words contains our secondary words and additional words which come from DDT or AR and often refer to tasks (such as manipulat*, regulat* or inspect*). Finally, we add 3D printing, which was in its infancy when DDT was written, and "labor", which often indicates that an innovation reduces labor costs. The most important keywords are those associated with "automat*" (Appendix A.2) and our main results are robust to only using these (Appendix A.6.2).

2.3 Identifying automation technology categories

We base our classification on the set of EPO patent applications from 1978 to 2018 with a description in English (1,538,370 patent applications) denoted Ω_{EPO}. Technological characteristics of patents are recorded in (generally, several) C/IPC codes. The C/IPC codes form a hierarchical classification system.[9] We now describe the aggregation level

[8] Doms, Dunne and Troske (1997) measure automation using the Survey of Manufacturing Technology (SMT) from 1988 and 1993 conducted by the US Census. The survey asked firms about their use of specific automation and information technologies. Acemoglu and Restrepo (2022) include imports of automation technology and associate specific HS-categories from Comtrade with automation technology.

[9] The IPC is the International Patent Classification and the CPC the Cooperative Patent Classification used by the USPTO and the EPO. The CPC is an extension of the IPC and contains around 250,000 codes in its most disaggregated form. The structure of the C/IPC classification is as follows: C/IPC "classes" have 3-digit codes (e.g. B25: "hand tools; portable power-driven tools; handles for hand implements; workshop equipment and manipulators"), "subclasses" have 4-digit codes (e.g. B25J:

that we use to classify patents.

Defining machinery C/IPC codes. First, we define "technology categories" as the highly disaggregated 6-digit C/IPC codes (e.g. B25J13). The co-occurrence of technology codes can also be informative about the characteristics of a patent. To capture this, we include pairs of 4-digit C/IPC codes in the definition of technology categories: for instance, a patent containing both codes B25J (manipulators) and B23K (a type of machine tools) belong to the technology category {B25J, B23K}. Finally, we include in the definition of technology categories the co-occurrence of 4-digit C/IPC codes with the 3-digit codes G05 or G06 (e.g. B25J with G05 or G06). The code G05 corresponds to "controlling; regulating" and G06 to "computing; calculating; counting" and Aschhoff et al. (2010) use these combinations to identify advanced manufacturing technologies.[10] The 6-digit codes alone identify 82% of our automation patents (see Appendix A.2.3).

Our keywords are best associated with automation in equipment. Accordingly, we restrict attention to C/IPC codes that belong to a group of technology fields which we call "machinery". Out of 34 technology fields (see Figure A.1), we focus on "machine tools", "handling", "textile and paper machines", and "other special machines" with some adjustments.[11] We classify pairs of 4-digit C/IPC codes or pairings of 4-digit C/IPC codes with G05 or G06 as machinery if at least one 4-digit code belongs to that field. This leaves us with 1010 6-digit codes, 1104 pairs of 4-digit codes, and 25 groupings of 4-digit codes with G05/G06 which form the set of machinery technology categories.

Defining automation C/IPC codes. We define a machinery patent as a patent that belongs to one of the machinery technology categories. We then denote MT_p the set of machinery technology categories associated with a patent p and T_p its text.[12] We define the function $k^{any}(T_p)$ which takes value 1 if any of the automation keywords is present in the text and 0 otherwise. For each machinery technology category t, we define the prevalence of automation keywords $s(t)$ as the share of patents containing at least

"manipulators; chambers provided with manipulation devices"), and main groups have 5 to 7 digit codes (e.g. B25J13: "controls for manipulators"). In the following, we refer to classes, subclasses, and main groups as 3-digit, 4-digit, and 6-digit codes respectively.

[10] To ensure that the set of patents available in Ω_{EPO} is sufficiently representative of a technology category, we restrict attention to categories containing at least 100 patents. We group 6-digit codes with the same 4-digit code and less than 100 patents into common artificial 6-digit codes.

[11] We make a few minor adjustments such as excluding weapons and ammunition and adding technology codes referring to "programme-control system". See Appendix A.2 for details.

[12] We use all C/IPC codes of the patent family associated with the EPO patent application p. See Section 2.4 for the definition of the patent family.

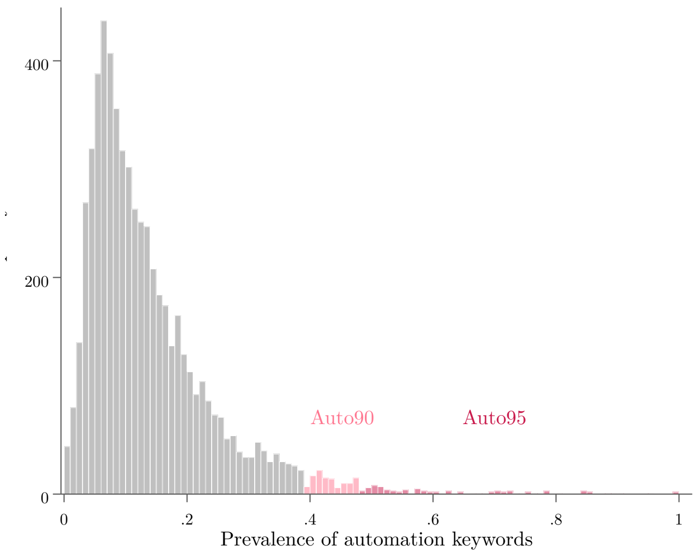

Figure 1: Prevalence of automation keywords for C/IPC 6-digit codes in machinery.

one of our keywords:

$$s\left(t\right)=\frac{\sum_{p\in\Omega_{EPO}}1_{t\in MT_p}k^{any}(T_p)}{\sum_{p\in\Omega_{EPO}}1_{t\in MT_p}}.$$

We similarly define the prevalence of specific keyword categories. We show that these measures are positively correlated for the main keywords, give examples of the prevalence measures in some C/IPC codes, and present additional statistics in Appendix A.2. We manually checked the C/IPC codes extensively and sampled patents from each category to ensure that the process delivered reasonable results.

We define automation technology categories as those with a prevalence measure above a threshold. Figure 1 shows the histogram of the prevalence of automation keywords for all C/IPC 6-digit codes in machinery. The distribution is skewed: most C/IPC codes have a low prevalence of automation keywords but a few codes have a very high value. We choose thresholds at the 90^{th} and 95^{th} percentiles of the distribution of the 6-digit code distribution (within machinery), which are given by 0.388 and 0.479, respectively, as our baselines.[13] Therefore, a technology category t belongs to the set of auto90 categories T^{90} if $s(t) > 0.396$ and to the set of auto95 category T^{95} if $s(t) > 0.480$. In Appendix A.2.4, we show that the technology categories with a high prevalence of automation keywords remain the same throughout the period considered. In particular, the correlation between the prevalence measures computed for the first half of the sample

[13]Choosing different thresholds is easy and we investigate how robust our results are in Section 4.5.

and the second half is 0.85.

2.4 Defining automation patents

We now proceed to classify automation patents. To do so, we use PATSTAT, which contains bibliographic information for almost the universe of patents. PATSTAT also allows us to identify patent families, a set of patent applications in different national or international patent offices representing the same innovation. For each patent family, we know the date of the first application (used as the year of an innovation), the corresponding patent offices, the identity of the applicants and the inventors, the number of citations received, and, importantly the C/IPC codes associated with the innovation.

We then define a patent family p in the PATSTAT dataset $\Omega_{PATSTAT}$ as an automation innovation if it belongs to at least one automation technology category. From now on, we refer to a patent family as a patent. That is, p is an auto95 patent if $\exists t_p \in MT_p$ such that $t_p \in T^{95}$, and similarly for an auto90 patent. Appendix A.2 provides additional statistics and gives examples of automation patents.

2.5 Comparison with other measures in the literature

A recent literature has emerged that uses patent data to identify automation technologies. Our approach – classifying technology categories rather than patents directly and using a dictionary method on the text of patents – is unique and we compare it here with two alternatives: Mann and Püttmann (2021)'s and Kogan et al. (2022)'s.[14]

Mann and Püttmann (2021) manually classify a set of patents as automation or not and then use machine-learning techniques to classify all patents. Relying on keywords instead of a training set presents several advantages. First, manually labeling patents is a difficult task that cannot be easily systematized and outsourced. Second, patents are technical descriptions of an innovation and do not primarily discuss its goal. Only a few words within the text are informative, so that the training set needs to be large. Third, our approach is more transparent, more easily replicable and modifiable, and leaves fewer degrees of freedom as we choose most of our keywords from the literature.

We compare our classifications in detail in Appendix A.2.6. The measures are positively correlated, though ours is generally more conservative. We classify 9.4% of the

[14]Bessen and Hunt (2007) directly use keywords to identify software patents. Webb (2020) similarly identifies patents in robotics, software and AI using keywords before matching them to occupations using machine-learning techniques.

common set of machinery patents as auto95 while they classify 29.8% of them as automation. 70% of our auto95 patents are classified as automation patents by Mann and Püttmann (to put this number in context, their algorithm has a 17% false negative error rate on their training set). Focusing on outlier technologies, we find that Mann and Püttmann classify a number of patents related to elevators and printing machines as automation, which we do not. This is in line with their definition of automation as "a device that carries out a process independently of human intervention", while we seek to identify innovations that replace workers in existing tasks.

In contemporaneous work, Kogan et al. (2022) and Autor et al. (2022) measure the distance between patent texts and the description of tasks in the Dictionary of Occupation Titles database to identify labor-displacing innovations. This approach has the advantages of easily matching innovations and the affected occupations and of being completely hands-off. However, as Kogan et al. (2022) point out, it captures not only automation innovations where workers are replaced by machines but also innovations where incumbent workers with newly obsolete skills are replaced by new workers.

Compared to these alternatives, our approach has the advantage of providing a measure that can be readily used by anyone on any patent data with C/IPC codes. Our approach also has some drawbacks: We rely on the existence of a set of well-identified automation technologies (robots, CNC, etc.), which is why we apply our method only to machinery patents. Heterogeneity within technology categories will also lead us to misclassify individual patents that are "exceptions" within their categories (i.e. non-automation patents in a category with many automation patents and vice versa).

2.6 Trends in automation innovations

Figure 2 plots the evolution of automation patent families. To restrict attention on innovations of sufficient quality and make the data more comparable across countries, we focus on patent families including patent applications in at least two countries, referred to as biadic patents. Several studies (e.g. De Rassenfosse et al., 2013, and Dechezleprêtre, Ménière and Mohnen, 2017) have shown that biadic patents are of higher quality than others.[15] Panel (a) shows that, globally, the share of automation patents in machinery

[15] We count applications and not-granted patents because certain patent offices, notably the Japanese, only formally grants a patent if the applicant requests an examination which they often only do when their rights are challenged. Further, biadic patents allow for better comparison across countries since several small patents typically cover the same large innovation in certain offices like the JPO but only one broad patent in others like the USPTO. To restrict attention to patent families of even higher

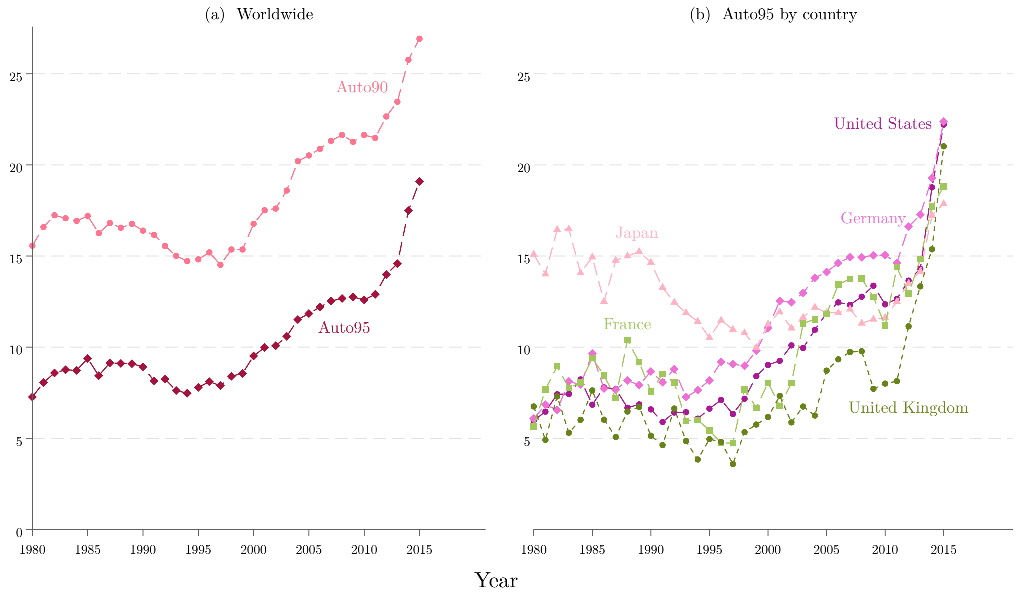

Figure 2: Share of automation patents in machinery for biadic families
Notes: Panel (a) reports the share of auto95 and auto90 patents worldwide; Panel (b) reports the share of auto95 patents in machinery by applicants' nationality.

declined slightly between the mid1980s (9.4% in 1985 for auto95) and the mid1990s (7.5% in 1994 for auto95) before increasing rapidly (reaching 19% in 2015 for auto95). Appendix Figure A.2 reports the raw numbers of auto90 and auto95 patents and their share out of total patents. Figure 2.b shows the trends for auto95 by applicant nationality. In the 1980s Japan had the highest share of automation patents in machinery, while Germany took that position from the 2000s.

2.7 Automation, routine tasks, skill composition and labor share

We now build a measure of automation at the sector level and relate it to changes in task and skill composition and factor shares. We do this in part to validate our classification of automation patents. We build on Autor et al. (2003) (hereafter ALM), who show that computerization was associated with a decrease in routine tasks at the sector level using U.S. data from 1960 to 1998. We report our main results here and details on the data construction and additional results in Appendix A.3.

quality, we carry out robustness checks where we use patent citations.

As ALM, we run sector level regressions of the type:

$$\Delta T_{jk} = \beta_0 + \beta_C \Delta C_j + \beta_{aut} aut_j + \varepsilon_j. \quad (1)$$

We focus here on the years 1980-1998. ΔT_{jk} represents the change in tasks of type k in sector j. We take this measure directly from ALM, who define 5 types of tasks: nonroutine analytic, nonroutine interactive, routine cognitive, routine manual, and nonroutine manual. ΔT_{jk} is measured as 10 times the annual within-sector change in task input, measured in percentiles of the 1960 task distribution. ΔC_j is ALM's measure of the change in computerization in sector j (available for the period 1984-1997). aut_j is our patent-based measure of the *use* of automation technologies in sector j. Since patenting is already a measure of knowledge flows, we do not first-difference this measure. We run similar regressions with changes in the skill ratio, the labor share, and employment as alternative outcome variables.

To build aut_j, we allocate patents in machinery to their sector of use, focusing on USPTO granted patents. Autor, Dorn, Hanson, Pisano and Shu (2020) match USPTO patents with firm-level data from Compustat, and provide detailed sectoral information on corporate patents. Using their data, we create a weighted concordance table from C/IPC 4-digit codes to 4-digit SIC industries that allows us to map patents to sectors of invention. We then combine this information with the 1997 capital flow table from the BEA to get the sector of use. The capital flow table is akin to an input-output table but it reports the flows in investment goods instead of intermediate inputs. Sectors that purchase a lot of capital from sectors innovating in machinery generally have high exposure to both automation and non-automation patents (defined here as pauto90 patents, i.e. machinery patents excluding auto90 patents). The correlation between the log counts of auto95 patents and pauto90 patents across sectors is 0.99 but the correlation between the ratio of auto95 patents over capital purchases and ratio of pauto90 patents over capital purchases drops to 0.76 (see Appendix Figure A.9).

For each sector of use j, we compute aut_j as the share of automation patents (auto95 in our baseline) among machinery patents from 1980-1998. Our measure captures whether the machinery patents used in a given sector are particularly intensive in automation technologies or not. We compute this statistic for the 133 sectors with machinery patents (our results are robust to excluding sectors with few machinery patents). Appendix Table A.1 reports summary statistics on this measure and our dependent variables: on average, the auto95 share is 7.5%. There is significant variation in the share

Table 2: Sectors with the highest and lowest shares of automation patents

Sectors with highest share of automated patents			Sectors with lowest share of automated patents		
Industry		Auto95	Industry		Auto95
756	Automotive services and repair shops	0.111	801	Bowling alleys, billiard and pool parlors	0.042
206	Household appliances (e.g., radio, TV, equipment)	0.106	100	Meat products	0.047
470	Water supply and irrigation	0.098	101	Dairy products	0.047
271	Iron and steal foundaries	0.097	102	Canned and preserved fruits and vegetables	0.047
130	Tobacco manufactures	0.093	110	Grain mill products	0.047
212	Misc. plastic products	0.093	111	Bakery products	0.047

Notes: This table lists the sectors with the highest and lowest share of auto95 patents out of all patents in machinery in 1980-1998. The industry codes and descriptions of the sectors correspond to the ind6090 industries described in ALM.

of automation patents across industries with a coefficient of variation of 17%. Exposure to automation is on average higher in manufacturing and we include a manufacturing dummy in our regressions. Interestingly, our automation measure is only weakly correlated with computerization, with a coefficient of 0.08 (and −0.16 when we weigh industries by employment). Table 2 lists the sectors with the highest and lowest shares of automation patents in machinery.[16]

Table 3: Effects of automation on tasks, skill composition, and the labor share

	Δ Routine cognitive		Δ Routine manual		Δ High/low skill workers		Δ Labor share (NBER)		Δ Labor share (BEA)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Share automation	-142.28^{***}	-142.96^{***}	-127.77^{***}	-127.11^{***}	3.81^{*}	3.51^{*}	-1.31^{**}	-0.86^{**}	-3.65^{*}	-3.59^{*}
(using industry)	(37.51)	(28.39)	(38.83)	(39.80)	(2.17)	(2.02)	(0.65)	(0.41)	(1.94)	(1.86)
Share automation		-22.13^{**}		-11.05		-0.21		-0.46		-0.21
(inventing industry)		(8.54)		(8.03)		(0.45)		(0.29)		(0.33)
Δ Computer use	-18.86^{***}	-20.49^{***}	-19.12^{***}	-21.79^{***}	0.98^{***}	0.92^{***}	0.25^{*}	0.27^{*}	0.12	0.15
(1984-1997)	(6.54)	(5.79)	(7.26)	(7.58)	(0.27)	(0.27)	(0.13)	(0.14)	(0.31)	(0.32)
Manufacturing	-1.70^{*}	-2.15^{**}	-0.02	-0.23	0.02	0.02			-0.05^{*}	-0.06^{*}
	(0.92)	(0.86)	(0.94)	(0.93)	(0.03)	(0.03)			(0.03)	(0.03)
R^2	0.29	0.38	0.21	0.23	0.21	0.19	0.18	0.27	0.22	0.23
Industries	133	126	133	126	133	126	56	56	60	60

Notes: This table shows the effect of automation technologies on tasks, skill composition, and the labor share. Each column is a separate OLS regression of the change in an industry outcome between 1980 and 1998 on the share of automation patents in machinery in 1980-1998, the annual percentage point change in industry computer use during 1984-1997, a dummy variable indicating the manufacturing sector, and a constant. In columns 1–2 the dependent variable is the change in routine cognitive tasks and in columns 3–4 the change in routine manual tasks, both measured as 10x the annual change in industry-level task input in centiles of the 1960 task distribution (see ALM). In columns 5–6 the dependent variable is the change in the ratio of high-skill workers (college graduates) over low-skill workers (others). In columns 7–8 and columns 9–10 the dependent variable is the change in the labor share in the NBER-CES manufacturing industry database, and in 60 aggregated industries from the BEA, respectively. As described in the text, the two automation share measures correspond to a different mapping between C/IPC codes and industries. Using industries allocates patents to their sector of use while innovating industry – added in columns 2, 4, 6, 8, and 10 – allocates patents to their sector of invention. The regressions are weighted by the mean industry share of total employment in FTEs in 1980 and 1998. Standard errors are clustered at the level of industry groups that have the same automation share by construction and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table 3 reports the results of regression (1) and Appendix Figure A.10 provides scatter plots. Columns (1) and (3) show a clear relationship: A 1 pp increase in the

[16] Some small sectors have the same share of automation patents because they map to the same C/IPC 4-digit codes. We cluster standard errors at the level of these 112 clusters of using industries.

automation share is associated with a 1.4 and 1.3 centiles decrease in routine cognitive and manual tasks per decade, respectively. This effect is larger than that of computerization (i.e. standardized beta coefficients are larger). At first sight, it may seem surprising that our measure of automation in machinery predicts a decline in routine cognitive tasks. However, ALM define routine cognitive tasks as the "adaptability to situations requiring the precise attainment of set limits, tolerances or standards". These correspond to inspection and control tasks that our automation machines may eliminate (Figure A.8 gives an example of such a machine). Metalworkers, for instance, are among the occupations with the highest intensity of routine cognitive tasks.

In Column (5), we instead use the change in the skill ratio (college graduates over all other workers) as the dependent variable. A 1 pp increase in the share of automation patents is associated with an increase in the skill ratio equal to one-third of its average increase over the 1980-1998 period. In Column (7), we use the change in the labor share for manufacturing industries (using the NBER Manufacturing Database). Industries with a high share of automation patents experience a decline in the labor share, consistent with the theoretical literature on automation. A 1 pp increase in the automation share is associated with 1.3 pp decline in the labor share. Interestingly, computerization does not have the same effect. We extend this analysis to non-manufacturing industries using more aggregate data on the labor share for 60 industries from the BEA. We still find a negative effect of automation (Column 9).

In even columns, we add a control for the share of automation patents invented by the sector. To allocate patents to the inventing sector, we omit the capital flow table step in the calculation of our automation variable. The coefficients on the automation share in the using sector remain stable. Moreover, the automation share in the using sector has a larger effect than the automation share in the inventing sector.[17]

Appendix A.3 contains additional robustness checks: We use biadic patents, auto90 patents, or an alternative concordance table between C/IPC codes and sectors from Lybbert and Zolas (2014). In all cases, we find a negative effect of the automation share on routine tasks, the skill ratio, and the labor share. We also look at the effect of automation on employment changes. We find a negative effect within manufacturing but no statistically significant effect for the economy as a whole.

To summarize, we have now classified machinery patents as automation or non-

[17] The standardized coefficients are larger for the using sector than the inventing sector (except in column (8)) as the s.d. for the share of automation patents in the using and inventing sectors are respectively 1.3% and 6%.

automation. Mapping C/IPC codes to sectors, we build a measure of automation at the sectoral level, which is more detailed than alternatives such as robotization. Doing that, we find that sectors that use automation technologies more intensely experience a decrease in routine tasks, an increase in the skill ratio, and a decline in the labor share.

3 Empirical Strategy and Data

We now move to our main empirical exercise and analyze the effect of labor cost shocks on automation innovations. Section 3.1 presents our empirical strategy, Sections 3.2 and 3.3 explain how we build our dataset, Section 3.4 describes our estimation equation and Section 3.5 shows summary statistics.

3.1 Empirical strategy

We motivate our empirical strategy with the business structure of the most prominent automation innovators. These are large firms that sell their automation equipment internationally to downstream firms. Automation equipment enables the replacement of low-skill workers with machines. It can also complement high-skill workers who program, operate, and maintain the machines. Therefore, the incentives for downstream producers to adopt automation technology are determined by labor costs in their local market. Higher low-skill labor costs for potential customers are associated with a larger market for automation machine producers, which in turn, would induce innovators to conduct more research on automation technologies.[18,19] Appendix A.4 presents a simple model that rationalizes this argument.

[18]Our conceptual argument is reflected in the business practices of large innovators. Siemens, the biggest innovator in our sample, is a very international company with 14% of its revenue in Germany in 2018. Its strongest growing division was the Digital Factory Division which provides a broad range of automation technology to manufacturers across the globe. The annual report (Siemens, 2018) uses a number of our keywords and describes how "The Digital Factory Division offers a comprehensive product portfolio and system solutions for automation technologies used in manufacturing industries, such as automation systems and software for factory automation, industrial controls and numerical control systems, motors, drives and inverters and integrated automation systems for machine tools and production machines...". The report is centrally interested in how "Changes in customer demand [for automation technology by downstream manufacturers] are strongly driven by macroeconomic cycles". It does not mention labor costs directly but uses euphemisms such as "increase competitiveness", "enhance efficiency", "improve cost position" and "streamline production". Siemens further discusses how such macroeconomic trends affect its *R&D* decisions.

[19]If automation innovations are internal to the firm, then the argument follows if one interprets the innovator's customers as the downstream production sites of the same firm.

Empirically, we aim to estimate by how much an increase in low-skill labor costs leads to an increase in automation innovations, and an increase in high-skill labor costs to a decrease in automation innovations. We focus on labor costs because they are the key factors that should affect automation innovations differently from non-automation innovations. In contrast, the total wage bill of low-skill workers would include their employment and a higher employment of low-skill workers may be associated with a more intensive use of non-automation machinery as well.

Ideally, we would measure the labor costs paid by the actual and potential customers of automation innovators. Such a measure would suffer from reverse causality, and we would need an instrument. A natural candidate would be a shift-share instrument. In the absence of direct data on the labor costs paid by innovators' customers, we directly use such a shift-share measure as a proxy. Our regression should therefore be seen as the reduced form of this instrumental approach.

More specifically, we measure the labor costs paid by the customer of an innovator as a weighted average of country-level labor costs where the weights reflect the market exposure of innovators. That is, we define the average low-skill $w_{L,i,t}$ and high-skill $w_{H,i,t}$ labor costs faced by firm i's customers as

$$w_{J,i,t} \equiv \sum_c \kappa_{i,c} w_{J,c,t} \text{ for } J \in \{L, H\}, \tag{2}$$

where $w_{L,c,t}$ ($w_{H,c,t}$) are the low-skill (high-skill) labor costs in country c at time t and $\kappa_{i,c}$ is the fixed weight of country c for firm i.[20] Similarly, we build controls for several macroeconomic variables such as labor productivity, GDP per capita, or the size of the manufacturing sector, which could also affect innovation.

With this shift-share measure, our identification strategy relies on how country-level shocks affect firms differently. We discuss identification extensively in Section 4.3. We now describe how we obtain country-level data (such as $w_{L,c,t}$) and firm data (including the weights $\kappa_{i,c}$).

[20] More precisely, innovation incentives depend on the expectation of future labor costs for automatable tasks. As we cannot measure expectations, we use current labor cost shocks as a proxy for shocks on expected future costs (in Section 4.5 we explicitly build a model for expected wages using an AR process). In addition, there are no good international occupational or task-level labor costs data. Since low-skill workers' tasks have been more intensely automated, we use low-skill labor cost as a proxy for the cost of automatable tasks. This proxy will be particularly good if labor markets are flexible across occupations within education groups or if labor shocks affect low-skill workers similarly across occupations. Otherwise, a noisy measure should result in a bias towards zero.

3.2 Country-level data

We use data on 41 countries. Most of our data come from the 2013 release of the World Input Output Tables (WIOD, Timmer et al. 2015), which contains information on hourly labor costs from 1995 to 2009 across educational attainment groups.[21] We focus on labor costs in manufacturing since our keywords largely come from the SMT (Survey of Manufacturing Technologies). Our results are robust to using labor costs in the entire economy. Although our data cover all labor costs, we refer to them as wages for simplicity. In the data, the low-skill workers are defined as having no high-school diploma or equivalent, whereas the high-skill workers have at least a college degree. Middle- and low-skill wages are highly correlated, and we can interpret our low-skill wage variable as reflecting both.

We calculate labor productivity in manufacturing as value added divided by hours. We gather exchange rate and GDP data from UNSTAT and compute the GDP gap to control for business cycles. We deflate all nominal values with the local PPI for manufacturing (indexed to 1995) and then convert everything into dollars using the average exchange rate for 1995, the starting year of our regressions. Appendix A.5.1 provides further details and summary statistics.

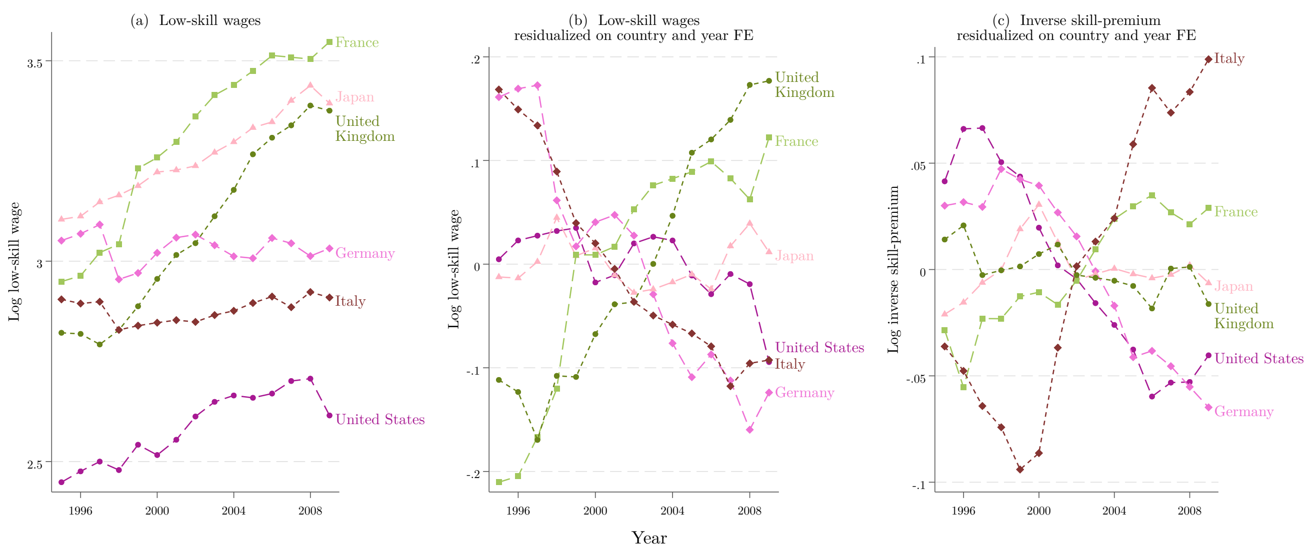

Figure 3: Country-level trends in low-skill wages and the inverse skill premium.

Notes: Panel (a) reports raw log low-skill wages; Panel (b) log low-skill wages residualized on country and year fixed effects; Panel (c) log inverse-skill premium residualized on country and year fixed effects.

[21]The data cover 40 countries, including all major markets (US, Japan, all EU countries of 2009, China, India, Brazil, Russia, etc.). We obtain similar data for Switzerland from the Swiss Federal Statistical Office.

Figure 3 uses the 6 countries with the largest average weights, and shows the log low-skill wages, the log low-skill wages residualized on country and year fixed effects, and the residualized log inverse skill premium. Trends vary markedly across countries providing a significant amount of variation in the data: for instance, the US sees a relative decrease in the inverse skill premium while France follows a non-monotonic pattern.[22]

3.3 Firm-level data

We now describe our firm-level data. To identify firms, we use Orbis Intellectual Property, which matches global patent data to the companies in Orbis (details in Appendix A.5.2). We then use PATSTAT to obtain all bibliographic information on firms' patents, including their C/IPC codes, which allows us to identify machinery and automation patents. We use this to build our dependent variable: the count of automation patents filed by a firm in a given year.

We use the firm's patent filing history as a proxy for its market exposure to measure the weights $\kappa_{i,c}$. This method builds on that of Aghion et al. (2016, henceforth ADHMV). Firms differ in their market exposure due to trade barriers, heterogeneous customer tastes, or various historical accidents. A patent grants its holder the exclusive right to commercially exploit an innovation in a specific country for a limited period. Inventors must file a patent in each country where they wish to protect their technology. Patenting is costly: a firm must hire lawyers, possibly translators, and pay filing costs. Therefore, inventors only seek patent protection in a country if they are confident about the potential market value of the technology (Eaton and Kortum, 1996). Indeed, empirical evidence suggests that inventors do not patent widely and indiscriminately, with the average invention being patented in only two countries (Dechezleprêtre et al., 2011).[23]

For each firm, we compute the fraction of its patents in machinery protected in each country c for which we have wage data, $\tilde{\kappa}_{i,c}$. We keep the weights fixed and compute them during the pre-sample period 1971-1994 to ensure that they are weakly exogenous.[24] We

[22]Appendix Figure A.12 shows more precisely the identifying variation taking into account the shift-share structure of our measure following Borusyak, Hull and Jaravel (2022)'s methodology.

[23]ADHMV verify that a method similar to ours accounts well for the sales distribution of major auto manufacturers. Coelli, Moxnes and Ulltveit-Moe (2020) carry out a more systematic exercise and verify that such a method accounts well for aggregate bilateral trade flows and firm exports across 8 country groups in a representative panel of 15,000 firms from 7 European countries (regressing patent weights on sales weights gives a coefficient of 0.89 with a s.e. of 0.008). In Appendix B.2, we also show that our patent weights correlate well with trade flows.

[24]This approach aligns with our goal of identifying the exogenous effect of an increase in wages on innovation. In reality, the exposure to different markets changes over time, in part in response to

restrict our attention to patent families with at least one citation (without self-citations) to exclude the lowest quality patents. See Appendix A.5.3 for details notably on how we treat European patents.

