Induced Automation: Evidence from Firm-level Patent Data

Antoine Dechezleprêtre David Hémous Morten Olsen

Carlo Zanella*

August 2022 (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, Morten Olsen, University of Copenhagen and Carlo Zanella, University of Zurich. 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 innovations? And if so, by how much? Automation technologies accelerate and political campaigns push for higher minimum wages. Economic theory, however, suggests that firms innovate more in automation technology when labor costs increase. This would affect the long-term effects of such policies. But, while the theoretical argument is well-understood, empirical evidence for induced automation is lacking. Current research faces two challenges: identifying automation innovations and finding exogenous variation in labor costs. Accordingly, our paper makes two contributions: we develop a new classification of automation innovations based on patent data, and isolate exogenous variation in labor costs from the innovating firms' perspective. We find that a 1% increase in low-skill wages induces between 2 and 5% more automation innovations.

For our classification, we aim to identify automation innovations that allow for the replacement of workers with equipment in some tasks. Our classification follows a two-step procedure: First, we classify technology categories in machinery (IPC and CPC codes) using patents' text and second, patents using their technology categories. This procedure leverages that the combined wording of many patents improves the signal of automation characteristics and permits classifying patents without text. The resulting classification covers a wide range of automation technologies. It is transparent and mappable to a detailed sectoral level. Our classification is broader than robots but stricter than others used in the literature (such as Mann and Püttmann, 2021). As a validation exercise, we reproduce the cross-sectoral analysis of Autor, Levy and Murnane (2003) but add a measure of the automation intensity of equipment based on our classification. We find that in the United States, sectors that use more automation-intensive equipment saw larger decreases in routine tasks.

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 innovating's firm perspective using a shift-share design. Automation innovators are often equipment manufacturers which sell their machines to downstream firms in various countries. Automation producers' incentives to innovate therefore 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 cost using data on innovating firms' international exposure and country-level labor costs. Country-level labor cost shocks can then be used as a source of variations to identify the effect of labor costs on automation innovations.

We carry out 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 geographical 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 biadic patents (that is, patents applied for in at least two countries). Our final sample covers the period 1997-2011 and contains 3,236 firms that account for 53.2% of global automation innovations. We run Poisson regressions employing multiple layers of firm, industry, and 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 depending on specification. In line with the capital-skill complementarity hypothesis (Krusell, Ohanian, Rios-Rull and Violante, 2000), increases in high-skill wages tend to reduce automation. We discuss our identification assumptions in the context of the recent shift-share literature. Borusyak, Hull and Jaravel (2022) argue that identification can be obtained from conditionally randomly assigned shocks. As domestic shocks could affect innovating firms through other channels than the labor costs paid by downstream firms, we include country-year fixed effects for the innovator's home country and further exclude the home country from the wage variables. In these specifications, foreign wage shocks are the source of identification. In addition, we carry out placebo regressions. We find that non-automation machinery innovations do not respond to wage shocks.

Because our main analysis is agnostic about the exact nature of the labor market shocks driving automation innovation, we complement it with two exercises. First, we build a measure of minimum wages for a subset of countries. We also find a positive effect of minimum wages on automation innovations. Second, we focus on a specific labor market shock, the Hartz reforms in Germany. The Hartz reforms 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 post the Hartz reforms. Finally, in a triple-difference exercise, we find that the reforms also decreased automation innovations relative to non-automation innovations.

We contribute to three literatures: on induced automation, on endogenous 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. We provide empirical support for this literature.

The existing empirical literature studying the effect of wages on technology adoption or innovation is limited.[1] A few papers show that labor market conditions affect adoption of labor-saving technology in agriculture (Hornbeck and Naidu, 2014, and Clemens, Lewis and Postel, 2018) and manufacturing (Lewis, 2011). Lordan and Neumark (2018) find that minimum wage hikes displace workers in automatable jobs and Fan, Hu and Tang (2020) that they induce Chinese firms to adopt industrial robots. In contrast, we focus on *innovation* not *adoption*. The distinction matters because i) the magnitudes of the effect of wages on innovation and adoption likely differ and ii) knowledge spillovers play a larger role for innovation than adoption.

The literature on induced automation innovation is scarcer. Acemoglu and Restrepo (2022) find a positive correlation in cross-country regressions between aging and patenting in robotics and numerical control, though they focus mainly on adoption. Our paper differs in at least three 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; and foremost, we conduct our analysis at the firm instead of the country-industry level. 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. In contrast, we exploit firm-level variation and focus on the effect of labor cost on global innovation.[2] Our

[1] 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, 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.

[2] Relatedly, Andersson, Karadja and Prawitz (2022) look at the effect of emigration to the US in the

contribution to the literature on induced automation is to show a new correlation at the firm level between downstream low-skill wages and automation innovation and to argue that this correlation reflects a causal effect from wages to automation innovations.

An extensive literature shows that the direction of innovation is endogenous in other contexts (e.g. Acemoglu and Linn, 2004, and Popp, 2002). We build on Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen (2016), who show that an increase in gas prices lead firms in the auto industry to engage more in clean and less in dirty innovations. We use a similar shift-share design as theirs and also measure firms' international exposure with patent weights. We advance their methodology in particular by introducing country-year fixed effects.[3]

Other researchers have built measures of automation with patent data. In contemporary work, Mann and Püttmann (2021) use machine-learing techniques to identify automation patents and Webb (2020) uses a dictionary approach similar to ours to identify robot, software, and artificial intelligence patents and links them to occupations. We compare our approaches below.[4] We classify technological codes, which are readily available, and not patents directly using text. Therefore, only our classification can be directly applied to other patent data.

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 the event study of the Hartz reforms. Section 6 concludes. The Appendix provides an analytical model, additional robustness checks, and details on our methodology.

19^{th}century in Sweden and find that more exposed municipalities experienced an increase in innovation (but 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).

[3]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.

[4]Recently, Autor, Salomons and Seegmiller (2021) and Kogan, Papanikolaou, Schmidt and Seegmiller (2022) also match patents to occupations to look at the effect of technology on labor.

2 Classifying Automation Patents

In this section, we develop our classification of automation patents. Then, we use the classification to build a measure of automation at the industry level and find that it predicts a decline in routine tasks, extending the analysis of Autor et al. (2003).

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 costs 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. Patent data have many advantages: they focus on the output of the innovative process, give the countries where inventions are being protected and are available at both the firm level and at a highly disaggregated technological level. They also have some drawbacks, such as a high heterogeneity in patent value. To tackle this problem, in our main analysis we focus on patents filed in at least two countries.

We use two patent 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 bibliographical information but not the text of close to the universe of patents. In these datasets, the technological characteristics of patents are recorded in technological codes (notably CPC and IPC codes, henceforth C/IPC codes, explained in footnote 8 below). Certain types of technologies, such as fossil fuel engines, can readily be identified to existing groupings of C/IPC codes. Such a grouping does not exist 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 technology categories as automation or not based on a cut-off. iv) We then 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 other researchers and we can use our technology category classification on patents without text and future patents.[5] Second, the C/IPC codes are by themselves informative of the characteristics of an innovation, including whether it relates to automation. Patents are written in varying styles. 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 this keyword is frequently 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 technological code corresponds to automation.[6]

Alternatively, we could have read and classified a subset of patents and then used machine-learning techniques to classify other patents or technology categories based on patent text. This is the procedure in Mann and Püttmann (2021), whose results we discuss in Section 2.3 and Appendix A.3. Relying on keywords instead of a training set of patents presents several advantages. First, manually classifying patents as automation is a difficult task that cannot be easily systematized and outsourced. Second, patents give technical descriptions of an innovation and do not primarily discuss its goal. Only a few words within the text are informative, and a machine-learning algorithm would require an extensive training set. Third, using a few keywords instead of a large training set makes our approach more transparent, easily replicable and modifiable, and leaves fewer degrees of freedom since we pick most of our keywords from the literature.

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 complement them with a few additional words as described below.[7] In fact, most of our search terms (for simplicity "keywords") correspond to the co-occurrence of

[5] 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.

[6] Our strategy follows the World Intellectual Property Organization (WIPO), which offers on its website a simple tool 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.

[7] 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.

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 and 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 (i.e. numerical control includes NC, numerically controlled and numeric control), ii) British/American spelling (i.e. labour/labor) and iii) hyphenated adjectives (i.e. computer aided / computer-aided design). "In the same sentence as secondary words" refers to at least one secondary word. We added words in italics, the others come from AR or DDT. See Appendix for details.

several words in the same sentence or patent or the repetition of these words a sufficient number of times. Table 1 describes the list of our search terms.

We have eight categories of keywords. Five of these, robot*, numerical control, computer-aided design and manufacturing, flexible manufacturing, and programmable logic control, are automation technologies in DDT or AR. 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. Furthermore, we add "automation" and "automatization". The stem "automat*" gathers too many false positives such as "automatic transmission". We resolve this in two ways: either we restrict attention to patents where the frequency is 5 or more or we combine automat* with our secondary words or other words that largely come from technologies described in DDT or AR and often describe tasks (such as manipulat*, regulat* or inspect*). We count patents where automat* and one of these words appear in the same sentence at least twice. 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*" (see Appendix A.2). Section 4.3 shows that our main results are robust to only using those.

2.3 Automation technology categories

Defining machinery C/IPC codes. We base our classification on EPO patent applications from 1978 to 2018 with a description in English (1,538,370 patent applications), which we denote Ω_{EPO}. To identify the technological characteristics of patents, we use their C/IPC codes. The C/IPC codes form a hierarchical classification system; most patents have several of them.[8] We define "technology categories" based on these codes, and use our keywords to classify technology categories as automation or not.

Specifically, we define technology categories using three C/IPC groupings. First, we use 6-digit C/IPC codes (e.g. B25J13). Second, we include pairs of 4-digit C/IPC codes (e.g. B25J and A61F) with the idea that the co-occurrence of technological codes can also be informative about the characteristics of a patent. Finally, we add the co-occurrence of 4-digit C/IPC codes with the 3-digit codes G05 or G06 (e.g. B25J together with G05 or G06). The code G05 corresponds to "controlling; regulating" and G06 to "computing; calculating; counting". Aschhoff et al. (2010) use these combinations to identify advanced manufacturing technologies. To ensure that the set of patents available in Ω_{EPO} is sufficiently representative of a technology category, we restrict attention to categories that contain at least 100 patents (we group 6-digit codes with the same 4-digit code and less than 100 patents in common artificial 6-digit codes). The 6-digit codes will identify close to 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 specific technological fields. There are 34 technological fields (see Figure A.1). We focus on "machine tools", "handling", "textile and paper machines", and "other special machines" with some adjustments, which we refer to as "machinery".[9] For pairs of 4-digit C/IPC codes or pairings of 4-digit C/IPC

[8]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: "manipulators; chambers provided with manipulation devices"), and main groups have 5 to 7 digit codes (e.g. B25J 9: "programme-controlled manipulators"). In the following, we refer to classes, subclasses, and main groups as 3-digit, 4-digit, and 6-digit codes respectively.

[9]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 con-

codes with G05 or G06 we classify them as machinery if at least a 4-digit code belongs to that field. This leaves us with 986 6-digit codes, 1104 pairs of 4-digit codes, and 25 groupings of 4-digit codes with G05/G06.

Defining automation C/IPC codes. We define a machinery patent as a patent which belongs to one of the machinery technological categories. We then denote MT_p the set of machinery technology categories associated with a patent p.[10] The combined set of machinery technology categories is $\mathcal{MT} = \cup_{p\in\Omega_{EPO}} MT_p$. A patent is also associated with a text T_p. For each keyword category (automat*, robot, CNC, etc.) we define functions $k^{automat*}(T_p)$, $k^{robot}(T_p)$, $k^{CNC}(T_p)$, etc. which take value 1 if one of the associated keywords is in the text and 0 otherwise. We define $k^{any}(T_p) = \max\{k^{automat*}(T_p), k^{robot}(T_p), k^{CNC}(T_p), ...\}$ which takes value 1 if any of the automation keywords are present. For all machinery technology category $t \in \mathcal{MT}$, we define the prevalence of automation keywords $s(t)$ as the share of patents containing at least 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 procedure delivered reasonable results and adjusted the keywords accordingly. Yet, we never modified the classification after carrying out our regressions.

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. It shows that most C/IPC codes have a low prevalence of automation keywords but a few codes have a very high value. As our baselines, 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 0.396 and 0.480,

tain 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 from the machine tools technological 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 qualitatively affect our results.

[10]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.

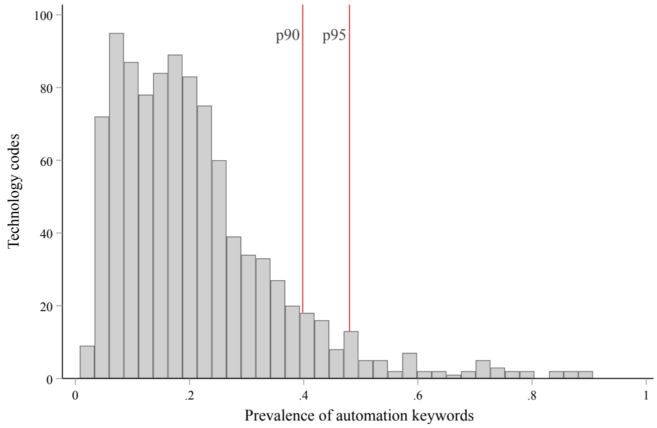

Figure 1: Prevalence of automation keywords for C/IPC 6-digit codes in machinery

respectively.[11] 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 and the second half is 0.85.

2.4 Automation patents

We now proceed to classify automation patents. To do so, we use PATSTAT, which contains bibliographical information for close to the universe of patents. PATSTAT further allows us to identify patent families, a set of patent applications across 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 slightly abuse language and 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 on how we identify automation patents, on the stability of our classification, and gives examples of automation patents.

[11] Choosing different thresholds is easy and we investigate how robust our results are in Section 4.3.

Comparison with Mann and Püttmann (2021). Mann and Püttmann (2021) also classify patents as automation and non-automation. Our approaches differ in three ways. First, they classify all patents while we focus on machinery. Second, they manually classify a training set and use machine learning to classify US patents in a given period, while we identify technology categories using a dictionary method. This way, we (or others) can classify any patent in machinery. Third, they define as automation "a device that carries out a process independently of human intervention", while we seek to identify innovations that replace workers in existing tasks. Therefore, they classify a number of patents related to elevators and printing machines as automation patents, which we do not. In Appendix A.3, we compare the two approaches in detail.[12]

2.5 Trends in automation innovations

To restrict attention to innovations of sufficient quality, we focus on patent families containing 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.[13]

Figure 2 plots the evolution of automation biadic patent families. Panel (a) shows that worldwide the share of automation patents in machinery slightly declined between the mid1980s (9.4% in 1985 for auto95) and the mid1990s (7.5% in 1994 for auto95) before increasing quickly (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. Initially, Japan's share of automation patents in machinery is the highest, but it declines through the 1980s and 1990s. It increases in the 2000s, yet is overtaken by other countries, in particular Germany which has the highest automation share in 2015.

[12]Bessen and Hunt (2007) also use keywords to identify software patents. Webb (2020) focuses on matching three technologies (robotics, software, and AI) to the occupations they might replace and similarly identifies the associated patents using keywords. We instead focus on all automation innovations in machinery, and classify technology categories first.

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

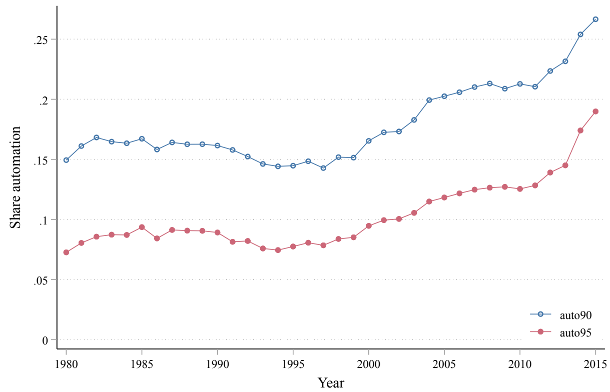

(a) Share of automation patents in machinery worldwide.

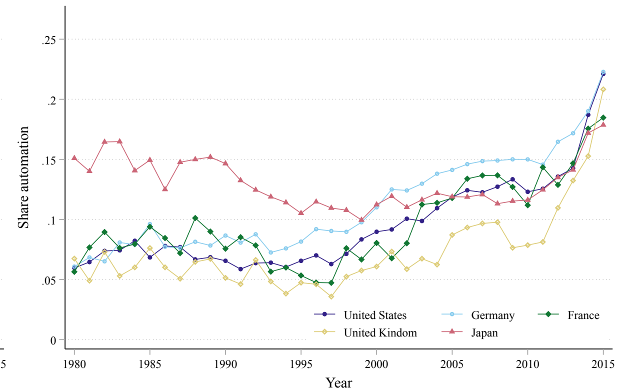

(b) Share of automation patents (auto95) in machinery by applicants' nationality.