The raw patent count indicates whether a firm intends to sell its products in a market but does not capture market size. A larger market attracts more firms, so the market size per firm does not grow 1 for 1 with country size. To account for this, we weigh each country c by $GDP_{0,c}^{0.35}$, where $GDP_{0,c}$ is the 5 year average GDP of country c at the end of the pre-sample period.[25] As a result, the weight of country c for firm i is:

$$\kappa_{i,c} = \frac{\tilde{\kappa}_{i,c} GDP_{0,c}^{0.35}}{\sum\limits_{c'} \tilde{\kappa}_{i,c'} GDP_{0,c'}^{0.35}}. \tag{3}$$

Following equation (2), we then combine the weights $\kappa_{i,c}$ with the macro variables presented in section 3.2 to build macro variables, including wages, at the level of firms' customers. We use 1971-1994 as a pre-sample period as PATSTAT's coverage is significantly better from the 1970s onward, and we prefer a long time period for our baseline measure. Importantly, the weights are stable over time. We show that our results are robust to alternative pre-sample periods and weighing schemes in Section 4.5.[26]

3.4 Estimation equation

We now describe how we estimate the effect of an increase in wages on automation innovations. We have a panel of firms with patent data and firm-level wage variables. Since our dependent variable is a count of patents, we use a Poisson specification. We assume that firms' innovation in automation follows:[27]

$$\begin{aligned} & \mathrm{E}\left(PAT_{Aut,i,t}\right) \\ = \ & \exp\left(\begin{array}{c} \beta_{w_L} \ln w_{L,i,t-2} + \beta_{w_H} \ln w_{H,i,t-2} + \beta_X X_{i,t-2} + \beta_{Ka} \ln K_{Aut,i,t-2} + \beta_{Ko} \ln K_{other,i,t-2} \\ +\beta_{Sa} \ln SPILL_{Aut,i,t-2} + \beta_{So} \ln SPILL_{other,i,t-2} + \delta_i + \delta_{j,t} \left(+\delta_{c,t}\right) \end{array}\right). \end{aligned} \tag{4}$$

changes in wages. Studying this response would be interesting but is beyond the scope of this paper.

[25] We use Eaton, Kortum and Kramarz (2011)'s study on French exports to compute the elasticity of the average export by firm with respect to the GDP of the destination country and find 0.35.

[26] We consider two alternative measures of low-skill wages where weights are based on 1971-1989 or 1985-1994. For the firms in our baseline regression sample, the correlation between these two wage variables is 0.86.

[27] For estimation, we use the ppmlhdfe command from Correia, Guimaraes and Zylkin (2020), which allows us to run Poisson regression models with high-dimensional fixed effects.

$PAT_{Aut,i,t}$ denotes the number of biadic automation patent families by firm i in industry j and country c with first application filed in year t. Automation patent families are the auto95 patents defined in Section 2. As mentioned in Section 2.6, we focus on biadic patent families, in line with our empirical strategy which relies on firms' exposure to international markets.

$w_{L,i,t}$ and $w_{H,i,t}$ are the average low-skill and high-skill manufacturing labor costs faced by the customers of firm i at time t defined in (2). $X_{i,t}$ represents a vector of macroeconomic controls: labor productivity in manufacturing or GDP per capita to capture technology, human capital shocks, or demand shocks in the customers' countries, and the GDP gap for business cycles fluctuations.

We include controls for knowledge stocks at the firm and country level. $K_{Aut,i,t}$ and $K_{other,i,t}$ denote the stocks of knowledge in automation and in all other technologies of firm i at time t. We compute these knowledge stocks using the perpetual inventory method. $SPILL_{Aut,i,t}$ and $SPILL_{other,i,t}$ similarly denote the stocks of external knowledge (spillovers) in automation and in other technologies to which firm i has access at time t. We compute these spillovers as a weighted average of country-level knowledge stocks, where the weights now reflect the location of the firms' inventors.[28] These controls serve two purposes. First, they ensure that we do not simply capture that some firms or countries are on different automation trends. Second, knowledge spillovers are an important characteristic of innovation processes and can amplify the short-run response of innovation to economic shocks over time. On the empirical side, Popp (2002) finds that including a measure of knowledge stocks is important to obtain correct estimates of energy prices on energy-saving innovations. On the theoretical side, including measures of the stock of knowledge is important for the calibration of macro models of directed technical change (Acemoglu, 2023).

δ_i are firm fixed effects such that our variation comes from how changes in wages affect changes in automation innovations.[29] $\delta_{j,t}$ are industry-year fixed effects. A firm's

[28]We use a depreciation rate of 15% when computing stocks at the firm or country level. The weights in the spillover variables correspond to the location of firms' innovators (obtained from PATSTAT) pre-sample in 1971-1994. When computing the log of stocks or spillovers, we replace 0's with 1's and add a dummy variable to indicate where stocks or spillovers are zero.

[29]We use the Hausman, Hall and Griliches (1984, HHG) method in our baseline specification to control for firm-level fixed effects. This is the count data equivalent to the within-group estimator. Technically, this method is inconsistent with equation (4) as it requires strict exogeneity and hence prevents the lagged dependent variable from appearing on the right-hand side (which it does here to a limited extent through the knowledge stock $K_{Aut,i,t-2}$). Nevertheless, we show in Section 4.5, that our coefficients of interest are not affected by this Nickell bias.

industry j is the manufacturing industry and corresponds to its 2-digit industry in Orbis. Appendix Table A.2 gives the distribution of firms and patents across the main industries in our sample. In some specifications, we include country-year fixed effects $\delta_{c,t}$, where the firm's country ("the home country") is defined as the country with the largest weight $\kappa_{i,c}$. We cluster standard errors at the firm level.[30]

We lag the right-hand side variables by 2 years in the baseline regressions for two reasons. First, the empirical literature for other types of innovation suggests a 2-year lag between R&D investment and the first results materialized by a patent application. Second, at the time of their R&D investment, innovators would use contemporaneous wages as a predictor of future wages. Section 4.5 explores alternative timing assumptions.

3.5 Baseline sample

We now describe the sample we rely on to estimate equation (4). Since wages are available for 1995-2009, our baseline datasets rely on firms that applied for at least one biadic automation patent between 1997 and 2011. These firms must also have at least one patent prior to 1995, so that we can compute the patent weights $\kappa_{i,c}$. We also exclude purely domestic firms (i.e., those that patented in only one country pre-sample), though our results are very similar if we include them. Our baseline sample for the auto95 measure corresponds to $3,255$ firms.

Appendix Table A.3 shows that our sample of firms covers a considerable share of worldwide automation innovations. Orbis' coverage is excellent: we can assign 84.7% of all biadic auto95 patent families in 1997-2011 to a firm. Moreover, most heavy patenters had already patented in at least two countries pre-sample: the firms of our sample account for 53.4% of all biadic auto95 patent families.

Appendix Table A.4 gives descriptive statistics on the number of automation patents per year and lists the sample's ten biggest automation innovators. The distribution of auto95 patents is strongly skewed: over the period 1997-2011, the median firm in the sample filed two auto95 patent applications, whereas the 99^{th} percentile filed 194. In our empirical analysis, we also look at the effect of wages on non-automation machinery patents (defined as pauto90 patents, i.e. machinery patents which are not auto90 patents). For that exercise, we restrict the sample further to firms that have at least one pauto90 patent. The average citations per patent is slightly higher for auto95 than

[30] Alternatively, we cluster at the country level in Appendix A.6 including following the approach of Cameron, Gelbach and Miller (2008) and discuss inference in the shift-share setting in Section 4.3.

Table 4: Descriptive statistics of the firms in our baseline regression

Average weights				Weight variation		
	(1)		(2)		(3)	(4)
					Total	*Foreign*
Largest country	0.47	United States	0.21			
Second largest	0.17	Germany	0.20	HHI country	13.3%	9.0%
Third largest	0.10	Japan	0.17	HHI country-year	0.9%	0.6%
Fourth largest	0.07	France	0.09	Mean pairwise corr.	0.08	0.13

Notes: This table presents summary statistics for the country weights of firms. Columns 1–3 report statistics for the total weights. Column 4 presents information on foreign weights (normalized to 1). Columns 1 and 2 report the average weights of the largest countries. Columns 3 and 4 report the Herfindahl-Hirschman Index (HHI) at the country and country-year level. The mean pariwise correlation is the average pairwise correlation between any two firms (column 3) or firms within a home country (column 4).

pauto90 patents (9.2 vs 7.6).

Table 4 provides descriptive statistics on the country weights for the firms in our sample. The largest country for a given firm has, on average, a weight of 0.47 (for the auto95 sample), and the second largest has a weight of 0.17. For regressions with country-year fixed effects, the latter is more relevant. The three countries with the largest weights on average are the United States, Germany, and Japan. Finally, we compute Herfindahl indices (HHI) of the weights. The HHI of country-level weights is 0.13. For regressions with country-year fixed effects, the relevant statistic is the HHI of the foreign weights (i.e. excluding the home country and renormalizing to 1), which is 0.09 at the country-level. With 15 years the HHI of country-year weights are 15 times smaller. The average pairwise correlation of weights is low: 0.08 (and 0.13 for foreign weights of firms from the same country).

Appendix Table A.5 reports standard deviations and a correlation matrix for the firm-level macroeconomic variables, residualized on firm and industry-year fixed effects. We still find significant variation in the residualized (log) low-skill wages, as the standard deviation is 0.03 (by comparison, the standard deviation is 0.1 when residualizing only on firm fixed effects). Appendix A.6.1 provides additional statistics computed at the level of the shock of our shift-share variable (see Appendix Table A.28).

4 Global Wages and Induced Automation

We present our results in three steps: First, we find a positive effect of low-skill wages on automation innovations. Second, we show that this effect does not exist for non-automation innovations in machinery. Third, we build on the recent shift-share literature (notably Borusyak, Hull and Jaravel, 2022, henceforth BHJ) and argue that the effect

Table 5: Baseline regressions: effect of wages on automation innovations (auto95)

Dependent variable	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.96***	2.71***	3.65***	2.26**	2.61**	3.70***	4.20***	5.32***	4.53**
	(0.80)	(0.85)	(0.96)	(1.00)	(1.14)	(1.28)	(1.33)	(1.56)	(1.79)
High-skill wage	−2.22***	−2.62***	−1.53*	−2.79***	−2.04*	−1.83*	−4.46***	−2.87*	−4.26***
	(0.73)	(0.79)	(0.82)	(0.96)	(1.07)	(1.06)	(1.32)	(1.47)	(1.41)
GDP gap	−3.74	−4.29	−2.17	4.47	5.42	6.93	−0.12	2.28	0.53
	(2.61)	(2.70)	(2.80)	(6.88)	(6.92)	(7.22)	(4.60)	(4.92)	(5.24)
Labor productivity		0.96			−1.73			−2.57	
		(0.92)			(1.77)			(1.60)	
GDP per capita			−1.91			−3.55*			−0.59
			(1.32)			(1.95)			(2.10)
Stock automation	−0.12***	−0.12***	−0.12***	−0.12***	−0.12***	−0.13***	−0.13***	−0.13***	−0.13***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Stock other	0.52***	0.52***	0.52***	0.52***	0.52***	0.52***	0.51***	0.51***	0.51***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Spillovers automation	0.60**	0.63**	0.76**	1.36***	1.35***	1.35***	1.33***	1.29***	1.33***
	(0.30)	(0.30)	(0.31)	(0.47)	(0.47)	(0.47)	(0.46)	(0.46)	(0.46)
Spillovers other	−0.20	−0.25	−0.34	−0.97***	−0.94***	−0.99***	−0.98***	−0.97***	−0.98***
	(0.22)	(0.22)	(0.24)	(0.36)	(0.36)	(0.36)	(0.35)	(0.35)	(0.35)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table presents the results of our baseline regression. The independent variables are lagged by two periods. Coefficients are estimated using conditional Poisson fixed effects regressions (HHG). All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

of low-skill wages on automation innovations is causal. We then discuss the magnitude of our effects and provide robustness checks.

4.1 Main results

Table 5 presents our baseline results. We run a panel analysis at the firm-level for the period 1997-2011 for the left-hand side and 1995-2009 for the right-hand side. We regress the count of automation patents in a firm on measures of the low- and high-skill labor costs faced by their potential customers. Columns (1)-(3) control for firm and industry-year fixed effects. An increase in the low-skill wage paid by the downstream producers of an innovating firm predicts an increase in automation innovation. The estimated coefficient is an elasticity, such that an increase of 1% in the low-skill wage is associated with between 2.7% and 3.6% more automation patents. In contrast, high-skill wages predict a decrease in automation innovation of roughly the same magnitude. The regressions also control for the business cycle (GDP gap), labor productivity in manufacturing (in Column (2)), or GDP per capita (in Column (3)) in the customers' countries. None of these macroeconomic controls have consistently significant effects. We find no evidence that firms specialize more in automation innovations following a successful innovation as the stock of automation knowledge at the firm level predicts fewer automation innovations in the future. In contrast, we find clear evidence of knowledge spillovers: firms exposed to more knowledge in automation technologies tend to undertake more automation innovations. As a result, the overall effect of an increase in low-skill wages on innovation is larger than its short-run effect (see Section 4.4).

Country-year fixed effects. Unobserved country-level shocks in the innovator's country may impact both wages and innovation by affecting the cost of innovation or the demand for automation equipment through other channels than downstream wages. For instance, a tax reform in Germany could affect both German low-skill wages and directly the incentive to innovate for German firms. Shocks that affect firms mainly through their home country can be captured with (home) country-year fixed effects. As discussed further in Section 4.3, our identifying assumption is then that *foreign* wages are exogenous to the firm's automation innovation given our controls. Columns (4)-(6) report the results with country-year fixed effects: we still obtain a positive effect of low-skill wages on automation innovations and a negative effect for high-skill wages with similar elasticities. In unreported regressions, we define home country using the headquarter location and find similar effects.

Foreign wages. Building on these results, we then decompose the macro variables into their home and foreign components and analyze the effect of foreign wages. We normalize foreign variables such that the coefficients can still be interpreted as elasticities.[31] Again, we find a positive effect of low-skill wages on automation innovation and a negative effect for high-skill wages (Columns (7)-(9)). Neither ADHMV nor other papers using their methodology include country-year fixed effects or focus on foreign variation.

[31] Specifically, we can decompose total low-skill wages $w_{L,i,t}$ as $w_{L,i,t} = \kappa_{i,D} w_{L,D,t} + \kappa_{i,F} w_{L,F,t}$, where $\kappa_{i,D}$ is the home weight, $w_{L,D,t}$ the home wage, $\kappa_{i,F} = 1 - \kappa_{i,D}$ the foreign weight and $w_{L,F,t}$ the average foreign wage. We use the normalized foreign (log) low-skill wage which is defined as $\frac{\kappa_{i,F} w_{L,F,0}}{w_{L,i,0}} \log w_{L,F,t}$. The ratio $\frac{\kappa_{i,F} w_{L,F,0}}{w_{L,i,0}}$ captures that more internationally exposed firms are more affected by foreign wages. We compute it at the beginning of the sample. As $d \log w_{L,i,t} = \frac{\kappa_{i,D} w_{L,D,0}}{w_{L,i,0}} d \log w_{L,D,t} + \frac{\kappa_{i,F} w_{L,F,0}}{w_{L,i,0}} d \log w_{L,F,t}$, an increase in the normalized foreign low-skill wage by 0.01 corresponds to an increase in total wages by 1%. We define normalized foreign high-skill wages, GDP per capita, and labor productivity similarly (GDP gap is already an average of logs so we simply multiply the average foreign GDP gap with $\kappa_{i,F}$).

As argued below, these fixed effects will generally be important for identification in such settings.

The elasticity of automation innovation with respect to low-skill wages range from 2.2 to 3.7 when we focus on total wages and somewhat larger, from 4.2 to 5.3 when we focus on foreign wages. This range does not depend on the inclusion of controls for for stocks, spillovers or the GDP gap (see Appendix Table A.6). To interpret the size of this elasticity, note that our analysis focuses on innovation with a high automation content and reflects the behavior of firms undertaking automation innovations.[32] We discuss the magnitude of our effect further in Section 4.4.

Auto90. Appendix Table A.7 reproduces Table 5 but for the auto90 measure of automation. The results are very similar, but the coefficients on low-skill wages tend to be of a smaller magnitude, consistent with auto95 being a stricter measure of automation.

Skill premium. The previous results suggest that the skill premium is a driver of automation innovations since the coefficients on low-skill and high-skill wages are of a similar magnitude but opposite signs. Appendix Table A.8 directly regresses automation innovation on the log of the inverse of the skill premium. The coefficient on the inverse skill premium is similar to that on low-skill wages in the previous specifications and significant at the 1% level in all specifications.

4.2 Non-automation innovation and the direction of innovation

Is the effect of wages on automation innovations specific to automation, or does it affect machinery patents in general? To answer this question, we now look at the non-automation innovations in machinery undertaken by the sample of firms in our baseline regressions. Specifically, we reproduce the regressions of Table 5 but for machinery innovations that are not auto90, denoted pauto90. We recompute the knowledge stocks and spillover variables for pauto90 innovations ("own") and for all innovations except pauto90 ("other"). Table 6 reports the results. The coefficients on low- and high-skill wages are much smaller and only significant in one specification without country-year fixed effects for low-skill wages.[33] Therefore, the *same* firms react differently following the same shocks in their automation and non-automation innovations.

[32] By comparison, the elasticities of clean and dirty patents wrt. fuel price in ADHMV are slightly smaller (between 0.5 and 3).

[33] We drop some firms from the sample of Table 5 because they do not have pauto90 patents during this period. The baseline results on auto95 innovations remain unchanged when restricting attention to the common subsample of Table 6.

Table 6: Effect of wages on non-automation innovations

Dependent variable	Placebo machinery (pauto90)								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	0.86	0.96	1.71*	0.30	0.50	0.85	0.98	1.63	1.02
	(0.72)	(0.78)	(0.89)	(0.97)	(1.03)	(1.29)	(1.53)	(1.64)	(1.77)
High-skill wage	−0.45	−0.29	0.29	−0.75	−0.34	−0.42	−1.44	−0.54	−1.42
	(0.82)	(0.78)	(0.83)	(1.15)	(1.21)	(1.18)	(1.57)	(1.75)	(1.68)
GDP gap	−2.15	−1.96	0.15	3.39	3.81	4.39	−0.16	1.16	−0.09
	(1.56)	(1.63)	(1.91)	(4.28)	(4.27)	(4.28)	(2.89)	(3.00)	(2.97)
Labor productivity		−0.38			−0.89			−1.43	
		(0.74)			(1.26)			(1.39)	
GDP per capita			−2.28*			−1.27			−0.07
			(1.33)			(1.90)			(1.73)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	42 703	42 703	42 703	42 570	42 570	42 570	42 570	42 570	42 570
Number of firms	2859	2859	2859	2856	2856	2856	2856	2856	2856

Notes: This table replicates our baseline regressions using placebo machinery innovations. Placebo machinery are innovations in machinery excluding auto90, denoted pauto90. The sample is restricted to firms having done an auto95 innovation in the sample period. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Spillover and stock variables are calculated with respect to the dependent variable (pauto90). Standard errors are clustered at firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

The pauto90 patents are a good placebo for the auto95 patents as they are both machinery patents and tend to target the same sectors (see Appendix Figure A.9). However, to ensure even greater comparability between the two sets of patents, we focus on a subset of pauto90 patents, "refined pauto90", which are technologically closer to auto95 patents. Specifically, we identify the set of 4-digit C/IPC machinery codes which contain at least one 6-digit auto95 code (i.e. a 6-digit machinery code for which the prevalence of automation keywords is in the top 5% percent of the distribution). Refined pauto90 patents are pauto90 patents which belong to one of these 4-digit C/IPC codes.[34] Columns (1)-(3) in Appendix Table A.9 show that the results are very similar to pauto90. In that Table, we also include the full sample of firms with pauto90 innovations (Columns (4)-(6)) and look at all machinery innovations excluding auto95 innovations (pauto95 in Columns (7)-(9)). Again, the coefficients on low- and high-skill wages are much smaller than for auto95 patents and insignificant.

Finally, we show that wages affect the direction, and not just the level, of innovation

[34] Refined pauto90 patents tend to be used by the same sectors as auto95 patents: the employment-weighted correlation between the ratio of auto95 patents over capital purchases and refined pauto90 over capital purchases across US sectors rises to 0.89 (versus 0.76 for all pauto90 patents).

in Appendix Table A.10: we regress the count of auto95 innovations controlling for the number of pauto90 patents. We find similar coefficients, so that our results are not driven by a general tendency for firms to innovate more.[35]

4.3 Shift-share structure and identification

The previous results establish a correlation between firms' automation innovations and the low-skill wages faced by their customers. We now argue that this correlation reflects a causal effect of an increase in low-skill wages on automation innovation.

Since our measure of wages has a shift-share structure, we draw on the recent literature that discusses the identifying assumptions in this type of setup. We interpret our results through the lens of BHJ. In the language of our setting, they show that the random assignment of wage shocks conditional on weights and controls can be sufficient for identification. The estimator is consistent if many country-year pairs are affected by weakly correlated shocks (we argue in Appendix A.6.1 that these conditions are met).[36]

Conditionally randomly assigned wage shocks. An important feature of our research design is that we include country-year fixed effects and focus on foreign macro variables. As a result, our regression coefficients do not reflect a spurious correlation between low-skill wages and innovation coming from domestic shocks such as domestic tax policy. Instead, our regression coefficients must reflect the effect of foreign shocks that affect the demand for automation machinery and are correlated with changes in wages. Further, since we control for high-skill wages, these foreign shocks must have asymmetric effects for low- and high-skill workers – a point further supported by the results on the skill premium. Wages themselves are an equilibrium outcome, and we can think of wage shocks as coming from four sources of variation: regulatory changes, labor supply shocks, customer demand shocks, and technology shocks. We discuss these in turn.

First and second, regulatory changes or labor supply shocks present an ideal source of variation. For example, the introduction of a minimum wage, demographic or education shocks are unlikely to affect automation innovations through any channel other than

[35] To handle 0's in the count of pauto90 patents we either use the arcsinh or replace 0's with 1's in the log count and add a dummy variable for any positive count. We also run a regression where we fix the coefficient on the log(pauto90) to 1, which is the equivalent of using the ratio of auto95 / pauto90 as a dependent variable.

[36] As shown in Table 4, the Herfindahl index for our foreign weights at the country level is 0.09 and 0.006 at the country-year level. In Appendix A.6.1, we argue that there is significant variation within countries (see also Figures 3 and A.12).

Table 7: Controls for demand effects and technology shocks

Dependent variable	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	3.19***	2.39***	2.57***	2.35**	2.19*	2.08*	4.09***	5.61***	7.00***
	(0.80)	(0.85)	(0.93)	(1.00)	(1.18)	(1.24)	(1.32)	(1.61)	(1.75)
High-skill wage	−1.65**	−2.35***	−2.35***	−2.05**	−1.96*	−1.14	−4.00***	−2.99**	−3.77**
	(0.75)	(0.76)	(0.82)	(1.01)	(1.07)	(1.01)	(1.42)	(1.50)	(1.48)
GDP gap	0.58	1.15	−4.06	9.82	8.65	7.29	2.35	1.29	4.06
	(3.04)	(2.84)	(2.69)	(6.99)	(6.88)	(6.83)	(5.65)	(5.21)	(5.00)
Labor productivity		1.99**	0.75		−0.63	−2.34		−3.02*	−5.27**
		(0.96)	(1.03)		(1.93)	(1.86)		(1.78)	(2.11)
Manufacturing size	−1.69**			−2.75***			−1.00		
	(0.67)			(1.07)			(1.16)		
Manufacturing size (low-skill weighted)		−1.85***			−1.75*			0.59	
		(0.61)			(0.98)			(1.02)	
Recent auto95 innovation			−1.16			−2.54**			1.25
			(0.76)			(1.26)			(0.93)
Recent other innovation			0.87*			1.58**			−0.46
			(0.48)			(0.77)			(0.80)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 574	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3239	3252

Notes: This table adds additional control variables. Manufacturing size denotes the log of weighted average of manufacturing value added in the customers' countries. Manufacturing size (low skill-weighted) weighs the value added of each manufacturing sector by the low-skill share in the total labor share in the US in 1995. Recent auto95 innovation and recent other innovation denote the log weighted averages of the flow of auto95 and other innovations in the last 3 years in the customers' countries. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. Columns 7–9 use the normalized foreign variables as defined in the text The normalized foreign manufacturing sizes and foreign innovations are defined similarly to normalized foreign low-skill wages. Standard errors are clustered at firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

an increase in labor costs. In principle, regulation or labor supply shocks could also affect the production costs of innovating firms and thus innovation. However, as long as production is concentrated in the home country, country-year fixed effects will absorb the effect.[37] In Section 5, we will focus on minimum wage changes and a specific labor-market shock, the Hartz reforms in Germany.

Third, foreign demand shocks for the customers of innovating firms may affect foreign manufacturing wages and thus automation innovation. But foreign demand shocks may also directly affect the demand for automation equipment and innovation, which could bias our wage coefficients. The asymmetry between low- and high-skill wage coefficients already rules out the possibility that our results are driven by skill-neutral demand shocks (which are also addressed by our controls for GDP gap and GDP per capita).

[37] If a seller of automation machinery serves a foreign market through local production instead of exporting, higher foreign low-skill wages in production would increase the price of machines and therefore bias our coefficient on low-skill wages toward 0.

However, sectoral demand shifts toward low-skill intensive sectors could raise low-skill wages (relative to high-skill wages) and may lead to an increase in auto95 innovations, if these innovations tend to target more low-skill intensive sectors (see Autor et al, 2022, for a model with such sectoral shifts). This is why we control for labor productivity in manufacturing in Table 5. Further, in Columns (1), (4) and (7) of Table 7, we control for the size of the manufacturing sector, computed as the other macro variables.[38] Our coefficients on low- and high-skill wages hardly change. Still, skill-biased sectoral demand shifts could also occur within manufacturing. We build a variable to capture these: We weigh time-varying sectoral value-added within manufacturing (for 46 sectors) at the country-level by the low-skill share of labor costs in each sector in the US in 1995. We then build this variable at the firm-level using a shift-share. Columns (2), (5), and (8) report the results: the coefficients on wages remain similar. The coefficients on these controls do not suggest a consistent effect for foreign customer demand shocks.[39]

Fourth, foreign technology shocks could also directly affect automation innovations. Skill-neutral technology shocks are already addressed by our control for labor productivity. In contrast, a recent period of higher-than-usual automation innovation will increase the skill premium, and could reduce future automation innovations through another channel than wages, for instance, if it affects the competitive landscape for machinery producers. To address this, we construct a measure of recent innovation analogous to that for low-skill wages: For each country, we compute the number of automation innovations applied for in the last three years. Then, we build firm-specific measures using the same patent weights as for wages (while spillover controls use inventor weights). We construct a similar control for other innovations. Columns (3), (6) and (9) of Table 7 report the results. Our coefficients on low-skill wages remain similar, and these controls do not show a consistent effect across specifications. Reverse causality would also manifest itself as a technology shock. The effect of a firm on its home market is already captured by country-year fixed effect, but a firm's own innovation may affect foreign wages if the firm is particularly large. In addition, to the above control, we also directly control for the stock of knowledge at the firm level and lag wages. Finally, reverse causality would, if anything, bias the low-skill wage coefficient downward.[40] Therefore skill-biased

[38]We remove the control for labor productivity in manufacturing since it is closely related to that control—though keeping it does not change the results. Controlling for the share (instead of the size) of manufacturing in GDP leads to similar results in unreported regressions.

[39]Offshoring is another form of foreign demand shocks. We show that our results are robust to a control for offshoring in Appendix A.6.2.

[40]An additional concern might come from low-skill human capital shocks (captured by $\gamma(i)$ in the

technology shocks do not seem to be driving our results.

To summarize, in the presence of country-year fixed effects and a control for high-skill wages, our regression coefficients must reflect the effect of foreign demand shocks for automation equipment producers with asymmetric effect on wages. As we have just seen, these foreign demand shocks are most likely regulatory shocks or labor supply shocks that allow us to identify a causal effect of wages on automation innovations. The stability of our coefficients to various controls can also be seen as a test of the exclusion restriction (BHJ, Aghion et al., 2022). Importantly, our coefficients on low-skill wages should be compared with those from regressions with the placebo innovations (reported in Table 6 and Appendix Table A.9). Should our result on the effect of low-skill wages on automation innovations come from a bias, then that bias would have to be absent for other types of machinery innovations undertaken by the same firms and for the same 4-digit C/IPC codes. Finally, Section 5 provides direct evidence from regulatory changes with quantitatively similar results.

Falsification tests and inference. We now turn to inference. Two potential issues arise: residual errors of firms with similar country distributions may be correlated (Adão, Kolesar and Morales, 2019) and our identification variations come from a limited set of country-year observations (BHJ). To address these in our Poisson setting, we implement a Monte Carlo simulation similar to those of Borusyak and Hull (2021). We base our simulation on the regressions of Table 5. For each country, we sample with replacement the full path of macro variables (wages, labor productivity, GDP per capita, and GDP gap) from the existing set of countries (for instance Germany will get the macro variables of Canada). Then, for each firm, we compute the firm-level macro variables as the weighted average of these new country-level variables keeping the original country weights. We keep the automation activity, the stocks of innovations and the spillover variables as in the data. We run the regressions, store the coefficients on low-skill and high-skill wages and repeat 4000 times. Table 8 reports the p-values of the original coefficients on low-skill wages and high-skill wages based on the simulated distribution of coefficients. The p-values are not markedly different from those in Table 5. In particular, the low-skill wage coefficients are significant at least at the 10% level (except in Column 6 with a p-value of 0.12) and at the 5% level when we focus on foreign wages. In the language of Adão et al. (2019), the set of controls absorbs most

model of Appendix A.4), which we cannot directly control for. However, a positive shock to low-skill human capital would be associated with higher wages and less automation innovation and would correspondingly also bias our estimates downwards.

Table 8: Falsification tests including Adão et al. (2019) s.e. bias

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.96***	2.71***	3.65***	2.26**	2.61*	3.70	4.20***	5.32**	4.53**
	[0.007]	[0.001]	[0.006]	[0.018]	[0.086]	[0.119]	[0.002]	[0.026]	[0.012]
	{0.006}	{0.000}	{0.002}	{0.002}	{0.008}	{0.015}	{0.001}	{0.002}	{0.002}
High-skill wage	−2.22***	−2.62*	−1.53***	−2.79	−2.04	−1.83*	−4.46	−2.87**	−4.26*
	[0.007]	[0.059]	[0.005]	[0.179]	[0.149]	[0.056]	[0.190]	[0.026]	[0.058]
	{0.004}	{0.066}	{0.001}	{0.068}	{0.058}	{0.002}	{0.015}	{0.002}	{0.000}
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity	No	Yes	No	No	Yes	No	No	Yes	No
GDP per capita	No	No	Yes	No	No	Yes	No	No	Yes
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table reproduces the baseline regression and performs two falsification tests. The first test addresses Adao et al. (2019)'s concern that in shift-share design, observations with similar weights distribution may have correlated errors, leading to an overrejection of the null hypothesis. For each country, we sample with replacement the entire path of macroeconomics variables (wages, labor productivity, GDP per capita, and GDP gap) from the existing set of countries. We then run the same regressions as in our baseline. We repeat 4000 times and compare the coefficients from the true regression with the distribution of coefficients from simulated regressions. The [ ] brackets report the p-values of our original coefficients. In a second exercise, we perform a similar exercise but instead of re-drawing the macro-variables, we re-draw firms' weights from the distribution of weights for firms in the same country. We run regressions on the simulated data and the brackets state the p-values. Significance levels at *10%, **5% and ***1% using the [ ] p-values.

of the country-specific shocks affecting the outcome variable and, consequently, no shift-share structure is left in the regression residuals. Appendix Figure A.3, panels a, b, c plot the distribution of coefficients for Columns (2), (5), and (8). We can also view this permutation exercise as a falsification test, and accordingly, the distributions are centered around 0.

Similarly, we perform an additional falsification test where instead of permuting the macro variables across countries, we permute the weights across firms from the same country. Specifically, for each firm, we keep the automation activity, the stocks of innovations and the spillover variables as in the data, but we sample (with replacement) their weights from the set of firms from the *same* country. That is, we may now attribute to Siemens the wages that, in reality, Bosch faces. We repeat the exercise 4000 times and Table 8 reports the p-values of the original coefficients on low-skill wages and high-skill wages based on the simulated distribution of coefficients. The p-values are similar or smaller than those of Table 5. Figure A.3 plots the distribution of coefficients we obtain for Columns (2), (5), and (8). The distributions are centered around 0 when we include country-year fixed effects (panels e and f). In panel d, we use the home country variation for identification and the mean coefficient is positive but significantly smaller. This exercise highlights that the relevant variation is between firms of the same country and not simply cross-country.

4.4 Magnitude of the effect

Having established that the effect of wages on automation innovations is likely causal, we now focus on the magnitude of the effects.