Figure 2: Share of automation patents in machinery for biadic families

2.6 Automation, routine tasks and skill composition

We now build a measure of automation at the industry level and relate it to changes in task and skill composition. We do this in part to validate our classification of automation patents. We build on Autor et al. (2003) (henceforth ALM), who show that computerization was associated with a decrease in routine tasks at the industry level on U.S. data from 1960 to 1998. We give our main results here and refer the reader to Appendix A.4 for details on the data construction and additional results.

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

$$\Delta T_{jk} = \beta_0 + \beta_C \Delta C_j + \beta_{aut} aut_j + \varepsilon_{j\tau}. \tag{1}$$

We focus here on the years 1980-1998. ΔT_{jk} represents the change in tasks of type k in industry 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-industry change in task input measured in percentile of the 1960 task distribution. ΔC_j is ALM's measure of the change of computerization in sector j (available for the period 1984-1997). aut_j is our patent-based measure of automation intensity in sector j. Since patenting is already a measure of the flow of knowledge, we do not first-difference this measure.

To construct 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, providing detailed sectoral information for corporate patents. We use their data to create a weighted concordance table from C/IPC 4-digit codes to 4-digit SIC industries. We use this mapping to allocate patents to sectors of invention. Then, we combine this information with the 1997 capital flow table from the BEA to get the sector of use. The capital flow table is similar to an input-output table but reports the flows in investment goods instead of intermediate inputs. For each sector j, we compute aut_j as the share of automation patents (auto95 in our baseline) among machinery patents applied for in 1980-1998. We compute this statistic for the 133 sectors with machinery patents (our results are robust to excluding sectors with few machinery patents). Interestingly, our automation measure auto95 is only weakly correlated with computerization with a coefficient of 0.08 (and -0.16 when we weigh industries by employment).

Table 2: Effect of the use of automation technologies on tasks and skill composition

	Δ Routine cognitive		Δ Routine manual		Δ High/low skill workers	
	(1)	(2)	(3)	(4)	(5)	(6)
Share automation	-155.29***	-159.70***	-128.32***	-129.88***	3.97*	3.64*
(using industry)	(37.30)	(25.62)	(34.89)	(36.82)	(2.15)	(1.91)
Share automation		-21.00***		-11.74		-0.22
(inventing industry)		(7.73)		(7.90)		(0.50)
Δ Computer use	-17.42***	-18.21***	-19.13***	-21.35***	0.97***	0.91***
(1984-1997)	(6.48)	(6.10)	(7.24)	(7.61)	(0.26)	(0.26)
R^2	0.27	0.36	0.21	0.24	0.21	0.19
Mean dependent variable	-2.50	-2.47	-2.27	-2.27	0.12	0.12
Observations	133	126	133	126	133	126

Notes: Each column represents a separate OLS regression of ten times the annual change in industry-level task input between 1980 and 1998, measured in centiles of the 1960 task distribution, on the share of automation patents in machinery, the annual percentage point change in industry computer use during 1984-1997, and a constant. Estimates are weighted by mean industry share of total employment in FTEs in 1980 and 1998. Robust standard errors are reported in parentheses. In columns 1–3 the dependent variable is the change in routine cognitive tasks, in columns 4–6 the change in routine manual tasks, and in columns 7–9 the change in the ratio of high-skill workers (college graduates) over low-skill workers (others). 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, and 6 – allocates patents to their sector of manufacturing. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 2 reports the results of regression (1) for routine cognitive tasks (Column 1) and routine manual tasks (Column 3). Appendix Figure A.10 provides scatter plots of the changes in routine tasks and the share of auto95 patents in machinery. There is a clear relationship: sectors with a higher share of automation patents experience a larger decline in routine cognitive and routine manual tasks. A 1 pp increase in the automation share is associated with a 1.6 and 1.3 centiles decrease in routine cognitive and manual tasks per decade. The standardized beta coefficients are larger for automation than for computerization: 2 vs 1.3 for routine cognitive and 1.7 vs 1.4 for routine manual tasks.[14]

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 render superfluous (Figure A.8 gives an example of such a machine). Metalworkers, for instance, are one of the occupations with the highest intensity in routine cognitive tasks.

In Column (5), we use the change in the ratio of high-skill workers (defined as college graduates) over low-skill workers (defined as all other workers) as the dependent variable instead. We find that sectors with a higher share of automation innovation experience a larger increase in the skill ratio.

In Columns (2), (4), and (6), we add a control for the share of automation patents invented by the industry. To allocate patents to the inventing sector, we simply omit the capital flow table step when computing our automation variable at the sectoral level. The coefficients on the automation share in the using sector remain similar to those in Columns (1), (3) and (5). In addition, the automation share in the using sector has a bigger effect than the automation share in the inventing sector.[15]

Appendix A.4 includes additional robustness checks: we use biadic patents, auto90 patents, or an alternative concordance table between C/IPC codes and sectors developed by Lybbert and Zolas (2014). In all cases, we find a negative effect of the automation share on routine tasks.

To summarize, we have now classified machinery patents as automation or non-automation. Importantly, given a mapping between C/IPC codes and sectors, this classification also delivers a measure of automation at a more detailed sectoral level than alternatives such as robotization. This measure is uncorrelated with computer use but is associated with a reduction in routine tasks and an increase in the skill ratio at the sectoral level.

[14] The employment-weighted standard deviation in the share of automation patents for the included industries is 1.3% and the mean 7.5%, while the standard deviation for computerization is 0.072. Routine tasks decline by 2.5 and 2.3 centiles per decade for these sectors.

[15] The standardized coefficients are larger for the using sector than the inventing sector as the s.d. for the share of automation patents in the using and inventing sectors are respectively 1.3% and 6.3%.

3 Empirical Strategy and Data

We now move to our main empirical exercise, which analyzes the effect of labor cost shocks on automation innovations. Section 3.1 presents our empirical strategy, Section 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 for our baseline sample. Section 4 discusses results and identification assumptions.

3.1 Empirical strategy

We motivate our empirical strategy with the business structure of the most prominent automation innovators. These are often large companies that sell their automation equipment internationally to downstream firms. Automation equipment allows for the replacement of low-skill workers with machines. It may also complement high-skill workers who program, operate, and maintain the machines. Therefore, the incentives of the downstream producers to adopt automation technology are determined by labor costs in their local market. Higher labor costs for potential customers are associated with a larger market for automation machine producers, which, in return, should induce innovators to undertake more research in automation technologies.[16,17] Appendix A.5 presents a simple model which rationalizes this argument.

Empirically, we aim to measure 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 would affect automation innovations differently from non-automation innovations (in contrast, for instance, with the market size of the downstream firms).

---

[16]For example, Siemens, the biggest innovator in our sample, had 31% of its workforce but only 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) 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...". Note that this sentence includes a lot of our keywords. The report is centrally interested in how "Changes in customer demand [for automation technology by downstream manufacturers] are strongly driven by macroeconomic cycles". Interestingly, the report never mentions "cost of labor" as a reason for automation but instead 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.

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

Ideally, we would measure the labor costs paid by automation innovators' actual and potential customers. 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 viewed as the reduced form of this instrumental approach.

More specifically, our measure of the labor cost paid by the customer of an innovator is 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 cost 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 $w_{L,c,t}$ (resp. $w_{H,c,t}$) is the low-skill (resp. high-skill) labor cost in country c at time t and $\kappa_{i,c}$ is the fixed weight of country c for firm i.[18] 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 this extensively in Section 4.2. We now describe how we obtain country-level data (such as $w_{L,c,t}$) and firm data (including the weights $\kappa_{i,c}$).

3.2 Macroeconomic data

We source country-level data primarily from the 2013 release of the World Input Output Tables (WIOD, Timmer et al. 2015). The database contains information on hourly labor costs from 1995 to 2009 across groups of educational attainment for the manufacturing sector in 40 countries, including all major markets (US, Japan, all EU countries of 2009, China, India, Brazil, Russia, etc.). We get similar data from the Swiss Federal Statistical

[18]To be more precise, innovation incentives depend on the expectation of future labor costs for automatable tasks, and ideally, we would measure these directly. We cannot measure expectations so we use current labor costs shocks as a proxy for shocks on expected future costs (see Section 4.3 for further discussion). In addition, there are no good international occupational or task-level labor costs data. Since low-skill and middle-skill workers are those whose 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 downward bias.

Office to add Switzerland, a large source of patents, to our analysis. For our baseline regressions, we focus on labor costs in manufacturing but check that 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, low-skill workers have no high-school diploma or equivalent and high-skill workers have at least a college degree. Middle- and low-skill wages are very highly correlated and we can interpret our low-skill wage variable as reflecting both.

From the same dataset, we calculate labor productivity in manufacturing as value added divided by hours and producer price indices (PPI for the whole economy and manufacturing). We gather exchange rate and GDP data from UNSTAT and compute the GDP gap to control for business cycles. For our baseline regressions, we deflate all nominal values by the local PPI for manufacturing (indexed to 1995), then convert everything into dollars using the average exchange rate for 1995, the starting year of our regressions. Appendix A.6.1 provides further details.

Appendix Table A.2 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. Appendix Figure A.3.a shows the log inverse skill premium in the 6 countries with the largest average weights. Trends in the skill premium vary markedly across countries, with non-monotonicities for some countries.[19]

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 with the companies in Orbis (Appendix A.6.2 details how we merge Orbis firms). We then use PATSTAT to obtain all bibliographical information about 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.

In the absence of sales data, we use the firm's history of patent filing as a proxy for its market exposure to measure the weights $\kappa_{i,c}$. This method follows and expands on that of Aghion et al. (2016, henceforth ADHMV). Firms differ in their market exposure because

[19]This figure only shows the raw data, the identifying variation, taking into account the fixed effects, can be seen in Appendix Figure A.12.

of trade barriers, heterogeneous tastes of customers, 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, and 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 apply for patent protection in a country if they are relatively confidant in the potential market value for the technology (Eaton and Kortum, 1996). Indeed, empirical evidence suggests that inventors do not patent widely and indiscriminately, with the average invention only patented in two countries (Dechezleprêtre et al., 2011).

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 they are weakly exogenous.[20] We restrict attention to patent families with at least one citation (excluding self-citations) to exclude the lowest quality patents. See Appendix A.6.3 for details.[21]

Although patenting indicates whether the firm intends to sell a that market, the raw patent count does not reflect 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.[22] As a result, the weight of country c for firm i is:

$$\kappa_{i,c} = \frac{\tilde{\kappa}_{i,c} GDP_{0,c}^{0.35}}{\sum_{c'} \tilde{\kappa}_{i,c'} GDP_{0,c'}^{0.35}}.$$

We then combine the weights $\kappa_{i,c}$ with the macro variables presented in section A.6.1 to build macro variables, including wages, at the level of the firms' customers along the lines of equation (2). We use 1971-1994 as a pre-sample period as PATSTAT's coverage

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

[21]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.

[22]Eaton, Kortum and Kramarz (2011) estimate the elasticity of French exports to the GDP of the destination country to be 1 and the elasticity of the number of French exporters to be 0.65. This gives an elasticity of the average export by firm of 0.35. ADHMV use a power of 1 on GDP instead of 0.35.

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.[23] We show that our results are robust to alternative pre-sample periods and weighing schemes in Section 4.3.

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 firm's innovation in automation follows:[24]

$$
\begin{aligned}
& PAT_{Aut,i,t} \\
= \quad & \exp\left(\begin{array}{c} \beta_{w_L} \ln w_{L,i,t-2} + \beta_{w_H} \ln w_{H,i,t-2} + \beta_X X_{i,t-2} + \beta_{Ka} \ln K_{Aut,i,t-2} + \beta_{Ko} \ln K_{other,i,t-2} \\ + \beta_{Sa} \ln SPILL_{Aut,i,t-2} + \beta_{So} \ln SPILL_{other,i,t-2} + \delta_i + \delta_{j,t} \, (+\delta_{c,t}) \end{array}\right) + \epsilon_{i,t}.
\end{aligned} \tag{3}
$$

$PAT_{Aut,i,t}$ denotes the number of biadic automation patent families by firm i with first application filed in year t. Automation patent families are the auto95 patents defined in Section 2. As mentioned in Section 2.5, we focus on biadic patent families to ensure that patents are of sufficiently good quality and more comparable across countries. Focusing on biadic patents is also consistent 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 wages (more generally labor costs) faced by the customers of firm i at time t defined in (2), deflated by the local manufacturing PPI. $X_{i,t}$ represents a vector of macroeconomic controls (labor productivity in manufacturing, GDP per capita, and GDP gap). Labor productivity captures technology or human capital shocks in the country where machines can be sold, GDP per capita similar shocks but also demand shocks and the GDP gap, business cycles fluctuations.

Following ADHMV, 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 other

[23] In Section 4.3, we consider an alternative measure 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 the two wage variables is 0.86.

[24] 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.

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 firms' inventors.[25] These controls ensure that we do not simply capture that some firms or countries are on different automation trends. In addition, knowledge spillovers are often an important characteristic of innovation processes and may amplify the short-run response of innovation to economic shocks over time.

δ_i are firm fixed effects so that we are looking at how changes in wages affect changes in automation innovations.[26] $\delta_{j,t}$ are industry-year fixed effects. The industry j of a firm is the industry of manufacturing and corresponds to its 2-digit industry in Orbis. Appendix Table A.1 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 is defined as the country with the largest weight $\kappa_{i,c}$. Finally, $\epsilon_{i,t}$ is an error term. In the baseline specification, we cluster standard errors at the firm level.

We lag the right-hand side variables by 2 years in the baseline regressions for two reasons. First, the empirical literature 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 predictive of future wages. Section 4.3 considers alternative timing assumptions.

3.5 Baseline sample

We now describe the firm sample we rely on to estimate equation (3). Since wages are available for 1995-2009, our baseline datasets rely on firms that applied for at least

[25]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.

[26]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 (3) 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 $K_{Aut,i,t-2}$). Yet, we show in Section 4.3, that our coefficients of interest are unsurprisingly not affected by Nickell's bias by either removing the stock control or by implementing the Blundell, Griffith and Van Reenen (1999) method, which uses the pre-sample average of the dependent variable to proxy for the fixed effect, in line with the patent literature.

Table 3: Descriptive statistics of the firms in our baseline regression

Sample	Auto95		Auto90			Auto95	Auto90
	(1)	(2)	(3)	(4)		(5)	(6)
Automation patents	Per year	1997-2011	Per year	1997-2011		Country weights	
Mean	0.56	11.98	0.64	13.70	Largest country	0.47	0.46
SD	3.58	54.21	3.94	62.31	Second largest	0.17	0.18
P50	0	2	0	2	US	0.21	0.21
P75	0	6	0	7	Japan	0.17	0.15
P90	1	20	1	22	Germany	0.20	0.21
P95	2	43	2	49	France	0.09	0.09
P99	10	194	11	215	UK	0.09	0.09
Number of firms	3236		4821				

Notes: Summary statistics for the firms used in our baseline regression. Columns (1), (2) and (5) consider the regression sample when the dependent variable is the count of auto95 patents. Columns (3), (4) and (6) consider the regression sample when the dependent variable is the count of auto90 patents. Columns (1) and (3) give statistics on the count of automation patents per year and Columns (2) and (4) for the whole period. Columns (5) and (6) give statistics on firms' country weights.

one biadic automation patent between 1997 and 2011. These firms must also have at least one patent before 1995, for us to compute weights on geographical coverage and on inventors' location. We further exclude wholly domestic firms (i.e. those which 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 $3,236$ 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 84.1% of all biadic auto95 patent families in 1997-2011 to a firm. Moreover, most heavy patenters had already patented in at least 2 countries pre-sample: the firms of our sample account for a disproportionate 53.2% of all biadic auto95 patent families.

Table 3 gives descriptive statistics on the number of automation patents per year and the country weights for the firms in our sample. The distribution of auto95 patents is strongly skewed: over the period 1997-2011, the median firm in the sample filed 2 auto95 patent applications, whereas the 99^{th} percentile filed 194. The largest country for a given firm has, on average, a weight of 0.47 (for auto95), and the second largest 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. Appendix Table A.4 lists our sample's ten biggest automation innovators.

Appendix Table A.5 gives standard deviations and a correlation matrix for the firm-level macroeconomic variables, residualized on firm and industry-year fixed effects.[27] We

[27]We note that the correlation between low-skill and middle-skill wages is very high. As a result, we

still find significant variation in the residualized (log) low-skill wages since the standard deviation is 0.03 (by comparison, the standard deviation is 0.1 when residualizing only on firm fixed effects). Appendix A.7 provides additional statistics computed at the level of the shock of our shift-share variable (see Appendix Table A.29).

4 Global Wages and Induced Automation

We present our results in three steps: First, we demonstrate a positive effect of low-skill wages on automation innovations. Second, we show that this effect does not exist for non-automation innovation in machinery. Third, we build on the recent shift-share literature (notably Borusyak, Hull and Jaravel, 2022) and argue that the effect of low-skill wages on automation innovations is causal. We then discuss additional results, notably on the minimum wage, and provide robustness checks.