High-skill wages. We found a large positive effect of high-skill wages when controlling for low-skill wages. This is consistent with a large literature on capital-skill complementarity (Krusell et al., 2000), with our model in Appendix A.4, and with the results in Section 2.7 on the skill ratio. More generally, papers studying the effect of automation technologies on employment outcomes frequently (though not always) document an increase in the skill ratio (see, among others, Graetz and Michaels, 2018, Humlum, 2021, and Boustan, Choi and Clingingsmith, 2022).

Comparison with the literature. We now compare our estimates to those of the two papers closest to ours: Lewis (2011) and Acemoglu and Restrepo (2022). Lewis (2011) shows that US manufacturing plants adopted fewer automation technologies when the local ratio of low- to middle-skill workers increased following immigration shocks. He measures automation with the SMT, from which we derived a large share of our keywords. He also measures the effect of an increase in the ratio of low- to middle-skill workers on the relative wages of high-, middle-, and low-skill workers. Combining these numbers, we can back out an elasticity of automation *adoption* with respect to the inverse skill premium of 3.6, very much in line with our coefficients. Deriving this number, however, requires making a number of assumptions and should be assessed accordingly (see Appendix A.6.3 for details).

Acemoglu and Restrepo (2022) study the effect of aging on robotics adoption and innovation across countries. They also report the effect of aging on blue-collar manufacturing wages across US commuting zones. Combining these results, we can back out an elasticity of robot adoption with respect to relative blue-collar manufacturing wages of 3.9 and an elasticity of innovation of 1.5 (see Appendix A.6.3 for details). Again, these "back-of-the-envelope" numbers should be taken with caution, and cannot substitute for our analysis which looks directly at the effect of wages on automation innovation. Still, it is reassuring that the magnitudes of the effects are similar.

Conceptually, and as emphasized in the introduction, it is important to distinguish the adoption response to a wage shock from the innovation response. In particular, calibrating directed technical change macro models requires estimates of the innovation response (Acemoglu, 2023). While a shock to adoption must die out over time (when all firms have adopted a given technology), a shock to innovation can build on itself. In the short-run, technology adoption is dominated by the increased use of existing technologies. In the medium run, however, the innovation response is a key determinant of the overall adoption response, while new adoption of existing technologies plays a smaller role. In the longer run, the innovation response also depends on knowledge spillovers.[41]

Simulation results. In the following, we analyze the economic magnitude of our coefficients and highlight the role played by spillovers. We do so by estimating the long-run effect of a change in the skill premium on tasks demand and the labor share that runs through endogenous automation innovation. To do so, we run a simulation in which we consider a uniform and permanent decrease in the skill premium by 10% between 1995 and 2009 (which is perhaps easier to interpret than a change in low-skill wages keeping high-skill wages constant). We run a regression that jointly estimates the effect of automation innovations (auto95) and other machinery innovations (pauto95). We then use our regression results to recompute the share of automation innovations in machinery over that period, both including spillovers and firms' knowledge stocks and not. Appendix A.6.4 details our procedure including the exact regression (Table A.41).

Figure 4 reports the results averaged over 500 simulations.[42] We first compute the direct effect of a decrease in the skill premium (keeping stocks and spillover variables constant) on the share of automation innovations in machinery. This is captured by the gap between the data curve and the data + direct effect curve. This gap reflects the elasticity of 2.51 of auto95 innovations with respect to the inverse skill premium (with an elasticity of 0.44 for pauto95). Next, we include the effect of updating firms' own innovation stocks in the curve "data + direct + stock", which slightly decreases the effect of low-skill wages reflecting the negative effect of the automation stock on auto95

[41]There can also be spillovers in the adoption of new technologies, though fundamentally they are of a different nature: spillovers in adoption are likely to hasten the diffusion of a given technology in the whole economy, while spillovers in innovation (such as the building-on-the-shoulders-of-giants externality) can lead to a different technology path. A rich literature studies spillovers in technology diffusion and adoption. These tend to be large for technologies exhibiting complementarities such as payment systems, but for the machinery technologies that we are studying, spillovers are likely to be mostly informational. Existing studies (Baptista, 2000, No, 2008, Bekes and Harasztozi, 2020) point toward spillovers with elasticities between 0.001 and 0.08, much smaller than the coefficients we find here (see Appendix A.6.3).

[42]The figure reports the share of automation patents for the firms in our regression sample. This differs from Figure 2 which reports the share of automation patents for all firms.

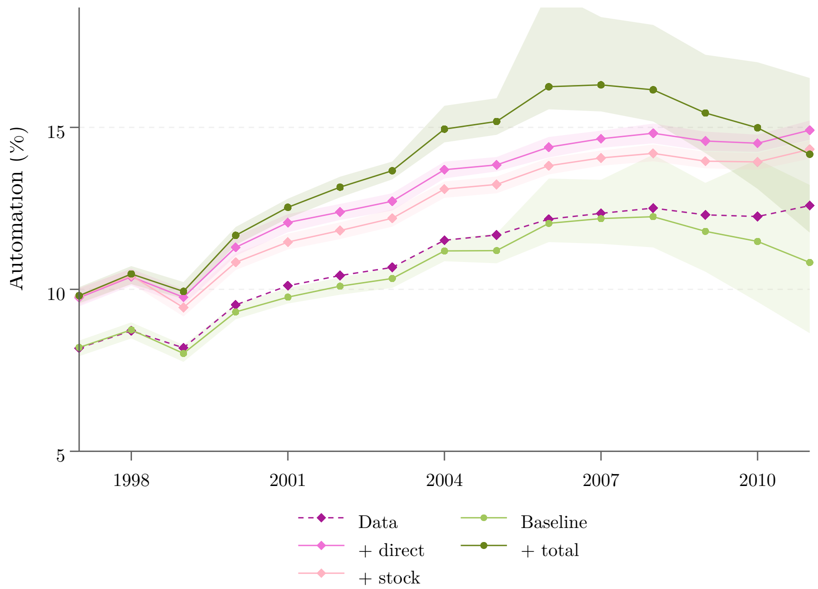

Figure 4: Simulation of a permanent and global 10% decrease in the skill premium on the share of automation innovations in machinery

Notes: We report the median share of automation patents in machinery across 500 simulations. The shaded areas depict the p25-p75 range. The effects are computed based on the regression results of Appendix Table A.41. The baseline curve reports the mean realization without any change in the skill premium, but when the knowledge spillovers are recomputed every period depending on the simulated number of innovation in the previous periods. The total effect of the skill premium shock is captured by the gap between the total effect curve and the baseline curve.

innovations and its positive effect on pauto95 innovations.

We then assess the importance of knowledge spillovers by recomputing the spillover variables for the auto95 and pauto95 innovations. This is not straightforward exercise because we need to predict not only the number of innovations but also their location. We start by simulating the baseline curve which is the median realization without a wage change (while the data series is one possible realization). The overall effect of an increase in the inverse skill premium on the share of automation innovation is then captured by the gap between the baseline curve and the baseline + total effect curve. Knowledge spillovers increase the elasticity of the share of automation patents with respect to the inverse skill premium. The average share of automation innovations in machinery between 1997 and 2011 increases by 3.3 pp, that is 1.3 pp more than the direct effect (for comparison, the share of automation innovations increased by 4.4 pp over the same time period).

Finally, we combine these effects with the results of Section 2.7 on the labor market outcomes of automation use (Columns (1), (3), (5) of Table 3). A 3.3 pp increase in the share of automation innovation is associated with a decline in routine cognitive tasks of 4.7 centiles, a decline in routine manual tasks of 4.2 centiles, and a decline in the labor share in manufacturing of 4.3 pp per decade (for comparison, routine cognitive, routine manual tasks and the labor share declined by 1.4 centiles, 1.3 centiles and 5.1 pp per decade, respectively, in the sectors considered over the period 1980-1998). Importantly, we stress that one *must not* interpret the results of this simulation as predictive, notably because a change in innovation will affect the skill premium. Nevertheless, this partial equilibrium analysis shows that the effect of the inverse skill premium on automation innovations is economically significant.

4.5 Additional Results and Robustness Checks

This section discusses additional results and robustness checks.

Shift-share design. BHJ show that, in our context, shift-share firm-level regressions are equivalent to weighted shock-level (i.e. country-year level) regressions. In Appendix A.6.1, instead of our Poisson regression, we consider a linear setting where such an equivalence result applies: we use the arcsinh of the count of automation patents as the dependent variable and replace our log of the average macro variables with the average of the logs. This linear setting allows us to give summary statistics on our shock variable and unpack the relationship between the inverse skill premium and automation in the data. Appendix Figure A.11 shows bin-scatter plots of the shock-level regressions of residualized automation measures on the inverse skill premium: the relationship appears linear and not driven by outliers. We also report how balanced our shocks are with respect to observables.

Goldsmith-Pinkham, Swan and Swift (2020) show that alternatively, identification in a shift-share design can be obtained if the weights are exogenous. In our context, this assumption is likely violated because firms' decision to innovate may be affected by other macro shocks in the destination countries, which would affect firms in proportion to the same weights. This is why we rely on BHJ. Nevertheless, we note that our weights appear predetermined and do not reflect firms' expectations of future wage growth: country-level growth rates in low- and high-skill wages between 1995 and 2000 have no predictive power on firm weights in 1995 (see Appendix Table A.11). In addition, we conduct robustness checks on our weights in Appendix Table A.12: We exclude automation patents from the weights; we use a longer lag between the period used to compute the weights and the regression period, either by computing weights only up to 1989 or by dropping the first 5 years of the regression; and drop the earlier years of the pre-sample period when computing the weights. Our results are robust in all cases.

We conduct additional exercises related to our shift-share setting in Appendix A.6.1. First, we show that no single country drives our results by sequentially excluding the six largest countries. Second, BHJ recommend considering other shock-level variables that may bias the results. We consider the effect of offshoring and the real interest rate. We also control for R&D costs by building firm-specific wage variables using weights based on the location of inventors instead of patent offices. Finally, we implement the correction suggested by Borusyak and Hull (2021) to address the non-linearity of logged shift-share measures. Our results are robust in all cases.

Timing and pre-trends. We now study different timing assumptions to assess whether there are pre-trends and because our choice of a 2-year lag is somewhat arbitrary. In Appendix Figure A.4, we look at alternative lags and leads for the dependent variables. We consider two specifications, both controlling for GDP gap, labor productivity, and country-year fixed effects. In Panel a, we look at total wages, corresponding to Column (5) of Table 5. In Panel b, we consider only foreign wages, corresponding to Column (8).[43] The 2-year lag delivers the highest coefficient for low-skill wages in both cases. This is in line with the empirical literature on induced innovation using patent data which often finds effects peaking with a 2-3 year lag (see e.g. ADHMV or Popp, 2002). A possible interpretation of this fast response is that firms may prioritize existing projects over starting new ones.[44]

Figure A.4 also looks at the effect of leads of wages on automation innovations. The early leads (up to 2 years) show significant effects for high-skill wages. This is not surprising: wages are auto-correlated and firms may anticipate shocks at short horizons. Importantly, though, we find no significant effect for longer leads, suggesting that there are no pre-trends (testing for such pre-trends is one of the recommendations of BHJ). Appendix Table A.13 runs a horse-race regression between 2-year lagged macro variables and macro variables with varying leads or lags. The coefficients on lead wages are always insignificant and the effect of 2-year lagged low-skill wages is always positive and significant except in one specification.

Additionally, innovators should only care about current wages insofar as they are

[43] We keep a lag of two periods for the stock variables; otherwise, the dependent variable would be included in the RHS in the lead and contemporaneous cases.

[44] In contrast, our regressions are unlikely to only capture the effect of patenting off-the-shelf inventions which already exist and have become commercially viable. First, Hall, Griliches and Hausman (1986) and Kaufer (1989) show patent applications to be timed closely to research expenditures because the first-to-file rule provides inventors with a strong incentive to patent as early as possible (Dechezleprêtre et al., 2017). Second, in that case, the largest effect of wages on patents would be contemporaneous.

predictive of future wages. In Appendix Table A.14, we compute predicted future wages at time $t-2$ based on an AR(1) process with country-specific trends instead of directly using lagged wages. The results are similar to our baseline.

Nickell's bias. Our regressions include the stock of automation innovations and may suffer from Nickell's bias. Appendix Table A.15 removes stocks or uses the standard method of Blundell, Griffith and van Reenen (1999), that proxies for the fixed effect with the firm's pre-sample average of the dependent variable. We obtain similar results.

Additional results. We derive more results in Appendix A.6.2. We control for middle-skill wages or firm-size year fixed effects. Then, we consider alternative specifications: we run long-difference regressions or cluster our baseline regressions at the country level. Finally, we look at alternative measures of firm-level wages (by pre-multiplying patents weights with other factors than $GDP^{0.35}$ or by converting macro variables in USD differently or using total wages) and firm-level innovations (by using citations-weighted measures of patents and subcategories of automation innovations).

5 Labor market reforms and induced automation

We now focus on two specific, identifiable labor shocks: minimum wage changes and the German Hartz reforms. This complements our main analysis, which was agnostic about the exact nature of the labor market shocks driving automation innovations.

5.1 Minimum wage

We compute the average minimum wage faced by the downstream customers of a specific firm selling automation equipment. We do this as a shift-share measure of country-level minimum wages exactly as we did for low-skill wages, with the caveat that we only have data for 22 countries instead of 41.[45] We then run panel regressions similar to our baseline where we replace the low-skill wage with the minimum wage. Table 9 reports the results. For regressions on total wages (Columns (1)-(6)), the coefficients on minimum wage are very similar to those on low-skill wages in previous tables, and for regressions on foreign wages (Columns (7)-(9)) the coefficients are positive but smaller and, in one

[45] We use data from the OECD. Importantly, not all countries have government-mandated minimum wages, and for some countries, we follow the literature and use sectorally bargained minimum wages. See details in Appendix A.5.1.

case, insignificant.[46] Overall, these results support our hypothesis that higher low-skill labor costs induce automation innovations.

Table 9: Effect of the minimum wage

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Minimum wage	2.14***	1.89***	2.17***	1.86**	2.00**	2.10*	2.27*	2.38*	1.24
	(0.63)	(0.64)	(0.79)	(0.89)	(0.94)	(1.07)	(1.19)	(1.26)	(1.43)
High-skill wage	−1.86***	−2.51***	−1.80**	−3.59***	−3.03**	−3.16**	−3.51**	−3.12*	−5.15***
	(0.67)	(0.79)	(0.84)	(1.03)	(1.26)	(1.43)	(1.38)	(1.86)	(1.86)
GDP gap	−2.49	−3.43	−2.38	7.43	8.15	8.35	3.16	3.61	−1.22
	(2.51)	(2.59)	(2.78)	(6.46)	(6.54)	(7.07)	(4.80)	(5.27)	(6.20)
Labor productivity		1.27			−1.02			−0.50	
		(0.79)			(1.50)			(1.63)	
GDP per capita			−0.11			−0.87			3.39
			(1.23)			(2.05)			(2.54)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 046	48 046	48 046	47 724	47 724	47 724	46 575	46 575	46 575
Number of firms	3252	3252	3252	3250	3250	3250	3167	3167	3167

Notes: This table replaces the low-skill wage with the minimum wage. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

5.2 Event study: the Hartz reforms in Germany

We now examine the effects of the Hartz reforms, a series of labor-market reforms in Germany, first drafted in 2002 and implemented between January 1st, 2003 and January 1st, 2005. In order to reduce unemployment and increase labor-market flexibility, the government reformed employment agencies, deregulated temporary work, offered wage subsidies for hard-to-place workers, reduced or removed social contributions for low-paid jobs, and reduced long-term unemployment benefits. Krause and Uhlig (2012), among others, have attributed an important role to the reforms in the remarkable performance of the German labor market since then, particularly in increasing labor supply and improving matching efficiency.

Such reforms would reduce the incentive to automate low-skill labor by decreasing labor costs, both directly and indirectly through an increase in labor supply and a

[46] A smaller coefficient is not surprising: First, we focus on manufacturing, where low-skill wages tend to be above the minimum wage. Second, the minimum wage captures only a portion of the labor costs. Third, the quality of the data is worse as we lose nearly half of our countries.

reduction in the expected cost of vacancies. The Hartz reforms are perhaps the most salient labor market reforms in a major country in our time period. This presents an ideal setting: The Hartz reforms are unlikely to have affected the direction of innovation in non-German firms through channels other than the German labor market and were the largest macroeconomic shock in Germany at the time. The reforms had a large and immediate effect. As soon as they were implemented in 2003, low-skill labor costs started stagnating while high-skill labor costs kept rising, leading to a sharp decline in the inverse skill-premium in Germany (see Appendix Figure A.5). In contrast, there is no such trend for the aggregate rest of the world.

We use an approach analogous to our main analysis, measuring innovation and firms' exposure to international markets. However, we exclude German firms as the Hartz reforms likely affected them through channels other than their customers' labor costs. We run the following regression over the years 1997–2014:

$$\mathrm{E}\left(PAT_{Aut,i,t}\right) = \exp\left(\beta_{DE,t} \cdot \delta_t \kappa_{i,DE} + \beta_{Ka} \ln K_{Aut,i,t-2} + \beta_{Ko} \ln K_{other,i,t-2} + \delta_i + \delta_{j,t} + \delta_{c,t}\right).$$

We keep a 2-year lag on the innovation stocks. As before, $PAT_{Aut,i,t}$ counts automation patents, $K_{Aut,i,t-2}$ and $K_{other,i,t-2}$ denote firm knowledge stocks, δ_i, $\delta_{j,t}$, and $\delta_{c,t}$ are firm, industry-year, and country-year fixed effects, respectively. $\kappa_{i,DE}$ is the fixed German weight of the firm; and δ_t is a set of year dummies (with 2005 the excluded year). $\beta_{DE,t}$ are the coefficients of interest. They state by how much more a firm exposed to Germany tends to file automation patents in a given year relative to 2005.

Figure 5.a reports the results. The coefficient of -2.71 in 2010 means that, on average, a firm with a German weight of 0.1 (the mean value is 0.104) had a 27.1% smaller increase in automation innovations between 2005 and 2010 than a firm with no German exposure. This aligns with our regression results: Between 2003 and 2008, the inverse skill-premium in Germany declined by 12.3% relative to the rest of the world. Using the elasticity of 2.5 of Column (4) in Table A.8, this would correspond to a decline in automation innovations of 30.8% between 2005 and 2010.

From 1999 to 2004, firms more exposed to Germany slightly increased their propensity to introduce automation innovations. As expected, the trend reversed between 2005 and 2009, consistent with the Hartz reform increasing labor supply from 2003 onward and a 2-year lag effect on innovation. From 2010 on, the coefficients increase again. This reversal suggests only temporary effects of the Hartz reform on the direction of innovation, or may be the result of the Great Recession.

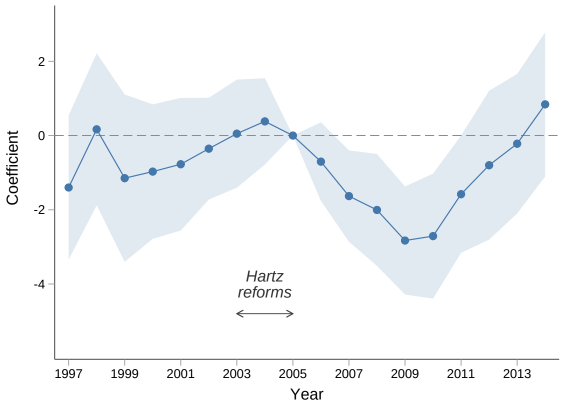

(a) Effect of German exposure on automation innovations

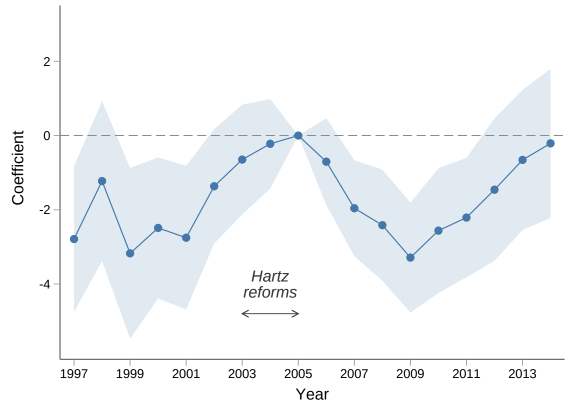

(b) Effect of German exposure on automation innovations relative to other machinery innovations.

Figure 5: Effect of German exposure on automation innovations.

Notes: Panel (a) reports coefficients on the interaction between the German weight and a set of year fixed effects in a Poisson regression of auto95 innovations controlling for a full set of fixed effects and firm innovation stocks with 2154 firms. Panel (b) reports coefficients on the triple interaction between the German weight, a dummy for auto95 innovations, and a set of year fixed effects in a regression of auto95 and other machinery innovations controlling for a full set of fixed effects, firm innovation stocks and the interaction between the German weight and a set of year fixed effects with 6690 firms. Standard errors are clustered at the firm level and the shaded areas represent 95% confidence intervals. The figure shows that the relative trend in automation innovation for firms more exposed to Germany reversed after the Hartz reforms.

We conduct a triple difference exercise to show that the trends above are specific to automation innovations. We compare automation innovations with non-automation machinery innovations by firms more or less exposed to Germany over time. Formally, we run the following regression:

$$\mathrm{E}\left(PAT_{k,i,t}\right) = \exp\left(\begin{array}{c} \beta_{DE,t} \cdot \delta_t \kappa_{i,DE} + \beta_{DE,t}^{aut} \cdot \delta_t \kappa_{i,DE} 1_{k=aut} + \beta_{Ka} \ln K_{Aut,i,t-2} \\ \beta_{Ka}^{aut} \ln K_{Aut,i,t-2} 1_{k=aut} + \beta_{Kp} \ln K_{Paut,i,t-2} + \beta_{Kp}^{aut} \ln K_{Paut,i,t-2} 1_{k=aut} \\ +\beta_{Ko} \ln K_{other,i,t-2} + \beta_{Ko}^{aut} \ln K_{other,i,t-2} 1_{k=aut} + \delta_{k,i} + \delta_{k,j,t} + \delta_{k,c,t} \end{array}\right). \tag{5}$$

k denotes the type of an innovation which is either auto95 or other machinery innovation (pauto95), $\delta_{k,i}$ represents a set of innovation type firm fixed effects, $\delta_{k,c,t}$ innovation type country-year fixed effects, $\delta_{k,j,t}$ innovation type industry year fixed effects and $1_{k=aut}$ is a dummy for an auto95 innovation. $K_{Paut,i,t}$ is the stock of other machinery innovations (pauto95) and $K_{other,i,t}$ the stock of non-machinery innovations. $\beta_{DE,t}^{aut}$ are the coefficients of interest. For each year, they measure how much exposure to Germany increases the relative propensity to introduce automation innovations compared to other forms of machinery innovations relative to 2005. The coefficients $\beta_{DE,t}$ measure the effect of

Table 10: Automation vs non-automation innovation and exposure to Germany: triple diff exercise

	Auto95 and pauto95			Auto95 and pauto90
	(1)	(2)	(3)	(4)
Time trend × auto95 dummy × German exposure × post	-1.06^{***}	-1.11^{***}	-1.09^{***}	-1.14^{***}
	(0.33)	(0.33)	(0.33)	(0.33)
Time trend × auto95 dummy × German exposure	0.50^{***}	0.47^{***}	0.46^{***}	0.47^{***}
	(0.18)	(0.18)	(0.18)	(0.18)
Time trend × German exposure × post			0.31	
			(0.22)	
Time trend × German exposure			-0.34^{**}	
			(0.15)	
Firm innovation stocks × innovation types	No	Yes	Yes	Yes
Year dummy × German exposure	Yes	Yes	No	Yes
Industry × year × innovation types FE	Yes	Yes	Yes	Yes
Country × year × innovation types FE	Yes	Yes	Yes	Yes
Firm × innovation types FE	Yes	Yes	Yes	Yes
Observations	76 136	76 136	76 136	74 124
Number of firms	5416	5416	5416	5278

Notes: This table shows that the effect of German exposure is specific to automation innovations. All regressions control for firm innovation types fixed effects, country-year-innovation types fixed effects, and industry-year-innovation types fixed effects. Innovation types are auto95 and pauto95 in Columns 1–3 and auto95 and pauto90 in Column 4. Column 2–4 control for innovation stocks lagged by two periods interacted with innovation types dummies. Column 3 controls for a linear time trend times the German exposure instead of yearly dummies times the German exposure. Throughout, German exposure is measured by the German weight. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

German exposure that is common to all machinery innovations. Figure 5.b reports the results: the pattern is, if anything, more pronounced than in Figure 5.a.

To formally test that the Hartz reform created a trend break, we replace the set of year fixed-effects δ_t in $\beta_{DE}^{aut} \cdot \delta_t \kappa_{i,DE} 1_{k=aut}$ in equation (5) with a time trend $t - 2005$ and a time trend interacted with a post 2005 dummy $(t - 2005)_{t>2005}$. We focus on the years 2000-2010 to have a panel centered on 2005 and to avoid the effects of the Great Recession on innovation. Table 10 reports the result. Column (2) corresponds exactly to this specification. We find a significant time trend in the effect of German exposure on the relative propensity to innovate in automation between 2000 and 2005. However, the trend sharply reverses in the following five years. Column (1) omits the controls for the stock variables. Column (3) replaces the flexible set of year dummies times German exposure, $\delta_t \kappa_{i,DE}$, by a time trend times German exposure and a time trend times German exposure post 2005. Finally, instead of looking at auto95 and pauto95 (i.e. all non-auto95 machinery innovations) innovation, Column (4) considers auto95 and pauto90 innovations (which we used as the default non-automation innovations in Table 6). In all cases, the trend break on automation innovations remains with a consistent magnitude. Overall, this section shows that, in line with our theory, the Hartz reforms reduced automation innovations by foreign firms highly exposed to Germany, both in absolute terms and relative to other types of machinery innovation.

6 Conclusion

In this paper, we identify automation patents and present evidence that equipment producers innovate more in automation technologies following increases in the low-skill labor costs of downstream firms. We develop a method to classify patents in machinery as automation or not, covering a broad range of technologies. We then use this classification to measure the use of automation technology by industry at a highly disaggregated level and find that our measure of automation predicts a decline in routine tasks, an increase in the skill ratio and a decrease in the labor share across US sectors.

Further, we use our classification to analyze labor market conditions' effect on machinery automation innovations. Using global data and firm-level variation, we find that automation innovations are highly responsive to changes in low-skill wages, with elasticities between 2 and 5, while an increase in high-skill wages decreases automation innovations. In contrast, other innovations in machinery by the same set of firms do not respond to changes in labor costs. To complement our analysis, we then focus on two policy-induced labor market shocks. We show that increases in the minimum wage lead to more automation innovations and that the Hartz reforms, aimed at reducing the effective cost of low-skill labor, induced a relative decrease in automation innovations by foreign firms with high exposure to Germany.

Our results highlight that labor market policies can generate technological responses, so that the long-term effects of these policies may differ from their short-term effects. Analyzing this feedback loop quantitatively would require the development of a macroeconomic model, which could be calibrated using our estimates. More generally, our estimates can discipline an emerging literature on automation and economic growth. In addition, future research could adapt our classification method to automation patents beyond machinery, and analyze how much the emergence of recent automation technologies for high-skill labor, such as AI, results from rising high-skill wages.

References

Acemoglu, Daron, "Distorted Innovation: Does the Market Get the Direction of Technology Right?," *American Economic Review Papers and Proceedings*, 2023.

— **and Amy Finkelstein**, "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector," *Journal of Political Economy*, 2008, *116* (5), 837–880.

— **and Joshua Linn**, "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *Quarterly Journal of Economics*, August 2004, *119* (3), 1049–1090.

— **and Pascual Restrepo**, "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

— **and** — , "Robots and Jobs: Evidence from US labor markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

— **and** — , "Demographics and Automation," *The review of economic studies*, 2022, *89* (1), 1–44.

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

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen, "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry," *Journal of Political Economy*, 2016, *214* (1), 1–51.

— **, Celine Antonin, Simon Bunel, and Xavier Jaravel**, "The Labor and Product Market Effects of Machines and Automation. Causal Evidence from French Firms and Industries," 2022.

— **, Roland Bénabou, Ralf Martin, and Alexandra Roulet**, "Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?," *American Economic Review: Insights*, forthcoming.

Andersson, David, Mournir Karadja, and Erik Prawitz, "Mass Migration and Technological Change.," *Journal of the European Economic Association*, 2022.

Aschhoff, Birgit, Dirk Crass, Katrin Cremers, Christoph Grimpe, Christian Rammer, Felix Brandes, Fernando Diaz-Lopez, Rosalinde Klein Woolthuis, Michael Mayer, and Carlo Montalvo, "European Competitiveness in Key Enabling Technologies," Technical Report, Centre for European Economic Research (ZEW) and TNO 2010.

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

— **, Caroline Chin, Anna Salomons, and Bryan Seegmiller**, "New Frontiers: The Origins and Content of New Work, 1940-2018," 2022.

— **, David Dorn, Gordon Hanson, Gary Pisano, and Pian Shu**, "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents," *American Economic Review: Insights*, 2020, *2* (3), 357–374.

— **, Frank Levy, and Richard Murnane**, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Baptista, Rui, "Do innovations diffuse faster within geographicalclusters?," *International Journal of Industrial Organization*, 2000, *18*, 515–535.

Bena, Jan and Elena Simintzi, "Machines could not compete with Chinese labor: Evidence from U.S. firms' innovation," 2019. mimeo.

Bessen, James and Robert Hunt, “An empirical look at software patents,” *Journal of Economics and Management Strategy*, 2007, *16* (1), 157–189.

—, **Maarten Goos, Anna Salomons, and Wiljan van der Berge**, “Automatic Reaction - What Happens to Workers at Firms that Automate ?,” 2019.

Békés, Gábor and Péter Harasztosi, “Machine imports, technology adoption, and local spillovers,” *Review of World Economy*, 2020, *156*, 343–375.

Blundell, Richard, Rachel Griffith, and John Van Reenen, “Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms,” *The Review of Economic Studies*, 1999, *66* (3), 529–554.

Borusyak, Kirill and Peter Hull, “Non-Random Exposure to Exogenous Shocks: Theory and Applications,” 2021.

—, —, **and Xavier Jaravel**, “Quasi-experimental Shift-share Research Designs,” *Review of Economic Studies*, 2022, *89* (1), 181–213.

Boustan, Leah Platt, Jiwon Choi, and David Clingingsmith, “Automation after the Assembly line: Computerized Machine Tools, Employment and Productivity in the United States,” 2022.

Cameron, A. Colin, Jonah Gelbach, and Douglas Miller, “Boostrap-based Improvements for Inference with Clustered Errors,” *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Clemens, Michael, Ethan Lewis, and Hannah Postel, “Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion,” *American Economic Review*, June 2018, *108* (6), 1468–87.

Coelli, Federica, Andreas Moxnes, and Karen Ulltveit-Moe, “Better, Faster, Stronger: Global Innovation and Trade Liberalization,” *The Review of Economics and Statistics*, 2020.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin, “Fast Poisson estimation with high-dimensional fixed effects,” *The Stata Journal*, 2020, *20* (1), 95–115.

Danzer, Alexander, Carsten Feuerbaum, and Fabian Gaessler, “Labor Supply and Automation Innovation,” June 2020. BGPE Discussion Paper No 198.

de Rassenfosse, Gaétan, Hélène Dernis, Dominique Guellec, Lucio Picci, and Bruno van Pottelsberghe de la Potterie, “The worldwide count of priority patents: A new indicator of inventive activity,” *Research Policy*, 2013, *42* (3), 720–737.

Dechezleprêtre, Antoine, Matthieu Glachant, Ivan Hascic, Nick Johnstone, and Yvan Ménière, “Invention and transfer of climate change-mitigation technologies: a global analysis.,” *Review of Environmental Economics and Policy*, 2011, *5*, 109–130.

—, **Yann Ménière, and Myra Mohnen**, “International patent families: from application strategies to statistical indicators,” *Scientometrics*, 2017, *111* (2), 793–828.

Doms, Mark, Timothy Dunne, and Kenneth Troske, “Workers, Wages, and Technology,” *Quarterly Journal of Economics*, 1997, *112* (1), 253–290.

Eaton, Jonathan and Samuel Kortum, “Trade in ideas: patenting and productivity in the OECD,” *Journal of International Economics*, 1996, *40*, 251–278.

—, —, **and Francis Kramarz**, “An anatomy of international trade: evidence from French firms,” *Econometrica*, 2011, *79* (5), 1453–1498.

Fan, Haichao, Yichuan Hu, and Lixin Tang, "Labor costs and the adoption of robots in China," *Journal of Economic Behavior & Organization*, 2021, *186*, 608–631.

Gaggl, Paul and Greg Wright, "A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive," *American Economic Journal: Applied Economics*, 2017, *9* (3), 262–94.

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

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

Hall, Bronwyn, Zvi Griliches, and Jerry Hausman, "Patents and R&D: Is There A Lag?," *International Economic Review*, 1986, *27* (2), 265–284.

Hausman, Jerry, Bronwyn Hall, and Zvi Griliches, "Econometric Models for Count Data with an Application to the patents-R&D Relationship," *Econometrica*, 1984, *52* (4), 909–938.

Hémous, David and Morten Olsen, "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," *American Economic Journal: Macroeconomics*, 2022, *14* (1), 179–223.