4.1 Main results

Table 4 presents our baseline results. Columns (1) to (3) control for firm and industry-year fixed effects. An increase in the low-skill manufacturing wage paid by the downstream producers of an innovating firm predicts an increase in automation innovation. The estimated coefficient is an elasticity, so 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 with a magnitude roughly similar to that of low-skill wages. 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 consistent significant effects. A higher stock of automation knowledge at the firm level predicts fewer automation innovations in the future—so that firms do not seem to specialize in automation technologies over time. The spillover coefficients indicate that firms exposed to more knowledge in automation technologies tend to undertake more automation innovations. The long-run effect of an increase in low-skill wages on innovation may therefore differ from its short-run effect (see Appendix A.8).

Country-year fixed effects. Columns (4) to (6) of Table 4 reproduce Columns (1) to (3) but add country-year fixed effects, where the country of a firm continues to be

will not look at the effect of middle-skill wages separately.

Table 4: 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	2.97***	2.72***	3.64***	2.24**	2.61**	3.64***	4.19***	5.30***	4.43**
	(0.80)	(0.85)	(0.96)	(1.01)	(1.14)	(1.28)	(1.34)	(1.57)	(1.80)
High-skill wage	-2.23***	-2.64***	-1.56*	-2.81***	-2.04*	-1.87*	-4.47***	-2.91**	-4.33***
	(0.73)	(0.80)	(0.82)	(0.97)	(1.08)	(1.07)	(1.32)	(1.48)	(1.42)
GDP gap	-3.80	-4.34	-2.26	4.56	5.53	6.95	0.04	2.40	0.50
	(2.62)	(2.71)	(2.81)	(6.87)	(6.90)	(7.21)	(4.59)	(4.91)	(5.24)
Labor productivity		0.96			-1.77			-2.53	
		(0.92)			(1.78)			(1.61)	
GDP per capita			-1.86			-3.45*			-0.42
			(1.32)			(1.97)			(2.12)
Stock automation	-0.12***	-0.12***	-0.12***	-0.12***	-0.12***	-0.12***	-0.12***	-0.13***	-0.12***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Stock other	0.51***	0.51***	0.51***	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	0.61**	0.64**	0.76**	1.36***	1.34***	1.35***	1.33***	1.29***	1.32***
	(0.30)	(0.30)	(0.31)	(0.47)	(0.47)	(0.47)	(0.46)	(0.46)	(0.46)
Spillovers other	-0.20	-0.25	-0.33	-0.97***	-0.93***	-0.99***	-0.97***	-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	–	–	–	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Number of firms	3236	3236	3236	3233	3233	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. 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 previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

the country with the largest weight. Unobserved country-level shocks in the innovator's country can 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 the incentive to innovate. Shocks that affect firms mainly through their home country can be captured through home country-year fixed effects. As further discussed in Section 4.2, our identification assumption is then that foreign wages are exogenous to the automation innovation of the firm, given our controls. 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, using the headquarters' location to define the home country gives similar results.

Foreign wages. Columns (7) to (9) go further and only consider the foreign component of wages and other macro variables. Specifically, we decompose total low-skill wages $w_{L,i,t}$ into their home and foreign components 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. With this specification, we can still interpret our coefficient as an elasticity on total wages. 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 (as GDP gap is already an average of logs, we directly interact the foreign variables with $\kappa_{i,F}$). Again, we find a positive effect of low-skill wages on automation innovation and a negative effect for high-skill wages. Neither ADHMV nor other papers using their methodology include country-year fixed effects or focus on foreign variation. As argued below, these will generally be important for identification in such settings.

Appendix Table A.6 reproduces similar regressions with fewer controls. Regardless of the control variables included, we find a very stable effect of low-skill wages on automation. The elasticities are between 2.2 and 3.7 when we focus on total wage and slightly larger, between 4.2 and 5.3, when we focus on foreign wages. To interpret the size of these elasticities, note that our analysis focuses on innovation with a high automation content and reflects the behavior of firms undertaking automation innovations.[28] Appendix A.8

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

Table 5: Effect of wages on non-automation innovations

Dependent variable	Placebo Machinery								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	0.87	0.96	1.74*	0.34	0.53	0.95	1.05	1.71	1.21
	(0.72)	(0.78)	(0.89)	(0.97)	(1.03)	(1.29)	(1.53)	(1.64)	(1.78)
High-skill wage	-0.47	-0.32	0.30	-0.72	-0.33	-0.35	-1.51	-0.59	-1.42
	(0.82)	(0.79)	(0.84)	(1.16)	(1.21)	(1.18)	(1.57)	(1.75)	(1.68)
GDP gap	-2.13	-1.96	0.22	3.40	3.80	4.51	-0.24	1.10	0.05
	(1.56)	(1.62)	(1.90)	(4.30)	(4.29)	(4.29)	(2.90)	(3.01)	(2.97)
Labor productivity		-0.33			-0.86			-1.45	
		(0.74)			(1.27)			(1.40)	
GDP per capita			-2.33*			-1.42			-0.26
			(1.32)			(1.91)			(1.74)
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
Observations	42538	42538	42538	42405	42405	42405	42405	42405	42405
Number of firms	2848	2848	2848	2845	2845	2845	2845	2845	2845

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. The sample is restricted to firms having done an auto95 innovation in the sample period. Placebo machinery are innovations in machinery excluding auto90, denoted pauto90. 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 previously defined. Spillover and stock variables are calculated with respect to the dependent variable (pauto90). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

runs a simulation to illustrate the macroeconomic effect of our coefficients.

Clustering level. In the baseline specification, we cluster at the firm level to account for auto-correlation in errors. As firms might also be affected by common country shocks, we cluster standard errors at the home country level in Appendix Table A.7. If anything, this tends to reduce the standard error on low-skill wages.[29] We discuss inference in the shift-share setting in Section 4.2.

Auto90. Appendix Table A.8 reproduces Table 4 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, in line with auto95 being a stricter measure of automation.

Non-automation innovations. 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 non-automation innovations in machinery. Specifically, we

[29] A potential explanation for the negatively correlated error terms is that a successful innovation by one firm captures the market and reduces the innovation of 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, Appendix Table A.7 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.

reproduce the regressions of Table 4 but for machinery innovations that are not auto90. We denote these pauto90. We restrict attention to the sample of firms included in the baseline regressions. We recompute knowledge stocks and spillover variables for these innovations ("own") and for all innovations except those ("other"). Table 5 reports the results. The coefficients on low-skill and high-skill wages are much smaller and only significant in one specification without country-year fixed effects for low-skill wages.[30]

Appendix Table A.9 shows additional placebo regressions, where we either include the full sample of firms undertaking pauto90 innovations or look at all machinery innovations excluding auto95 innovations (pauto95). Again, the coefficients on low- and high-skill wages are not significant. The table also shows regressions of the number of auto95 patents controlling for the number of non-automation innovations (pauto90) patents. This control ensures that our results are not driven by a general tendency for firms to innovate or patent more conditional on innovating. Our coefficients of interest remain unaffected.

The placebo regressions suggest that the effect of wages on innovation is specific to automation innovation. These results validate both our measure of automation and our empirical approach. Suppose our results were explained by another reason than a causal link from low-skill wages to automation innovation. In that case, that alternative reason should also not lead to a comovement between low-skill wages and other innovations in machinery by the *same* firms.

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. Table 6 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 previous specifications and significant at the 1% level in all specifications.

4.2 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.

[30]We drop some firms from the sample of Table 4 because they do not have pauto90 patents during this period. Needless to say, the baseline results on auto95 innovations remain unchanged when restricting attention to the common subsample of Table 5.

Table 6: 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	2.52***	2.68***	2.53***	2.53***	2.39***	2.63***	4.39***	4.20***	4.37***
	(0.70)	(0.70)	(0.70)	(0.89)	(0.88)	(0.89)	(1.28)	(1.25)	(1.27)
GDP gap	-4.12	-4.40*	-4.14	4.77	5.15	5.50	-0.02	0.66	0.40
	(2.59)	(2.61)	(2.61)	(6.79)	(6.73)	(6.85)	(4.60)	(4.64)	(4.68)
Labor productivity		1.03			-1.21			-0.59	
		(0.64)			(1.10)			(0.73)	
GDP per capita			0.04			-1.62			-0.33
			(0.71)			(1.14)			(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	–	–	–	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Number of firms	3236	3236	3236	3233	3233	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. Standard errors are clustered at the firm-level and reported in parentheses. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). 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 previously defined. In these columns, GDP gap, GDP per capita and labor productivity also correspond to their normalized foreign values. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Conditionally randomly assigned wage shocks. Since our measure of wages has a shift-share structure, we rely on the recent literature that discusses the identifying assumptions in this type of set-up. We interpret our results through the lens of Borusyak, Hull and Jaravel (2022). In the language of our setting, they show that the random assignment of wage shocks conditional on weights and our controls can be sufficient for identification. The inference is valid if many country-year pairs are affected by weakly correlated shocks (we argue that these conditions are met in Appendix A.7).[31]

Wages are an equilibrium outcome. So, how can wage/labor costs shocks be conditionally randomly assigned in our context? We include country-year fixed effects and focus on foreign wages. Additionally, our analysis controls for high-skill wages and finds that the skill premium largely drives automation innovation. As such, we are foremost interested in foreign shocks that affect low-and high skill wages differently. We can think of wage shocks as coming from four sources of variation: changes in regulation, labor supply shocks, demand shocks, and technology shocks. We discuss these in turn.

Changes in regulation or labor supply shocks in manufacturing present an ideal source of variation. The introduction of a minimum wage, demographic or education shocks, or

[31]The Herfindahl index for our weights at the country level is 0.13 and, therefore, with 15 years, 0.009 at the country-year level. For foreign weights, these numbers are 0.09 and 0.006, respectively. In Appendix A.7, we argue that there is significant variation within countries.

Table 7: Including additional controls

Dependent variable	Auto95									
	Domestic and foreign					*Foreign*				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Low-skill wage	1.99**	2.07*	3.09***	2.56**	2.47**	4.14***	6.94***	5.19***	5.28***	6.96***
	(1.01)	(1.25)	(1.17)	(1.16)	(1.18)	(1.34)	(1.76)	(1.53)	(1.56)	(1.88)
High-skill wage	-2.49**	-1.14	-1.24	-1.87*	-2.26**	-4.38***	-3.80**	-2.97**	-2.74*	-2.99*
	(0.97)	(1.01)	(1.01)	(1.07)	(1.15)	(1.32)	(1.49)	(1.47)	(1.46)	(1.73)
GDP gap	6.68	7.32	5.68	6.20	5.31	1.62	4.15	3.08	3.05	3.60
	(6.81)	(6.82)	(6.89)	(7.00)	(6.85)	(4.61)	(4.99)	(5.33)	(4.88)	(5.45)
Labor productivity		-2.36	-2.93*	-1.66	-1.62		-5.19**	-2.11	-2.76*	-3.65**
		(1.86)	(1.70)	(1.79)	(1.79)		(2.11)	(1.56)	(1.59)	(1.74)
Manufacturing size	-0.49***					-0.54***				
	(0.19)					(0.20)				
Recent auto95 innovation		-2.51**					1.24			
		(1.26)					(0.93)			
Recent other innovation		1.56**					-0.47			
		(0.78)					(0.80)			
Offshoring			11.65**					-1.87		
			(5.47)					(4.55)		
Long-term interest rate				0.08					-0.03	
				(0.11)					(0.06)	
Low-skill wage (iw)					-0.00					0.05
					(0.47)					(0.55)
High-skill wage (iw)					0.27					-0.23
					(0.37)					(0.46)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47453	47453	47453	47158	46693	47453	47453	47453	47060	35248
Number of firms	3233	3233	3233	3209	3181	3233	3233	3233	3205	2413

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at firm-level and reported in parentheses. All columns include firm, industry-year, and country-year fixed effects. Manufacturing size denotes the log of weighted averages of manufacturing value added in the customer's countries. Recent auto95 innovation, recent other innovation, offshoring and long-term interest rate similarly denote the log weighted averages of respectively auto95 innovations in the last 3 years, other innovations in the last 3 years, the share of foreign value added in the gross value added in manufacturing, and 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. Columns 6-10 use the normalized foreign variables previously defined. Normalized foreign manufacturing size, recent innovation variables and offshoring are defined similarly to normalized foreign low-skill wages; normalized foreign long-term interest rate is defined like normalized foreign GDP gap. Columns 4 and 9 restrict attention to countries for which interest rates are available. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

shifts in labor demand in non-manufacturing sectors, for instance, are unlikely to affect automation innovations through any other channel than an increase in labor costs and can be a source of conditionally randomly assigned wage shocks. In principle, regulation or labor supply shocks could also affect the production costs of innovating firms and thereby innovation. However, as long as production is concentrated in the home country, country-year fixed effects will absorb the effect.[32] In Section 5, we will focus on a specific labor-market shock, the Hartz reforms in Germany.

[32] If a firm 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.

Foreign demand shocks, in contrast, can directly affect both foreign manufacturing wages and the demand for automation equipment and innovation. To address this issue, we have already included several macro controls: GDP gap, GDP per capita, and labor productivity in manufacturing.[33] In Table 7, Columns (1) and (6), we additionally control for the size of the manufacturing sector, computed as the weighted average of country-level manufacturing value added. Our coefficients on low- and high-skill wages remain very similar. Manufacturing size itself has a significant but small negative effect.[34]

Our control for labor productivity addresses foreign technology shocks if they are not skill-biased. Nevertheless, one may still be concerned by skill- (or unskill-) biased foreign technology shocks. For instance, a recent period of higher than usual automation innovation might leave both wages and the incentive for further innovation low, creating a spurious positive correlation. To address this, we construct a measure of recent innovation analogous to that of the low-skill wages: for each country we compute the number of automation innovations (from our set of firms or others) applied for in the last three years and build firm-specific measures. We build a similar control for other innovations. Columns (2) and (7) 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.[35]

A related but distinct issue is that of reverse causality. Distinct because reverse causality concerns the effect of firms' own innovations on wages. As a result, this issue is largely addressed by country-year fixed effects: a shock that leads German firms to introduce more automation innovations will lower German wages but is unlikely to strongly affect non-German wages. In addition, we include a lag between automation innovations and wages and control for past automation innovations in the form of the knowledge stocks at the firm level.

To summarize, we consider that conditional on high-skill wages, macro-controls, and the set of fixed effects; our low-skill wages can be considered as good randomly assigned (or similarly, that the skill premium is as good as randomly assigned conditional on

[33] A foreign demand shock should increase low- but also high-skill wages and the demand for all machines. It will be associated with a higher GDP gap and higher GDP per capita. If the shock occurs in manufacturing, it will lead to an increase in the manufacturing sector and, if value added increases more than employment, higher labor productivity.

[34] 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.

[35] An additional concern might come from low-skill human capital shocks (captured by $\gamma(i)$ in the model of Appendix A.5), 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 bias our estimates downwards.

macro-control and the fixed effects). The stability of our coefficients to various controls suggests that once one controls for high-skill wages, the shocks that identify an effect of low-skill wages on automation innovations are primarily the labor supply and regulation shocks mentioned earlier. The stability of our coefficients can also be seen as a test of the exclusion restriction (Borusyak et al., 2022, Aghion et al., 2022).

Alternative explanations. Borusyak et al. (2022) recommend considering other shock-level variables that may bias results. Accordingly, we control for offshoring, the real interest rate, and inventor-located weighted wages. Increased offshoring in the foreign country might reduce both wages and the willingness to buy automation technology. We construct a measure of offshoring at the country level based on the methodology of Timmer et al. (2014): the share of foreign value added in the gross value added in manufacturing. Then, as for other variables, we build the firm-specific value and control for it in Columns (3) and (8) of Table 7. The real interest rate covaries with the business cycle and is potentially an important determinant of the cost of purchasing equipment. Columns (4) and (9) control for the real yield on 10-year government bonds.[36] 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 in Columns (5) and (10) of Table 7. These regressions provide an additional placebo test, treating firms with the same macroeconomic shocks but weighing them differently. Across the specifications, our coefficients on total and foreign low-skill wages remain largely stable.

Placebo. Perhaps most importantly, our coefficients on low-skill wages should be compared to those from regressions with the placebo innovations. In Table 5, we reported regressions with non-automation innovations in machinery and found persistently little effect from low-skill wages on innovation. Therefore, if our result on the effect of low-skill wages on automation innovations came from a bias, then that bias would have to be absent for other types of machinery innovations undertaken by the same firms.

Shift-share checklist. Borusyak et al. (2022) 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.7, we consider a linear setting for which such an equivalence result applies: we use arcsinh of the count of automation patents as the dependent variable and replace our log of average macro variables with the average of logs. The linear setting allows us to give summary statistics on our shock variable

[36] 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.

Table 8: Monte-Carlo simulation to address Adão et al. (2019) s.e. bias

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.97**	2.72***	3.64***	2.24**	2.61*	3.64	4.19***	5.30**	4.43**
	[0.015]	[0.003]	[0.008]	[0.020]	[0.077]	[0.118]	[0.003]	[0.030]	[0.018]
High-skill wage	-2.23***	-2.64*	-1.56***	-2.81	-2.04	-1.87**	-4.47	-2.91**	-4.33*
	[0.010]	[0.069]	[0.005]	[0.179]	[0.152]	[0.048]	[0.184]	[0.023]	[0.060]
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity	–	Yes	–	–	Yes	–	–	Yes	–
GDP per capita	–	–	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
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Firms	3236	3236	3236	3436	3436	3436	3436	3436	3436

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). P-values are reported in brackets. Columns 1–3 include firm and industry-year fixed effects. Columns 4–9 add country-year fixed effects. Columns 7–9 use the normalized foreign macro variables previously defined. All regressions include controls for stocks and spillovers. The p-values are computed by sampling with replacement the entire path of macroeconomic variables for each firm with 4000 draws. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

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. In addition, we show that a single country does not drive our results by sequentially excluding the six largest countries.

Adão, Kolesar and Morales (2019) show that shift-share design applications tend to over-reject the null. In our application, a problem arises when the residual errors of firms with similar country distributions are correlated, and it is not solved by standard clustering. To address this issue 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 4. Specifically, for each firm, we keep the automation activity, the stocks of innovations, the spillover variables, and the distribution of country weights based on actual data. Then, 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. For each sample, we compute firm-level macro variables as the weighted average of these new country-level variables. 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 4. 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 soaks up most country-specific shocks affecting the outcome variable and, consequently, no shift-share structure is left in the regression residuals. Figure A.4 plots the distribution of coefficients we obtain for Columns 2, 5, and 8.[37]

4.3 Additional Results and Robustness Checks

In this section, we discuss additional results and robustness checks.

Table 9: Effect of the minimum wage

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Minimum wage	2.12***	1.86***	2.12***	1.84**	1.99**	2.04*	2.33*	2.43*	1.22
	(0.63)	(0.64)	(0.79)	(0.89)	(0.94)	(1.07)	(1.20)	(1.26)	(1.44)
High-skill wage	-1.88***	-2.54***	-1.87**	-3.66***	-3.08**	-3.30**	-3.61***	-3.25*	-5.38***
	(0.67)	(0.79)	(0.84)	(1.03)	(1.26)	(1.44)	(1.38)	(1.87)	(1.87)
GDP gap	-2.56	-3.51	-2.55	7.48	8.22	8.25	3.25	3.66	-1.48
	(2.51)	(2.59)	(2.79)	(6.46)	(6.53)	(7.07)	(4.79)	(5.27)	(6.22)
Labor productivity		1.30			-1.04			-0.46	
		(0.79)			(1.50)			(1.63)	
GDP per capita			-0.01			-0.73			3.66
			(1.23)			(2.07)			(2.57)
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
Observations	47767	47767	47767	47436	47436	47436	46287	46287	46287
Number of firms	3233	3233	3233	3231	3231	3231	3148	3148	3148

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. 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 previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Minimum wage. In Table 9, we look at a particular type of labor market regulation, namely the minimum wage. We have data for 22 countries instead of 41.[38] Replacing low-skill wages with the minimum wage, we find a positive effect on automation innovations.

[37]Borusyak and Hull (2021) show that a regression based on a logged shift-share measure may be biased due to the non-linearity of the log function. We implement their correction to remove this potential bias in Appendix A.7.

[38]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.6.1. We do not use the minimum wage as an instrument for low-skill wages

The coefficients are similar to those on low-skill wages for regressions on total wages (Columns (1)-(6)) but smaller and, in one case, insignificant for regressions on foreign wages. This is not surprising: First, we focus on manufacturing, where low-skill wages tend to be substantially above the minimum wage. Second, the minimum wage only captures part of the labor costs. Third, we lose nearly half of our countries.

Macroeconomic magnitude. To illustrate the macroeconomic magnitude of our coefficients and the effect of spillovers and stock variables, we run a simulation in Appendix A.8 where we uniformly and permanently decrease the global skill premium by 10%. This increases the share of automation innovations in machinery by 4.8 p.p. over 1997-2011, with 2.7 p.p. coming from the adjustment of stocks and spillovers. We combine this estimate with the coefficients from our industry-level analysis of Section 2.6, specifically Columns (1) and (4) in Table A.27. We find that the 10% increase in the skill premium would lead to a decline in routine cognitive tasks of 7.5 centiles and a decline in routine manual tasks of 6.2 centiles over a decade (for comparison, routine cognitive and manual tasks declined at 2.5 and 2.3 centiles per decade in the sectors considered).

Timing and pre-trends. In Appendix Figure A.5, we look at alternative lags (and leads) for the dependent variables.[39] 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 4. In Panel b, we only consider foreign wages, corresponding to Column 8. The 2-year lag delivers the highest coefficient 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 among many ADHMV or Popp, 2002). A possible interpretation of this fast response is that firms may prioritize existing automation projects over starting new projects.[40]

Figure A.5 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.

because it would be inconsistent: if low-skill wages are endogenous, then high-skill wages are likely endogenous too, so we would need a second instrument.

[39]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.

[40]In contrast, it is unlikely that our regressions only capture the effect of patenting off-the-shelf inventions which already exist within the firm 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 in the R&D process (Dechezleprêtre et al., 2017). Second, if that were the case, then the largest effect of wages on patents should be contemporaneous.

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 Borusyak, Hull and Jaravel, 2022).

Additionally, innovators should only care about current wages insofar as they are predictive of future wages. In Appendix Table A.10, 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.

Long-difference. For most of our regressions, we follow the large patent literature and rely on the Poisson estimator, which best handles the count data nature of our dependent variable. In Appendix Table A.11, 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 time 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 some specifications, 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.

Innovation types. In Appendix Table A.12, we look at other definitions or subcategories of automation innovations in regressions with foreign wages. The results are robust to excluding the codes that we added to the definition of the machinery technological field listed in footnote 9. Though the coefficients are a bit 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.

Pre-determined weights. 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, as argued in Section 4.2, firms' decision to innovate may be affected by other macro shocks in the destination countries, the exposure to which would be captured by our weights. This is why we rely on Borusyak et al. (2022). Nevertheless, we note that our weights are pre-determined and do not reflect firms' expectations of future wage growth. Appendix Table A.13 shows that country-level growth rates in low- and high-skill wages between 1995 and 2000 have no predictive power on firm weights in 1995. Appendix Table A.14 shows that our results are robust to excluding automation patents from the weights. We also 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.[41]

Additional robustness checks. Appendix Table A.14 looks at alternatives to premultiplying our patents weights with $GDP^{0.35}$: with no multiplication, multiplying by GDP, or with total payment to low-skill workers raised to the power of 0.35, $(w_L L)^{0.35}$. These weights may better measure the potential market for technology that automates low-skill work. The results remain similar.

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) instead, which proxies for the fixed effect with the firm's pre-sample average of the dependent variable. We obtain similar results.

Appendix Table A.16 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 technological field and year of application. The results are weaker with total wages and country-year fixed effects but are very similar to the case without weighing patents in our preferred specification with foreign wages and country-year fixed effects.[42]

Firms of different sizes may be on different trends in automation innovation. In Appendix Table A.17, 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. Appendix Table A.18 shows that our results (using foreign wages and country-year fixed effects) are robust to using different deflators, converting in USD yearly or replacing manufacturing wages with total wages. As mentioned previously, Appendix A.7 contains several exercises linked to our shift-share setting including removing large countries sequentially.

[41]The table also shows that the results are robust to dropping the earlier years from the weights.

[42]This reflects in part that the number of citations is quite right-skewed. Once it is winsorized at the 99^{th} percentile, the t-stats for Columns (4) and (5) rise from 1.2 and 1.27 to 1.43 and 1.56.

5 Event study: the Hartz reforms in Germany

We now focus on a specific exogenous labor shock, namely the German Hartz reforms. This complements our main analysis, which was agnostic about the exact nature of the labor market shocks driving automation innovations. The Hartz reforms were a series of labor-market reforms in Germany first designed 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 given the reforms an important role in the remarkable performance of the German labor market since, in particular, in increasing labor supply and improving matching efficiency.

Such reforms should reduce the incentive to automate low-skill labor by both directly and indirectly decreasing labor costs through an increase in labor supply and a reduction in the expected cost of vacancies. These are perhaps the most salient labor market reforms in a major country during our time period and are an ideal setting for us. They are unlikely to have affected the direction of innovation in non-German firms through channels other than the German labor market and were the major macroeconomic shock in Germany at the time. Furthermore, they had a large and immediate effect: Appendix Figure A.3.b shows that the inverse skill-premium in Germany started to decline as soon as the Hartz reforms were implemented in 2003 while it was flat beforehand. In contrast, there is no such trend for the aggregate rest of the world.

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

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

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

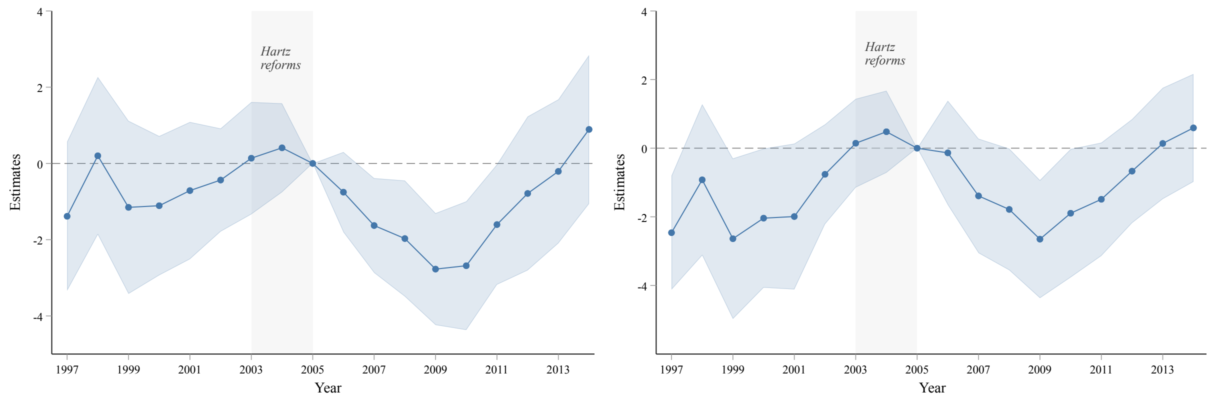

(a) Effect of German exposure on automation innovations

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

Figure 3: Effect of German exposure on automation innovations. 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 Poisson 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.

tends to file automation patents in a given year relative to 2005.

Figure 3.a reports the results. The coefficient of -2.68 in 2010 means that, on average, a firm with a German weight of 0.1 (the mean value is 0.104) had a 26.8% 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 6, this would correspond to a decline in automation innovations of 30.8% between 2005 and 2010.

From 2000 to 2004, firms more exposed to Germany increased their propensity to introduce automation innovations. As expected, the trend reversed between 2006 and 2009, consistent with the Hartz reform increasing labor supply from 2003 onward and decreasing the incentive to introduce automation innovations from 2005. From 2010, the coefficients increase again. This reversal may suggest only temporary effects of the Hartz reform on the direction of innovation, or it was the result of the Great Recession. Since Germany was less affected than other countries by the recession starting in 2008, the relative cost of labor may have risen, leading to a relative increase in automation innovation 2 years later.

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:

$$PAT_{k,i,t} = \exp\left(\begin{array}{c} \beta_{DE,t}\cdot\delta_t\kappa_{i,DE} + \beta_{DE,t}^{aut}\cdot\delta_t\kappa_{i,DE}1_{k=aut} + \beta_{Ka}\ln K_{Aut,i,t-2} \\ \beta_{Ka}^{aut}\ln K_{Aut,i,t-2}1_{k=aut} + \beta_{Kp}\ln K_{Paut,i,t-2} + \beta_{Kp}^{aut}\ln K_{Paut,i,t-2}1_{k=aut} \\ +\beta_{Ko}\ln K_{other,i,t-2} + \beta_{Ko}^{aut}\ln K_{other,i,t-2}1_{k=aut} + \delta_{k,i} + \delta_{k,j,t} + \delta_{k,c,t} \end{array}\right) + \epsilon_{k,i,t}. \tag{4}$$

k denotes the type of an innovation which is either auto95 or other machinery innovation (pauto95), δ_k represents a set of innovation type dummies, $\delta_{k,i}$ 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 interests. 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 German exposure common to all machinery innovations. Figure 3.b reports the results: the pattern is, if anything, more pronounced than in Figure 3.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 (4) 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 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

Table 10: 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***	-1.10***	-1.09***	-1.12***
	(0.33)	(0.33)	(0.33)	(0.33)
Time trend × auto95 dummy × German exposure	0.50***	0.47***	0.46***	0.47***
	(0.19)	(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	–	Yes	Yes	Yes
Year dummy × German exposure	Yes	Yes	–	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	76026	76026	76026	74091
Number of firms	5415	5415	5415	5279

Notes: The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. 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 (all other machinery innovations) 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.

pauto90 innovations (which we used as the default non-automation innovations in Table 5). 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 innovation of 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 firms respond to increases in downstream firms' low-skill labor costs with an increase in automation innovations. We develop a method to classify patents in machinery as automation or not, covering a broad range of technologies. Then, we use this classification to measure the use of automation technology by industry at a highly disaggregated level and find that our automation measure predicts a decline in routine tasks across US sectors. Future research could adapt our classification method to automation patents beyond machinery. Such an extension would allow for an analysis of automation in the service industry or automation of high-skill tasks through Artificial Intelligence.

Further, we use our classification to analyze labor market conditions' effect on machinery automation innovations. Relying on global data, we find that automation innovations are very responsive to changes in low-skill wages with elasticities between 2 and 5. Exploiting the German Hartz reforms, we find a relative decrease in automation innovations by foreign firms with high exposure to Germany after the reform. Though using different variations in the data, both exercises emphasize that automation innovations are much more responsive to changes in labor costs than other innovations.

These results suggest that policies that increase labor costs for low-skill workers, such as increases in the minimum wage, will induce innovations that replace them. Therefore, endogenous technological change is likely to reduce the costs of such policies for the overall economy, as well as limit the welfare gains of these policies for low-skill workers. Our paper provides a building block toward estimating the extent to which a policy-induced increase in low-skill wages could be undone through innovation over time.

References

Acemoglu, D. and Linn, J. (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics*, 119(3):1049–1090.

Acemoglu, D. and Restrepo, P. (2018). The race between machine and man: Implications of technology for growth, factor shares and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Acemoglu, D. and Restrepo, P. (2022). Demographics and automation. *The review of economic studies*, 89(1):1–44.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2022). The labor and product market effects of machines and automation: Causal evidence from french firms and industries.

Aghion, P., Bénabou, R., Martin, R., and Roulet, A. (forthcoming). Environmental preferences and technological choices: Is market competition clean or dirty? *American Economic Review: Insights*.

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 214(1):1–51.

Andersson, D., Karadja, M., and Prawitz, E. (2022). Mass migration and technological change. *Journal of the European Economic Association*.

Aschhoff, B., Crass, D., Cremers, K., Grimpe, C., Rammer, C., Brandes, F., Diaz-Lopez, F., Klein Woolthuis, R., Mayer, M., and Montalvo, C. (2010). European competitiveness in key enabling technologies. Technical report, Centre for European Economic Research (ZEW) and TNO.

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

Autor, D., Dorn, D., Hanson, G., Pisano, G., and Shu, P. (2020). Foreign competition and domestic innovation: Evidence from u.s. patents. *American Economic Review: Insights*, 2(3):357–374.

Autor, D., Levy, F., and Murnane, R. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Autor, D., Salomons, A., and Seegmiller, B. (2021). New frontiers: The origins and content of new work, 1940-2018.

Bena, J. and Simintzi, E. (2019). Machines could not compete with Chinese labor: Evidence from U.S. firms' innovation. mimeo.

Bessen, J., Goos, M., Salomons, A., and van der Berge, W. (2019). Automatic reaction - what happens to workers at firms that automate ?

Bessen, J. and Hunt, R. (2007). An empirical look at software patents. *Journal of Economics and Management Strategy*, 16(1):157–189.

Blundell, R., Griffith, R., and Van Reenen, J. (1999). Market share, market value and innovation in a panel of British manufacturing firms. *The Review of Economic Studies*, 66(3):529–554.

Borusyak, K. and Hull, P. (2021). Non-random exposure to exogenous shocks: Theory and applications.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Cameron, A. C., Gelbach, J., and Miller, D. (2008). Boostrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427.

Clemens, M., Lewis, E., and Postel, H. (2018). Immigration restrictions as active labor market policy: Evidence from the Mexican bracero exclusion. *American Economic*

Review, 108(6):1468–87.

Coelli, F., Moxnes, A., and Ulltveit-Moe, K. (2020). Better, faster, stronger: Global innovation and trade liberalization. *The Review of Economics and Statistics*.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1):95–115.

Danzer, A., Feuerbaum, C., and Gaessler, F. (2020). Labor supply and automation innovation. BGPE Discussion Paper No 198.

de Rassenfosse, G., Dernis, H., Guellec, D., Picci, L., and van Pottelsberghe de la Potterie, B. (2013). The worldwide count of priority patents: A new indicator of inventive activity. *Research Policy*, 42(3):720–737.

Dechezleprêtre, A., Glachant, M., Hascic, I., Johnstone, N., and Ménière, Y. (2011). Invention and transfer of climate change-mitigation technologies: a global analysis. *Review of Environmental Economics and Policy*, 5:109–130.

Dechezleprêtre, A., Ménière, Y., and Mohnen, M. (2017). International patent families: from application strategies to statistical indicators. *Scientometrics*, 111(2):793–828.