Hornbeck, Richard and Suresh Naidu, "When the levee breaks: Black migration and economic development in the American South," *American Economic Review*, 2014, *104(3)*, 963–990.

Humlum, Anders, "Robot Adoption and Labor Market Dynamics," 2021. mimeo.

Kaufer, Erich, *The Economic of the Patent System*, New York: Rootledge, 1989.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, "Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations," 2022. mimeo.

Krause, Michael and Harald Uhlig, "Transitions in the German labor market: Structure and crisis," *Journal of Monetary Economics*, 2012, *59* (1), 64–79.

Krusell, Per, Lee Ohanian, José-Víctor Ríos-Rull, and Giovanni Violante, "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 2000, *68* (5), 1029–1053.

Lewis, Ethan, "Immigration, Skill Mix and Capital Skill Complementarity," *Quarterly Journal of Economics*, May 2011, *126 (2)*, 1029–1069.

Lordan, Grace and David Neumark, "People versus machines: The impact of minimum wages on automatable jobs," *Labour Economics*, 2018, *52* (40-53).

Lybbert, Travis and Nikolas Zolas, "Getting patents and economic data to speak to each other: An 'Algorithmic Links with Probabilities' approach for joint analyses of patenting and economic activity," *Research Policy*, 2014, *43*, 530–542.

Mann, Katja and Lukas Puettmann, "Benign Effects of Automation: New Evidence from Patent Texts," *The Review of Economic and Statistics*, August 2021, pp. 1–45.

No, Joung, "Cities and Growth: Knowledge Spillovers in the Adoption of Advanced Manufacturing Technologies," 2008. The Canadian Economy in Transition Series.

Noailly, Joëlle and Roger Smeets, "Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data," *Journal of Environmental Economics and Management*, 2015, *72*, 15–37.

Popp, David, "Induced Innovation and Energy Prices," *The American Economic Review*, 2002, *92* (1), 160–180.

San, Shmuel, "Labor Supply and Directed Technical Change: Evidence from the Abrogation of the Bracero Program in 1964," *American Economic Journal: Applied Economics*, 2023, *15* (1), 136–63.

Siemens, "Annual Report," Technical Report 2018.

Timmer, Marcel, Abdul Azeez Erumban, Bart Los, Robert Stehrer, and Gaaitzen de Vries, "Slicing up Global Value Chains," *Journal of Economic Perspectives*, 2014, *28* (2), 99–118.

__ , Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries, "An Illustrated User Guide to the World Input-Output Database: the Case of Global Automotive Production," *Review of International Economics*, 2015, *23* (3).

Voth, Hans-Joachim, Bruno Caprettini, and Alex Trew, "Fighting for Growth: Labor Scarcity and Technological Progress During the British Industrial Revolution," 2022.

Webb, Michael, "The Impact of Artificial Intelligence on Labor Market," 2020.

Zeira, Joseph, "Workers, Machines, and Economic Growth," *Quarterly Journal of Economics*, 1998, *113* (4), 1091–1117.

Online Appendix

Table of Contents

A Main Appendix

A.1 Additional Figures and Tables

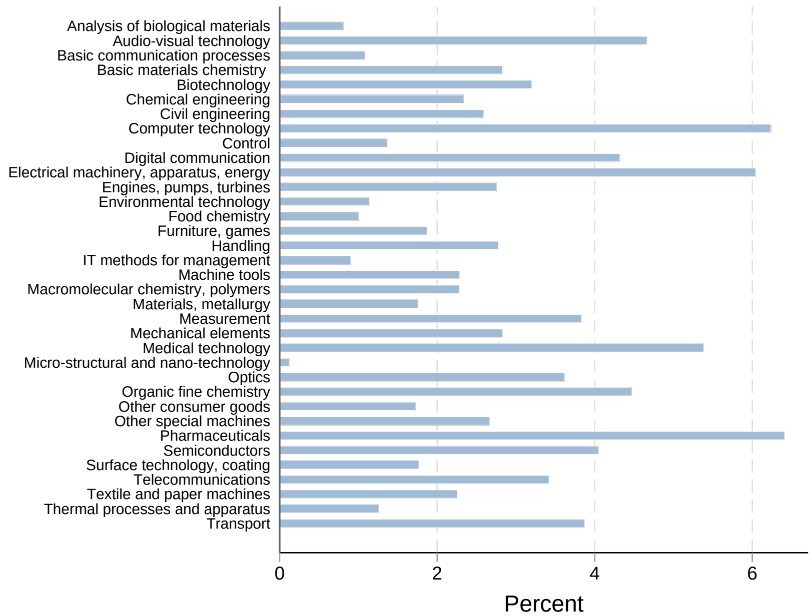

Figure A.1: Share of biadic patent applications in the different technical fields in 1997-2011.

Figure A.2: Trends in automation patents.

Notes: Panel (a) reports the share of automation patents (auto90 or auto95) in machinery out of total patents according to the auto90 and auto95 definitions. Panel (b) reports the raw number of automation patents (auto90 or auto95) worldwide. We restrict attention to biadic families.

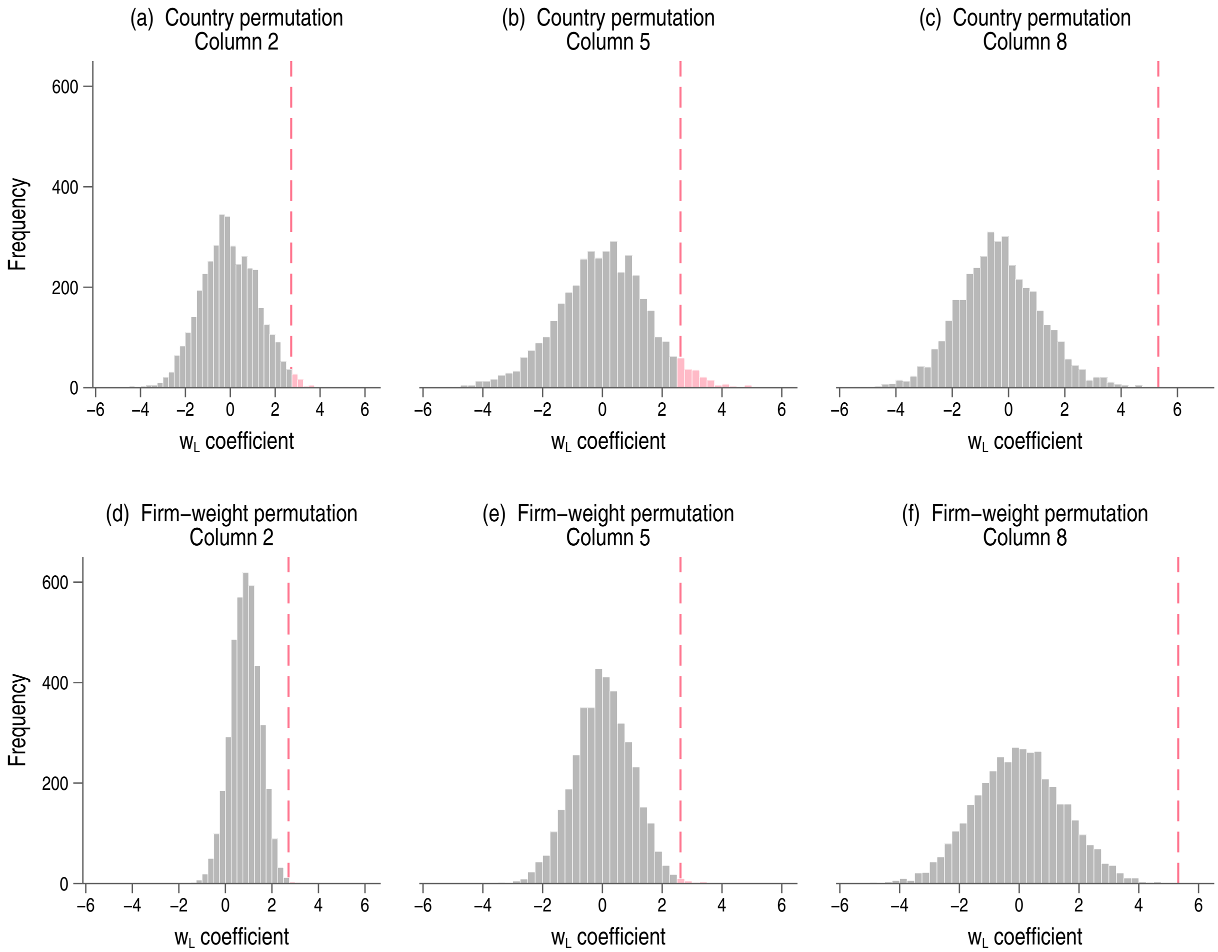

Figure A.3: Distribution of coefficients in Monte-Carlo simulations.

Notes: We run Monte-Carlo simulations where for each country, we sample with replacement the entire path of macroeconomics variables (wages, labor productivity and GDP gap) from the existing set of countries. We then re-run our regressions 4000 times. Panels a), b) and c) report histograms on the distribution of low-skill wage coefficients. The vertical red lines correspond to the coefficients of the true regressions. We then carry a symmetric exercise, where for each firm, we sample with replacement the set of country-weights from the existing set of firms within the same country. We re-run our regressions 4000 times and panels d), e) and f) report histograms on the distribution of low-skill wage coefficients. Each panel corresponds to a different column in Table 8.

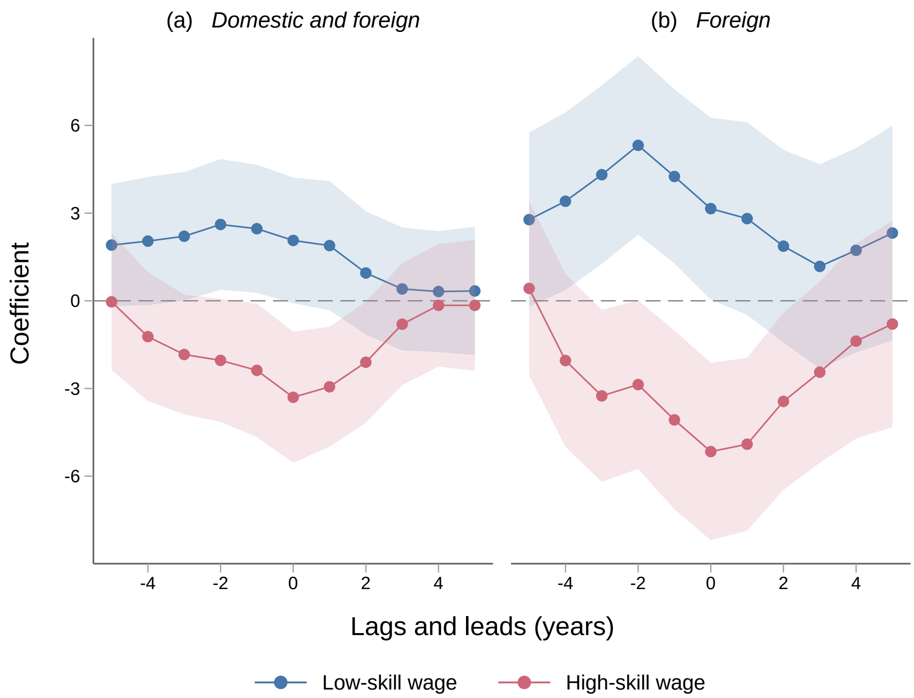

Figure A.4: Lag and leads.

Notes: This figure reports regression coefficients on low-skill and high-skill wages at different lags and leads. Each panel and each year corresponds to a different Poisson regression of auto95 innovations on wages, GDP gap, labor productivity, stocks, spillovers, firm fixed effects, industry-year fixed effects, and country-year fixed effects. Explanatory variables are computed at year t + the year marked on the x-axis except the stocks for which we keep the same lag of 2 years throughout. Panel a consider the total macroeconomic variables while Panel b looks at the normalized foreign variables previously defined. The shaded area represent 95% confidence interval, standard errors are clustered at the firm level. Panel a, year -2 corresponds to Column 5 of our baseline Table 5, and Panel b, year -2 corresponds to Column 8. The leads test for the presence of pre-trends.

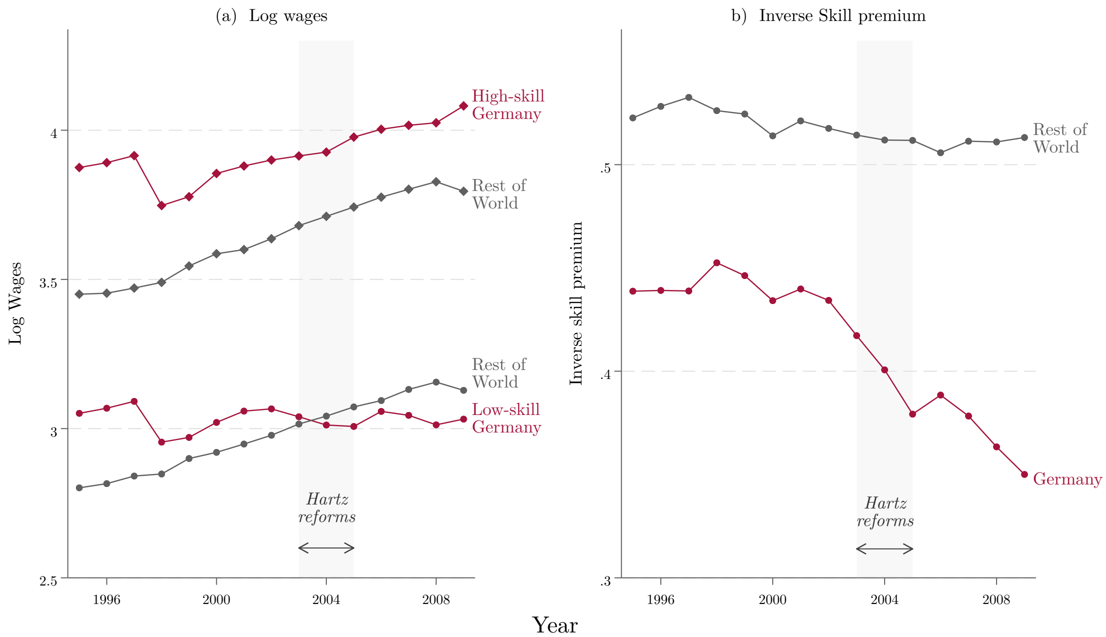

Figure A.5: Effect of the Hartz reforms on labor costs and the inverse skill premium.
Notes: Panel a) shows log low-skill and high-skill labor costs (denoted wages) in Germany and in the rest of the world. Panel b) shows the inverse skill premium. The rest of the world series is computed as a weighted average using the weights (excluding Germany) of the firms included in the regression of Figure 5.a

Table A.1: Summary statistics on the industry level regressions

	Mean	SD	Min	P10	P50	P90	Max	N
Share automation (using industry)	0.075	0.013	0.042	0.059	0.079	0.088	0.111	133
Share automation (inventing industry)	0.081	0.060	0.011	0.027	0.076	0.166	0.382	126
Δ Computer use (1984-1997)	0.192	0.072	-0.159	0.104	0.187	0.280	0.412	133
Δ Routine cognitive	-2.493	4.216	-21.667	-8.286	-2.710	4.020	9.666	133
Δ Routine manual	-2.308	4.336	-23.283	-9.330	-1.435	3.073	12.516	133
Δ High/low skill workers	0.123	0.176	-0.105	-0.003	0.070	0.318	1.132	133
Δ Labor share (NBER manufacturing)	-0.093	0.063	-0.230	-0.179	-0.084	-0.040	0.035	56
Δ Labor share (BEA)	-0.046	0.121	-0.616	-0.191	-0.015	0.045	0.327	60

Notes: This table shows summary statistics for the variables in our industry level regression. Share automation (using industry) represents the share of automation patents among machinery patents used by an industry. Share automation (inventing industry) represents the share of automation patents among machinery patents invented by an industry. Patents are USPTO granted patents over the years 1980-1998. Δ Computer use is the change in computer per-employee between 1997 and 1984. Δ routine cognitive, routine manual and high/low skill workers denote changes in these variables between 1980-1998. Δ labor share (NBER manufacutring) is the change in payroll / value added in the NBER-CES manufacuring industry database. Δ Labor Share (BEA) is the change in total compensation / value added in 60 aggregated industries. Industries are weighed by mean industry employment in 1980 and 1998.

Table A.2: Industry of innovators

	Industry	Share auto95 (%)	Share firms (%)
20	Manufacture of chemicals and chemical products	2.13	3.41
25	Manufacture of fabricated metal products, except machinery and equipment	1.18	4.42
26	Manufacture of computer, electronic and optical products	23.20	7.62
27	Manufacture of electrical equipment	9.45	2.89
28	Manufacture of machinery and equipment n.e.c.	24.36	21.20
29	Manufacture of motor vehicles, trailers and semi-trailers	5.30	3.53
30	Manufacture of other transport equipment	4.57	1.17
46	Wholesale trade, except of motor vehicles and motorcycles	1.32	3.29
64	Financial service activities, except insurance and pension funding	1.69	0.98
72	Scientific research and development	2.04	2.37
	Other industries	12.99	26.82
	No information on industry	11.79	22.27

Notes: The table reports the industry of patenting firms included in our baseline regression with industry-year fixed effects at the NACEv2 division level, and the share of biadic auto95 families for each industry. Industries representing less than 1% of patents are summed up in the 'Other industries' category.

Table A.3: Coverage of the regression sample

	Applications	Families	Biadic Families	Firms
Patstat 1997-2011	432 095	179 954	61 699	–
Matched with Orbis	348 342	140 707	52 331	4251
Firms in sample	206 959	86 030	33 025	3255

Notes: This table reports the number of auto95 patent applications, families, biadic families, and firms (that do at least one auto95 biadic innovation) for the time period 1997-2011 for three different samples based on PATSTAT: the whole sample, the sample of firms observed in ORBIS and the sample of firms included in our baseline regression.

Table A.4: Descriptive statistics on innovation

(a) Top 10 auto95 innovators in our sample

Company	Auto95's in 1997-2011
Siemens Aktiengesellschaft	1781
Honda Motor Co., Ltd.	815
Fanuc Co.	779
Samsung Electronics Co., Ltd.	718
Robert Bosch Gmbh	673
Mitsubishi Electric Co.	669
Tokyo Electron Limited	583
Murata Machinery, Ltd.	502
Kabushiki Kaisha Toshiba	491
Panasonic I.P.M Co., Ltd.	460

Notes: This table reports the 10 firms with the highest number of biadic auto95 patents in our baseline sample.

(b) Summary statistics on auto95 and pauto90 innovation

Sample	Baseline		Restricted	
	Auto95		Auto95	Pauto90
	(1)	(2)	(3)	(4)
Number of patents	Yearly	1997-2011	1997-2011	1997-2011
Mean	1	12	13	83
SD	4	54	57	313
P50	0	2	2	15
P75	0	6	7	49
P90	2	20	24	166
P95	3	43	50	335
P99	14	194	200	1184
Average citations received in 5 years	9.4		9.2	7.6
Number of firms	3255		2859	

Notes: This table presents summary statistics for the firms' patenting activity. Columns 1 and 2 show statistics for the baseline regression sample. Columns 3 and 4 describe the restricted sample in which we include non-automation machinery (pauto90) patents. Average citations are calculated as the average number of citations received by a patent within 5 years after the application. The firms are the non-domestic firms that patent at least once before 1995 and during the sample period 1997-2011.

Table A.5: Summary statistics on the firm-level macro variables

	Low-skill wage	Middle-skill wage	High-skill wage	GDP gap	GDP per capita	Labor productivity
Low-skill wage	1.000					
Middle-skill wage	0.942	1.000				
High-skill wage	0.609	0.750	1.000			
GDP gap	−0.063	−0.051	−0.032	1.000		
GDP per capita	0.709	0.804	0.732	0.114	1.000	
Labor productivity	0.674	0.736	0.772	0.039	0.668	1.000
Standard deviation	0.032	0.029	0.034	0.004	0.026	0.026

Notes: This table shows the correlation of residuals for the auto95 baseline regression sample, controlling for firm and year-industry fixed effects. The last row shows the standard deviation of the residual variables.

Table A.6: Baseline regressions with fewer controls

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	3.42^{***}	2.65^{***}	3.00^{***}	2.74^{***}	2.65^{***}	2.25^{**}	4.69^{***}	4.20^{***}	4.20^{***}
	(0.75)	(0.76)	(0.79)	(0.98)	(0.76)	(1.00)	(1.32)	(1.31)	(1.33)
High-skill wage	-1.57^{**}	-1.51^{**}	-2.20^{***}	-2.72^{***}	-1.51^{**}	-2.81^{***}	-4.94^{***}	-4.50^{***}	-4.46^{***}
	(0.68)	(0.65)	(0.72)	(0.92)	(0.65)	(0.96)	(1.38)	(1.32)	(1.31)
Stock automation		-0.11^{***}	-0.12^{***}		-0.11^{***}	-0.12^{***}		-0.12^{***}	-0.13^{***}
		(0.03)	(0.03)		(0.03)	(0.03)		(0.03)	(0.03)
Stock other		0.52^{***}	0.51^{***}		0.52^{***}	0.52^{***}		0.50^{***}	0.51^{***}
		(0.04)	(0.04)		(0.04)	(0.04)		(0.04)	(0.04)
Spillovers automation			0.58^{*}			1.35^{***}			1.33^{***}
			(0.29)			(0.47)			(0.46)
Spillovers other			-0.19			-0.97^{***}			-0.98^{***}
			(0.22)			(0.36)			(0.35)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	48 091	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3255	3252	3252	3252	3252

Notes: This table shows our baseline regressions with fewer controls. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Significance levels at *10%, **5%, ***1%.

Table A.7: Auto90 innovations

Dependent variable	Auto90								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.35^{***}	2.07^{***}	3.31^{***}	1.71^{**}	1.73^{*}	2.82^{***}	3.30^{***}	3.88^{***}	3.90^{***}
	(0.66)	(0.68)	(0.79)	(0.82)	(0.89)	(1.06)	(1.13)	(1.32)	(1.45)
High-skill wage	-1.96^{***}	-2.46^{***}	-0.92	-1.80^{**}	-1.75^{*}	-1.06	-3.80^{***}	-2.95^{**}	-3.45^{***}
	(0.60)	(0.65)	(0.66)	(0.81)	(0.92)	(0.86)	(1.17)	(1.30)	(1.23)
GDP gap	-3.61^{*}	-4.29^{**}	-1.20	3.77	3.84	5.66	-0.30	0.93	0.87
	(2.09)	(2.14)	(2.24)	(5.25)	(5.33)	(5.43)	(3.26)	(3.52)	(3.69)
Labor productivity		1.15			-0.13			-1.35	
		(0.73)			(1.30)			(1.33)	
GDP per capita			-2.72^{**}			-2.72^{*}			-1.07
			(1.06)			(1.49)			(1.56)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	72 196	72 196	72 196	71 905	71 905	71 905	71 905	71 905	71 905
Number of firms	4857	4857	4857	4854	4854	4854	4854	4854	4854

Notes: This table shows our baseline regression using a weaker measure of automation (auto90). All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Stock and spillover variables are calculated with respect to the dependent variable (auto90). Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.8: Effect of the inverse skill premium on auto95 innovations

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill / High-skill wages	2.51***	2.67***	2.51***	2.53***	2.39***	2.63***	4.38***	4.20***	4.36***
	(0.69)	(0.69)	(0.69)	(0.88)	(0.88)	(0.88)	(1.27)	(1.24)	(1.26)
GDP gap	−4.06	−4.35*	−4.08	4.67	5.03	5.39	−0.17	0.49	0.25
	(2.58)	(2.61)	(2.60)	(6.80)	(6.75)	(6.86)	(4.61)	(4.64)	(4.69)
Labor productivity		1.03			−1.16			−0.58	
		(0.64)			(1.10)			(0.73)	
GDP per capita			0.03			−1.62			−0.33
			(0.71)			(1.13)			(0.89)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table shows the effect of the skillpremium on automation innovations. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. Columns 7–9 compute the normalized foreign (log) inverse skill premium as the difference between the normalized (log) foreign low-skill wages and the normalized (log) foreign high-skill wages as defined in the text. In these columns, GDP gap, GDP per capita and labor productivity also correspond to their normalized foreign values. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.9: Additional regressions with non-automation patents

Dependent variable	Pauto90 refined			Pauto90			Pauto95		
	Dom. and Fgn.		*Fgn.*	*Dom. and Fgn.*		*Fgn.*	*Dom. and Fgn.*		*Fgn.*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	0.97	1.17	2.05	0.72	0.32	0.98	0.97	0.48	1.56
	(0.90)	(1.21)	(1.70)	(0.59)	(0.77)	(1.21)	(0.75)	(0.99)	(1.61)
High-skill wage	−1.18	0.82	1.33	−0.20	−0.35	−0.61	−0.43	−0.45	−0.76
	(0.85)	(1.32)	(1.93)	(0.56)	(0.86)	(1.26)	(0.73)	(1.17)	(1.71)
GDP gap	−3.05	0.39	−2.47	−3.03**	1.35	0.36	−2.04	3.62	0.86
	(2.13)	(4.35)	(3.31)	(1.35)	(3.39)	(2.34)	(1.57)	(4.13)	(2.85)
Labor productivity	1.42*	−1.62	−3.01*	−0.12	0.03	−0.88	−0.13	−0.59	−1.15
	(0.80)	(1.51)	(1.67)	(0.60)	(0.96)	(1.01)	(0.70)	(1.21)	(1.34)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Observations	35 678	35 500	35 500	149 250	149 015	149 015	44 094	43 971	43 971
Number of firms	2399	2397	2397	9990	9987	9987	2951	2948	2948

Notes: This table presents additional regressions using non-automation innovations. In columns 1–3 the dependent variable is refined pauto90 (non-auto90 machinery patents that list at least one 4-digit C/IPC code containing a 6-digit code classified auto95), and the sample is restricted to the firms in the baseline auto95 regressions. In columns 4–6 the dependent variable is pauto90 (machinery patents excluding auto90) but the sample is unrestricted. In columns 7–9 the dependent variable is pauto95 (machinery patents excluding auto95), and the sample is again restricted to the firms in the baseline auto95 regression. All columns include firm and industry-year fixed effects, Columns 2, 3, 5, 6, 8 and 9 add country-year fixed effects. In Columns 3, 6, and 9 the macroeconomic variables are the normalized foreign variables as defined in the text. Stocks and spillovers are defined in terms of the respective dependent variable. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.10: Wages and the direction of innovation

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.12***	2.13***	2.01**	2.39**	2.40**	2.59**	4.81***	4.80***	5.06***
	(0.73)	(0.73)	(0.80)	(0.98)	(0.98)	(1.04)	(1.34)	(1.34)	(1.48)
High-skill wage	−2.13***	−2.14***	−1.94***	−2.11**	−2.12**	−2.23**	−2.89**	−2.92**	−3.04**
	(0.66)	(0.66)	(0.72)	(0.97)	(0.98)	(1.08)	(1.32)	(1.33)	(1.52)
GDP gap	−2.37	−2.37	−2.25	2.48	2.68	−0.81	3.78	3.69	5.01
	(2.25)	(2.27)	(2.29)	(5.54)	(5.59)	(5.30)	(4.20)	(4.20)	(5.20)
Labor productivity	0.89	0.89	0.80	−1.44	−1.41	−1.42	−1.99	−1.94	−1.95
	(0.84)	(0.84)	(0.93)	(1.61)	(1.61)	(1.71)	(1.40)	(1.41)	(1.56)
Arcsinh pauto90	0.51***			0.51***			0.51***		
	(0.02)			(0.02)			(0.02)		
Log pauto90		0.48***	1.00		0.48***	1.00		0.48***	1.00
		(0.02)			(0.02)			(0.02)	
Any pauto90		0.42***	0.09*		0.43***	0.10*		0.42***	0.09*
		(0.05)	(0.05)		(0.05)	(0.05)		(0.05)	(0.05)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table shows regressions with a control for non-automation machinery innovations (pauto90). Columns 1, 4, and 7 control for the arcsinh of pauto90 patent flow. Columns 2, 5, and 8 control for log pauto90 and a dummy variable indicating at least 1 pauto90 innovation. Columns 3, 6, and 9 constrain the coefficient on log pauto90 to 1. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.11: Predicting weights using subsequent wages

	Weight			Foreign weight		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in low-skill wages, 1995-2000	−0.14	−0.26	−0.13	−0.10	−0.31	−0.33
	(0.12)	(0.28)	(0.29)	(0.11)	(0.26)	(0.30)
Growth in high-skill wages, 1995-2000		0.13	0.01		0.20	0.23
		(0.24)	(0.27)		(0.21)	(0.24)
Patent weighted	No	No	Yes	No	No	Yes
Observations	133 455	133 455	133 455	130 200	130 200	130 200
Firms	3255	3255	3255	3255	3255	3255

Notes: This table shows OLS regressions of firm-level weights on country growth rates for low-skill and high-skill wages between 1995 and 2000. Columns 3 and 6 weigh observations by the number of auto95 patents between 1997 and 2011. In columns 4–6, the dependent variable is the the foreign weight component only. Standard errors are clustered at the country-level. Significance levels at *10%, **5%, ***1%.