Doms, M., Dunne, T., and Troske, K. (1997). Workers, Wages, and Technology. *Quarterly Journal of Economics*, 112(1):253–290.

Eaton, J. and Kortum, S. (1996). Trade in ideas: patenting and productivity in the OECD. *Journal of International Economics*, 40:251–278.

Eaton, J., Kortum, S., and Kramarz, F. (2011). An anatomy of international trade: evidence from French firms. *Econometrica*, 79(5):1453–1498.

Fan, H., Hu, Y., and Tang, L. (2021). Labor costs and the adoption of robots in china. *Journal of Economic Behavior & Organization*, 186:608–631.

Gaggl, P. and Wright, G. (2017). A short-run view of what computers do: Evidence from a UK tax incentive. *American Economic Journal: Applied Economics*, 9(3):262–94.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *The American Economic Review*, 110(8):2586–2624.

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

Hall, B., Griliches, Z., and Hausman, J. (1986). Patents and r&d: Is there a lag? *International Economic Review*, 27(2):265–284.

Hausman, J., Hall, B., and Griliches, Z. (1984). Econometric models for count data with an application to the patents-R&D relationship. *Econometrica*, 52(4):909–938.

Hémous, D. and Olsen, M. (2022). The rise of the machines: Automation, horizontal innovation and income inequality. *American Economic Journal: Macroeconomics*, 14(1):179–223.

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

Kaufer, E. (1989). *The Economic of the Patent System*. New York: Rootledge.

Kogan, L., Papanikolaou, D., Schmidt, L., and Seegmiller, B. (2022). Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations. mimeo.

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

Krusell, P., Ohanian, L., Ríos-Rull, J.-V., and Violante, G. (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 68(5):1029–1053.

Lewis, E. (2011). Immigration, skill mix and capital skill complementarity. *Quarterly Journal of Economics*, 126 (2):1029–1069.

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

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.

Mann, K. and Puettmann, L. (2021). Benign effects of automation: New evidence from patent texts. *The Review of Economic and Statistics*, pages 1–45.

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

Popp, D. (2002). Induced innovation and energy prices. *The American Economic Review*, 92(1):160–180.

Siemens (2018). Annual report. Technical report.

Timmer, M., Dietzenbacher, E., Los, B., Stehrer, R., and de Vries, G. J. (2015). An illustrated user guide to the world input-output database: the case of global automotive production. *Review of International Economics*, 23(3).

Timmer, M., Erumban, A. A., Los, B., Stehrer, R., and de Vries, G. (2014). Slicing up global value chains. *Journal of Economic Perspectives*, 28(2):99–118.

Webb, M. (2020). The impact of artificial intelligence on labor market.

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

Online Appendix

Table of Contents

A Main Appendix

A.1 Additional Figures and Tables

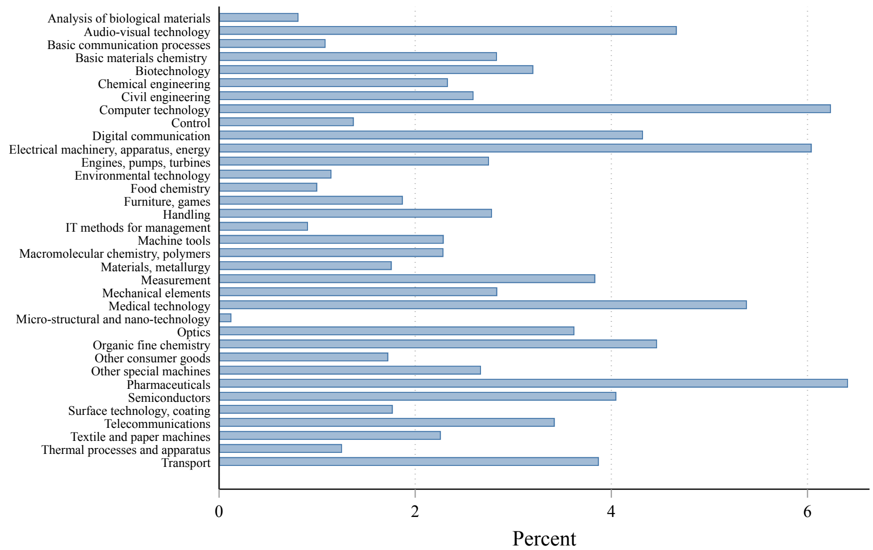

Figure A.1: Share of biadic patent applications in the different technical fields in 1997-2011.

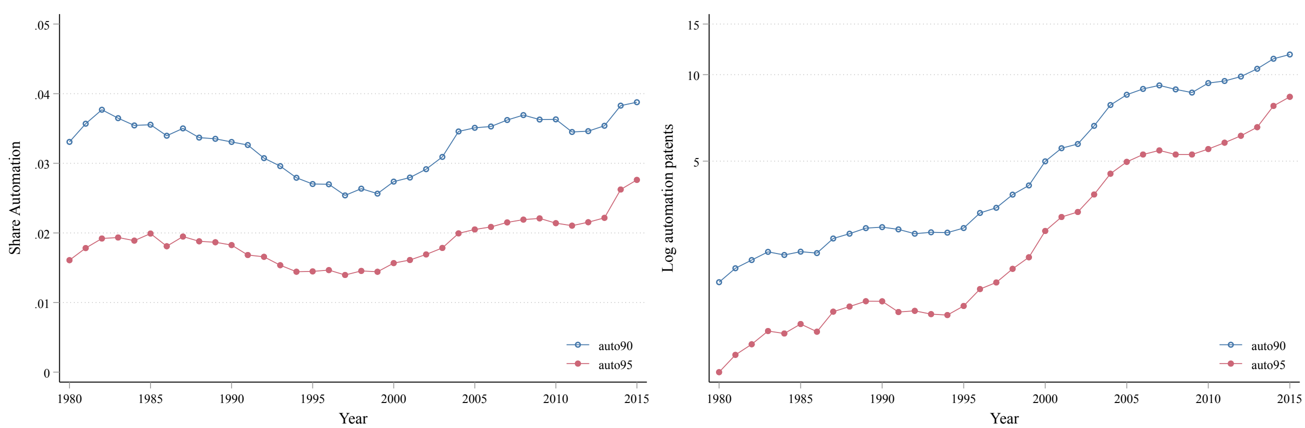

(a) Share of automation patents in machinery out of total patents according to the auto90 and auto95 definitions.

(b) Number of automation patents worldwide according to the auto90 and auto95 definitions.

Figure A.2: Trends in automation (for biadic applications)

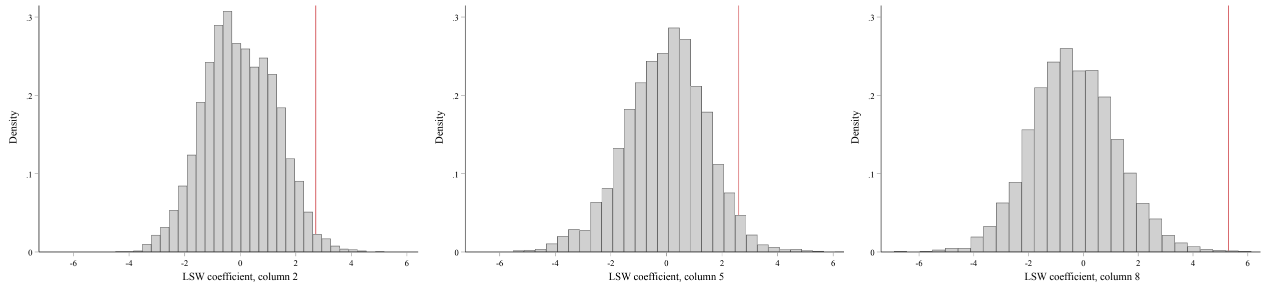

(a) Column 2 of Table 8 **(b)** Column 5 of Table 8 **(c)** Column 8 of Table 8

Figure A.4: Distribution of coefficients in Monte-Carlo simulations. 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. The figure reports histograms on the distribution of low-skill wage coefficients. The vertical red lines correspond to the coefficients of the true regressions. Each panel corresponds to a different column in Table 8.

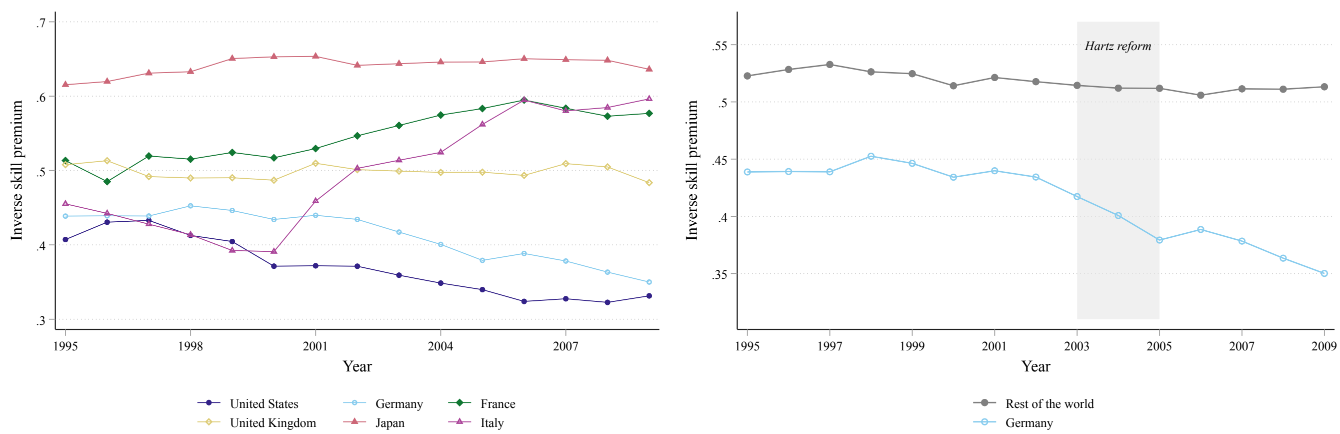

(a) Evolution of the inverse skill premium in the 5 most important countries of our main regressions.

(b) Evolution of the inverse skill premium in Germany and the rest of the world. Note: 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 3.a.

Figure A.3: Trends in the inverse skill premium.

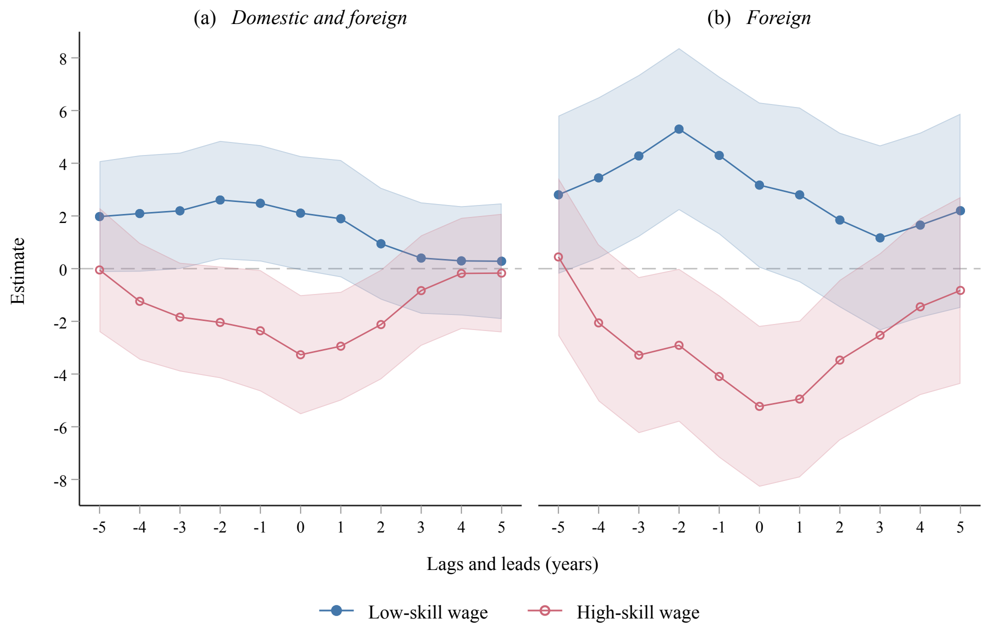

Figure A.5: Lag and leads. This figure reports regression coefficients on low-skill and high-skill wages at different lags and leads. Each panel and each year corresponds to a different Poisson regression of auto95 innovations on wages, GDP gap, labor productivity, stocks, spillovers, firm fixed effects, 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 4, and Panel b, year -2 corresponds to Column 8. The leads test for the presence of pre-trends.

Table A.1: Industry of innovators

	Industry	Share auto95 (%)	Share firms (%)
20	Manufacture of chemicals and chemical products	2.14	3.43
25	Manufacture of fabricated metal products, except machinery and equipment	1.18	4.42
26	Manufacture of computer, electronic and optical products	23.26	7.66
27	Manufacture of electrical equipment	9.47	2.9
28	Manufacture of machinery and equipment n.e.c.	24.29	21.11
29	Manufacture of motor vehicles, trailers and semi-trailers	5.32	3.55
30	Manufacture of other transport equipment	4.58	1.17
46	Wholesale trade, except of motor vehicles and motorcycles	1.32	3.31
64	Financial service activities, except insurance and pension funding	1.68	0.99
72	Scientific research and development	2.05	2.38
	Other industries	12.96	26.83
	No information on industry	11.75	22.22

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.2: 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

Note: 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.

Table A.3: Coverage of the regression sample

	Applications	Families	Biadic Families	Firms (with auto95 bia)
Patstat 1997-2011	430783	179025	60941	-
Matched with Orbis	347242	139538	51250	4231
Firms in sample	206313	85371	32397	3236

Notes: This table reports the number of auto95 patent applications, families, biadic families and firms 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: Top 10 auto95 innovators in our sample

Company	Number of biadic auto95 patents in 1997-2011
Siemens Aktiengesellschaft	1781
Honda Motor Co., Ltd.	815
Fanuc Co.	779
Samsung Electronics Co., Ltd.	718
Mitsubishi Electric Co.	669
Robert Bosch GmbH	663
Tokyo Electron, Ltd.	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 most auto95 patent families in our baseline sample.

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	0.608	0.749	1.000			
GDP gap	-0.063	-0.051	-0.032	1.000		
GDP per capita	0.709	0.805	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***	3.01***	2.72***	2.65***	2.24**	4.67***	4.19***	4.19***
	(0.76)	(0.76)	(0.80)	(0.98)	(0.76)	(1.01)	(1.33)	(1.32)	(1.33)
High-skill wage	-1.56**	-1.51**	-2.21***	-2.72***	-1.51**	-2.83***	-4.94***	-4.51***	-4.47***
	(0.68)	(0.65)	(0.73)	(0.93)	(0.65)	(0.97)	(1.39)	(1.33)	(1.32)
Stock automation		-0.11***	-0.12***		-0.11***	-0.12***		-0.11***	-0.12***
		(0.03)	(0.03)		(0.03)	(0.03)		(0.03)	(0.03)
Stock other		0.51***	0.51***		0.51***	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***			-0.97***
			(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	–	–	–	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47812	47812	47812	47453	47812	47453	47453	47453	47453
Number of firms	3236	3236	3236	3233	3236	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. 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 previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.7: Baseline regressions for auto95 with country-level clustering

	Auto95								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.97	2.72	3.64	2.24	2.61	3.64	4.19	5.30	4.43
	(0.70)	(0.77)	(1.11)	(0.73)	(0.55)	(1.59)	(0.86)	(1.65)	(1.79)
	[0.000]	[0.000]	[0.001]	[0.002]	[0.000]	[0.022]	[0.000]	[0.001]	[0.013]
	{0.027}	{0.000}	{0.001}	{0.039}	{0.054}	{0.061}	{0.016}	{0.022}	{0.005}
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	–	Yes	–	–	Yes	–	–	Yes	–
GDP per capita	–	–	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
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Firms	3236	3236	3236	3233	3233	3233	3233	3233	3233

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. The account for few clusters, the { } brackets report cluster-bootstrapped p-values following Cameron et. al (2008).

Table A.8: Auto90 innovations

Dependent variable	Auto90								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.33^{***}	2.06^{***}	3.29^{***}	1.69^{**}	1.72^{*}	2.80^{***}	3.26^{***}	3.83^{***}	3.87^{***}
	(0.67)	(0.69)	(0.79)	(0.83)	(0.90)	(1.07)	(1.14)	(1.34)	(1.47)
High-skill wage	-1.95^{***}	-2.44^{***}	-0.91	-1.79^{**}	-1.73^{*}	-1.05	-3.73^{***}	-2.88^{**}	-3.37^{***}
	(0.60)	(0.66)	(0.67)	(0.82)	(0.93)	(0.87)	(1.18)	(1.31)	(1.24)
GDP gap	-3.61^{*}	-4.27^{**}	-1.21	3.68	3.77	5.58	-0.32	0.92	0.89
	(2.09)	(2.15)	(2.25)	(5.28)	(5.36)	(5.47)	(3.27)	(3.54)	(3.71)
Labor productivity		1.12			-0.15			-1.36	
		(0.73)			(1.31)			(1.35)	
GDP per capita			-2.72^{**}			-2.73^{*}			-1.10
			(1.06)			(1.49)			(1.57)
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
Observations	71656	71656	71656	71367	71367	71367	71367	71367	71367
Number of firms	4821	4821	4821	4818	4818	4818	4818	4818	4818

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. 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 previously defined. Stock and spillover variables are calculated with respect to the dependent variable (auto90). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

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

Dependent variable	Pauto90			Pauto95			Auto95		
	Dom. and Fgn.		*Fgn.*	*Dom. and Fgn.*		*Fgn.*	*Dom. and Fgn.*		*Fgn.*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	0.73	0.33	1.01	0.95	0.49	1.60	2.11***	2.37**	4.75***
	(0.59)	(0.77)	(1.21)	(0.76)	(1.00)	(1.61)	(0.73)	(0.99)	(1.35)
High-skill wage	-0.22	-0.35	-0.62	-0.44	-0.43	-0.81	-2.15***	-2.13**	-2.94**
	(0.56)	(0.86)	(1.26)	(0.74)	(1.18)	(1.71)	(0.66)	(0.98)	(1.33)
GDP gap	-3.06**	1.34	0.38	-2.03	3.49	0.77	-2.55	2.13	3.83
	(1.35)	(3.39)	(2.33)	(1.57)	(4.16)	(2.87)	(2.24)	(5.54)	(4.19)
Labor productivity	-0.11	0.02	-0.91	-0.11	-0.58	-1.14	0.89	-1.46	-1.90
	(0.60)	(0.96)	(1.01)	(0.71)	(1.22)	(1.35)	(0.84)	(1.62)	(1.41)
Arcsinh pauto90							0.51***	0.51***	0.51***
							(0.02)	(0.02)	(0.02)
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
Observations	149580	149345	149345	43809	43686	43686	47812	47453	47453
Number of firms	10012	10009	10009	2932	2929	2929	3236	3233	3233

Notes: The independent variables are lagged by two periods. Standard errors are clustered at the firm-level and reported in parentheses. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). In columns 1–3 the dependent variable is pauto90 (machinery patents excluding auto90). In columns 4–6 the dependent variable is pauto95 (machinery patents excluding auto95), and the sample is restricted to the firms in the baseline auto95 regression. In columns 7-–9 the dependent variable is auto95 innovation and we control for contemporaneous placebo innovations, defined as the arcsinh of pauto90 patents. 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 previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.10: Predicted wages

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.44***	1.84**	2.46***	1.64*	1.56	1.65*	3.82***	4.24***	3.81***
	(0.82)	(0.82)	(0.82)	(0.94)	(1.02)	(0.94)	(1.30)	(1.41)	(1.31)
High-skill wage	-2.78***	-4.75***	-2.83***	-3.31***	-3.55**	-3.32***	-4.52***	-3.56**	-4.51***
	(0.83)	(1.08)	(0.83)	(1.04)	(1.42)	(1.04)	(1.33)	(1.53)	(1.34)
GDP gap	-4.40*	-3.77	-4.45*	4.67	4.66	4.68	-0.13	0.74	-0.10
	(2.61)	(2.56)	(2.61)	(6.80)	(6.81)	(6.80)	(4.55)	(4.59)	(4.58)
Labor productivity		2.85***			0.35			-1.59	
		(0.94)			(1.57)			(1.50)	
GDP per capita			0.14			0.03			-0.01
			(0.11)			(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	–	–	–	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Number of firms	3236	3236	3236	3233	3233	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. 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 previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.11: 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	0.98***		0.58		0.61		1.81***	1.34**	1.80**
	(0.31)		(0.40)		(0.60)		(0.48)	(0.62)	(0.90)
Δ High-skill wage	-0.98***		-1.19***		-1.62***		-1.58***	-1.80***	-2.82***
	(0.30)		(0.44)		(0.63)		(0.46)	(0.66)	(0.93)
Δ Low-skill / High-skill wages		0.98***		0.83**		1.08**			
		(0.25)		(0.33)		(0.52)			
Δ GDP gap	-1.54	-1.54	-0.68	-0.27	-1.01	-0.19	-2.54*	-2.95	-1.61
	(1.15)	(1.13)	(2.38)	(2.33)	(1.80)	(1.63)	(1.54)	(3.38)	(2.56)
Δ Labor productivity	-0.00	0.00	0.41	-0.17	0.86	0.06	-0.02	0.13	0.96
	(0.42)	(0.29)	(0.65)	(0.43)	(0.61)	(0.27)	(0.62)	(0.97)	(0.93)
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
Observations	32360	32360	32330	32330	32330	32330	21710	21690	21690
Number of firms	3236	3236	3233	3233	3233	3233	2171	2169	2169

Notes: Estimation is done by OLS. Standard errors are clustered at the firm-level and reported in parentheses. 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 the 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. Columns 1, 2, and 7 include industry-year fixed effects, while 3, 4, and 8 and include industry-year and country-year fixed effects. In Columns 3, 4, and 9 the macroeconomic variables are the normalized foreign variables previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.13: 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	–	–	Yes	–	–	Yes
Observations	132676	132676	132676	129440	129440	129440
Firms	3236	3236	3236	3236	3236	3236

Notes: 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. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.12: 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	5.30^{***}	5.42^{***}	3.53^{***}	8.97^{***}	6.13^{***}	6.16^{*}	7.49^{***}	1.68	-1.56
	(1.57)	(1.62)	(1.32)	(3.04)	(1.99)	(3.39)	(2.54)	(4.80)	(3.05)
High-skill wage	-2.91^{**}	-1.42	-2.11	-1.14	-2.13	-0.10	-3.06	6.49	1.75
	(1.48)	(1.63)	(1.32)	(2.95)	(1.80)	(3.12)	(2.37)	(6.12)	(3.61)
GDP gap	2.40	0.74	1.97	9.61	4.17	4.83	1.22	-1.69	-1.17
	(4.91)	(4.58)	(2.85)	(6.30)	(4.48)	(7.99)	(6.79)	(13.10)	(9.68)
Labor productivity	-2.53	-3.87^{**}	-1.78	-8.49^{***}	-4.53^{**}	-7.38^{***}	-5.70^{**}	-8.37	-1.03
	(1.61)	(1.71)	(1.22)	(2.50)	(1.76)	(2.83)	(2.25)	(5.50)	(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	47453	45838	97449	22517	48032	15049	23268	6476	13617
Number of firms	3233	3144	6544	1595	3272	1096	1632	508	1001

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. All regressions include firm fixed effects, industry-year and country-year fixed effects. 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 threshold used to delimit patents which is the 90th or the 80th percentile of the distribution of automation keywords for 6 digit C/IPC codes. The macroeconomic variables are the normalized foreign variables previously defined. Stocks and spillovers are computed with respect to the dependent variable. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.14: Alternative weights

	Auto95						
Weights	1971–1989	1985–1994	start 2000	pauto95	GDP^0	GDP^1	$(w_L \cdot L)^{0.35}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Foreign:							
Low-skill wage	5.41^{***}	5.12^{***}	6.64^{***}	5.55^{***}	4.15^{***}	6.15^{***}	5.30^{***}
	(1.92)	(1.52)	(2.11)	(1.72)	(1.40)	(1.71)	(1.54)
High-skill wage	-3.46^{**}	-1.39	-3.03	-2.94^{*}	-3.62^{***}	-3.26^{**}	-3.56^{***}
	(1.70)	(1.55)	(2.05)	(1.66)	(1.35)	(1.63)	(1.35)
GDP gap	0.85	3.41	0.69	7.62^{*}	-2.13	-0.76	-0.50
	(4.15)	(4.81)	(3.91)	(4.10)	(3.67)	(3.89)	(3.76)
Labor productivity	-2.48	-3.87^{**}	-4.80^{***}	-2.61^{*}	-1.60	-1.93	-2.23
	(1.79)	(1.62)	(1.78)	(1.55)	(1.44)	(1.59)	(1.57)
Stock automation	-0.13^{***}	-0.12^{***}	-0.31^{***}	-0.12^{***}	-0.12^{***}	-0.12^{***}	-0.12^{***}
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Stock other	0.60^{***}	0.53^{***}	0.49^{***}	0.58^{***}	0.51^{***}	0.51^{***}	0.51^{***}
	(0.05)	(0.05)	(0.06)	(0.05)	(0.04)	(0.04)	(0.04)
Stocks and spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33752	43262	25854	44672	47318	47457	47338
Number of firms	2319	2949	2624	3057	3230	3231	3234

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. All regressions include firm, country-year and industry-year fixed effects. Firms' country weights for the macroeconomic variables are computed over the period 1971–1989 in column 1; and over the period 1985–1994 for Column 2. Columns 3–7 use the baseline pre-sample period of 1971–1994. Column 3 restricts the sample to the years 2000–2009. Column 4 uses weights calculated using pauto95 patents applications (i.e., machinery patents excluding auto95); Column 5 does not adjust for GDP in the computation of the weights; Column 6 uses GDP instead of $GDP^{0.35}$ to adjust for country size and Column 7 replaces GDP with total low-skilled payment wL in the baseline formula. In all columns the macroeconomic variables are the normalized foreign variables previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

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***	2.68**	2.57**	4.80***	3.86***
	(0.80)	(0.78)	(1.07)	(1.02)	(1.46)	(1.39)
High-skill wage	-2.55***	-1.16	-2.22**	-1.74*	-2.76**	-2.17
	(0.78)	(0.80)	(1.02)	(1.00)	(1.40)	(1.47)
GDP gap	-4.32	-3.02	4.95	6.31	1.85	0.87
	(2.77)	(3.46)	(7.04)	(7.31)	(4.97)	(5.24)
Labor productivity	0.85	0.49	-1.48	-1.15	-1.92	-0.91
	(0.90)	(0.98)	(1.69)	(1.44)	(1.50)	(1.50)
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	–	–	Yes	Yes	Yes	Yes
Estimator	HHG	BGVR	HHG	BGVR	HHG	BGVR
Observations	47812	47812	47453	47453	47453	47453
Number of firms	3236	3236	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. 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 previously defined. Standard errors are clustered at the firm-level and reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.16: Citations-weighted patents

Dependent variable	Citations-weighted auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.03**	1.74	3.05***	1.27	1.64	3.35**	3.61**	4.52**	3.93*
	(1.00)	(1.11)	(1.14)	(1.23)	(1.46)	(1.53)	(1.71)	(1.87)	(2.23)
High-skill wage	-2.28**	-2.81***	-1.11	-3.15**	-2.22*	-1.72	-4.37***	-3.22*	-4.19**
	(0.96)	(0.98)	(1.09)	(1.30)	(1.32)	(1.42)	(1.61)	(1.88)	(1.74)
GDP gap	-2.95	-3.65	-0.38	0.66	1.97	4.41	-0.40	1.60	0.24
	(3.23)	(3.42)	(3.32)	(7.90)	(8.06)	(8.07)	(5.23)	(5.64)	(5.88)
Labor productivity		1.22			-2.06			-1.92	
		(1.22)			(2.29)			(1.82)	
GDP per capita			-2.98*			-5.15**			-0.56
			(1.63)			(2.41)			(2.66)
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
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Number of firms	3236	3236	3236	3233	3233	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. Patents are citations-weighted: we add to each patent the number of citations received within 5 years normalized by technological field and year of application. 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 previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.17: Firm bin size - year fixed effects

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	3.12***	2.84***	3.63***	2.37**	2.78**	3.71***	4.45***	5.71***	4.68***
	(0.79)	(0.85)	(0.96)	(0.99)	(1.13)	(1.27)	(1.32)	(1.56)	(1.78)
High-skill wage	-2.40***	-2.85***	-1.89**	-2.89***	-2.02*	-2.01*	-4.79***	-3.03**	-4.66***
	(0.72)	(0.78)	(0.81)	(0.95)	(1.08)	(1.05)	(1.33)	(1.48)	(1.42)
GDP gap	-2.83	-3.46	-1.67	4.46	5.55	6.75	-0.12	2.50	0.33
	(2.72)	(2.82)	(2.90)	(6.77)	(6.82)	(7.11)	(4.66)	(4.93)	(5.28)
Labor productivity		1.09			-2.00			-2.85*	
		(0.91)			(1.78)			(1.63)	
GDP per capita			-1.42			-3.28*			-0.41
			(1.34)			(1.99)			(2.10)
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	–	–	–	Yes	Yes	Yes	Yes	Yes	Yes
Observations	47812	47812	47812	47453	47453	47453	47453	47453	47453
Number of firms	3236	3236	3236	3233	3233	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. Standard errors are clustered at the firm-level and reported in parentheses. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). 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 defined previously. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A.18: 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	5.16***	4.48***	5.12***	5.85**	5.39***
	(1.54)	(1.43)	(1.96)	(2.79)	(2.06)
High-skill wage	-2.63*	-3.66**	-2.56*	-2.53	-3.42
	(1.40)	(1.43)	(1.49)	(2.34)	(2.30)
GDP gap	2.60	1.52	2.52	1.09	0.33
	(4.85)	(4.91)	(4.91)	(4.50)	(4.64)
Labor productivity	-2.71*	-1.39	-2.70*	-3.63	-3.01
	(1.54)	(1.57)	(1.64)	(3.10)	(2.93)
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	47453	47453	47453	47453	47453
Number of firms	3233	3233	3233	3233	3233

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. All regressions include firm fixed effects, industry-year fixed effects and country-year fixed effects. 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. In all columns, the macroeconomic variables are the normalized foreign variables previously defined. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

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 technological 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.

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

Table A.19: Summary statistics on the prevalence of keywords

	IPC/CPC 6 digit					IPC4 + (G05 or G06)					IPC4 pairs			
Share	All	Robot	Automat*	CNC	Labor	All	Robot	Automat*	CNC	Labor	All	Robot	Automat*	CNC
Mean	0.21	0.04	0.11	0.03	0.06	0.53	0.15	0.32	0.11	0.10	0.19	0.05	0.09	0.02
SD	0.15	0.09	0.10	0.06	0.04	0.19	0.18	0.11	0.17	0.04	0.16	0.10	0.10	0.05
P25	0.11	0.01	0.04	0.00	0.03	0.40	0.07	0.27	0.01	0.07	0.08	0.01	0.03	0.00
P50	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
P75	0.27	0.05	0.15	0.02	0.08	0.64	0.16	0.40	0.16	0.11	0.23	0.04	0.11	0.01
P90	0.40	0.09	0.25	0.06	0.11	0.78	0.36	0.43	0.38	0.15	0.37	0.09	0.22	0.04
P95	0.48	0.14	0.30	0.13	0.13	0.86	0.44	0.45	0.55	0.16	0.52	0.15	0.31	0.08
P99	0.76	0.60	0.46	0.33	0.19	0.90	0.83	0.60	0.57	0.18	0.84	0.59	0.45	0.22

Notes: This table computes summary statistics on the share of patents with any automation keywords, robot keywords, automat* keywords, CNC or labor keywords for each type of technological categories (6 digit codes, pairs of 4 digit codes and combinations of ipc4 codes with G05 or G06) v

Table A.19 gives summary statistics on the prevalence of automation keywords across technology categories in machinery, $p(t)$, as well as 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 or 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 also more common as the prevalence of automat* is significantly higher than that of the other keywords. Nevertheless, the difference narrows somewhat in the right tail: the 95^{th} percentile for 6-digit codes is 30% for automat* and 14% and 13% for robot and CNC. In fact, the thresholds (5 and 2) used in the definition of the automat* keywords were chosen such 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.20: Correlation between the main prevalence measures

Keywords	Automat	Robot	CNC	Labor
Automat	1.000			
Robot	0.379	1.000		
CNC	0.210	0.205	1.000	
Labor	0.394	0.224	0.085	1.000

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

Table A.20 shows the correlation between the prevalence of the 4 mains keyword 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.

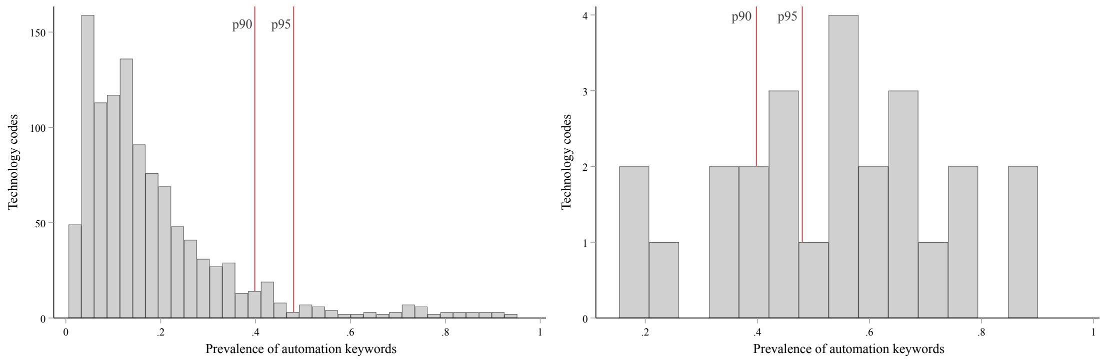

(a) For all pairs of C/IPC 4-digit codes within machinery

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

Figure A.6: Histograms of the prevalence of automation keywords. These 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 technological categories.