Table A.12: Alternative weights

Dependent variable	Auto95							
Weight robustness	Pauto95		1971–1989		1985–1994		start 2000	
	Dom. and fgn. (1)	*Fgn.* (2)	*Dom. and fgn.* (3)	*Fgn.* (4)	*Dom. and fgn.* (5)	*Fgn.* (6)	*Dom. and fgn.* (7)	*Fgn.* (8)
Low-skill wage	2.60**	2.19*	2.86**	3.44**	5.52***	5.32***	5.15***	6.69***
	(1.17)	(1.28)	(1.16)	(1.46)	(1.71)	(1.91)	(1.51)	(2.11)
High-skill wage	−2.07*	−2.39**	−1.20	−2.03	−2.86*	−3.39**	−1.37	−2.99
	(1.07)	(1.21)	(1.14)	(1.62)	(1.65)	(1.69)	(1.54)	(2.04)
GDP gap	−3.06	3.80	3.92	−0.84	7.15*	0.94	3.32	0.55
	(5.70)	(6.72)	(6.62)	(6.76)	(4.08)	(4.14)	(4.82)	(3.90)
Labor productivity	−0.36	0.20	−3.13*	0.45	−2.65*	−2.46	−3.90**	−4.84***
	(1.63)	(1.89)	(1.77)	(2.16)	(1.54)	(1.78)	(1.61)	(1.77)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44 936	33 959	43 548	26 020	44 936	33 959	43 548	26 020
Number of firms	3075	2333	2968	2640	3075	2333	2968	2640

Notes: This table uses alternative weights to compute firm's macroeconomic variables. In Columns 1–2 the firm's country weights are calculated using pauto95 patents (machinery patents excluding auto95). Columns 2–4 compute the weights over the period 1971–1989 and Columns 5–6 over the period 1985–1994. Columns 7–8 use the baseline pre-sample period of 1971–1994 to compute weights but restrict the regression sample to the years 2000–2009. In columns 2, 4, 6, and 8 the macroeconmic variables are the normalized foreign variables as described in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.13: Horse-race regressions between 2 year lags and other lags / leads

Dependent variable	Auto95								
Macrovars lag j	-6	-5	-4	-3	-2	-1	0	1	2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A. Domestic and Foreign									
Low-skill wage (L2)	3.37***	2.68**	2.53**	2.51**	2.61**	8.14***	4.25**	3.37**	2.06
	(1.21)	(1.29)	(1.25)	(1.27)	(1.14)	(2.23)	(1.91)	(1.68)	(1.72)
Low-skill wage Lj	−0.62	−0.01	0.07	−0.05		−4.83**	−1.03	−0.02	−0.21
	(0.45)	(0.52)	(0.52)	(0.65)		(1.99)	(1.65)	(1.60)	(1.70)
High-skill wage (L2)	−3.19***	−1.94*	−2.43**	−2.97**	−2.04*	−1.99	−0.86	−2.03	−1.52
	(1.18)	(1.08)	(1.15)	(1.38)	(1.07)	(1.95)	(1.76)	(1.83)	(1.85)
High-skill wage Lj	0.56	−0.67	0.61	0.50		−0.49	−2.13*	−1.16	0.02
	(1.07)	(1.26)	(1.28)	(1.54)		(1.56)	(1.28)	(1.28)	(1.61)
Observations	47 741	47 741	47 741	47 741	47 741	43 160	38 835	34 749	30 816
Number of firms	3252	3252	3252	3252	3252	3148	3050	2958	2862
Panel B. Foreign									
Low-skill wage (L2)	5.69***	5.19***	5.13***	4.67***	5.32***	10.62***	8.19***	7.02***	5.00**
	(1.60)	(1.65)	(1.60)	(1.69)	(1.56)	(2.84)	(2.50)	(2.44)	(2.42)
Low-skill wage Lj	−0.47	0.15	0.16	0.57		−4.92**	−2.44	−1.37	−0.38
	(0.57)	(0.61)	(0.64)	(0.91)		(2.37)	(2.17)	(2.32)	(2.47)
High-skill wage (L2)	−3.31**	−2.20	−2.34	−1.53	−2.87*	−2.62	−1.53	−2.76	−3.11
	(1.65)	(1.66)	(1.73)	(1.83)	(1.47)	(2.07)	(1.89)	(1.94)	(2.30)
High-skill wage Lj	0.45	−1.24	−0.67	−1.76		−1.44	−3.61**	−2.70	−0.20
	(1.20)	(1.49)	(1.50)	(1.76)		(1.93)	(1.74)	(1.69)	(2.02)
Observations	47 341	47 429	47 539	47 651	47 741	43 160	38 835	34 749	30 816
Number of firms	3241	3243	3246	3250	3252	3148	3050	2958	2862
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table runs a horserace regressions between different lag of wages. The wages variables are included twice: lagged by two periods (as in the baseline) and shifted as indicated by lag j in the header. All columns include controls for labor productivity and the business cycle, firm and industry-year fixed effects, and country-year fixed effects. In Panel B, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.14: Predicted wages

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.45^{***}	1.84^{**}	2.47^{***}	1.66^{*}	1.55	1.66^{*}	3.83^{***}	4.24^{***}	3.82^{***}
	(0.81)	(0.82)	(0.82)	(0.93)	(1.02)	(0.93)	(1.30)	(1.40)	(1.30)
High-skill wage	-2.78^{***}	-4.75^{***}	-2.82^{***}	-3.29^{***}	-3.59^{**}	-3.30^{***}	-4.51^{***}	-3.57^{**}	-4.50^{***}
	(0.82)	(1.07)	(0.83)	(1.03)	(1.40)	(1.03)	(1.33)	(1.52)	(1.34)
GDP gap	-4.34^{*}	-3.71	-4.39^{*}	4.58	4.57	4.59	-0.28	0.56	-0.24
	(2.60)	(2.56)	(2.60)	(6.81)	(6.82)	(6.81)	(4.55)	(4.59)	(4.59)
Labor productivity		2.86^{***}			0.45			-1.55	
		(0.94)			(1.56)			(1.49)	
GDP per capita			0.14			0.02			-0.02
			(0.10)			(0.12)			(0.14)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table uses predicted wages as main RHS variables. We estimate for each country an AR(1) process with time trends for wages, labor productivity, and GDP per capita. We then use the estimated process to predict with the information available at time t-2 the average values between the years t+2 and t+7, which are in turn the independent variables in these regressions. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Table A.15: Addressing Nickell's bias

	Auto95					
	Domestic and foreign				*Foreign*	
	(1)	(2)	(3)	(4)	(5)	(6)
Low-skill wage	2.67^{***}	2.26^{***}	2.70^{**}	2.60^{**}	4.83^{***}	3.91^{***}
	(0.79)	(0.78)	(1.06)	(1.02)	(1.45)	(1.39)
High-skill wage	-2.54^{***}	-1.14	-2.22^{**}	-1.74^{*}	-2.72^{*}	-2.16
	(0.77)	(0.79)	(1.01)	(1.00)	(1.40)	(1.46)
GDP gap	-4.26	-2.93	4.83	6.15	1.67	0.70
	(2.76)	(3.45)	(7.06)	(7.31)	(4.98)	(5.24)
Labor productivity	0.84	0.47	-1.46	-1.14	-1.96	-0.94
	(0.90)	(0.98)	(1.68)	(1.44)	(1.50)	(1.49)
Stock automation	No	Yes	No	Yes	No	Yes
Stock other	Yes	Yes	Yes	Yes	Yes	Yes
Spillovers	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	Yes	Yes	Yes	Yes
Estimator	HHG	BGVR	HHG	BGVR	HHG	BGVR
Observations	48 091	48 091	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3252	3252	3252	3252

Notes: This table addresses potential Nickell's bias. The coefficients are estimated with conditional Poisson regressions fixed-effects (HHG) in columns 1, 3, and 5. In columns 2, 4, and 6, the coefficients are estimated with Poisson regressions where the firm fixed effects are replaced by the pre-sample mean, following Blundell, Griffith and Van Reenen (1999, BGVR). All columns include firm and industry-year fixed effects. Columns 3–6 add country-year fixed effects. In Columns 5 and 6 the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

A.2 Appendix on the classification of automation patents

This Appendix provides additional information on our classification of automation patents in machinery. First, we report details on our approach not contained in the main text in Appendix A.2.1. Then, we show additional statistics at the technology category level in Appendix A.2.2 and at the patent level in Appendix A.2.3. Appendix A.2.4 shows that our classification is stable. Finally, Appendix A.2.5 gives the prevalence of automation keywords for a few technology categories and examples of automation patents.

A.2.1 Additional details on our classification

We derived the exact list of keywords in Table 1 after experimenting extensively with variations around them and looking at the resulting classification of technology categories and the associated patents. Relative to the original list of technologies given in the Survey of Manufacturing Technologies (Doms, Dunne and Troske, 1997), we did not include keywords related to information network, as these seem less related to the automation of the production process and the patents containing words such as "local area network" do not appear related to automation. We also did not count all laser patents as they are not all related to automation—but we obtain patents related to automation using laser technologies thanks to our other keywords. Furthermore, the Y section of the CPC classification is organized differently from the rest and is only designed to provide additional information. As a result, we ignore Y codes.

As mentioned in the text, we focus on the technology fields: "machine tools", "handling", "textile and paper machines", and "other special machines" with a few adjustments. First, we exclude F41 and F42, which correspond to weapons and ammunition and are in "other special machines". Moreover, we include B42C which corresponds to machines for book production and B07C which corresponds to machines for postal sorting as both correspond to equipment technologies and contain 6-digit codes with a high prevalence of automation keywords. We further include the 6-digit codes G05B19 and G05B2219, which correspond to "programme-control systems" and contain many computer numerically controlled machine tool patents without C/IPC codes from the machine tools technology field. Finally, we include the 6-digit code B62D65 which deals with engine manufacturing (though the rest of the B62D code deals with the vehicle parts themselves). We verify that these additional codes do not affect our results.

A.2.2 Statistics on the classification at the technology category level

Table A.16: Summary statistics on the prevalence of keywords

	C/IPC6					C/IPC4 + (G05 or G06)					C/IPC4 pairs				
	All	Robot	Automat*	CNC	Labor	All	Robot	Automat*	CNC	Labor	All	Robot	Automat*	CNC	Labor
Mean	0.21	0.04	0.11	0.02	0.06	0.53	0.15	0.32	0.11	0.09	0.18	0.04	0.09	0.02	0.02
SD	0.14	0.08	0.09	0.06	0.04	0.19	0.18	0.11	0.17	0.04	0.16	0.10	0.10	0.05	0.05
25th	0.10	0.01	0.04	0.00	0.03	0.40	0.07	0.27	0.01	0.07	0.08	0.01	0.02	0.00	0.00
50th	0.18	0.02	0.09	0.00	0.05	0.54	0.10	0.32	0.03	0.10	0.14	0.02	0.05	0.00	0.00
75th	0.27	0.05	0.15	0.02	0.08	0.64	0.16	0.40	0.15	0.11	0.23	0.04	0.11	0.01	0.01
90th	0.39	0.09	0.24	0.06	0.10	0.78	0.36	0.43	0.38	0.15	0.37	0.09	0.22	0.04	0.04
95th	0.48	0.14	0.29	0.13	0.13	0.86	0.44	0.45	0.55	0.16	0.52	0.15	0.31	0.08	0.08
995th	0.75	0.36	0.44	0.33	0.18	0.90	0.83	0.60	0.57	0.18	0.84	0.59	0.45	0.22	0.22

Notes: This table computes summary statistics on the share of patents with any automation keywords, robot keywords, automat* keywords, CNC keywords or labor keywords for each type of technological categories (6-digit C/IPC codes, pairs of 4-digit C/IPC codes and combinations of 4-digit C/IPC codes with G05 or G06) within machinery with at least 100 patents.

Table A.16 gives summary statistics on the prevalence of automation keywords across technology categories in machinery, $p(t)$, and the prevalence of the 4 main subgroups of keywords: automat*, robot, numerical control (CNC) and labor. The 95^{th} and 90^{th} percentile for the prevalence of automation keywords for 6-digit codes in machinery define the thresholds used to categorize auto95 and auto90 patents. The distributions are quite similar for the C/IPC 6-digit codes and for pairs of IPC 4-digit codes and shifted to the right for combinations of C/IPC 4-digit codes with G05/G06 (see also the histograms below). All prevalence measures are right-skewed, particularly for 6-digit codes and 4-digit pairs, and even more for the robot and CNC patents. The automat* keywords are more frequently used than the other keywords but the difference narrows in the right tail: the 95^{th} percentile for 6-digit codes is 29% for automat* and 14% and 13% for robot and CNC. In fact, we chose the thresholds (5 and 2) used in the definition of the automat* keywords so that the distributions of the prevalence measures are somewhat comparable. The right tails of the distribution are similar for the prevalence of the robot and CNC keywords.

Table A.17: Correlation between the main prevalence measures

Keywords	Automat	Robot	CNC	Labor
Automat	1.000			
Robot	0.383	1.000		
CNC	0.215	0.206	1.000	
Labor	0.391	0.225	0.090	1.000

Notes: This table shows the correlation between the prevalence of the main keywords, computed for C/IPC 6-digit codes.

Table A.17 shows the correlation between the prevalence of the 4 mains keyword

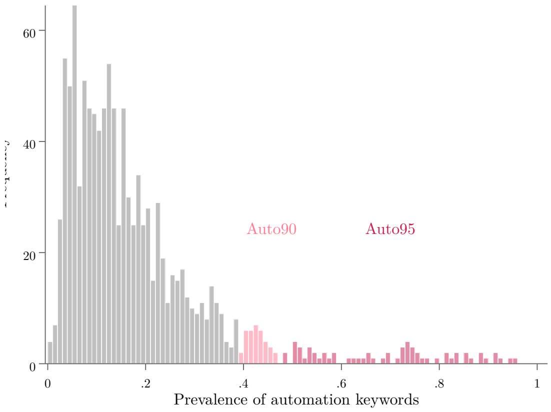

(a) For all pairs of C/IPC 4-digit codes within machinery

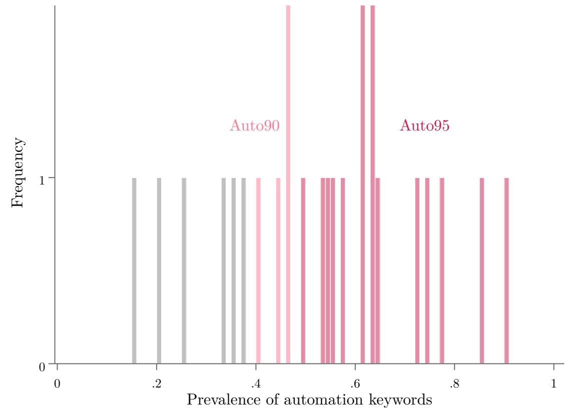

(b) For combinations of IPC4 in machinery with G05 G06

Figure A.6: Histograms of the prevalence of automation keywords.
Notes: We only include technology categories with at least 100 patents. The p90 and p95 lines, based on the 6-digit distribution, mark the thresholds used to define auto90 and auto95 technology categories.

categories (automat*, robot, CNC and labour) for 6-digit C/IPC codes. These measures are positively correlated with a coefficient above 0.2 in all cases except CNC and labour. The broadest category, automat*, is the one with the highest correlation coefficients.

Figure A.6.a gives the histograms of the prevalence of automation keywords for machinery technology categories which are pairs of C/IPC 4-digit codes. The histograms are very similar to those of C/IPC 6-digit codes in Figure 1. Figure A.6.b shows the histograms for all combinations of machinery C/IPC 4-digit codes with G05 or G06. The distribution is considerably shifted to the right. This is in line with expectations as G05 proxies for control and G06 for algorithmic, two set of technologies which have been used heavily in automation. There are, however, many fewer combination of these types, and accordingly fewer patents can be characterized as automation innovations this way. Overall, we classify 51 6-digit codes, 15 combination of 4-digit codes with GO5/GO6 and 63 pairs of 4-digit codes as auto95.

A.2.3 How are auto90 and auto95 patents identified?

Given that our classification procedure is relatively complex, we assess here which features dominate. To do so, we focus on biadic patent families in 1997-2011, the set of innovations which we use for our main regressions. There are 61,699 auto95 biadic patent families and 106,538 auto90 ones. Table A.18.a gives the share of biadic patents which are identified through a C/IPC 6-digit code, a pair of 4-digit codes or a combination of

Table A.18: Identification of automation technology categories

(a) Type of C/IPC codes identifying auto90 and auto95 patents

IPC codes / Patents	Auto90	Auto95
Matches C/IPC6	82.1%	83.4%
Matches C/IPC4 pair	41.3%	41.8%
Matches C/IPC4 - G05/G06 combination	16.1%	22.7%

Notes: This table shows the share of innovations classified as automation innovation through 6-digit C/IPC codes, 4-digit C/IPC pairs or 4-digit C/IPC - G05/G06 pairs. The statistics are computed on biadic patents from 1997-2011.

(b) Auto patents and subcategories of automation innovations

Sources / Patents	Auto80	Auto90	Auto95
Auto80	100.0%	100.0%	100.0%
Automat*80	36.4%	54.1%	72.2%
CNC80	5.0%	8.3%	13.2%
Robot80	12.1%	20.0%	34.3%
Auto90	60.7%	100.0%	100.0%
Automat*90	22.2%	36.5%	58.1%
CNC90	2.1%	3.4%	5.8%
Robot90	7.8%	12.8%	22.1%
Auto95	35.2%	57.9%	100.0%
Automat*95	3.3%	5.4%	9.3%
CNC95	1.5%	2.5%	4.4%
Robot95	6.5%	10.8%	18.6%

Notes: This table shows the share of innovations classified as automation innovation through 6-digit C/IPC codes, 4-digit C/IPC pairs or 4-digit C/IPC - G05/G06 pairs. The statistics are computed on biadic patents from 1997-2011.

4-digit code with G05/G06 (the shares sum up to more than 100% since patents may be identified as automation innovations in several ways). 6-digit codes are the most relevant since they identify more than 80% of either auto90 or auto95 patents alone.

Similarly, one may wonder which keywords are the most important in identifying automation patents. To assess that, we define robot95 patents as patents which contain a technology category with a prevalence of "robot" keywords above the threshold used to define auto95 (namely 0.480). Therefore, those patents are a subset of the auto95 patents. We define CNC85, automat*95, robot90, CNC90, automat*90, robot80, CNC80 and automat*80 similarly. The other keywords are much less common. Table A.18.b reports the share of auto95, auto90 and auto80 patents which belong to each subcategory. "Automat*" is the most important keyword: 72% of auto95 patents are also automat*80 patents. "Robot" matters as well with 34% of auto95 patents which are robot80 and 19% which are even robot95 (more than automat*95). CNC does not matter much: only 13% of auto95 patents are CNC80.

A.2.4 Stability of the classification

To assess the stability of our classification, we redo exactly the same exercise but instead of using EPO patents from 1978 to 2017, we restrict attention to EPO patents from the first half of the sample (1978-1997), the second half (1998-2017) or the period of our

Table A.19: Correlation between the prevalence of automation keywords for different periods

Prevalence of automation keywords by period				
Keywords	1978-2017	1997-2011	1978-1997	1998-2017
1978-2017	1.000			
1997-2011	0.960	1.000		
1978-1997	0.913	0.858	1.000	
1998-2017	0.973	0.981	0.849	1.000

Notes: Correlation between the prevalence of the main keywords, computed for C/IPC 6-digit codes.

Table A.20: Confusion table for different classification periods

Classification periods		First half 1978-1997		Second half 1998-2017		Regression period 1997-2011		Total
		Yes	No	Yes	No	Yes	No	
Baseline 1978-2017	Yes	51 812	9887	55 820	5879	54 021	7678	61 699
	No	7698	3 118 139	5041	3 120 796	5550	3 120 287	3 125 837
	Total	59 510	3 128 026	60 861	3 126 675	59 571	3 127 965	3 187 536

Notes: This table classifies all biadic patent families from 1997-2011 as auto95 or not using EPO patents from different time periods. Our baseline measure uses all patents from 1978-2017, while the other measures use patents from the first half of the sample, the second half, or the regression period time.

main regression analysis (1997-2011). There is a modest increase in the share of patents with automation keywords within each technology category. The share of patents with an automation keyword increases on average from 0.191 in the first half of the sample to 0.216 in the second half. Nevertheless, the ranking of codes is remarkably stable as shown in Table A.19 which reports the correlations of the prevalence measures for the different time periods.

Further, focusing on the same set of biadic patent families in 1997-2011, Table A.20 shows confusion tables on the classification of patents as auto95 according to each of the classification period. Regardless of the time period used, the set of automation patents stays roughly the same. In particular, 87.6% of the baseline auto95 patents are still auto95 if we run the classification over the years 1997-2011. This common set of patents then represent 90.7% of all biadic patents classified as auto95 patents when using the period 1997-2011 instead of the full sample.

A.2.5 Examples

To better illustrate our approach, we now give a few examples. First, Table A.21 shows a few 6-digit C/IPC codes in machinery with their prevalence of automation keywords $p(t)$, their rank according to that measure and the prevalence of the most important sub-

Table A.21: Examples of 6-digit C/IPC codes in machinery

Code	Description	# Patents	Any	Rank	Robot	Automat*	CNC	Labor
High Prevalence Codes								
B25J5	Manipulators mounted on wheels or on carriages	504	0.91	1	0.87	0.27	0.01	0.10
B25J9	Programme-controlled manipulators	2809	0.86	4	0.78	0.29	0.29	0.08
B23Q15	Automatic control or regulation of feed movement, cutting velocity or position of tool or work	591	0.79	7	0.09	0.36	0.36	0.06
A01J7	Accessories for milking machines or devices	395	0.77	9	0.62	0.52	0.52	0.10
G05B19	Programme-control systems	7133	0.70	17	0.22	0.39	0.39	0.08
B65G1	Storing articles, individually or in orderly arrangement, in warehouses or magazines	1064	0.58	30	0.18	0.46	0.46	0.11
Low Prevalence Codes								
B23P6	Restoring or reconditioning objects	613	0.26	267	0.07	0.06	0.05	0.09
A01B63	Lifting or adjusting devices or arrangements for agricultural machines or implements	264	0.24	307	0.01	0.20	0.00	0.04
B66D3	Portable or mobile lifting or hauling appliances	215	0.13	678	0.02	0.07	0.00	0.06

Notes: This table reports the prevalence of automation keywords for examples of 6-digit C/IPC codes. 'Any' is the share of patents with any of the keywords. 'Rank' is the rank of the code among 986 6-digit C/IPC codes in machinery with at least 100 patents. 'Robot' , 'Automat*', 'CNC' and 'labor' are the shares of patents with at least one keyword from these categories.

categories (automat*, robots, CNC, and labor). C/IPC codes associated with robotics (B25J) have the highest prevalence numbers (91% for B25J5). There are also codes associated with machine tools at the top of the distribution such as B23Q15 and codes associated with devices used in the agricultural sector such as A01J7. The last three C/IPC codes are examples with a low prevalence of automation keywords: machine-tools and processes for repairing or reconditioning objects (B23P6), devices typically mounted on tractors (A01B63), and lifting or hauling appliances such as hoists (B66D3), which do not replace workers in new tasks. The table also shows that the different sub-measures do not capture the same technologies: the robotic codes are ranked highly thanks to the prevalence of "robot" keyword, B23Q15 thanks to its CNC prevalence, and B65G1 thanks to its "automat*" prevalence.

Figure A.7 shows an automated storage cabinet patent. We classify it as automation because it contains the 6-digit code B65G 1 which has a high prevalence measure (0.58, see Table A.21). This patent itself contains several keywords: a sentence with the words "automatic" and "storing," and another sentence with "robot". Figure A.8 shows an automation patent of a similar storage cabinet that belongs to the same C/IPC code but does not contain any keywords and still describes a labor-saving innovation. Appendix B.1 provides more examples.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) **EP 2 604 550 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
01.10.2014 Bulletin 2014/40

(21) Application number: **10855839.6**

(22) Date of filing: **12.08.2010**

(51) Int Cl.:
B65G 1/137 *(2006.01)* ***B65G 1/04*** *(2006.01)*
B66F 9/07 *(2006.01)* ***A47B 96/02*** *(2006.01)*
B65G 1/08 *(2006.01)*

(86) International application number:
PCT/ES2010/070549

(87) International publication number:
WO 2012/020149 (16.02.2012 Gazette 2012/07)

(54) **AUTOMATIC PLANT FOR STORING AND DISPENSING GOODS**
AUTOMATISCHE ANLAGE ZUR AUFBEWAHRUNG UND AUSGABE VON WAREN
INSTALLATION AUTOMATIQUE POUR STOCKER ET DISTRIBUER DES PRODUITS

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO SE SI SK SM TR

(43) Date of publication of application:
19.06.2013 Bulletin 2013/25

(73) Proprietor: **Automatismos Y Montajes Industriales J. Martin, S.L.**
47012 Valladolid (ES)

(72) Inventors:
- **MARTÍN DE PABLO, Francisco Javier**
  E-47012 Valladolid (ES)
- **GONZÁLEZ LÓPEZ, Isabel**
  E-47012 Valladolid (ES)

(74) Representative: **Ungria López, Javier**
c/o UNGRIA Patentes y Marcas, S.A.,
Avda. Ramon y Cajal, 78
28043 Madrid (ES)

(56) References cited:

EP-A1- 2 113 473	**CH-A5- 680 434**
DE-A1- 4 336 885	**DE-A1- 4 339 055**
DE-A1- 19 635 396	**DE-A1- 19 724 378**
DE-U1- 20 021 440	**US-A- 3 782 565**
US-A1- 2010 168 910	

Description

OBJECT OF THE INVENTION

[0001] The present invention, as expressed in the wording of this specification, relates to an automatic plant for storing and dispensing goods, essentially applicable to the pharmaceutical sector, although it is also applicable to any other sector needing to store and dispense different small-sized goods.

[0002] The products are stored in principle in modular shelves, which may be inclined or not, shelves that are part of characteristic modular shelving units that also configure an elongated shelving structure in the longitudinal direction.

[0003] Based on this premise, the essence of the invention is based on characteristic modular horizontal guides along which respective modular subsets (robots) move, for the loading and unloading of products with respect to the shelves of the modular shelving units, modular horizontal guides that can easily adapt to the required length of the elongated structure of shelving units, so that both loading and unloading subsets have a horizontal translation movement parallel to said elongate structure of shelving units and a vertical movement to access the different levels of the shelves where the products are stored.

Figure A.7: Example of an automation patent

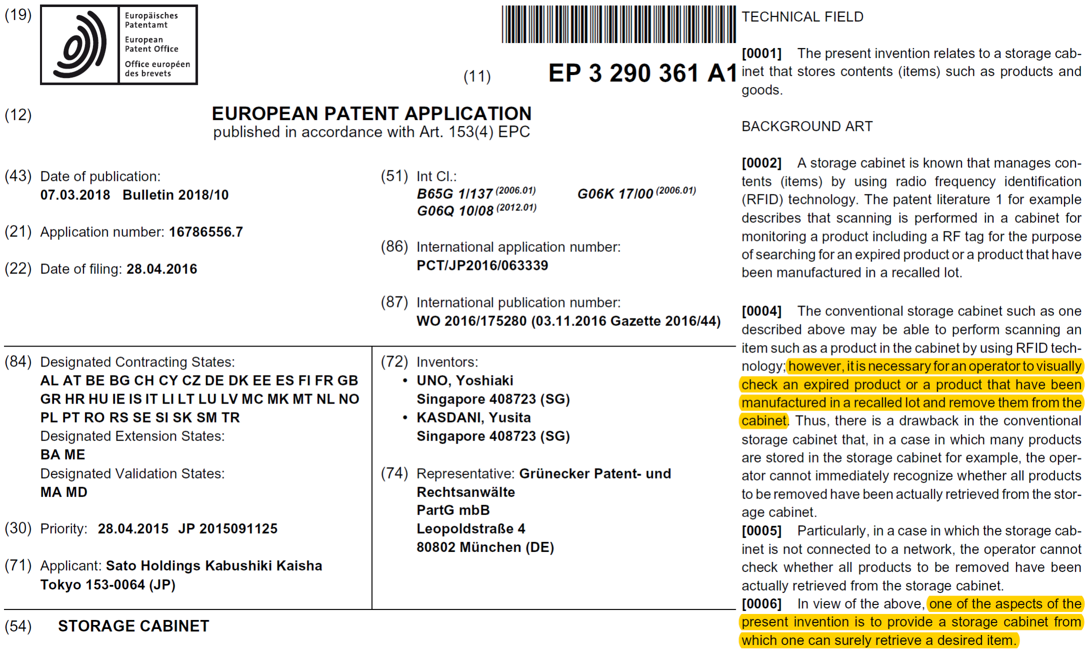

(19) Europäisches Patentamt European Patent Office Office européen des brevets

(11) EP 3 290 361 A1

(12) EUROPEAN PATENT APPLICATION
published in accordance with Art. 153(4) EPC

(43) Date of publication:
07.03.2018 Bulletin 2018/10

(21) Application number: 16786556.7

(22) Date of filing: 28.04.2016

(51) Int Cl.:
B65G 1/137 (2006.01) *G06K 17/00* (2006.01)
G06Q 10/08 (2012.01)

(86) International application number:
PCT/JP2016/063339

(87) International publication number:
WO 2016/175280 (03.11.2016 Gazette 2016/44)

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME
Designated Validation States:
MA MD

(30) Priority: 28.04.2015 JP 2015091125

(71) Applicant: Sato Holdings Kabushiki Kaisha
Tokyo 153-0064 (JP)

(72) Inventors:
- UNO, Yoshiaki
Singapore 408723 (SG)
- KASDANI, Yusita
Singapore 408723 (SG)

(74) Representative: Grünecker Patent- und Rechtsanwälte
PartG mbB
Leopoldstraße 4
80802 München (DE)

(54) STORAGE CABINET

TECHNICAL FIELD

[0001] The present invention relates to a storage cabinet that stores contents (items) such as products and goods.

BACKGROUND ART

[0002] A storage cabinet is known that manages contents (items) by using radio frequency identification (RFID) technology. The patent literature 1 for example describes that scanning is performed in a cabinet for monitoring a product including a RF tag for the purpose of searching for an expired product or a product that have been manufactured in a recalled lot.

[0004] The conventional storage cabinet such as one described above may be able to perform scanning an item such as a product in the cabinet by using RFID technology; however, it is necessary for an operator to visually check an expired product or a product that have been manufactured in a recalled lot and remove them from the cabinet. Thus, there is a drawback in the conventional storage cabinet that, in a case in which many products are stored in the storage cabinet for example, the operator cannot immediately recognize whether all products to be removed have been actually retrieved from the storage cabinet.

[0005] Particularly, in a case in which the storage cabinet is not connected to a network, the operator cannot check whether all products to be removed have been actually retrieved from the storage cabinet.

[0006] In view of the above, one of the aspects of the present invention is to provide a storage cabinet from which one can surely retrieve a desired item.

Figure A.8: Example of an automation patent without keywords

A.2.6 Comparison with Mann and Püttmann (2021)

In this section, we compare our classification of automation patents with that of Mann and Püttmann (2021, henceforth MP). We first show that our classifications are correlated though ours is generally stricter than theirs. Then, we focus on outlier technologies to understand where the differences come from.

We considered the 737,711 US machinery patents (according to our definition) of MP and classified them as auto95 or not. We have a lower share of automation patents (9.4% for auto95) than MP who have 29.8%. 70% of our auto95 patents are classified as automation patents by MP (to analyze this number, it is useful to note that their algorithm has a 17% false negative error rate on the training set), while we classify 22% of their automation patents as auto95 (see Table A.22). Therefore, our measure of automation is generally stricter than theirs although it is not a perfect subset.

To get a sense of where our classifications differ the most, we look for outlier C/IPC codes: we compute the difference between our prevalence measure and their share of automation patents and look at the codes with the highest and lowest values (focusing on codes with at least 100 patents in both their dataset and our EPO dataset). Table A.23 lists the 6 codes with the largest positive difference among auto95 codes, which correspond to codes that we more strongly identify as automation than MP do, and the

Table A.22: Confusion table for MP's and our classification

Machinery patents		MP Automation Yes (%)	MP Automation No (%)	Total (%)
DHOZ Automation	Yes (%)	6.6	2.8	9.4
	No (%)	23.2	67.4	90.6
	Total (%)	29.8	70.2	100.0

Notes: This table reports the shares of machinery patents that we (auto95) or Mann and Puettmann classify as automation. The sample is the set of US patents analyzed by Mann and Puettmann.

6 codes with the largest (in absolute value) negative difference among non-auto90 codes, which correspond to codes that MP more strongly identify as automation than we do.[47] Three of the codes with a high difference belong to the manipulator subclass (B25J): joints (B25J17), gripping heads (B25J15) and accessories of manipulators (B25J19). MP classify a large share of these patents as automation but our prevalence number is even higher. In their definition of automation patents, MP specify that they exclude innovations which only refer to parts of a machine. This accounts for some of the patents in these codes that they do not classify as automation. D01H9 corresponds to "arrangements for replacing or removing bobbins, cores, receptacles, or completed packages at paying-out or take-up stations" for textile machines. The share of automation patents in MP is low at 38%, however their "raw share" (computed before they exclude certain patents) is quite high at 71%. The excluded patents are not chemical or pharmaceutical patents (as emphasized in the paper), but belong to the "other" technology field (according to the Hall-Jaffe-Trajtenberg classification). The same situation occurs for B65B2210 (which is about packaging machines) where their raw automation score is actually at 63% and the patents excluded by MP are not chemical. B23P23 is a machine tool subclass (specifically "Machines or arrangements of machines for performing specified combinations of different metal-working operations not covered by a single other subclass") which often involves CNC technologies.

The non-auto90 codes where MP find a high share of automation patents but for which we have a comparatively low prevalence measure are of two types. Among the top 6, half are in the subclass B66B which corresponds to elevators and the other half are in the subclass B41J which corresponds to typewriters and printing machines. In fact, the

[47] We identify outliers using our prevalence measure at the 6-digit level instead of our share of automation patents because by construction, our share of automation patents is 100% for all auto95 codes so doing so would mask some of the underlying heterogeneity in our approaches. Table A.23 reports the share of auto95 patents for each code for clarity. Codes with a low prevalence score still feature some auto95 patents since a patent in a code with a low prevalence score can also have an auto95 code.

Table A.23: Outliers 6-digit C/IPC codes in the comparison between our measure and MP's measure

Code	Simplified description	DHOZ Keyword prevalence	DHOZ Share auto95
Positive outliers among auto95 codes			
B25J17	Manipulators (joints)	0.84	1.00
D01H9	Textile machines (arrangements for replacing or removing various elements)	0.62	1.00
B65B2210	Manipulators (gripping heads)	0.48	1.00
B25J15	Metal working machines (specified combinations n.e.c)	0.71	1.00
B23P23	Manipulators (accessories)	0.67	1.00
B25J19	'B33Y70$_descr'$	0.89	1.00
Negative outliers among non-auto90 codes			
B66B2201	Control systems of elevators	0.19	0.01
B66B3	Elevators (signalling and indicating device applications)	0.19	0.03
B41J23	Typerwriters / printing machines (power drive)	0.08	0.11
B66B1	Elevators (control systems)	0.16	0.02
B41J19	Typerwriters / printing machines (characters and line spacing mechanisms)	0.14	0.04
B41J5	Typerwriters / printing machines (controlling character selection)	0.21	0.09

Notes: This table lists the 6 auto95 codes with the largest positive difference between the prevalence of automation keywords ba sification and the share of automation patents according to MP in their data; and the 6 non-auto90 codes with the largest neg between the two measures. We additionaly list the share of patents classified auto95 according to our definition. We restrict att with at least 100 patents in both datasets.

first 32 6-digit C/IPC codes belong to either B66B, B41J or the subclass B65H which is about handling thin or filamentary material and also involves patents associated with printing machines. It is not surprising that our classifications differ for these types of innovation, since they do correspond to processes performed independently of human action (in line with MP's criterion); yet elevators and printers do not (or at least, no longer) replace humans in existing tasks.

A.3 Reproducing ALM

We detail how we build the variables used in Section 2.7 and provide further results.

A.3.1 Data for the ALM exercise

Except for the automation and labor share measures, we take the variables directly from ALM. We refer the reader to that paper for a detailed explanation. The task measures are computed using the 1977 *Dictionary of Occupational Titles* (DOT) which measure the tasks content of occupations. Occupations are then matched to industries using the Census Integrated Public Micro Samples 1% extracts for 1960, 1970, and 1980 (IPUMS) and the CPS Merged Outgoing Rotation Group files for 1980, 1990, and 1998 (MORG). The task change measure at the industry level reflects changes in occupations holding the task content of each occupation constant, which ALM refer to as the extensive margin. Since tasks measures do not have a natural scale, ALM convert them into percentile values corresponding to their rank in the 1960 distribution of tasks across sectors. Therefore, the employment-weighted means of all tasks measure across sectors in 1960 is 50. Our analysis starts in 1980 and drops a few sectors but we keep the original ALM measure to facilitate comparison. As in ALM, the dependent variable in Table 3 corresponds to 10 times the annualized change in industry's tasks inputs. Computerization ΔC_j is measured as the change per decade in the percentage of industry workers using a computer at their jobs between 1984 and 1997 (estimated from the October Current Population Survey supplements). For all regressions, observations are weighed by the employment share in each sector.