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 61,497 auto95 biadic patent families and 104,886 auto90 ones. Table A.21.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 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.21.b reports the share of auto95, auto90 and auto80 patents which belong to each subcategory. "Automat*" is the most important keyword: 71% 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.

Table A.21: Identification of automation technology categories

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

IPC codes / Patents	Auto90	Auto95
Matches ipc6	82.1%	83.3%
Matches ipc4 pair	16.4%	22.8%
Matches ipc4 - G05/G06 combination	40.7%	41.9%

Notes: Share of innovations classified as automation innovation through ipc6 codes, ipc4 pairs or ipc4 - G05/G06 pairs. Statistics 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	35.8%	53.6%	71.2%
CNC80	5.0%	8.4%	13.3%
Robot80	12.2%	20.3%	34.4%
Auto90	59.9%	100.0%	100.0%
Automat*90	10.6%	17.7%	27.1%
CNC90	1.8%	2.9%	5.0%
Robot90	7.6%	12.6%	21.5%
Auto95	35.2%	58.6%	100.0%
Automat*95	3.3%	5.5%	9.3%
CNC95	1.6%	2.6%	4.4%
Robot95	6.6%	10.9%	18.6%

Notes: Share of auto95 (auto90 and auto80, respectively) innovations which are also classified as automat*80/90/95, CNC80/90/95, and robot80/90/95 innovations. Statistics computed on biadic patents from 1997-2011.

Table A.22: 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	0.958	1.000		
1978-1997	0.906	0.851	1.000	
1998-2017	0.963	0.972	0.851	1.000

Notes: Correlation matrix for the prevalence of automation keywords by C/IPC 6-digit codes in machinery using EPO patents over different time periods. We exclude catch-all categories made at the 4-digit level.

Table A.23: Confusion table for different classification periods

Confusion Matrix		Auto95 based on the 1978-1997 classification		Auto95 based on the 1998-2017 classification		Auto95 based on the 1997-2011 classification		Total
		Yes	No	Yes	No	Yes	No	
Auto95 based on	Yes	51,243	10,254	55,290	6,207	52,027	9,470	61,497
the 1978-2017	No	4,378	3,121,661	5,243	3,120,796	5,752	3,120,287	3,126,039
classification	Total	55,621	3,131,915	60,533	3,127,003	57,779	3,129,757	3,187,536

Notes: This table classifies all biadic patent families from 1997-2011 as auto95 or not, but using EPO patents from different time periods to classify technological categories as automation or not. 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.

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 main regression analysis (1997-2011). There is a modest increase in the share of patents with automation keywords within each technology category. At the C/IPC 6-digit level in machinery, the share of patents with an automation keyword increases on average from 0.19 in the first half of the sample to 0.21 in the second half. Nevertheless, the ranking of codes is remarkably stable as shown in Table A.22 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.23 shows confusion tables on the classification of patents as auto95 according to each of the classification period. Regardless of the time period used the number of automation patents stays roughly constant. In particular, 84.6% 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% 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.24 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 subcategories (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

Table A.24: Examples of 6-digit C/IPC codes in machinery

Code	Description	# Patents	Any	Rank	Robot	Automat*	CNC	Labor
				High Prevalence				
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	0.07	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	0.65	0.06
A01J7	Accessories for milking machines or devices	395	0.77	9	0.62	0.52	0	0.1
G05B19	Programme-control systems	7133	0.70	17	0.22	0.39	0.25	0.08
B65G1	Storing articles, individually or in orderly arrangement, in warehouses or magazines	1064	0.58	30	0.18	0.46	0.01	0.11
				Low Prevalence				
B23P6	Restoring or reconditioning objects	613	0.26	262	0.07	0.06	0.05	0.09
A01B63	Lifting or adjusting devices or arrangements for agricultural machines or implements	264	0.24	301	0.01	0.20	0	0.04
B66D3	Portable or mobile lifting or hauling appliances	215	0.13	665	0.02	0.07	0.01	0.06

Notes: Prevalence of automation keywords for a few 6 digit C/IPC codes. "Any" is the share of patents with any of the keywords. "Rank" is the rank of the code among 1009 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.

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.24). This patent itself contains several keywords: a sentence with the words "automatic" and "storing," and another sentence with "robot".

(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

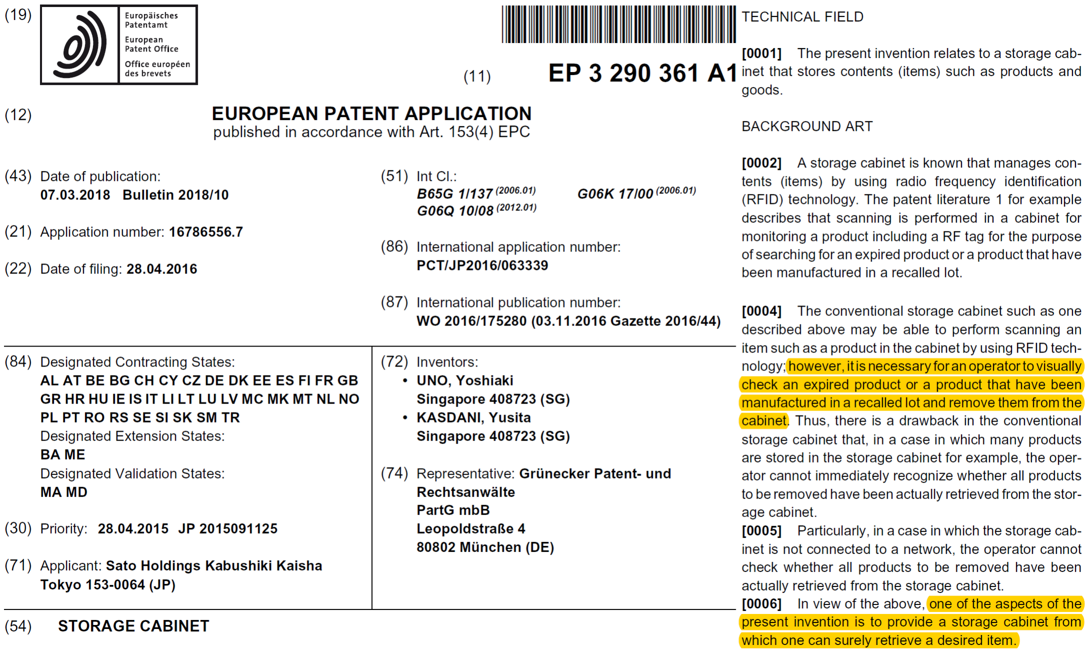
(19) Europäisches Patentamt European Patent Office Office européen des brevets

(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

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.

A.3 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 machinery patents (according to our definition) of MP and classified them as auto95 or not. We have a lower share of automation patents (18% for auto90 and 9.9% for auto95) than MP who have 31%. 71.5% 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 22.9% of their automation patents as auto95. Therefore, our measure of automation is generally stricter than theirs although it is not a perfect subset.

To facilitate comparison, we compute the share of automation patents at the C/IPC 6-digit level according to their classification and compare this number with our measure

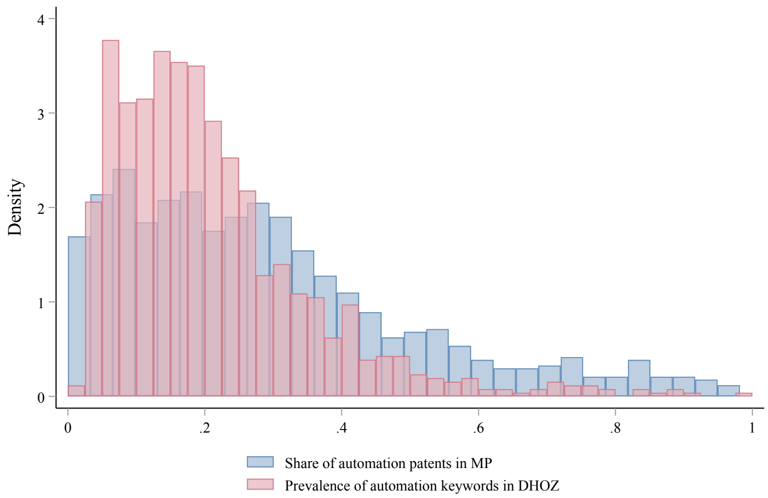

Figure A.9: Histograms of the share of automation patents in MP and of the prevalence of automation keywords in this paper at the 6-digit level in machinery.

of the prevalence of automation keywords. The correlation between these two measures is high (at 0.59). Figure A.9 shows the histograms of the two distributions. Our prevalence measure is more skewed (with a kurtosis of 7 versus 3.5), and as such, it more clearly identifies a small set of outliers among 6-digit 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.25 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 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. 3 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 0.38, however their "raw share" (computed before they exclude certain patents) is quite high at 0.71. The excluded patents are not chemical or pharmaceutical patents (as emphasized

Table A.25: Outliers 6-digit C/IPC codes in the comparison between our measure and MP's measure

Code	Simplified description	Prevalence of automation keywords (DHOZ)	Share of automation patents (MP)
	Positive outliers among auto95 codes		
B25J17	Manipulators (joints)	0.84	0.54
D01H9	Textile machines (arrangements for replacing or removing various elements)	0.62	0.38
B25J15	Manipulators (gripping heads)	0.71	0.50
B23P23	Metal working machines (specified combinations n.e.c)	0.67	0.46
B25J19	Manipulators (accessories)	0.89	0.69
B33Y70	3D printing materials	0.52	0.32
	Negative outliers among non-auto90 codes		
B66B2201	Control systems of elevators	0.19	0.97
B66B3	Elevators (signalling and indicating device applications)	0.19	0.92
B41J23	Typerwriters / printing machines (power drive)	0.08	0.82
B66B1	Elevators (control systems)	0.16	0.89
B41J19	Typerwriters / printing machines (characters and line spacing mechanisms)	0.14	0.84
B41J5	Typerwriters / printing machines (controlling character selection)	0.21	0.91

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

in the paper), but belong to the "other" technological field (according to the Hall-Jaffe-Trajtenberg classification). 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. Finally, B33Y70 refers to materials adapted for additive manufacturing (e.g. 3D printing). These are typically inputs for automation machines but not full machines themselves, which may explain why they are excluded by MP. Regardless, very few of our auto95 patents are identified through the 3D printing keyword.

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 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.4 Reproducing ALM

We detail how we build the variables used in Section 2.6 and provide further results.

A.4.1 Data for the ALM exercise

Except for the automation 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 2 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.

Table A.26: 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 code and description		Auto95	Industry code and description		Auto95
756	Automotive services and repair shops	0.110	801	Bowling alleys, billiard and pool parlors	0.042
206	Household appliances (e.g., radio, TV, equipment)	0.106	100	Meat products	0.046
470	Water supply and irrigation	0.098	102	Canned and preserved fruits and vegetables	0.046
271	Iron and steal foundaries	0.096	110	Gain milk products	0.046
130	Tobacco manufactures	0.093	112	Sugar and confectionary products	0.046
212	Misc. plastic products	0.093	101	Dairy products	0.046

Notes: The share is the share of auto95 (95th percentile threshhold) patents out of all patents in machinery in 1980-1998 in the respective sector. The industry codes and descriptions are SIC 1987.

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.[43] We dropped commodities associated with the construction sector which are structures. Combining the two BEA tables, we obtain an investment flow table at the industry level. We combine that table with the C/IPC to industry of manufacturing table previously derived to get an C/IPC to industry of use table mapping 656 4-digit C/IPC codes into 966 SIC industries.

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

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

patents over the time period 1980-1998 and take the ratio to be our measure of automation intensity. Table A.26 shows the sectors with the highest and lowest shares of automation patents in machinery.

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.

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 "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. 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.4.2 Additional results

We now provide a few additional results which complements those in the main text. 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 Table 2 Columns (1), (3) and (5); the only difference being that the figure does not control for computerization.

We carry a number of robustness checks in Table A.27. In Columns (1), (4) and (7), we compute the share of automation patents using only granted USPTO patents which are also biadic. The results are similar to those in Table 2 though less precise for the skill ratio. In Columns (2), (5) and (8), we use the share of auto90 patents in machinery

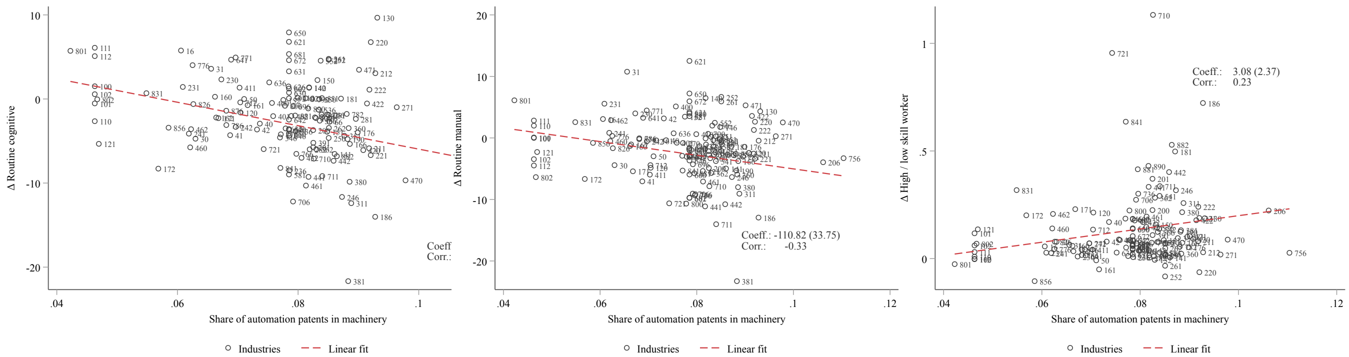

(a) Change in routine cognitive tasks and automation intensity

(b) Change in routine manual tasks and automation intensity

(c) Change in skill composition and automation intensity

Figure A.10: Scatter plots of routine tasks and skill composition changes and automation intensity (auto 95) in 1980-1998 in the United States.

Table A.27: Changes in routine task intensity and different measures of sectoral automation

	Δ Routine cognitive			Δ Routine manual			Δ High/low skill workers		
	Biadic	Auto90	Lybbert and Zolas	Biadic	Auto90	Lybbert and Zolas	Biadic	Auto90	Lybbert and Zolas
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Share automation	-134.82***	-80.84***	-20.69***	-103.76***	-58.13***	-10.97**	2.67	1.86	0.45*
	(28.23)	(19.85)	(5.74)	(32.38)	(18.29)	(4.80)	(1.88)	(1.16)	(0.23)
Δ Computer use	-20.03***	-17.56**	-19.34*	-20.88***	-18.74**	-13.09	0.99***	0.96***	0.58*
(1984-1997)	(7.29)	(7.21)	(10.43)	(7.79)	(7.80)	(8.60)	(0.26)	(0.26)	(0.29)
R^2	0.23	0.22	0.39	0.17	0.15	0.22	0.17	0.18	0.28
Mean dependent variable	-2.5	-2.5	-2.7	-2.27	-2.27	-1.35	.12	.12	.1
Observations	133	133	71	133	133	71	133	133	71

Notes: Each column represents a separate OLS regression of ten times the annual change in industry-level task input between 1980 and 1998, measured in centiles of the 1960 task distribution, on the share of automation patents in machinery, the annual percentage point change in industry computer use during 1984-1997, and a constant. Estimates are weighted by mean industry share of total employment in FTEs in 1980 and 1998. Robust standard errors are reported in parentheses. In columns 1–3 the dependent variable is the change in routine cognitive tasks, in columns 4–6 the change in routine manual tasks, and in columns 7–9 the change in the ratio of high-skill workers (college graduates) over low-skill workers (others). Biadic uses only biadic auto95 patents, Auto90 defines automation patents as auto90 patents. In both cases, patents are allocated to their sector of use. Lybbert and Zolas uses auto95 patents and allocates patents using a concordance table between C/IPC codes and industries from Lybbert and Zolas (2014). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

to measure automation in the sector of use. The results are similar but with smaller coefficients than in the regressions using auto95, in line with auto95 being a stricter measure of automation. In Columns (3), (6), and (9), 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 0.089, the standardized coefficients are relatively similar.

In unreported regressions, we also find that our results are robust to considering the different time periods analyzed by ALM: the 1970s, 1980s and 1990-1998.

A.5 A Simple Model

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

$$y_i = h_{1,i}^{1-\beta}\left(\gamma\left(i\right)l_i + \alpha\left(i\right)\nu^{\nu}(1-\nu)^{1-\nu}x_i^{\nu}h_{2,i}^{1-\nu}\right)^{\beta}. \tag{5}$$

$\gamma(i)$ is the productivity of low-skill workers, $\alpha(i)$ is an index which takes the value 0 for non-automated intermediates and 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, for each non-automated intermediate i, there is an innovator. The innovator creates a machine specific to intermediate i with probability λ if she spends $\theta\lambda^{\psi+1}Y/(\psi+1)$ units of the manufacturing 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 manufacturing 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:

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

This expression is increasing in the low-skill wage w_L and decreasing in the high-skill wage w_H with a magnitude which is larger 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 one 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.6 Data Appendix for the main analysis

This Appendix provides details on the data and the variable construction for our main analysis.