To map patents to sectors we proceed in 4 steps. First, we build a mapping between C/IPC 4-digit codes and the SIC sector that holds the patent (inventing sector). To do that, we use Autor et al. (2020) who match 72% of domestic USPTO corporate patents to firms in Compustat. This allows us to assign a 4-digit SIC sector to this subset of patents. We match the USPTO patents to our patent family data from PATSTAT, which we use to get the full set of C/IPC codes of the family. We then restrict attention to granted patents in machinery applied for in the period 1976-2010. Each patent family for which we have a sector creates a link between its C/IPC codes and that sector. We weigh that link inversely to the number of 6-digit C/IPC codes in the patent. Counting these connections allows us to build a weighted concordance table between 656 4-digit C/IPC codes and 397 SIC codes (at different levels of aggregation), where the industries refer to the industry of invention / manufacturing.

Second, to obtain the sector of use we rely on the 1997 "investment by using industries" table from the BEA (at the most disaggregated level, 180 commodities for 123 industries) which gives the flows of investment from commodities to industry available at www.bea.gov/industry/capital-flow-data. Since machines are a capital input, this is the appropriate equivalent of a standard IO table. Beforehand, we assign commodities to industries using the 1997 make table at the detailed level from the BEA (available at www.bea.gov/industry/historical-benchmark-input-output-tables) which gives the commodities produced by each industry.[48] We dropped commodities associated with the

[48]Since our industries are in SIC 1987, we use concordance tables from the IO industries to NAICS 1997 provided by the BEA and then the weighed concordance table between NAICS 1997 and SIC 1987 from David Dorn's website https://www.ddorn.net/data.htm which we complete with a concordance table from the Census available here (www.census.gov/eos/www/naics/concordances/concordances.html). To generate weights in the mapping between IO industries and NAICS 1997 and to disaggregate the NAICS industries from the capital flow table, we use CBP data from 1998

construction sector which are structures. Combining the two BEA tables, we obtain an investment flow table at the industry level. We then combine that table with the table mapping C/IPC to industry of manufacturing in order to obtain a mapping between C/IPC codes and (932 SIC) industries of use.

Third, we allocate patent families fractionally to their C/IPC 4-digit codes and use the previous table to assign them to an industry of use in the SIC classification (having restricted attention to the C/IPC codes which appear in the table). Fourth, we use a concordance table from the US Census Bureau from SIC industries to the Census industries from 1990 (ind90) given by Scopp (2003) and ALM concordance table from ind90 to consistent Census industries (ind6090) in order to allocate patents to their industry of use in ALM's classification.

Finally, for each sector, we compute the sums of automation patents and machinery patents over the time period 1980-1998 and take the ratio to be our measure of automation intensity.

To compute the share of automation patents in machinery according to the industry of manufacturing / invention, we proceed as above but skip step 3 with the investment flow table. Once patents are assigned to a SIC industry of manufacturing, we use the same concordance tables to assign patents to an ind6090 industry of manufacturing.

We source our labor share data from the NBER manufacturing database and the BEA. In the NBER manufacturing database, we calculate the labor share as total payroll / value-added and apply the concordance procedure described in step 3 above to go from the 4-digit SIC industries to the consistent Census industries. The database is limited to industries in the manufacturing sector. The BEA provides labor share data for more aggregate SIC industries for the whole economy. We calculate the labor share as total compensation / value-added and build a crosswalk from the 4-digit SIC level to these more aggregate industries to map our patents.

Finally, in robustness checks, we also use an alternative mapping from patents to sectors based on Lybbert and Zolas (2014) who provide a concordance table between IPC codes at the 4-digit level and NAICS 1997 6-digit industry codes. The concordance table is probabilistic (so that each code is associated with a sector with a certain probability). The Lybbert and Zolas concordance tables are derived by matching patent texts with industry descriptions, and as such they cannot *a priori* distinguish between sector of use and industry of manufacturing. We checked, however, that patents associated with

(https://www.census.gov/data/datasets/1998/econ/cbp/1998-cpb.html).

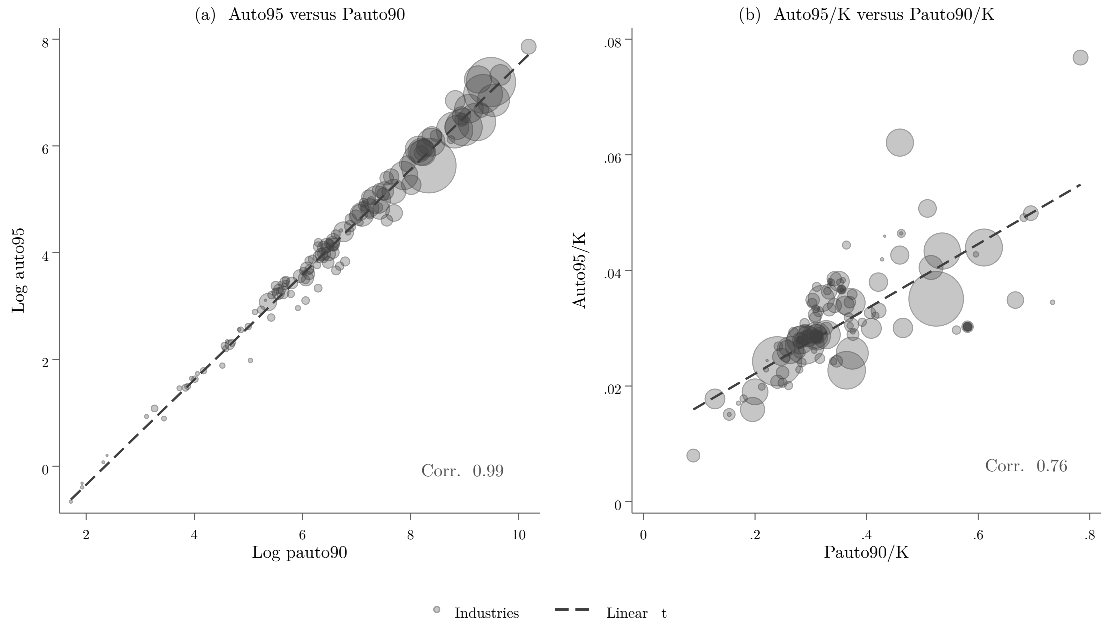

Figure A.9: Correlation between counts of auto95 and pauto90 patents at the sectoral level.
Notes: Panel (a) shows the log counts and Panel (b) shows counts scaled by capital purchases. Sectors are employment-weighted.

"textile and paper machines" for instance are associated with the textile and paper sectors and not with the equipment sector. Therefore, we think of this mapping as rather corresponding to the using sector as well for our set of technologies. In addition, it has the advantage of providing a much more direct mapping between C/IPC codes and industries. We attribute patents to sectors fractionally in function of their C/IPC codes. To assign patents to the consistent Census industry codes used by ALM, we first use a Census concordance table (https://www.census.gov/topics/employment/industry-occupation/guidance/code-lists.html) to go from NAICS 1997 to Census industry codes 1990, and then again use ALM concordance table.

A.3.2 Additional results

We now provide a few additional results which complements those in the main text. As discussed in the text, machinery patents tend to be used by the same sectors whether they are automation or non-automation patents. Figure A.9.a shows the (employment-weighted) correlation between the log of auto95 patents and the log of pauto90 (i.e. non-automation) patents across US sectors. The very strong correlation reflects our procedure which allocates patents according to capital purchases by sector. To remove this partly

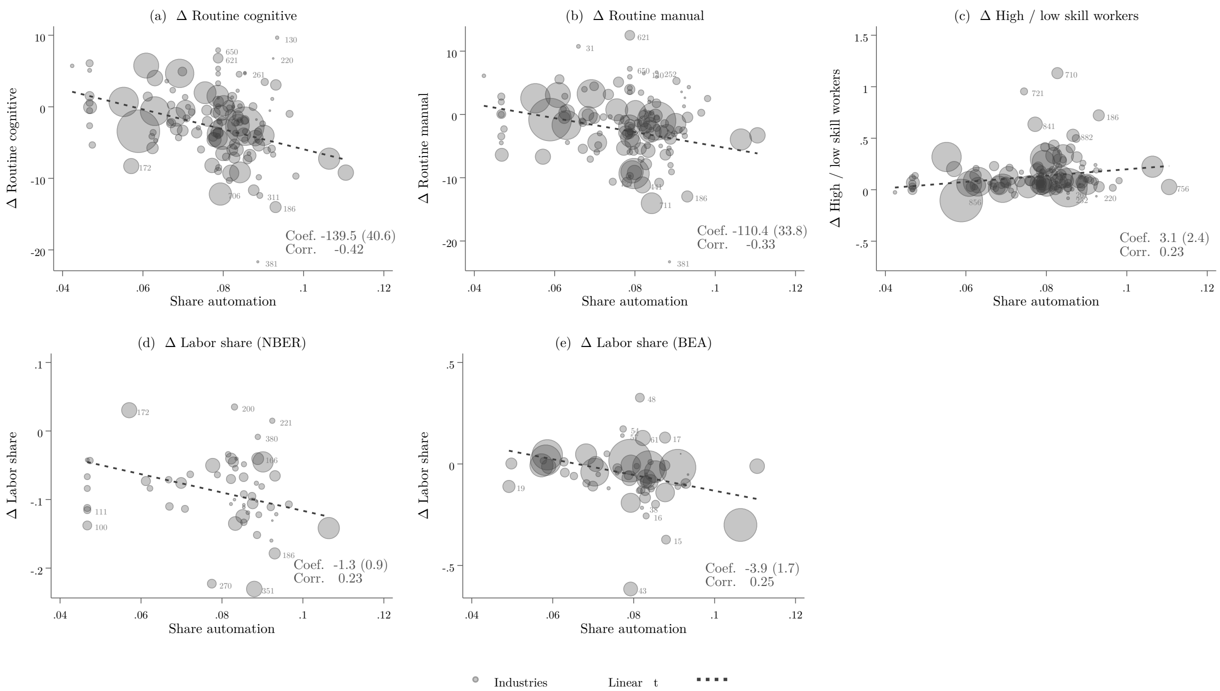

Figure A.10: Scatter plots of changes in routine tasks, skill composition, and the labor share versus the share of automation patents (auto95) in machinery patents used by the industry in 1980-1998.

mechanical effect, Figure A.9.b shows the correlation between the ratio of auto95 patents over capital purchases and pauto90 patents over capital purchases. There is still a substantial correlation 0.76, showing that automation and non-automation patents tend to be used by the same sectors even controlling for the amount of capital purchased. Nevertheless, the sectoral variation is sufficient to enable us to look at the effect of the share of automation among machinery patents across sectors.

Figure A.10 shows scatter plots of the change in routine tasks and skill composition and the share of automation patents in 1980-1998. This figure shows the raw data underlying the regressions in Columns (1), (3), (5), (7) and (9) of Table 3 – but the figure does not control for computerization or the manufacturing dummy.

We carry a number of robustness checks in Table A.24. In Columns (1), (4), (7) and (10) we compute the share of automation patents using only granted USPTO patents which are also biadic. The results are similar to those in Table 3 though less precise for the skill ratio. In Columns (2), (5), (8) and (11), we use the share of auto90 patents in machinery to measure automation in the sector of use. The results are similar but with smaller coefficients than in the regressions using auto95 (and less precise for the

Table A.24: Robustness checks for the sectoral analysis

	Δ Routine cognitive			Δ Routine manual			Δ High/low skill workers			Δ Labor Share (NBER)		
	Biadic (1)	Auto90 (2)	Lybbert and Zolas (3)	Biadic (4)	Auto90 (5)	Lybbert and Zolas (6)	Biadic (7)	Auto90 (8)	Lybbert and Zolas (9)	Biadic (10)	Auto90 (11)	Lybbert and Zolas (12)
Share automation	-120.82^{***}	-69.67^{***}	-23.10^{***}	-102.55^{***}	-58.55^{***}	-13.48^{**}	2.45	1.72	0.74^{**}	-1.19^{*}	-0.73^{*}	-0.27^{**}
	(27.51)	(20.22)	(4.85)	(35.70)	(20.81)	(5.68)	(1.85)	(1.18)	(0.30)	(0.62)	(0.37)	(0.11)
Δ Computer use (1984-1997)	-21.12^{***}	-18.41^{**}	-13.45	-20.90^{***}	-18.58^{**}	-7.53	1.01^{***}	0.96^{***}	0.42	0.24^{*}	0.26^{**}	0.23
	(7.29)	(7.44)	(8.93)	(7.81)	(7.81)	(8.38)	(0.26)	(0.27)	(0.28)	(0.13)	(0.13)	(0.14)
Manufacturing	-1.70^{*}	-1.20	-1.66	-0.07	0.34	-1.65^{*}	0.03	0.01	0.02			
	(0.93)	(1.02)	(1.65)	(0.94)	(1.03)	(0.96)	(0.03)	(0.03)	(0.02)			
R^2	0.26	0.23	0.40	0.17	0.15	0.32	0.17	0.18	0.43	0.18	0.19	0.27
Industries	133	133	71	133	133	71	133	133	71	56	56	56

Notes: This table provides robustness checks for the effect of automation technologies on tasks, skill composition, and the labor-share. Columns 1, 4, 7, and 10 use biadic auto95 patents: that is, patents applied for in at least two countries. Columns 2, 5, 8, and 11 define automation patents as auto90 patents. In both cases, patents are allocated to their sector of use. Columns 3, 6, 9, and 12 use auto95 patents (as in the baseline) but allocate patents using a concordance table between C/IPC codes and industries from Lybbert and Zolas (2014). The regressions are weighed by the mean industry share of total employment in FTEs in 1980 and 1998. Standard errors are clustered at the level of industry groups that have the same automation share by construction and reported in parentheses. Significance levels at *10%, **5%, ***1%.

skill ratio), in line with auto95 being a stricter measure of automation. In Columns (3), (6), (9) and (12) we instead map patents to sectors based on a concordance table from Lybbert and Zolas (2014) between 4-digit C/IPC codes and sectors. This method has the advantage of mapping more directly patents to sectors but cannot distinguish between manufacturing and using sectors. We still find that sectors with a high share of automation patents experienced a decline in routine tasks. The coefficients are smaller, but given that the standard deviation of the share of automation patents in that case is 0.086, the standardized coefficients are relatively similar.

Finally, in Table A.25, we look at the effect of the share of automation patents on total employment and employment by skill type. Panel A looks at all industries. As already seen in Table 3, automation is associated with a relative decrease in low-skill employment compared to high-skill labor. The effect on low-skill employment is negative but non-significant and the effect on total employment is closer to 0 (as there is a positive non-significant effect on high-skill employment). The results are clearer in the manufacturing sector, where an increase in automation is associated with a significant decrease in both low-skill and total employment.

A.4 A Simple Model

We incorporate the business features described in 3.1 into a simple model built on Hémous and Olsen (2022). A final good is produced with a continuum of intermediate inputs according to the Cobb-Douglas production function $Y = \exp\left(\int_0^1 \ln y\,(i)\,di\right)$, where $y(i)$ denotes the quantity of intermediate input i. The final good is the numéraire. Each

Table A.25: Changes in employment and automation

	Δ Log employment		Δ Log high-skilled		Δ Log low-skilled	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. All industries						
Share automation	−2.31	−2.02	3.15	3.46	−4.44	−4.11
(using industry)	(3.45)	(3.69)	(3.42)	(3.98)	(3.98)	(4.03)
Share automation		0.99		1.05		1.11*
(inventing industry)		(0.62)		(0.88)		(0.61)
Δ Computer use	1.45*	1.56*	1.40**	1.52**	0.96	1.08
(1984-1997)	(0.80)	(0.84)	(0.68)	(0.72)	(0.82)	(0.85)
R^2	0.10	0.12	0.08	0.11	0.07	0.10
Mean dependent variable	−2.50	−2.50	0.12	0.12	−2.27	−2.27
Industries	133	133	132	132	133	133
Panel B. Manufacturing industries						
Share automation	−4.67***	−4.67**	−1.36	−2.41	−6.17***	−5.90**
(using industry)	(1.54)	(2.29)	(1.91)	(2.66)	(1.77)	(2.42)
Share automation		0.01		1.06		−0.28
(inventing industry)		(1.34)		(1.17)		(1.38)
Δ Computer use	1.37***	1.37***	2.01***	1.97***	1.06**	1.07**
(1984-1997)	(0.50)	(0.51)	(0.56)	(0.57)	(0.52)	(0.52)
R^2	0.14	0.14	0.15	0.16	0.14	0.14
Mean dependent variable	−4.26	−4.26	0.14	0.14	−2.62	−2.62
Industries	58	58	57	57	58	58

Notes: This table shows the effect of automation technologies on employment. Each column represents a separate OLS regression of the change in log employment between 1980 and 1998 on the share of automation patents in machinery, the annual percentage point change in industry computer use during 1984-1997, and a constant. Panel A considers all industries. Panel B focuses on industries in manufacturing. In columns 1–2 the dependent variable is the change in log employment, in columns 2–3 the change in log employment of high-skilled workers (college graduates), and in columns 3–4 the change in log employment of low-skilled workers (others). The two automation share measures correspond to a different mapping between C/IPC codes and industries. Using industries allocates patents to their sector of use while innovating industry – added in columns 2,4, and 6 – allocates patents to their sector of invention. The regressions are weighed by the mean industry share of total employment in FTEs in 1980 and 1998. Standard errors are clustered at the level of industry groups that have the same automation share by construction and reported in parentheses. Significance levels at *10%, **5%, ***1%.

intermediate input is produced competitively with high-skill labor ($h_{1,i}$ and potentially $h_{2,i}$), low-skill labor, l_i, and potentially machines, x_i, according to:

$$y_i = h_{1,i}^{1-\beta}\left(\gamma\left(i\right)l_i + \alpha\left(i\right)\nu^{\nu}(1-\nu)^{1-\nu}x_i^{\nu}h_{2,i}^{1-\nu}\right)^{\beta}. \tag{6}$$

$\gamma(i)$ is the productivity of low-skill workers, $\alpha(i)$ is an index equal to 0 for non-automated intermediates and to 1 for automated intermediates and ν and β are parameters in $(0,1)$. Machines are specific to the intermediate input i. If a machine is invented, it is produced monopolistically 1 for 1 with the final good so that the monopolist charges a price $p_x(i) \geq 1$. At the beginning of the period, a potential innovator has the opportunity to create a specific machine for each non-automated intermediate i. She does so with probability λ if she spends $\theta\lambda^{\psi+1}Y/(\psi+1)$ units of the final good with $\psi > 0$.

For an automated intermediate input ($\alpha(i) = 1$), the downstream producer is indifferent between using low-skill workers or machines together with high-skill workers in production whenever $w_H^{\nu}p_x^{1-\nu} = w_L/\gamma(i)$. Therefore, the machine producer is in Bertrand competition with low-skill workers. As a machine costs 1, the machine producer charges a price $p_x(i) = \max\{(w_L/\gamma(i))^{\frac{1}{1-\nu}}\, w_H^{-\frac{\nu}{1-\nu}}, 1\}$ such that machines are used if $w_L/\gamma(i) > w_H^{\nu}$. Since the final good is produced according to a Cobb-Douglas production function, we get $p(i)y(i) = Y$ for all intermediates. We can then derive the profits of the machine producer as $\pi_i^A = \max\left(1-(\gamma(i)/w_L)^{\frac{1}{1-\nu}}\, w_H^{\frac{\nu}{1-\nu}}, 0\right)\nu\beta Y$.

In turn, at the beginning of the period, the potential innovator solves $\max \lambda\pi_i^A - \theta\lambda^{\psi+1}Y/(\psi+1)$, giving the equilibrium innovation rate $\lambda = \left[\pi_i^A/(\theta Y)\right]^{1/\psi}$. As a result, the number of automation innovations is equal to:

$$Aut = \left(\frac{\nu\beta}{\theta}\right)^{1/\psi}\int_0^1 \left(1-\alpha\left(i\right)\right)\left[\max\left(\left(1-\left(\frac{\gamma(i)}{w_L}\right)^{\frac{1}{1-\nu}} w_H^{\frac{\nu}{1-\nu}}\right), 0\right)\right]^{1/\psi} di.$$

This expression is increasing in the low-skill wage w_L and decreasing in the high-skill wage w_H with a larger magnitude for a lower ψ. Intuitively, the incentive to replace low-skill workers with machines (and high-skill workers) increases with low-skill wages, leading to a higher demand for machines. The reverse holds for high-skill wages. An upward shift in low-skill worker productivity, $\gamma(i)$, also reduces the number of automation innovations. Our empirical analysis aims at computing $\partial \ln Aut/\partial \ln w_L$.

To contrast automation with other types of innovations, assume that the production of an intermediate takes place according to:

$$y_i = (q_i m_i)^\delta h_{1,i}^{1-\beta-\delta} \left(\gamma(i)\, l_i + \alpha(i)\, \nu^\nu (1-\nu)^{1-\nu} x_i^\nu h_{2,i}^{1-\nu}\right)^\beta,$$

where m_i denotes non-automation "Hicks" machines with quality q_i. Hicks machines are also produced 1 for 1 with the final good. Each period a potential innovator may improve on the available quality of Hicks machines for intermediate i by a factor μ by investing in R&D. If she spends $\theta_m \lambda_m^{\psi+1} Y/(\psi+1)$ units of the final good, she is successful with probability λ_m. In that case, the innovator becomes the monopolistic provider of Hicks machine i under the pressure of a competitive fringe which has access to the previous technology, and the technology diffuses after one period. Otherwise, the good is produced competitively.

The previous analysis on automation innovations remains identical. A successful Hicks innovator can charge a mark-up μ leading to profits $\pi_i^H = (1-\mu^{-1})\,\delta Y$. The innovation rate is then $\lambda_m = \left[(1-\mu^{-1})\,\delta/\theta_m\right]^{1/\psi}$, so that the number of Hicks innovations is a constant given by λ_m. In contrast to automation innovations, the number of non-automation innovations is independent of low- or high-skill wages.

A.5 Data Appendix for the main analysis

Here, we provide details on the data and the variable construction for our main analysis.

A.5.1 Macroeconomic variables

Our main source of macroeconomic variables is the *World Input Output Database (WIOD)* from Timmer et al. (2015) which contains information on hourly wages (low-skill, middle-skill and high-skill) for the manufacturing sector and the total economy from 1995 to 2009 for 40 countries. It also contains information on GDP deflators and PPIs, both for manufacturing and for the whole economy. They employ the ISCED skill-classification, where categories 1+2 denote low-skill (no high-school diploma in the US) 3+4 denote middle-skill (high-school but not completed college) and 5+6 denote high-skill (college and above). Switzerland is not included in the WIOD database and we use data on skill-dependent wages, productivity growth and price deflators obtained directly from *Federal Statistical Office of Switzerland.*

We add data from *UNSTAT* on exchange rates and GDP (and add Taiwan from the *Taiwanese Statistical office*). We calculate the GDP gap as the deviations of log

GDP from HP-filtered log GDP using a smoothing parameter of 6.25. To compute the offshoring variable we follow Timmer et al. (2014) and compute the share of foreign value added in manufacturing from the WIOD 2013 (except for Switzerland where we use the 2016 release and assign to the years 1995-1999 the same value as in 2000). For the nominal interest rate, we use the yield on 10-year government bonds with data from the OECD for AT AU BE CA CH DE DK ES FI FR GB IE IT JP NL PT SE US and from the IMF for KR GR LU.

The primary data source for the hourly minimum wage data is *OECD Statistics.*[49] For the US, we use data from FRED for state minimum wages and calculate the nation-level minimum wage as the weighed average of the state-by-state maximum of state minimum and federal minimum wages, where the weight is the manufacturing employment in a given state. Further, the UK did not have an official minimum wage until 1999. Before 1993, wage councils set minimum wages in various industries (see Dickens, Machin and Manning, 1999). We compute an employment-weighed industry average across manufacturing industries and use the 1993 nominal value for the four years in our sample (1995-1998) with no minimum wage. Finally, Germany did not have a minimum wage during the time period we study. Instead, we follow Dolado et al. (1996) and use the collectively bargained minimum wages in manufacturing which effectively constitute law once they have been implemented. These data come from personal correspondence with Sabine Lenz at the *Statistical Agency of Germany.*

Table A.26 shows that low-skill and high-skill wages differ considerably across countries and that the skill premium also varies for countries of similar development level. For instance, between 1995 and 2009, the skill premium in the United States rose from 2.46 to 3.02 but slightly declined in Belgium from 1.56 to 1.46.

A.5.2 Merging Orbis firms

For our analysis, we need to decide the level at which R&D decision are undertaken. Orbis IP links patent data to companies. For companies in the same business group, R&D decisions could happen at the group level, though treating a group as one agent is

[49]Not all countries have government-imposed hourly minimum wages. Spain, for instance, had a monthly minimum wage of 728 euros in 2009. To convert this into hourly wage we note that Spain has 14 "monthly" payments a year. Further, workers have 6 weeks off and the standard work week is 38 hours. Consequently we calculate the hourly minimum wages as monthly minimum wage$\times 14/\left[(52-6)\times 38\right]$, which in 2009 is 5.83 euros per hour. We perform similar calculations, depending on individual work conditions, for other countries with minimum wages that are not stated per hour: Belgium, Brazil, Israel, Mexico, Netherlands, Poland and Portugal.

Table A.26: Low-skill wages and the skill premium in manufacturing for selected countries

Country	Low-skill wages (1995$)		High-skill wages (1995$)		Skill-premium (HSW/LSW)	
	1995	2009	1995	2009	1995	2009
India	0.19	0.28	0.89	1.38	4.79	4.98
Mexico	0.89	0.61	3.46	2.56	3.90	4.21
Bulgaria	1.29	0.71	4.27	1.60	3.32	2.25
United States	11.57	13.67	28.42	41.23	2.46	3.02
Belgium	29.50	41.89	45.98	61.24	1.56	1.46
Sweden	19.92	42.16	34.44	55.92	1.73	1.33
Finland	23.41	43.63	28.10	63.71	1.20	1.46

Notes: Wages data, taken from WIOD. The table shows manufacturing low-skill and high-skill wages (technically labor costs) deflated by (manufacturing) PPI and converted to USD using average 1995 exchange rates. Skill-premium is the ratio of high-skill to low-skill wages. The table shows the three countries with the lowest low-skill wages in 2009, the three with the highest and the US.

often too aggressive (as subsidiaries might be in different sectors). Therefore, for firms within the same business group, we normalize company names by removing non-firm specific words such as country names or legal entity types and then merge firms with the same normalized name. All other firms are treated as separate entities. E.g., Siemens S.A., Siemens Ltd. or Belgian Siemens S.A. are merged, but Primetals Technologies Germany Gmbh which belongs to the same group remains a separate entity.

A.5.3 Firm-level patent weights

We give further details on the firm level patent weights. As mentioned in the text, we only count patents in machinery because some of the biggest innovators in automation technologies are large firms which produce a wide array of products with different specialization patterns across industries. Further, we exclude firms which have more than half of their patents in countries for which we do not have wage information.

In Europe, firms can apply both at national patent offices and at the EPO, in which case they still need to pay a fee for each country where they seek protection. We count a patent as being protected in a given European country if it is applied for either directly in the national office or through the EPO. In addition, we take the following steps in order to deal with EP patents. We assign EP patents to countries when they enter into the national phase. A firm's untransferred EP patents are assigned using information on where that firm previously transferred its EP patents. If a firm does not have any already transferred EP patents, we assign the patent based on a firm's direct patenting history in EPO countries. Untransferred EP patents that are still left are assigned to countries based on the EPO-wide distribution of transfers. We also drop a firm if more than half of its patents are EP patents assigned using the EPO-wide distribution.

Finally, as mentioned in the text, we only count patents in families with at least one (non self-) citation. Including all patents generally increases the weight of the country with the most patents, in line with the finding that poor quality patents tend to be protected in fewer countries. However, further increasing the threshold from 1 to more citations does not significantly change the distribution of weights.

A.6 Additional results and robustness checks for the main analysis

This Appendix presents robustness checks linked to our shift-share set-up (Appendix A.6.1), other robustness checks (Appendix A.6.2), details on the comparison of our estimates with estimates in the literature found in Section 4.4 (Appendix A.6.3), and finally details on the simulation exercise presented in Section 4.4 (Appendix A.6.4).

A.6.1 Shift-share analysis

We present a number of additional results related to our shift-share set-up. We first do a "shock-level" analysis as recommended by BHJ, then we show that our results do not depend on a single country and include additional shock-level controls, finally, we address Borusyak and Hull (2021)'s concern regarding the use of a nonlinear shift-share.

Shock-level regressions. BHJ show that identification in a shift-share setting can be obtained from conditionally randomly allocated shocks. Key to their argument is an equivalence result between what in our context would be a linear firm-level regression and a linear regression run at the level of the shocks (country-year). They advise practitioners to run the shock-level regression and to provide several statistics showing that there are enough variations in the shocks, that there are sufficiently many shocks, and how the shocks correlate with other variables.

To follow their approach we need to turn to a linear setting. To do that, we first replace our dependent variables which are defined as log of averages with average of logs. In addition, it is easier to map our analysis with theirs if we consider a single shock. Therefore, given the previous results showing that low- and high- skill wages often have coefficients of opposite magnitude, we directly look at the effect of the inverse skill premium. We define it here as:

$$ISP_{i,t} \equiv \sum_c \kappa_{i,c} \ln\left(\frac{w_{L,c,t}}{w_{H,c,t}}\right). \tag{7}$$

Table A.27: From firm-level to shock level regressions

Dependent variable	Auto95				
	Firm-level		Country-level		
	(1)	(2)	(3)	(4)	(5)
Low-skill / High-skill wages	2.49***	0.40***	0.40***	0.33**	0.37***
	(0.86)	(0.15)	(0.07)	(0.16)	(0.07)
Labor productivity				−0.31	
				(0.50)	
GDP gap				−0.32	
				(1.82)	
Estimator	Poisson	Linear (arcsinh)	Linear (arcsinh)	Linear (arcsinh)	Linear (arcsinh)
Stocks and spillovers	Yes	Yes	Yes	Yes	No
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	47 741	48 780	615	615	615
Firms / Countries	3252	3252	41	41	41

Notes: This table reports shock-level equivalent regressions. The coefficients are estimated with conditional Poisson fixed effect regressions (HHG) in column 1 and OLS in columns 2–5. The dependent variable in columns 2–5 is the arcsinh transformation of auto95 innovations. Standard errors are reported in parentheses. Standard errors are clustered at the firm-level in columns 1 and 2 and country-level clustered in columns 3–5. Columns 3–5 run equivalent shock-level regressions following Borusyak, Hull and Jaravel (2022, BHJ) (see text for details). All regresions include firm fixed effects, industry-year fixed effects and country-year fixed effects. Significance levels at *10%, **5%, ***1%.

We also define the other macro variables (GDP per capita, labor productivity, etc) as average of logs. Second, we switch from a Poisson estimator to a linear one where we use arcsinh of the count of patents as a dependent variables (the arcsinh is approximately linear for low values and approximately log for higher values which allows us to deal with 0s). That is we replace (4) with:

$$\begin{array}{rc} & \operatorname{arcsinh}\left(PAT_{Aut,i,t}\right) \\ = & \begin{array}{c} \beta_{ISP} ISP_{L,i,t-2} + \beta_X X_{i,t-2} + \beta_{Ka} \ln K_{Aut,i,t-2} + \beta_{Ko} \ln K_{other,i,t-2} \\ + \beta_{Sa} \ln SPILL_{Aut,i,t-2} + \beta_{So} \ln SPILL_{other,i,t-2} + \delta_i + \delta_{j,t} + \delta_{c,t} + \epsilon_{i,t} \end{array} . \end{array} \tag{8}$$

Finally, we focus this analysis on total wages (with country-year fixed effects) since this set-up is more easily transcribed in the BHJ framework.

Table A.27 shows the results. Columns (1) and (2) report regressions at the firm-level. In Column (1), we only replace the previous definition of the inverse skill premium (the difference between the log average of low- and high-skill wages) with that of equation (7). We control for firm, industry-year and country-year fixed effects, stocks and spillovers but not for any other macro variables in order to focus on the direct effect of the shock in consideration. We obtain a coefficient much in line with those of Table A.8. Column (2) runs a linear regression at the firm level as in (8). We obtain a similar result – the magnitude is smaller as the range of variations for arcsinh is smaller than for log.

Column (3) follows the BHJ approach and runs a shock-level regression. That is,

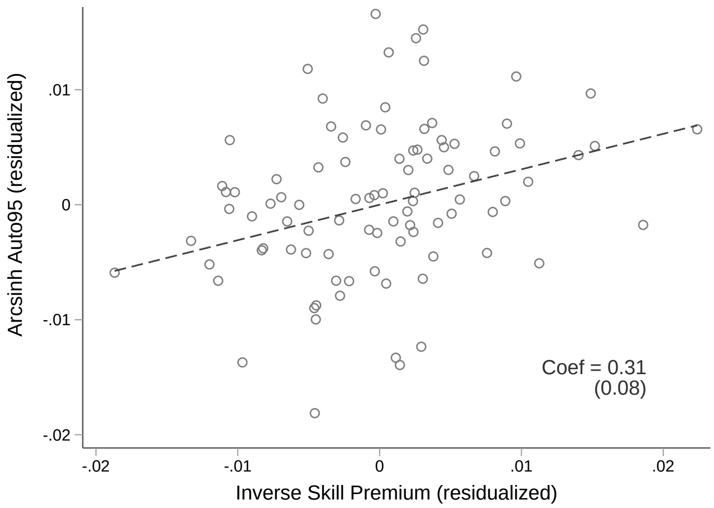

Figure A.11: Bin-scatter plot of the shock-level regression.
Notes: This figure shows bin-scatter plot regressions of automation on the inverse skill premium. We residualize both arcsinh(auto95) and the inverse skill premium on firm, industry-year and country-year fixed effects and on stocks and spillover variables. We then compute weighted average of the residuals at the shock (i.e. country-year) level following BHJ. We then group observation in 100 bins of the inverse skill premium.

we first residualize our automation measure on our controls (fixed effects, stocks and spillovers) and similarly residualize the inverse skill premium measure. We then compute a weighted average of the residualized automation measure at the country-year level, where, for each country, we weigh each firm-year observation by the firm-country weight $\kappa_{i,c}$. We then run a linear regression of that average measure of automation on the inverse skill premium at the country-year level. Each country-year observation is weighed by its average weight at the firm level. As demonstrated by BHJ, we get exactly the same coefficient. Column (4) adds controls for labor productivity in manufacturing and Column (5) removes the controls for stocks and spillovers so that the only controls are the fixed effects. While the original regression looks at the effect of a weighted average of wages on firms' innovations, this "shock-level" regression inverts the relationship and looks at the effect of wages on a weighted average of firms' innovations. It is important to realize that this does not mean that our original shift-share approach would simply mean re-weighing firm-level variables to run a country-level regression. Our measure of automation innovation arcsinh $(PAT_{Aut,i,t})$ is first residualized on country-year fixed effects, so that we remove the average contribution of domestic firms to automation innovation when we run the shock level regression.[50]

[50] As already mentioned, we run this analysis at the level of the inverse skill premium because this allows us to keep track of only one shock. In addition, regressions with arcsinh and separate low- and high- skill wages do not show a significant effect for low-skill wages when we use the full sample. This

Table A.28: Shock-level summary statistics

	(1)	(2)	(3)	(4)
Mean	-0.78	0	0	0
Standard deviation (%)	36.4	2.1	0.9	1.0
Interquartile range (%)	55.7	2.9	1.0	1.0
Residualizing on ...				
F fixed effect	–	Yes	Yes	Yes
IY+CY fixed effects	–	–	Yes	Yes
Stocks/Spillovers	–	–	–	Yes

Notes: This table reports summary statistics on the log inverse skill premium weighted by the average country weight in our regression sample as in Borusyak, Hull and Jaravel (2022). The log inverse skill premium is residualized on firm fixed effects (Columns 2, 3 and 4), industry-year and country-year fixed effects (Columns 3 and 4) and stocks and spillovers (Column 4).

To unpack our regression results, Figure A.11 shows a bin-scatter plot of the residualized measures of automation and the inverse skill premium at the country-year level. The figure corresponds to the regression of Column (5) in Table A.27 which only controls for fixed effects. We group observations in 100 bins of equal weights. The overall relationship between automation and the inverse skill-premium does not seem to be driven by outliers or specific parts of the inverse skill premium distribution.

Shock-level summary statistics. Table A.28 reports summary statistics on the shock-level regressions. The standard-deviation of the shock, namely the log inverse skill premium residualized on firm, industry-year and country-year fixed effects is 0.9%. This is a significant amount of variation given that the standard deviation of the log inverse skill premium residualized only on firm fixed effects (i.e. only taking away level differences across countries) is 2.1% (see also the distribution in Figure A.11 and Table A.5).

Table 4 reports that the HHI of weights are 0.13 for total weights and 0.09 for foreign weights at the country level and therefore 0.009 and 0.006 at the country-year level. The "true" level of variation depends on how much variation there actually is in the time dimension for a given country. To assess this, Figure 3.c shows the evolution of the inverse skill premium for the 6 countries with the largest average weights residualized on country and year fixed effects. Figure A.12 does the same thing but residualizes the log inverse skill premium on the full set of fixed effects, stocks and spillovers (i.e. as in Column 3 of Table A.27). The two figures look overall similar: there is a significant amount of variation both across and within countries. Of course, the inverse skill premium is correlated from year to year, but after a few years, the correlation is much weaker. We

is due to the difference in functional forms between the arcsinh and log. We recover our original result when we focus on firms with at least 2 patents over the full time period. This result is exactly in line with our long-difference regressions that also use arcsinh (see Appendix Table A.35).

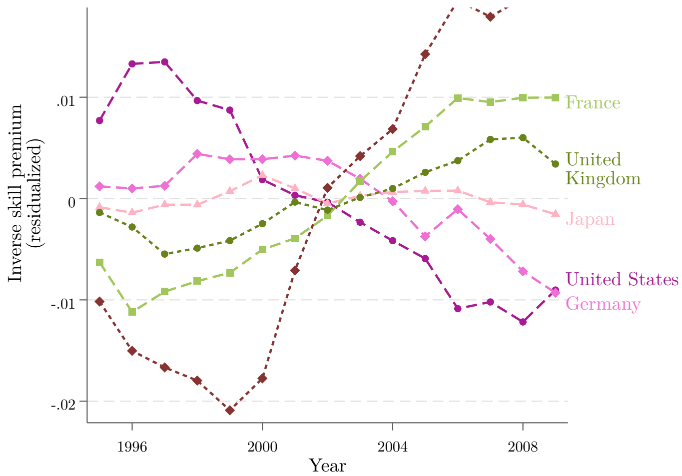

Figure A.12: Residualized inverse skill premium in the 6 most important countries.
Notes: This figure reports our identifying shocks, namely the log inverse skill premium residualized on firm fixed effects, industry-year and country-year fixed effects, stocks and spillovers variable and aggregated at the country level following BHJ's methodology.

find no correlation between the log skill premium and its fifth lag, so loosely speaking one may consider that we have at least 3 "separate observations" for each country.

Shock-level balance tests. In Table A.29, we look at the balance of our shocks against observables (offshoring is defined below). We regress the macro variables on the log inverse skill premium at the country-year level. All variables are residualized on our full set of fixed effects, stocks and spillovers, and observations are weighted following the BHJ procedure. The only macro variables that are significantly correlated with the skill premium are the recent innovation variables (there is also a significant coefficient for low-skill weighted manufacturing size but the effect is small). More automation innovations are associated with a higher skill premium as one would expect. This is also true for all other innovations – which include non machinery innovations such as innovations in IT, for instance. Table 7 shows that controlling for recent innovations does not affect the effect of wages on automation innovations in our central regressions.

Excluding one country at the time. Next, we check whether our results are driven by a specific country. We go back to our original firm-level Poisson regressions. We successively remove the six largest countries by average weight (US, JP, DE, GB, FR, IT, and ES). Excluding a country means that we treat it like the home country when computing normalized foreign wages. We control for the weight of the excluded country times a year dummy. Table A.30 reports the results (with foreign wages). The coefficient on low-skill wages always remains positive and significant.[51]

[51] Goldsmith-Pinkham, Sorkin and Swift (2020) suggest carrying out a similar exercise by excluding

Table A.29: Shock balance tests

	Estimate (1)	(SE) (2)
GDP Gap	0.00	(0.01)
Labor Productivity	−0.23	(0.17)
GDP per capita	0.04	(0.19)
Manufacturing size	−0.07	(0.10)
Manufacturing size (low-skill weighted)	-0.21^{*}	(0.12)
Offshoring	0.01	(0.03)
Recent auto95 innovation	-1.00^{***}	(0.39)
Recent other innovation	-1.34^{**}	(0.67)
Stocks and spillovers	Yes	
Fixed effects	F+IY+CY	
Number of country-years	615	

Notes: This table reports coefficients from separate regressions of country-level observables on the log inverse skill premium. The respective independent variables are residualized on firm, industry-year, and country-year fixed effects. Standard errors are reported in Column 2 and clustered at the country-level. Significance levels at *10%, **5%, ***1%.

Table A.30: Excluding one country at the time

	Auto95							
Excluded country	None	US	DE	JP	GB	FR	IT	ES
Average weight		0.21	0.20	0.17	0.09	0.09	0.03	0.03
	(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Foreign:								
Low-skill wage	5.32^{***}	5.75^{***}	3.84^{***}	3.67^{***}	4.95^{***}	3.62^{**}	5.53^{***}	5.10^{***}
	(1.56)	(1.70)	(1.40)	(1.34)	(1.34)	(1.50)	(1.48)	(1.53)
High-skill wage	-2.87^{*}	-2.55^{*}	−1.74	−1.58	−0.82	−2.16	-4.61^{**}	−2.46
	(1.47)	(1.46)	(1.31)	(1.31)	(1.35)	(1.33)	(1.92)	(1.50)
GDP gap	2.28	2.23	3.33	2.55	3.09	1.91	1.98	1.97
	(4.92)	(5.13)	(5.63)	(3.96)	(4.90)	(5.06)	(5.22)	(4.98)
Labor productivity	−2.57	-3.99^{**}	-2.56^{*}	−1.76	-3.61^{**}	−1.87	−1.15	-2.75^{*}
	(1.60)	(1.68)	(1.38)	(1.49)	(1.60)	(1.49)	(1.65)	(1.58)
Excluded country weight × year dummy	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47 741	46 984	47 272	47 562	47 333	47 681	47 606	47 670
Number of firms	3252	3200	3218	3240	3225	3248	3243	3247

Notes: This table excludes one country at the time. Column 0 reproduces the baseline regression with normalized foreign wages. Columns 1–7 exclude the country in the column header in addition to the domestic country when computing the normalized foreign macroeconomic variables. Additionally, columns 1–7 control for the weight of the excluded country times year dummies. The average weight in the header reports the average country weight for the firms in the sample of column 1. All columns include firm, industry-year and country-year fixed effects. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Additional controls. BHJ also recommend considering other shock-level variables

countries with a large Rotemberg weight. Rotemberg weights require a linear shift-share instrument. When wages are computed as average of logs, the six countries with the largest Rotemberg weights are the UK, FR, SE, DE, US, and BE. Our results are also robust to excluding Belgium and Sweden.

Table A.31: Including additional controls

Dependent variable	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.75***	2.90***	2.70***	3.09***	2.56**	2.48**	5.22***	5.30***	6.95***
	(0.86)	(0.88)	(0.92)	(1.17)	(1.15)	(1.18)	(1.52)	(1.55)	(1.87)
High-skill wage	−2.37***	−2.44***	−3.06***	−1.24	−1.86*	−2.25*	−2.93**	−2.69*	−2.90*
	(0.74)	(0.80)	(0.92)	(1.00)	(1.06)	(1.15)	(1.46)	(1.45)	(1.73)
GDP gap	−4.96*	−3.15	−4.08	5.57	6.07	5.21	2.96	2.94	3.50
	(2.76)	(2.75)	(2.70)	(6.90)	(7.02)	(6.87)	(5.33)	(4.88)	(5.45)
Labor productivity	0.70	0.49	1.06	−2.89*	−1.62	−1.58	−2.14	−2.80*	−3.72**
	(0.89)	(0.98)	(0.92)	(1.69)	(1.78)	(1.78)	(1.55)	(1.59)	(1.73)
Offshoring	4.15			11.62**			−1.88		
	(2.62)			(5.45)			(4.53)		
Long-term interest rate		−0.05			0.09			−0.03	
		(0.07)			(0.11)			(0.06)	
Low-skill wage (iw)			−0.14			−0.00			0.05
			(0.45)			(0.46)			(0.54)
High-skill wage (iw)			0.45			0.25			−0.23
			(0.39)			(0.37)			(0.46)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	47 670	47 329	47 741	47 446	46 981	47 741	47 348	35 485
Number of firms	3255	3228	3204	3252	3228	3200	3252	3224	2429

Notes: This table tests three alternative explanations. Offshoring denotes the log weighted averages of the share of foreign value added in gross value added in manufacutring. Long-term interest rate denotes the real yield on 10-year government bonds. Low-skill wages (iw) and high-skill wages (iw) compute log weighted averages of wages in the countries where the firm's inventors are located. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. Columns 7–9 use the normalized foreign variables as defined in the text. Low-skill wage (iw) and high-skill wage (iw) in Column 9 are still the total wages. Normalized offshoring is defined similarly to normalized foreign low-skill wages; normalized foreign long-term interest rate is defined like normalized foreign GDP gap. Standard errors are clustered at firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

that may bias results. Increased offshoring in the foreign country might reduce both wages and the willingness to buy automation technology. We measure offshoring at the country level as the share of foreign value-added in the gross value-added in manufacturing (Timmer et al., 2014) and compute it at the firm-level as the other macro variables. The real interest rate may be an important determinant of the cost of purchasing equipment and we control for the real yield on 10-year government bonds.[52] Labor costs could affect inventing firms through their R&D costs. We re-build our firm-specific wage variables using weights based on the location of inventors instead of patent offices and control for these inventor-location-weighted wages. Table A.31 reports the results, our coefficients on total and foreign low-skill wages remain largely stable.

Borusyak and Hull (2021). Borusyak and Hull (2021) show that a regression

[52] We obtain data for 21 countries (AT AU BE CA CH DE DK ES FI FR GB GR IE IT JP KR LU NL PT SE US) from the IMF and the OECD and deflate nominal yields using the manufacturing PPI. We compute the variable at the firm-level using patent weights for these 21 countries only.

using a logged shift-share measure may be biased due to the non-linearity of the log function. Table A.27 already shows firm-level regressions with a linear independent variable (the average of log inverse skill premium). Table A.32 implements Borusyak and Hull (2021)'s suggested correction in our default specification to remove the potential bias.[53] The results remain very similar.

Table A.32: Borusyak and Hull (2021)'s correction

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.35***	2.21***	3.83***	1.61*	2.21**	4.24***	5.18***	5.33***	3.43**
	(0.77)	(0.84)	(0.97)	(0.95)	(1.10)	(1.27)	(1.42)	(1.51)	(1.67)
High-skill wage	−1.99***	−2.22***	−0.80	−2.73***	−1.41	−1.49	−3.72***	−3.53**	−3.78***
	(0.71)	(0.77)	(0.81)	(0.96)	(1.06)	(1.05)	(1.26)	(1.58)	(1.25)
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity	No	Yes	No	No	Yes	No	No	Yes	No
GDP per capita	No	No	Yes	No	No	Yes	No	No	Yes
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table replicates the baseline regression applying the correction suggested by Borusyak and Hull (2021). We sample with replacement the entire path of log macroeconomic variables (wages, labor productivity, GDP per capita, and GDP gap) for each firm with 4000 draws, take the average value, and subtract it from the original macroeconomic variable. Significance levels at *10%, **5%, ***1%.

A.6.2 Other results and robustness checks

This Appendix presents a number of additional results. We first include additional control variables, second we consider alternative specifications (long-differences and different clustering) and third we look at alternative measures of firm-level wages and innovation.

Middle-skill wages. Lewis (2011) focuses on the effect of the low- to middle-skill ratio on the adoption of automation technologies. Table A.33 looks at the effect of middle-skill wages on automation innovations. A clear pattern emerges: low-skill wages always have a positive and significant effect, while middle-skill wages have a positive effect in regressions without low-skill wages but a negative effect otherwise. This is also in line with Graetz and Michaels (2018) who find that robots decrease the share of low-skill labor and increase the share of both high and middle-skill labor (and in contrast

[53]The correction consists in rescaling the original variables as follows: We sample with replacement the entire path of macroeconomic variables for each firm. We take the average across many draws and remove it from the original macroeconomic variables.

Table A.33: Middle-skill wages

Dependent variable	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	5.88^{***}		4.18^{***}	5.72^{***}		4.44^{**}	8.89^{***}		7.83^{**}
	(1.44)		(1.36)	(2.08)		(2.10)	(3.13)		(3.15)
Middle-skill wage	-5.01^{***}	2.78^{***}	-2.08	-4.45^{*}	2.80^{*}	-2.54	-5.63	5.02^{***}	-3.52
	(1.53)	(1.06)	(1.63)	(2.37)	(1.44)	(2.59)	(3.51)	(1.86)	(3.70)
High-skill wage		-3.06^{***}	-2.14^{**}		-2.56^{**}	-1.50		-3.39^{**}	-2.16
		(0.92)	(0.88)		(1.17)	(1.18)		(1.57)	(1.54)
GDP gap	-3.35	-4.74^{*}	-4.20	5.95	5.35	5.26	1.86	3.02	1.30
	(2.68)	(2.68)	(2.71)	(6.73)	(6.92)	(6.94)	(5.05)	(4.98)	(5.20)
Labor productivity	-0.07	1.29	0.97	-2.94^{*}	-1.24	-1.78	-2.99^{**}	-2.29	-2.13
	(0.88)	(0.91)	(0.91)	(1.62)	(1.72)	(1.77)	(1.45)	(1.67)	(1.61)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table reports the effect of middle-skill wages. All columns include firm and industry-year fixed effects. Columns 4–6 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

with the literature on IT which tends to finds more negative effects for middle-skill workers). Nevertheless, we prefer not to over-emphasize these results because low- and middle-skill wages are strongly correlated (see Table A.5).

Firm-size. Firms of different sizes may be on different trends in automation innovation. In Table A.34, we group firms into four bins according to their number of automation patents in 1995 and allow for bin-year fixed effects. We find similar results.

Table A.34: Firm bin size - year fixed effects

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	3.11***	2.83***	3.63***	2.38**	2.78**	3.75***	4.45***	5.73***	4.78***
	(0.79)	(0.84)	(0.96)	(0.98)	(1.12)	(1.26)	(1.31)	(1.55)	(1.77)
High-skill wage	−2.38***	−2.83***	−1.86**	−2.87***	−2.01*	−1.96*	−4.78***	−2.98**	−4.59***
	(0.71)	(0.78)	(0.81)	(0.95)	(1.08)	(1.04)	(1.32)	(1.48)	(1.41)
GDP gap	−2.79	−3.41	−1.60	4.39	5.46	6.74	−0.28	2.39	0.34
	(2.72)	(2.82)	(2.89)	(6.78)	(6.83)	(7.11)	(4.66)	(4.93)	(5.28)
Labor productivity		1.08			−1.96			−2.91*	
		(0.91)			(1.77)			(1.62)	
GDP per capita			−1.45			−3.36*			−0.58
			(1.33)			(1.97)			(2.08)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bin × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table controls for the size of the firms. Firms are classified into five bins by the stock of total patents in 1995 with 25th, 50th, 75th, and 95th percentiles as four thresholds. All columns include firm, industry-year and bin-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Long-difference. We now turn to alternative specifications. For most of our analysis, we follow the large patent literature and rely on a panel setting using the Poisson estimator, which best handles the count data nature of our dependent variable. In Table A.35, we conduct a long-difference estimation. To allow for zeros in the number of patents, we use the arcsinh transformation and construct ten 5-year overlapping differences from our 15 years of data. Columns (1)-(6) focus on firms that patented at least once over the period considered (now 1995-2013), mirroring what a Poisson regression would do. We find a positive effect of low-skill wages and a negative effect of high-skill wages – though, in one specification, the positive effect of low-skill wages is non-significant. The inverse skill premium, however, always has a positive and significant effect. The diminished significance of low-skill wages reflects the noisy behavior of one-time patenters and the difference in functional forms between the log function and arcsinh for low patent counts. Columns (7)-(9) restrict attention to firms that have patented at least twice and recover the same results as in our Poisson regressions. These results suggest that automation responds to medium-run changes in wages.

Table A.35: Five-year difference estimation

Dependent variable	Δ Arcsinhauto95								
Firm restriction	At least one auto95 innovation						At least two auto95 innovations		
	Domestic and Foreign				*Foreign*		*Dom. and Fgn.*		*Fgn.*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ Low-skill wage	1.14***		0.81*		1.00		2.14***	1.71**	2.45**
	(0.36)		(0.47)		(0.67)		(0.55)	(0.72)	(1.01)
Δ High-skill wage	−1.05***		−1.40***		−2.05***		−1.67***	−2.13***	−3.57***
	(0.31)		(0.45)		(0.68)		(0.47)	(0.68)	(0.99)
Δ Low-skill / High-skill wages		1.09***		1.08***		1.49**			
		(0.28)		(0.38)		(0.59)			
Δ GDP gap	−0.82	−0.86	0.87	1.00	−0.31	0.32	−1.39	0.96	−0.03
	(1.04)	(1.04)	(1.93)	(1.93)	(1.65)	(1.53)	(1.41)	(2.75)	(2.33)
Δ Labor productivity	−0.39	−0.32	0.14	−0.34	0.88	0.06	−0.67	−0.19	1.06
	(0.38)	(0.30)	(0.60)	(0.45)	(0.64)	(0.33)	(0.56)	(0.90)	(0.96)
Spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	Yes	Yes	Yes	Yes	No	Yes	Yes
Observations	32 550	32 550	32 520	32 520	32 520	32 520	21 890	21 870	21 870
Number of firms	3255	3255	3252	3252	3252	3252	2189	2187	2187

Notes: This table conducts five-year difference regressions. Estimation is done by OLS for the years t=2000-2009. The dependent variable is the difference between the arcsinh of the sum of yearly auto95 patents in t to t+4 and the arcsinh of the sum of yearly auto95 patents in t-5 to t-1. All independent variables are the sum of yearly counterparts from t-4 to t. Columns 1–6 focus on firms that have at least patented once in 1995–2013 while columns 7–9 restrict attention to firms that patented at least twice in 1995–2013. All columns include industry-year fixed effects. Columns 3–6 and 8–9 add country-year fixed effects. In Columns 3, 4, and 9 the macroeconomic variables are normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Clustering level. In the baseline specification, we cluster at the firm level to account for auto-correlation in errors. Firms that share similar weight distributions may be affected by common shocks. The best way to address this issue is through the Monte-Carlo simulations of Table 8. As an alternative, we cluster standard errors at the home country level in Table A.36. If anything, this tends to reduce the standard error on low-skill wages. A potential explanation for the negatively correlated error terms is that a successful innovator may capture the market thereby discouraging innovation by its competitors. In addition, standard errors may overstate confidence levels if the number of clusters is small or the size distribution of clusters is skewed. To address this, Table A.36 also includes p-values for low-skill wages using the BDM bootstrap-t approach of Cameron, Gelbach and Miller (2008). All coefficients of interest remain significant.

Different weights. We now turn to different measures of firm-level wages. First, we look at alternatives to pre-multiplying patent weights with $GDP^{0.35}$ (see equation (3)) in Table A.12. We either use patent weights directly, or multiply them by GDP, or by total payment to low-skill workers raised to the power of 0.35, $(w_L L)^{0.35}$. These latter weights may better measure the potential market for technology that automates

Table A.36: Baseline regressions for auto95 with country-level clustering

	Auto95								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.96***	2.71***	3.65***	2.26***	2.61***	3.70**	4.20***	5.32***	4.53**
	(0.68)	(0.75)	(1.09)	(0.72)	(0.55)	(1.57)	(0.84)	(1.64)	(1.76)
	[0.000]	[0.000]	[0.001]	[0.002]	[0.000]	[0.019]	[0.000]	[0.001]	[0.010]
	{0.018}	{0.000}	{0.001}	{0.036}	{0.057}	{0.065}	{0.015}	{0.016}	{0.007}
High-skill wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity	No	Yes	No	No	Yes	No	No	Yes	No
GDP per capita	No	No	Yes	No	No	Yes	No	No	Yes
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table reproduces the baseline table using different inference procedures. The standard errors in parentheses are clustered at country-level (instead of firm-level). The [ ] brackets report the associated p-values. To account for few clusters, the { } brackets report cluster-bootstrapped p-values following Cameron et. al (2008). Significance levels at *10%, **5%, ***1%.

low-skill work. The results remain similar.

Table A.37: Alternative weights

Dependent variable	Auto95					
Weight market size adj.	GDP^0		GDP^1		$(w_L \cdot L)^{0.35}$	
	Dom. and fgn. (1)	*Fgn.* (2)	*Dom. and fgn.* (3)	*Fgn.* (4)	*Dom. and fgn.* (5)	*Fgn.* (6)
Low-skill wage	2.74**	3.62***	2.93***	4.20***	6.10***	5.29***
	(1.10)	(1.19)	(1.12)	(1.39)	(1.70)	(1.53)
High-skill wage	−3.45***	−2.46**	−2.92***	−3.62***	−3.19**	−3.49***
	(1.06)	(1.06)	(1.03)	(1.34)	(1.62)	(1.34)
GDP gap	−5.95	1.40	−3.81	−2.43	−0.92	−0.62
	(5.15)	(5.21)	(5.38)	(3.65)	(3.89)	(3.75)
Labor productivity	0.77	0.23	−0.28	−1.61	−1.92	−2.27
	(1.53)	(1.45)	(1.58)	(1.42)	(1.59)	(1.55)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47 597	47 730	47 631	47 597	47 730	47 631
Number of firms	3249	3250	3253	3249	3250	3253

Notes: This table varies the market size adjustment in the firm's country weights. Columns 1–2 do not adjust for GDP in the computation of the weights, Columns 3–4 use GDP instead of $\text{GDP}^{0.35}$ to adjust for country size and Columns 5–6 replace GDP with total low-skilled payment $w_L * L$ in the baseline formula. All regressions include firm, country-year and industry-year fixed effects. In columns 2, 4, and 6 the macroeconomic variables are the normalized foreign variables as described in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Different deflators and wages. Second, we look at other macro measures of wages (our baseline regressions use manufacturing wages deflated by local PPI and converted in USD with the 1995 exchange rate). Table A.38 shows that our results (with foreign wages and country-year fixed effects) are robust to converting in USD yearly or in another year (2005), using a GDP deflator or replacing manufacturing wages with total wages.

Table A.38: Robustness to total wages and different deflators

Dependent variable	Auto95				
Sector	Manufacturing			Total	
Deflator	Manufacturing PPI, conversion in 2005	US manufacturing PPI, conversion every year	GDP deflator, conversion in 1995	Manufacturing PPI, conversion in 1995	US manufacturing PPI, conversion every year
	(1)	(2)	(3)	(4)	(5)
Foreign:					
Low-skill wage	5.18***	4.50***	5.18***	5.91**	5.40***
	(1.54)	(1.42)	(1.95)	(2.78)	(2.05)
High-skill wage	−2.59*	−3.60**	−2.53*	−2.50	−3.38
	(1.39)	(1.43)	(1.48)	(2.33)	(2.30)
GDP gap	2.48	1.42	2.38	0.95	0.20
	(4.86)	(4.91)	(4.92)	(4.51)	(4.65)
Labor productivity	−2.75*	−1.46	−2.73*	−3.67	−3.05
	(1.53)	(1.56)	(1.64)	(3.08)	(2.92)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	47 741	47 741	47 741	47 741	47 741
Number of firms	3252	3252	3252	3252	3252

Notes: This table shows robustness to different wage conversions. Columns 1–3 use manufacturing wages and columns 4 and 5 total wages. In column 1, macroeconomic variables are deflated with the local manufacturing PPI and converted to USD in 2005. In Columns 2 and 5 they are converted to USD every year and deflated with the US manufacturing PPI. In Column 3, macroeconomic variables are deflated with the local GDP deflator and converted to USD in 1995. In Column 4, macroeconomic variables are deflated with the local manufacturing PPI and converted to USD in 1995. All regressions include firm fixed effects, industry-year fixed effects and country-year fixed effects. In all columns, the macroeconomic variables are the normalized foreign variables as defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Citations. Finally, we look at other measures of innovation. Table A.39 investigates whether our results are robust to focusing on patents of higher quality and weighs patents by citations. We add to each patent the number of citations received within 5 years normalized by technology field, patent office and year of application, and winsorized at the 75^{th} percentile. We find similar coefficients as in the baseline, which shows that our results are not driven by low-quality innovations.[54]

[54] If we do not winsorize the patent counts at the 75^{th} percentile, we lose significance in columns (4) and (5). The number of citations is quite right-skewed and one possible interpretation is that conditional on R&D investment, whether an innovation turns out to be of very high quality is largely random. This dampens the effect of low-skill wages on (non-winsorized) citations-weighted patents.

Table A.39: Citations-weighted patents

Dependent variable	Citations-weighted auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.50***	2.27**	3.22***	1.79*	2.21*	3.30**	4.04***	5.16***	4.28**
	(0.85)	(0.91)	(1.03)	(1.07)	(1.25)	(1.37)	(1.39)	(1.56)	(1.80)
High-skill wage	−2.25***	−2.64***	−1.52*	−3.00***	−2.10*	−1.98*	−4.79***	−3.23**	−4.65***
	(0.82)	(0.86)	(0.88)	(1.08)	(1.12)	(1.17)	(1.34)	(1.54)	(1.47)
GDP gap	−3.38	−3.89	−1.73	3.28	4.42	5.90	−1.61	0.77	−1.15
	(2.62)	(2.70)	(2.78)	(6.62)	(6.61)	(6.92)	(4.41)	(4.71)	(5.01)
Labor productivity		0.90			−2.06			−2.53	
		(0.95)			(1.85)			(1.61)	
GDP per capita			−2.01			−3.75*			−0.42
			(1.36)			(2.06)			(2.19)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48 091	48 091	48 091	47 741	47 741	47 741	47 741	47 741	47 741
Number of firms	3255	3255	3255	3252	3252	3252	3252	3252	3252

Notes: This table weighs patents by citations. We add to each auto95 patent the number of citations received within 5 years normalized by technological field, patent office, and year of application, and winsorized at the 75th percentile. All columns include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. In Columns 7–9 the macroeconomic variables are the normalized foreign variables defined in the text. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Innovation types. We look at other definitions or subcategories of automation innovations in regressions with foreign wages in Table A.40. The results are robust to excluding the codes that we added to the definition of the machinery technology field listed in footnote 11. Though the coefficients are smaller, they are also robust to using the laxer auto80 definition of automation innovations. Subcategories of automation innovations are defined by re-classifying codes according to the prevalence of each category of automation keywords. We find large effects of low-skill wages on automat* and robot patents; but no significant effect on CNC patents, for which the sample size is smaller.

A.6.3 Computing automation elasticities from the literature

In this Appendix, we explain how we compute the elasticities reported in Section 4.4. Lewis (2011) identifies low-skill workers as high-school dropouts and middle-skill workers as high-school graduate, which does not align with our analysis. Nevertheless, he estimates that a 1 point increase in the ratio of low- to middle-skill workers decreases the number of technologies adopted by 7.75 (Table V, column 2), decreases $\ln(w_L/w_M)$ by 0.199 (Table VIII, column 2) and increases $\ln(w_H/w_M)$ by 0.474 (Table VIII, column 5), so that $\ln(w_L/w_H)$ decreases by -0.673. The mean number of adopted technologies

Table A.40: Innovation categories

Dependent variable	Auto95	AutoX95	Auto80	Automat*90	Automat*80	Robot90	Robot80	CNC90	CNC80
	(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Foreign:									
Low-skill wage	5.32^{***}	5.42^{***}	3.56^{***}	7.48^{***}	5.94^{***}	5.92^{*}	7.49^{***}	-1.48	-1.56
	(1.56)	(1.62)	(1.32)	(2.10)	(1.96)	(3.32)	(2.54)	(4.08)	(3.05)
High-skill wage	-2.87^{*}	-1.42	-2.16	-2.49	-2.09	0.51	-3.06	5.52	1.75
	(1.47)	(1.63)	(1.32)	(1.90)	(1.77)	(3.01)	(2.37)	(5.58)	(3.61)
GDP gap	2.28	0.62	1.91	8.31^{*}	4.23	6.03	1.22	-1.69	-1.17
	(4.92)	(4.60)	(2.85)	(4.93)	(4.42)	(8.15)	(6.79)	(12.03)	(9.68)
Labor productivity	-2.57	-3.87^{**}	-1.78	-5.41^{***}	-4.49^{**}	-7.65^{***}	-5.70^{**}	-4.43	-1.03
	(1.60)	(1.71)	(1.22)	(1.82)	(1.75)	(2.81)	(2.25)	(5.19)	(3.25)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47 741	45 928	97 705	32 424	48 900	15 831	23 268	7080	13 617
Number of firms	3252	3150	6561	2244	3329	1151	1632	547	1001

Notes: This table analyzes the effect of wages on different automation innovation categories. AutoX95 excludes the C/IPC codes which we added when defining the machinery technological field. Auto80 lowers the threshold to define automation innovation to the 80th percentile of the C/IPC 6-digit distribution. Automat*90 and Automat*80 only count words associated with automat. Robot90 and Robot80 only count words associated with robot. CNC90 and CNC80 words associated with CNC. 90 and 80 refer to the thresholds used to define the corresponding technology categories, which are the 90th and 80th percentile of the distribution of automation keywords for 6-digit C/IPC codes. The macroeconomic variables are the normalized foreign variables as defined in the text. Stocks and spillovers are computed with respect to the dependent variable. All regressions include firm fixed effects, industry-year, and country-year fixed effects. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

is 3.09, while the mean change in the ratio of low- to middle-skill workers is -0.03 (Table I). From this, we can back an elasticity of automation adoption with respect to the inverse skill premium of $[\ln(3.09 - 7.75 \times 0.03) - \ln(3.09)] / [0.673 \times 0.03] = 3.6$.

Acemoglu and Restrepo (2022) measure aging as the predicted change in the ratio of above 56 to below 56 workers between 1995 and 2025. They find that aging leads to an increase in the log ratio of robot imports over all intermediate imports of 1.96 (Table 4, column 3) and an increase in the log number of robotics over all patents of 0.75 (Table 5, column 3). They also report that aging between 1990 and 2015 is associated with a relative increase of blue-collar manufacturing wages compared to average wages of 0.418 across US commuting zones (Table A.20, Panel B, column 4). Taking ratios and adjusting for the different time lengths gives an elasticity of $\frac{1.96}{0.418}\frac{25}{30} = 3.9$ for adoption and $\frac{0.75}{0.418}\frac{25}{30} = 1.5$ for innovation.

Finally, we report on elasticities in the adoption of new technologies in footnote 41. Baptista (2000) studies the adoption of CNC machines in the UK. He estimates the effect of the number of previous adopters in an area on the hazard rate of adoption. Using the coefficient from their Table 3 and the mean number of adopters from Table 2, one gets that a 1% increase in the number of local adopters reduces the hazard rate of adoption by 0.08%. No (2008) looks at the adoption of advanced manufacturing technologies in Canada and reports elasticities with respect to the number of previous adopters in similar industries between 0.0012 and 0.0015 (their Table 3 and 4). Finally, Bekes and Harasztozi (2020) shows that Hungarian firms are more likely to import specific machines when a nearby peer already imports the same machine. Combining the coefficient of Table 7 (0.003—coefficients in the table are multiplied by 100) with the probability that there is a peer (0.2 from their Table 4) and an average hazard rate of importing of 1%, we get an elasticity of 0.06.

A.6.4 Macroeconomic interpretation of the regression coefficients

This section provides details on the simulation results of Section 4.4. Table A.41 shows the exact regression that supports our simulations. We jointly estimate the effect of the inverse skill premium on auto95 and pauto95 innovations (without restricting attention to the sample of firms of the baseline regression). This requires that we compute separately the stocks and spillovers of auto95 innovations, pauto95 innovations and non-machinery innovations. We also include quadratic terms for the knowledge spillovers.[55]

Recomputing the spillover variables involves two complications. First, our model applies only to the number of innovations, not their location. To allocate innovations to countries, we assign the simulated innovations proportionally to contemporaneous inventor weights of the firms (while the spillover variables are computed using predetermined inventor weights). These contemporaneous weights reflect the distribution of where firms' innovators are located in the respective year (or the closest year if there's no patenting).

Second, our regression dataset does not include all firms with biadic innovations but our spillover variables are computed using country-level stocks of biadic innovations. To account for this, we assume that out-of-sample firms respond similarly to in-sample firms. When assigning simulated innovations to countries, we increase the innovations by those of out-of-sample firms so that the ratio of in-sample to out-of-sample innovations in that country-year remains the same as in the data. We make this adjustment for countries with at least 10 in-sample machinery patents.

[55] The coefficients on knowledge spillovers in log linear regressions are greater than 1 leading to an explosive behavior. Coefficients on the knowledge spillover squares are significant which justifies the inclusion of the square terms. We also use ln(1+) to compute stocks and spillovers in this exercise. This has no effect on the regression results but ensures a more stable behavior in the simulations.

Table A.41: Regression supporting the simulation of Figure 4

Dependent variable	Auto95	Pauto95
	(1)	(2)
Low-skill / High-skill wages	2.51***	0.44
	(0.69)	(0.53)
Stock automation	−0.15***	0.13***
	(0.05)	(0.03)
Stock non-automation	0.34***	0.27***
	(0.06)	(0.03)
Spillovers automation	2.24**	−1.01
	(0.96)	(0.63)
Spillovers automation squared	−0.10*	0.04
	(0.06)	(0.04)
Spillovers non-automation	4.44*	4.42***
	(2.32)	(1.50)
Spillovers non-automation squared	−0.20*	−0.13
	(0.12)	(0.08)
GDP gap	Yes	Yes
Non-machinery stock	Yes	Yes
Non-machinery spillovers	Yes	Yes
Non-machinery spillovers squared	Yes	Yes
Firm fixed effects	Yes	Yes
Industry × year fixed effects	Yes	Yes
Country × year fixed effects	No	No
Observations	48 091	155 183
Number of firms	3255	10 382

Notes: This table shows regressions of automation (column 1) and non-automation machinery innovations (column 2) on the inverse skill-premium, the GDP gap, and firm-level stock and spillover variables. We consider automation, non-automation, and non-machinery stocks and spillovers separately and include squared spillovers. Stocks and spillovers are computed as log(1+). The regressions include firm and year-industry fixed effects. Standard errors are clustered at the firm-level and reported in parentheses. Significance levels at *10%, **5%, ***1%.

Even without any change in the skill premium, the noise in the Poisson process means that the exact number of patents in each country can vary from one simulation to the next. If the spillover variables are kept as in the data, the average effect of this noise is null, and the average simulation (with no change in the skill premium) looks exactly like the data series. However, when the spillover variables at time t are updated to reflect the simulated innovations in the years before $t-2$, the predicted number of innovations at t may be different fron that in the data. This is why the baseline curve in Figure 4 differs from the data series, especially toward the end of the sample. And this is also why the total effect of the change in the skill premium should be computed as the difference between the baseline + total effect curve and the baseline curve. Figure 4 displays the median simulation but the mean looks similar.

Online Appendix References

Dolado, J., Kramarz, F., Machin, S., Manning, A., Margolis, D., and Teulings, C. (1996). The Economic Impact of Minimum Wages in Europe. *Economic Policy* 11 (23).

Dickens, R., Machin, S., and Manning, A. (1999). The Effects of Minimum Wages on Employment: Theory and Evidence from Britain. *Journal of Labor Economics,* 17 (1).

Lybbert, T. and Zolas, N. (2014). Getting patents and economic data to speak to each other: An 'Algorithmic Links with Probabilities' approach for joint analyses of patenting and economic activity. *Research Policy,* 43: 530-542.

Scopp, T. (2003). The relationship between the 1990 Census and Census 2000 Industry and Occupation Classification Systems. US Census Bureau.

B Supplemental material

B.1 Additional examples

We provide a few additional examples of automation and non-automation patents. Figure B.1 shows the example of a robot with a patent containing the IPC code B25J9. The patent describes a multi-axis robot with a plurality of tools which can change the working range of each arm. This essentially increases the flexibility of the robot. Figure B.2 shows an automation innovation used in the dairy industry. The patent contains the code A01J7 which is a high automation code (see Table A.21). It describes a system involving a robotic arm to disinfect the teats of cows after milking. The patent argues that this reduces the need for human labor and therefore saves costs. Figure B.3 describes an automated machining device – yet another example of a high automation innovation – which contains the code B23Q15 (a high automation code described in Table A.21). The devices features a built-in compensation system to correct for errors thereby reducing the need for a "labor-intensive adjustment process". Figure B.4 describes another high automation patent belonging to the same IPC code as well as to G05B19. This is also a machining device. The patent explains that innovations in machining have aimed at making the process as automated as possible by involving some feedback mechanism (as in the previous older patent). This invention aims at better predicting the machining requirements in the first place.

In contrast, Figure B.5 describes a low automation innovation in machinery (none of the codes are above the 90th percentile in the 6-digit C/IPC distribution). The innovation relates to a "conveying belt assembly for a printing device"', which is about the circulation of paper in the printing machine. This innovation does not directly involve automation. Similarly Figure B.6 describes a winch to raise and lower people, another low-automation innovation in machinery. This innovation seems rather low-skill labor complementary as its goal is to enable workers to move in a plurality of directions. Finally, Figure B.7 describes a harvester (which also counts as a machinery innovation since the code A01B63 belongs to other special machinery). This is also a low-automation innovation as its goal is to ensure that the harvester can both operate in the field and travel on roads.

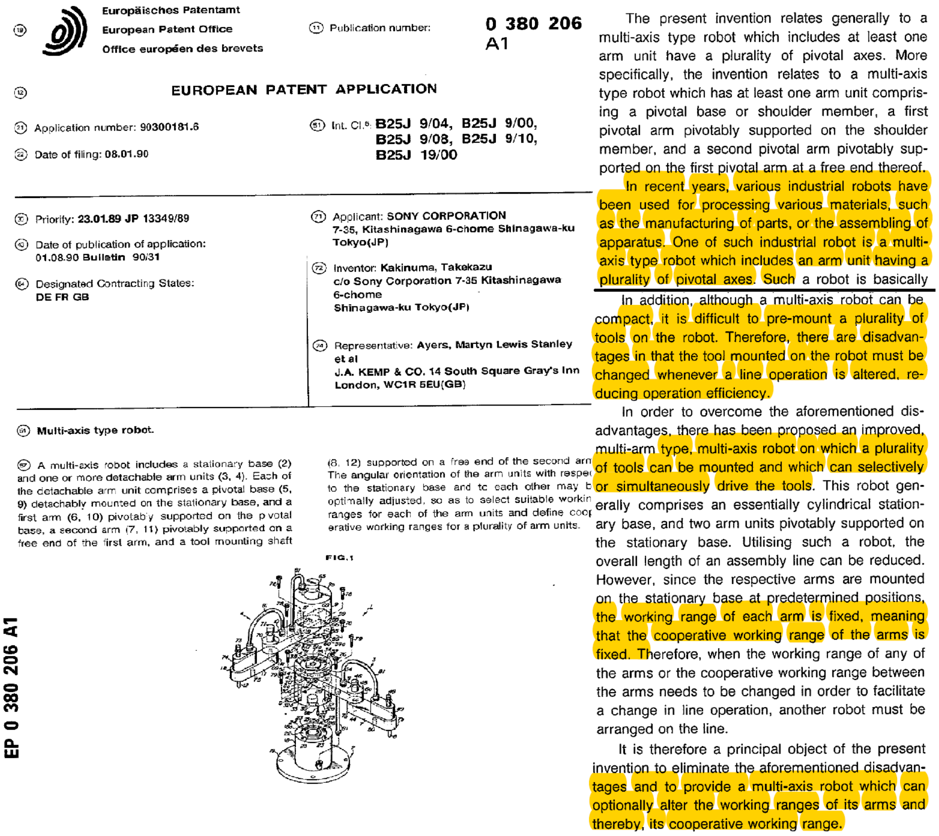

Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number: **0 380 206 A1**

(12) **EUROPEAN PATENT APPLICATION**

(21) Application number: 90300181.6

(22) Date of filing: 08.01.90

(51) Int. Cl.5: **B25J 9/04, B25J 9/00, B25J 9/08, B25J 9/10, B25J 19/00**

(30) Priority: **23.01.89 JP 13349/89**

(43) Date of publication of application: 01.08.90 Bulletin 90/31

(84) Designated Contracting States: DE FR GB

(71) Applicant: SONY CORPORATION
7-35, Kitashinagawa 6-chome Shinagawa-ku Tokyo(JP)

(72) Inventor: Kakinuma, Takekazu
c/o Sony Corporation 7-35 Kitashinagawa 6-chome
Shinagawa-ku Tokyo(JP)

(74) Representative: Ayers, Martyn Lewis Stanley et al
J.A. KEMP & CO. 14 South Square Gray's Inn London, WC1R 5EU(GB)

(54) **Multi-axis type robot.**

(57) A multi-axis robot includes a stationary base (2) and one or more detachable arm units (3, 4). Each of the detachable arm unit comprises a pivotal base (5, 9) detachably mounted on the stationary base, and a first arm (6, 10) pivotably supported on the pivotal base, a second arm (7, 11) pivotably supported on a free end of the first arm, and a tool mounting shaft (8, 12) supported on a free end of the second arm The angular orientation of the arm units with respec to the stationary base and to each other may b optimally adjusted, so as to select suitable workin ranges for each of the arm units and define coop erative working ranges for a plurality of arm units.

FIG.1

EP 0 380 206 A1

The present invention relates generally to a multi-axis type robot which includes at least one arm unit have a plurality of pivotal axes. More specifically, the invention relates to a multi-axis type robot which has at least one arm unit comprising a pivotal base or shoulder member, a first pivotal arm pivotably supported on the shoulder member, and a second pivotal arm pivotably supported on the first pivotal arm at a free end thereof.

In recent years, various industrial robots have been used for processing various materials, such as the manufacturing of parts, or the assembling of apparatus. One of such industrial robot is a multi-axis type robot which includes an arm unit having a plurality of pivotal axes. Such a robot is basically

In addition, although a multi-axis robot can be compact, it is difficult to pre-mount a plurality of tools on the robot. Therefore, there are disadvantages in that the tool mounted on the robot must be changed whenever a line operation is altered, reducing operation efficiency.

In order to overcome the aforementioned disadvantages, there has been proposed an improved, multi-arm type, multi-axis robot on which a plurality of tools can be mounted and which can selectively or simultaneously drive the tools. This robot generally comprises an essentially cylindrical stationary base, and two arm units pivotably supported on the stationary base. Utilising such a robot, the overall length of an assembly line can be reduced. However, since the respective arms are mounted on the stationary base at predetermined positions, the working range of each arm is fixed, meaning that the cooperative working range of the arms is fixed. Therefore, when the working range of any of the arms or the cooperative working range between the arms needs to be changed in order to facilitate a change in line operation, another robot must be arranged on the line.

It is therefore a principal object of the present invention to eliminate the aforementioned disadvantages and to provide a multi-axis robot which can optionally alter the working ranges of its arms and thereby, its cooperative working range.

Figure B.1: Example of a high automation patent: an industrial robot

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) **EP 3 300 593 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
04.04.2018 Bulletin 2018/14

(21) Application number: **17198024.6**

(22) Date of filing: **12.08.2011**

(51) Int Cl.:
A01K 1/12 *(2006.01)* ***A01J 5/00*** *(2006.01)*
A01J 7/04 *(2006.01)* ***A01J 5/003*** *(2006.01)*

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

(30) Priority: **31.08.2010 US 378871 P**
28.04.2011 US 201113095963

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
11746122.8 / 2 611 285

(71) Applicant: **Technologies Holdings Corporation**
Houston, TX 77019 (US)

(72) Inventors:
- **HOFMAN, Henk**
NL-8531 PC Lemmer (NL)
- **VAN DER SLUIS, Peter, William**
8271 PP IJsselmuiden (NL)
- **GROENSMA, Yep**
8441 CA Ca Heerenveen (NL)

(74) Representative: **Moore, Derek**
Jensen & Son
366-368 Old Street
London EC1V 9LT (GB)

Remarks:
This application was filed on 24-10-2017 as a divisional application to the application mentioned under INID code 62.

(54) **METHOD AND AUTOMATED SYSTEM FOR APPLYING DISINFECTANT TO THE TEATS OF DAIRY LIVESTOCK**

SUMMARY OF THE INVENTION

[0003] According to embodiments of the present disclosure, disadvantages and problems associated with previous systems supporting dairy milking operations may be reduced or eliminated.

[0004] In certain embodiments, a system for applying disinfectant to the teats of a dairy livestock includes a carriage mounted on a track, the carriage operable to translate laterally along the track. The system further includes a robotic arm including a first member pivotally attached to the carriage such that the first member may rotate about a point of attachment to the carriage, a second member pivotally attached to the first member such that the second member may rotate about a point of attachment to the first member, and a spray tool member pivotally attached to the second member such that the spray tool member may rotate about a point of attachment to the second member. The system further includes a controller operable to cause at least a portion of the robotic arm to extend between the hind legs of a dairy livestock such that a spray tool of the spray tool member is located at a spray position from which the spray tool may discharge an amount of disinfectant to the teats of the dairy livestock.

[0005] Particular embodiments of the present disclosure may provide one or more technical advantages. For example, certain embodiments of the present disclosure may provide an automated system for applying disinfectant to the teats of dairy livestock. Additionally, certain embodiments of the present disclosure may minimize overspray, thereby reducing the volume of the disinfectant needed. By reducing the need for human labor and reducing the volume of disinfectant used, certain embodiments of the present disclosure may reduce the cost associated with applying disinfectant to the teats of dairy livestock in certain dairy milking operations. Furthermore, the use of the automated system of the present disclosure in conjunction with a rotary milking platform may increase the throughput of the milking platform, thereby increasing the overall milk production of the milking platform.

Figure B.2: Example of a high automation patent: a milking robot

⑲ Europäisches Patentamt
European Patent Office
Office européen des brevets

⑪ Publication number: **0 412 635 A2**

⑫ **EUROPEAN PATENT APPLICATION**

㉑ Application number: **90305164.7**

㉒ Date of filing: **14.05.90**

㉛ Int. Cl.5: **B23Q 15/16**, B23Q 15/18

㉚ Priority: **10.08.89 US 391929**

㊸ Date of publication of application: **13.02.91 Bulletin 91/07**

㊹ Designated Contracting States: **DE ES FR GB IT**

㉑ Applicant: **CINCINNATI MILACRON INC.**
4701 Marburg Avenue
Cincinnati Ohio 45209(US)

㉒ Inventor: **Wood, David B. III**
106 Sherwood Green Court
Mason, Ohio 45040(US)

㉔ Representative: **Carpmael, John William Maurice et al**
CARPMAELS & RANSFORD 43 Bloomsbury Square
London, WC1A 2RA(GB)

㉔ **High production machining device.**

TECHNICAL FIELD

This invention relates to a high-productivity, twin-spindle turning center featuring a built-in compensation system to correct for processing errors, and, more particularly, to an improved two-spindle machining device having a built-in tool compensation system which provides for individual process control for each spindle.

Heretofore, the industry has attempted to address the problems of these inherent errors by measuring resulting parts and assigning offset errors which can be compensated for by providing adjustable tool blocks, or by undertaking tedious shimming operations of the tools themselves. Often a machinist had no other choice but to average the errors between the two tools, and attempt to adjust the tools and/or tool blocks to compensate. Once these initial errors were reduced sufficiently as a result of such labor-intensive adjustment procedures, it was often necessary to slow the turning process down to reserve tool life and, thereby, delay the tedious process of replacing worn tools as long as possible. Such compromise directly undermined productivity levels, and the process of averaging errors does not generally yield part accuracies which are competitive with the quality of parts made on single-spindle machines, let alone achieving the higher level of accuracy demanded in this industry.

Consequently, heretofore, there has not been available a reliable, low-cost, built-in tool compensating system for lathe machines. Moreover, compensation systems previously available could not effectively provide a multi-spindle machine tool wherein individual process control for each spindle was possible. While multi-spindle machines have been available for quite some time, there has not been presented a compensation system which can consistently maintain high production rates on each spindle in a relatively simple and efficient manner.

Figure B.3: Example of a high automation patent: an automated machining device

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 913 229 B1

(12) EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent: **19.01.2005 Bulletin 2005/03**

(21) Application number: **98907226.9**

(22) Date of filing: **13.03.1998**

(51) Int Cl.[7]: **B23Q 15/00**, G05B 19/4093

(86) International application number: **PCT/JP1998/001074**

(87) International publication number: **WO 1998/041357 (24.09.1998 Gazette 1998/38)**

(54) **MACHINING PROCESSOR**
PROZESSOR FÜR MASCHINELLE BEARBEITUNG
PROCESSEUR D'USINAGE

(84) Designated Contracting States: **AT CH DE FR GB IT LI SE**

(30) Priority: **15.03.1997 JP 8219497**

(43) Date of publication of application: **06.05.1999 Bulletin 1999/18**

(73) Proprietor: **MAKINO MILLING MACHINE CO. LTD. Meguro-ku, Tokyo (JP)**

(72) Inventors:
- **YOSHIDA, Jun-Makino Milling Machine Co., Ltd. Kanagawa 243-0308 (JP)**
- **KAWANA, Akira Makino Milling Machine Co., Ltd. Kanagawa 243-0308 (JP)**
- **INOUE, Shinichi Makino Milling Machine Co., Ltd. Kanagawa 243-0308 (JP)**
- **HISAKI, Tatsuya Makino Milling Machine Co., Ltd. Kanagawa 243-0308 (JP)**

(74) Representative: **Bibby, William Mark Mathisen, Macara & Co., The Coach House, 6-8 Swakeleys Road Ickenham Uxbridge UB10 8BZ (GB)**

(56) References cited:

EP-A- 0 753 805	JP-A- 1 205 954
JP-A- 2 178 711	JP-A- 3 251 907
JP-A- 3 294 146	JP-A- 4 283 047
JP-A- 4 284 507	JP-A- 5 077 138
JP-A- 6 102 923	JP-A- 6 119 029
JP-A- 6 138 929	JP-A- 6 170 694
JP-A- 8 132 332	JP-A- 62 140 741
JP-A- 62 241 635	JP-U- 5 008 604
US-A- 4 837 703	

[0001] The present invention relates to a control apparatus for a machine tool and a machining system comprising the control apparatus and a machine tool wherein, by supplying a raw workpiece and inputting data regarding a machining profile of a final product (hereinafter referred to as machining profile data), the workpiece to be machined is machined according to the machining profile data so that a final product can be fabricated.

Background Art

[0002] In the conventional method of machining a workpiece by a NC machine tool, the first step is to prepare a drawing representing the profile of a product to be machined. A programmer determines the machining steps from the drawing and creates a NC program manually or by an automatic programming unit. An operator inputs the NC program into the NC machine tool while, at the same time, setting up the workpiece on the NC machine tool manually or by using an automatic workpiece changer. Then, the cutting tool to be used is preset, and the amount of tool offset is defined. The cutting tool is then mounted in the tool magazine of the NC machine tool. After that, the NC program is executed thereby to machine the workpiece and fabricate a product. Various inventions have hitherto been developed with the aim of automating these steps as far as possible and reflecting the know-how accumulated by programmers and operators on the machining steps.

[0008] These conventional techniques are based on the architecture of securing a high accuracy and a high production efficiency by feedback correction of the machining conditions, but not intended to realize a high-accuracy, high-efficiency machining process by predicting machining requirements and determining a tool path and machining conditions based on the prediction.

[0010] An object of the present invention is to provide a machine tool control apparatus and a machining system including the control apparatus and a machine tool, in which an intended product can be automatically machined at high efficiency while meeting the precision requirements in response to only profile data on the product to be finished and data on the workpiece to be machined.

Figure B.4: Example of a high automation patent: another automated machining device

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 2 990 363 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication: **02.03.2016 Bulletin 2016/09**

(21) Application number: **15181736.8**

(22) Date of filing: **20.08.2015**

(51) Int Cl.: ***B65H 5/02*** (2006.01) ***B65H 7/10*** (2006.01) ***B41J 11/00*** (2006.01)

(84) Designated Contracting States: **AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR**
Designated Extension States: **BA ME**
Designated Validation States: **MA**

(30) Priority: **26.08.2014 EP 14182318**

(71) Applicant: **OCE-Technologies B.V. 5914 CA Venlo (NL)**

(72) Inventor: **ALBERS, Antonius G.H. 5914 CA Venlo (NL)**

(74) Representative: **Cornelissen, Leandra Océ-Technologies B.V. Corporate Patents P.O. Box 101 5900 MA Venlo (NL)**

(54) **CONVEYING BELT ASSEMBLY FOR A PRINTING DEVICE**

Description

[0001] The present invention relates to a conveying belt assembly for a printing device, a method for controlling the position of an endless conveyor belt, and the use of a conveying belt assembly.

[0002] In printing devices conveying belts are used to transport a sheet of paper through the printing device. The sheet of paper transported through the printing device requires high accuracy in control of its position.

[0003] The present invention has as its object to provide a conveying belt assembly for a printing device, which conveying belt assembly allows very accurate control of the position of the conveyor belt or its conveying part.

Figure B.5: Example of a low automation patent: a printer

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) **EP 1 452 478 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
01.09.2004 Bulletin 2004/36

(51) Int Cl.7: **B66D 3/22**, B66D 3/26, A61G 7/10

(21) Application number: **03004482.0**

(22) Date of filing: **28.02.2003**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LI LU MC NL PT SE SI SK TR
Designated Extension States:
AL LT LV MK RO

(71) Applicant: **ERGOLET A/S**
4220 Korsor (DK)

(72) Inventor: **Hjort, Mogens**
4220 Korsor (DK)

(74) Representative:
van Walstijn, Bartholomeus Gerard G.
Walstijn Intellectual Property ApS
Parkovsvej 3
2820 Gentofte (DK)

(54) **A winch for raising and lowering persons**

[0001] The present invention relates to a winch for raising and lowering persons, comprising a housing provided with a first attachment member, a first opening formed in the housing substantially opposite to the first attachment member, an electric motor coupled to the input of a reduction gearing, a reel component coupled to the output of the reduction gearing, and a flexible elongated traction member connected to the reel component for winding and unwinding the traction member for raising and lowering a person. Further, the invention relates to the use of a winch according to the invention as a ceiling lift. The invention also relates to a ceiling lift assembly, comprising an overhead rail with at least one carriage guided therein, the carriage being provided with an attachment member, a winch provided with at least one attachment member on the winch housing and the winch comprising a flexible elongated traction member with an attachment member on its free end and a spreader bar with an attachment member.

[0004] Against this background, it is an object of the present invention to provide a winch of the kind referred to initially, which overcomes or at least reduces the above mentioned problems by allowing it to operate in a plurality of orientations. This object is achieved in accordance with claim 1 by providing a winch of said kind with the housing having a second opening so that the traction member can be guided through the first opening or through the second opening.

[0005] Thus, it becomes possible to operate the winch in more orientations.

Figure B.6: Example of a low automation patent: a winch

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) **EP 1 226 745 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
31.07.2002 Bulletin 2002/31

(51) Int Cl.7: **A01B 63/00**, A01B 73/00

(21) Application number: **02075380.2**

(22) Date of filing: **28.01.2002**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **29.01.2001 NL 1017208**

(71) Applicant: **Frans Vervaet B.V.**
4521 PE Biervliet (NL)

(72) Inventors:
- **Vervaet, Edwin Joseph Germain**
4521 PT Biervliet (NL)
- **Poppe, Bertus Marinus**
4365 NG Meliskerke (NL)
- **Vervaet, Robin Richard**
4521 PE Biervliet (NL)

(74) Representative: **Eveleens Maarse, Pieter**
Arnold & Siedsma,
Advocaten en Octrooigemachtigden,
Sweelinckplein 1
2517 GK Den Haag (NL)

(54) **Harvester**

Description

[0001] The invention relates to an agricultural machine provided with at least one pair of wheels and at least one wheel for performing operations on the land.

[0002] Such agricultural machines are generally known.

[0003] In order to perform operations on the land, it is important for stability that the wheels are placed far apart, while for travel without performing operations it is important that the wheels are placed closer together to improve the quality of travel.

[0004] The object of the invention is to provide a machine which can meet both requirements.

Figure B.7: Example of a low automation patent: a harvester

B.2 Validating our weights approach

We compare our firm-level weights to bilateral trade flows and show that they are strongly correlated. The first step is to compute patent-based weights at the country level. For this exercise (and this exercise only), we define the home country D of a firm based on the location of its headquarters according to the country code of its

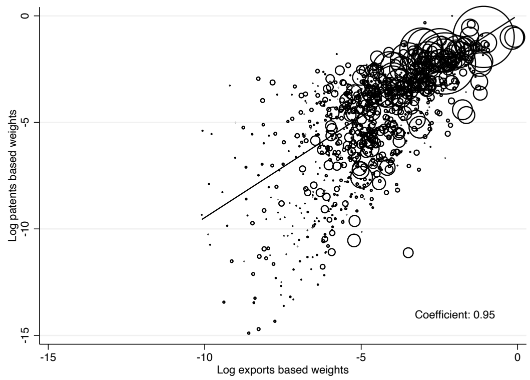

(a) All countries

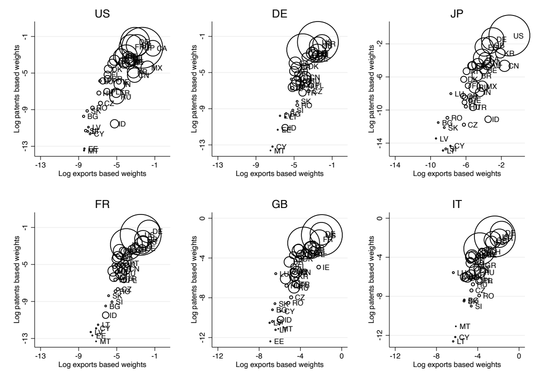

(b) Trade from the 6 largest countries

Figure B.8: Bilateral patent flows and trade flows in machinery.
Notes: Panel (a) plots log patent based weights, which are a weighted average of the destination country's weights in the (foreign) patent portfolio of firms from the origin country, against export shares in machinery over the years 1995-2009. The size of each circle represents the product of the GDP of both countries, which is used as a weight in the regression. Panel (b) focuses on the weights from the listed countries and observations are weighted by the GDP of the partner country.

identifier in the Orbis database. For firms which we merged, we keep the country code of the largest entity by biadic machinery patents in 1997-2011. We compute the foreign weights for each firm i by excluding the home country. Therefore, the foreign weight for country $c \neq D$ for firm i is given by $\kappa_{i,c}/(1-\kappa_{i,D})$ (recall that these weights are computed based on patenting from 1971 to 1994). We then build the foreign patent-based weight in country c for country D as a weighted average of the foreign weights in country c of the firms from country D, where each firm is weighted according to the number of machinery biadic patents in 1997-2011.

The second step is to build similar weights based on exports. To do that, we collect sectoral bilateral trade flow from UN Comtrade data between between 1995 and 2009 for 40 countries (Taiwan is not included in the data). To obtain trade flows in machinery, we use the Eurostat concordance table between 4-digit IPC codes and 2 or 3-digits NACE Rev 2 codes (van Looy, Vereyen, and Schmoch, 2014): this concordance table matches IPC codes to the industry of manufacturing. The concordance table assigns a unique industry to each IPC code. Then, for each industry, we compute the share of biadic patents over the period 1995-2009 that are in machinery according to our definition.[56] This gives us a machinery weight for each industry code and each country. We then multiply sectoral trade flows (after having aggregated the original data to the NACE

[56]To do that we use a fractional approach: each patent is allocated NACE sectoral weights (and machinery weights) depending on the share of IPC codes associated with a NACE sector or machinery.

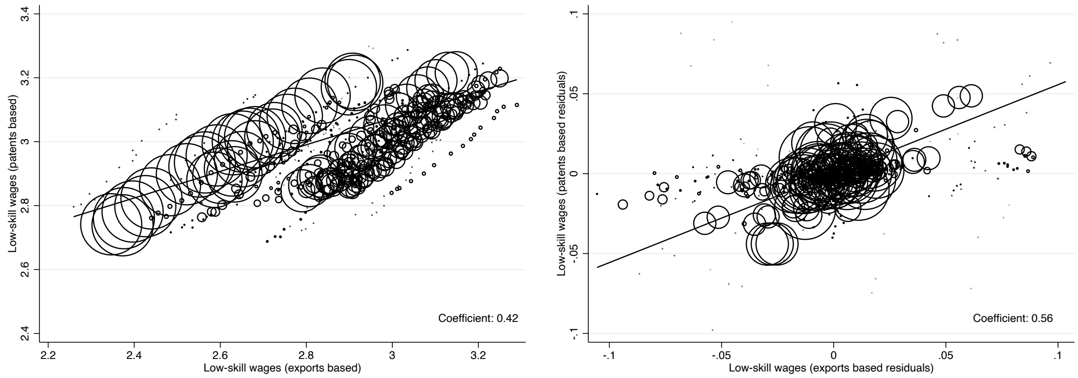

(a) Low-skill wages. **(b)** Residualized low-skill wages

Figure B.9: Foreign low-skill wages for each country computed either with patent-based weights or with trade-based weights.

Notes: Wages are computed for the years 1995-2009. Panel (a) plots log foreign low-skill wages using either patent-based weights or trade-based weights. Panel (b) plots the residuals of foreign wages according to both methods controlling for country and year fixed effects. Observations are weighted by the number of biadic machinery patents by firms from the country over the years 1997-2011.

Rev 2 codes used in the concordance table) by this weight to get bilateral trade in machinery. We then compute the export share in machinery across destinations. We compute trade based weights for each year in 1995-2009 and take the average (there are a few missing observations for 1995).

Figure B.8 plots the patent-based weights against the trade-based weights. Panel (b) focuses on a few origin countries while Panel (a) plots all countries together. We find a strong correlation between the two measures with a regression coefficient of 0.94 (when observations are weighted by the trade flow in 1996).

Figure B.9 goes further and compares low-skill wages computed with either sets of weights. For each country, we compute "foreign low-skill wages" as a weighted average of foreign wages where the weights are either the patent-based weights or the trade-based weights derived above. Foreign wages are deflated with the local PPI and converted in USD in 1995 as in our main analysis. Panel (a) then reports foreign log low-skill wages according to both types of weights in 1995-2009 and finds that they are strongly correlated. Panel (b) reports the same foreign log low-skill wages but taking away country and year fixed effects. The regression coefficient is 0.56, when observations are weighed by the number of machinery patents in the country between 1997 and 2011.

Overall, this exercise shows that there is tight relationship between our patent-based weights and (future) trade flows, suggesting that we can use these patent-based weights as proxies for firms' markets exposure.

References

van Looy, B., Vereyen, C., and Schmoch, U. (2014). Patent Statistics: Concordance IPC V8 - NACE REV.2. Eurostat.