A.6.1 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 further contains information on GDP deflators and PPIs, both for manufacturing and for the whole economy. They employ the ISCED skill-classification, where category 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 denotes high-skill (college and above). Switzerland is not included in the WIOD database and we add data on skill-dependent wages, productivity growth and price deflators using data obtained directly from *Federal Statistical Office of Switzerland.*

We supplement this data with 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.*[44] For the US, we use data from FRED for state minimum wages and calculate the nation-level

[44] 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.

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*.

A.6.2 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 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 in our regressions.

A.6.3 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.7 Shift-share analysis

This appendix presents a number of additional results related to our shift-share set-up. We first do a "shock-level" analysis as recommended by Borusyak et al. (2022, henceforth BHJ), then we show that our results do not depend on a single country and 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:

$$ISP_{i,t} \equiv \sum_c \kappa_{i,c} \ln\left(\frac{w_{L,c,t}}{w_{H,c,t}}\right). \tag{6}$$

Table A.28: 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***	0.32**	0.36***
	(0.87)	(0.15)	(0.08)	(0.16)	(0.07)
Labor productivity				-0.32	
				(0.50)	
GDP gap				-0.30	
				(1.88)	
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	47453	48495	615	615	615
Firms / Countries	3233	3233	41	41	41

Notes: The independent variables are lagged by two periods. 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 regressions include firm fixed effects, industry-year fixed effects and country-year fixed effects. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

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 (3) with:

$$\begin{aligned} & \text{arcsinh}\left(PAT_{Aut,i,t}\right) \qquad (7) \\ = \quad & \begin{array}{c} \beta_{ISP} ISP_{L,i,t-2} + \beta_X X_{i,t-2} + \beta_{Ka} \ln K_{Aut,i,t-2} + \beta_{Ko} \ln K_{other,i,t-2} \\ +\beta_{Sa} \ln SPILL_{Aut,i,t-2} + \beta_{So} \ln SPILL_{other,i,t-2} + \delta_i + \delta_{j,t} + \delta_{c,t} + \epsilon_{i,t} \end{array}. \end{aligned}$$

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.28 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 (6). 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 6. Column (2) runs a linear regression at the firm level as in (7). We obtain a similar result – the magnitude is smaller as the range of variations for arcsinh is smaller than for the log function.

Column (3) follows the BHJ approach and runs a shock-level regression. That is, 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 weighted 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.[45]

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.28 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.29 reports summary statistics on the shock-level regressions following BHJ's recommendation. In line with Table A.5 and the distribution shown in Figure A.11, 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

[45]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 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.11).

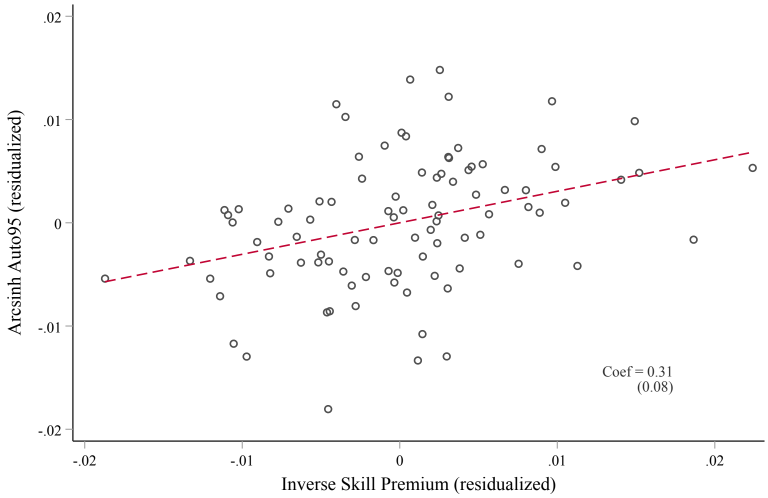

Figure A.11: Bin-scatter plot of the shock-level regression. We residualized both arcsinh(auto95) and the inverse skill premium on firm, industry-year and country-year fixed effects and on stocks and spillover variables. We then compute weighted average of the residuals at the shock (i.e. country-year)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

Table A.29: Shock-level summary statistics

	(1)	(2)	(3)	(4)
Panel A: Descriptive statistics on the inverse skill premium				
Mean	-0.78	0	0	0
Standard deviation (%)	36.4	2.1	0.9	1.0
Interquartile range (%)	55.7	2.9	1.1	1.0
Residualizing on				
F fixed effect		Yes	Yes	Yes
IY+CY fixed effects			Yes	Yes
Stocks/Spillovers				Yes

Panel B: Herfindahl–Hirschman index of weights

	Total weights	Foreign weights
Country	0.133	0.090
Country and year	0.009	0.006

Notes: Panel A reports descriptive statistics for the log inverse skill premium weighted by the average country weight in our regression sample as in Borusyak et al. (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). Panel B reports the Herfindahl–Hirschman index of weights at the country and country-year level for both the total weights and foreign weights (normalized to sum up to 1).

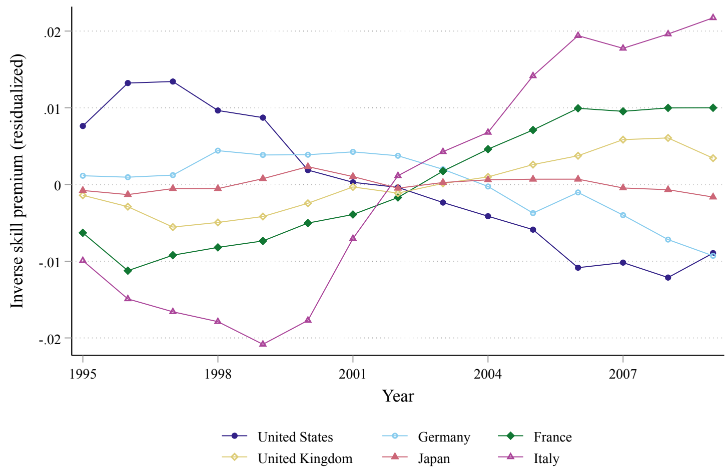

Figure A.12: Residualized inverse skill premium in the 6 most important countries. This figure reports our identifying shocks: namely the log inverse skill premium residualized on firm fixed effects, industry-year and country-year fixed effects, stocks and spillovers variable and aggregated at the country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

the log inverse skill premium residualized only on firm fixed effects (i.e. only taking away level differences across countries) is 2.1%. The Herfindahl index of the weights at the country level is 0.13 (as reported in Table 3, the largest weight is the US weight with a value of 0.21). With 15 years, however, the Herfindahl index of the weights at the country-year level is 15 times smaller (at 0.009). For foreign weights, the Herfindahl index are 0.09 at the country-level 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.

Figure A.3.a shows the evolution of the inverse skill premium for the 6 countries with the largest average weights in the raw data. 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.28). 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 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.30, we look at the balance of our shocks against observables. 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,

Table A.30: Shock balance tests

	Estimate (1)	SE (2)
GDP Gap	0.00	(0.01)
Labor Productivity	-0.22	(0.17)
GDP per capita	0.04	(0.19)
Manufacturing size	-0.11	(0.10)
Offshoring	0.01	(0.03)
Recent auto95 innovation	-1.01***	(0.38)
Recent other innovation	-1.35**	(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.

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. More automation innovations are associated with a higher skill premium as one would expect. This holds true for all other innovations—note that these include non machinery innovations such as innovations in computers, 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 also include the weight of the excluded country times a year dummy as a control. Table A.31 reports the results (with foreign wages and controlling for labor productivity). The coefficient on low-skill wages always remains positive and significant.[46]

[46] Goldsmith-Pinkham, Sorkin and Swift (2020) suggest carrying out a similar exercise by excluding countries with a large Rotemberg weight. Rotemberg weights require a linear shift-share instrument, we check that 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: 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	5.30***	5.61***	3.80***	3.59***	4.95***	3.55**	5.47***	5.02***
	(1.57)	(1.70)	(1.41)	(1.34)	(1.34)	(1.51)	(1.48)	(1.54)
High-skill wage	-2.91**	-2.42*	-1.77	-1.56	-0.77	-2.15	-4.61**	-2.39
	(1.48)	(1.47)	(1.32)	(1.32)	(1.36)	(1.33)	(1.93)	(1.51)
GDP gap	2.40	2.37	3.44	2.50	3.19	2.03	2.05	2.09
	(4.91)	(5.08)	(5.63)	(3.95)	(4.90)	(5.05)	(5.22)	(4.97)
Labor productivity	-2.53	-4.00**	-2.52*	-1.73	-3.64**	-1.86	-1.12	-2.77*
	(1.61)	(1.68)	(1.39)	(1.50)	(1.60)	(1.49)	(1.66)	(1.58)
Excluded country weight × year dummy	–	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	47453	46677	46984	47274	47045	47393	47318	47382
Number of firms	3233	3181	3199	3221	3206	3229	3224	3228

Notes: The independent variables are lagged by two periods. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Standard errors are clustered at the firm-level and reported in parentheses. All columns include firm, industry-year and country-year fixed effects. In column 0, the macroeconomic variables are the normalized foreign variables previously defined. 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. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Borusyak and Hull (2021). Borusyak and Hull (2021) show that a regression using a logged shift-share measure may be biased due to the non-linearity of the log function. Table A.28 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.[47] The results remain very similar.

[47]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.32: Borusyak and Hull (2021)'s correction

	Auto95								
	Domestic and foreign						*Foreign*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.44***	2.32***	3.84***	1.57	2.23**	4.06***	5.24***	5.45***	3.44**
	(0.78)	(0.85)	(0.96)	(0.96)	(1.11)	(1.26)	(1.45)	(1.52)	(1.68)
High-skill wage	-2.07***	-2.27***	-0.94	-2.76***	-1.28	-1.60	-3.81***	-3.54**	-3.87***
	(0.72)	(0.78)	(0.80)	(0.95)	(1.05)	(1.03)	(1.25)	(1.57)	(1.24)
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity		Yes			Yes			Yes	
GDP per capita			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
Observations	47783	47783	47783	47424	47424	47424	47424	47424	47424
Number of firms	3234	3234	3234	3231	3231	3231	3231	3231	3231

Notes: The independent variables are lagged by two periods. Standard errors are clustered at the firm-level and reported in parentheses. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). All columns include firm and industry-year fixed effects. Columns 4—9 add country-year fixed effects. Columns 7-–9 use the normalized foreign macro variables previously defined. All regressions include controls for stocks and spillovers. The macroeconomic variables in log, low-skill wages, high-skill wages, GDP per capita and labor productivity are adjusted following the correction suggested in Borusyak and Hull (2021): we sample with replacement the entire path of macroeconomic variables for each firm with 1000 draws, take the average value and subtract it from the original macroeconomic variable. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

A.8 Macroeconomic interpretation of the regression coefficients

In this section, we analyze the economic magnitude of our regression coefficients by estimating the effect of a change in the skill premium on tasks demand that runs through automation. To do that, we can combine our regressions results with the results of Section 2.6. Two issues prevent us from simply directly multiplying coefficients: First, our regressions control for the stock of knowledge of firms and the knowledge spillovers that they are potentially subject to – both change as a change in wages affect (all) firms' innovation decisions. Second, in Section 2.6, we looked at the effect over a decade of changes in the share of automation patents over total machinery patents, so we need to jointly estimate the effect on automation innovations and other machinery innovations.

Therefore, we run a simulation where we consider a uniform and permanent decrease in the skill premium by 10% between 1995 and 2009 in all countries. We use our regression results to recompute the share of automation innovations in machinery over that period. Importantly, we stress that one *must not* interpret the result of this simulation as predictive, notably because a change in innovation should in turn affect the skill premium. Yet, our analysis could be used to calibrate a model which predicts that the direction of innovation reacts to changes in the skill premium. We focus on a changes in the skill premium as it is easier to interpret than a change in low-skill wages keeping high-skill wages constant.

Specifically, we consider the regression results reported in Panel a of Figure A.13. Given our goal of computing changes in the share of automation innovations in machinery, this regression differs slightly from the ones in the paper. We regress both auto95 innovations and all other machinery innovations (pauto95) on the inverse of the skill premium, the GDP gap, stock and spillover variables and firm and industry-year fixed effects and we consider separately the stocks and spillovers of auto95 innovations, pauto95 innovations and all other innovations.[48]

Figure A.13 reports the results averaged over 1000 simulations (using the median gives similar results).[49] 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 counterfactual (direct effect) curve. This gap reflects the elasticity of 2.36 of auto95 innovations with respect to the inverse skill premium (with an elasticity of 0.26 for other machinery innovations). Taking into account the response of firms' own innovation stocks slightly decreases the effect of low-skill wages reflecting the negative effect of the automation stock on auto95 innovations and its positive effect on other machinery innovations.

We then assess the importance of knowledge spillovers by recomputing the spillover variables for the auto95 innovations and other machinery innovations (but not the non-machinery innovations). This involves two complications. First, our model only applies to the number of innovations and not their location. To allocate innovations to countries, we assign the simulated innovations proportionally to the firm's inventor weights (used to construct the spillover variables). Second, firms in our sample account for only 53.2% of all biadic innovations in 1997-2011. We assume that the other firms respond similarly, so that when we assign simulated innovations to countries, we increase innovations by out-of-sample firms to keep the ratio of in-sample to out-of-sample innovations constant.

The overall effect of an increase in the inverse skill premium is then captured by the gap between the baseline curve and the counterfactual one. The baseline curve and the data series differ because the baseline is an average while the data series is only one

[48] For technical reasons, we also use $ln(1+)$ for spillovers and stocks instead of replacing 0's with 1's and adding a dummy for 0 stocks or spillovers as in the baseline regressions.

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

Dependent variable	Auto95	Pauto95
	(1)	(2)
Low-skill / High-skill wages	2.36***	0.26
	(0.68)	(0.50)
GDP gap	-4.51*	-2.98**
	(2.69)	(1.35)
Stock automation	-0.15***	0.12***
	(0.05)	(0.03)
Stock non-automation	0.34***	0.28***
	(0.06)	(0.03)
Stock other	0.34***	0.25***
	(0.06)	(0.04)
Spillovers automation	1.04***	-0.14
	(0.36)	(0.21)
Spillovers non-automation	1.12*	2.24***
	(0.61)	(0.38)
Spillovers other	-1.67**	-1.89***
	(0.74)	(0.49)
Firm fixed effects	Yes	Yes
Industry × year fixed effects	Yes	Yes
Observations	47812	155183
Number of firms	3236	10382

Notes: The independent variables are lagged by two periods. Standard errors are clustered at the firm-level and reported in parentheses. The coefficients are estimated with conditional Poisson fixed effects regressions (HHG). Both regressions include firm and year-industry fixed effects. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

(a) Supporting regression

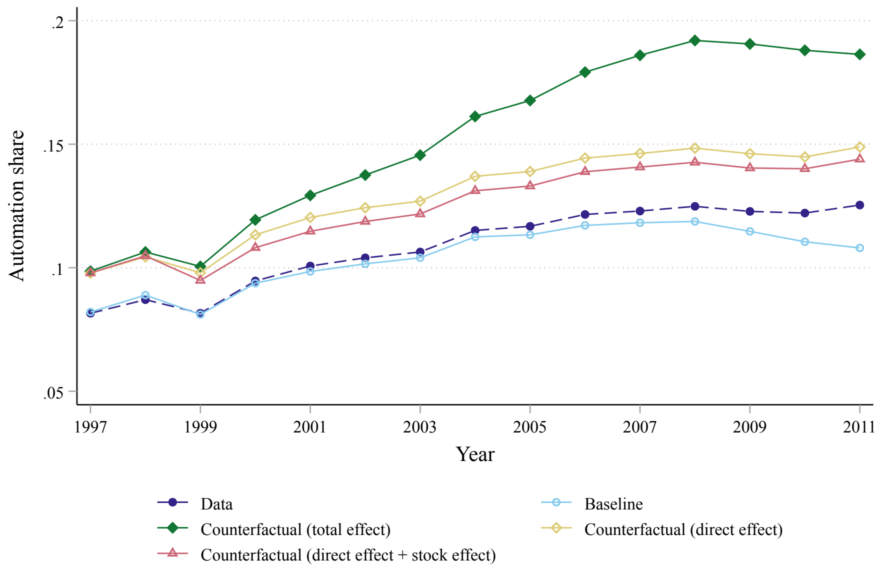

(b) Simulation result

Figure A.13: Simulation of a permanent and global 10% decrease in the skill premium on the share of automation innovations in machinery possible realization. Knowledge spillovers increase the overall elasticity of the share of automation patents with respect to low-skill wages. The average share of automation innovations in machinery between 1997 and 2011 increases by 4.8 p.p. from 10.4% to 15.3%. This is 2.7 p.p. more than the direct effect. This 4.8 p.p. increase can be compared to the 4.4 p.p. increase in the data over the same time period. As mentioned in Section 4.3, we can then combine these effects with the results of Section 2.6, and obtain that this 4.8 p.p. increase in the share of automation innovation would be associated with a decline in routine cognitive tasks of 7.5 centiles and a decline in routine manual tasks of 6.2 centiles. Though one should not interpret these numbers as causal, they indicate that the effect of the skill premium on automation innovations is economically significant.

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.24). 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.24). 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.

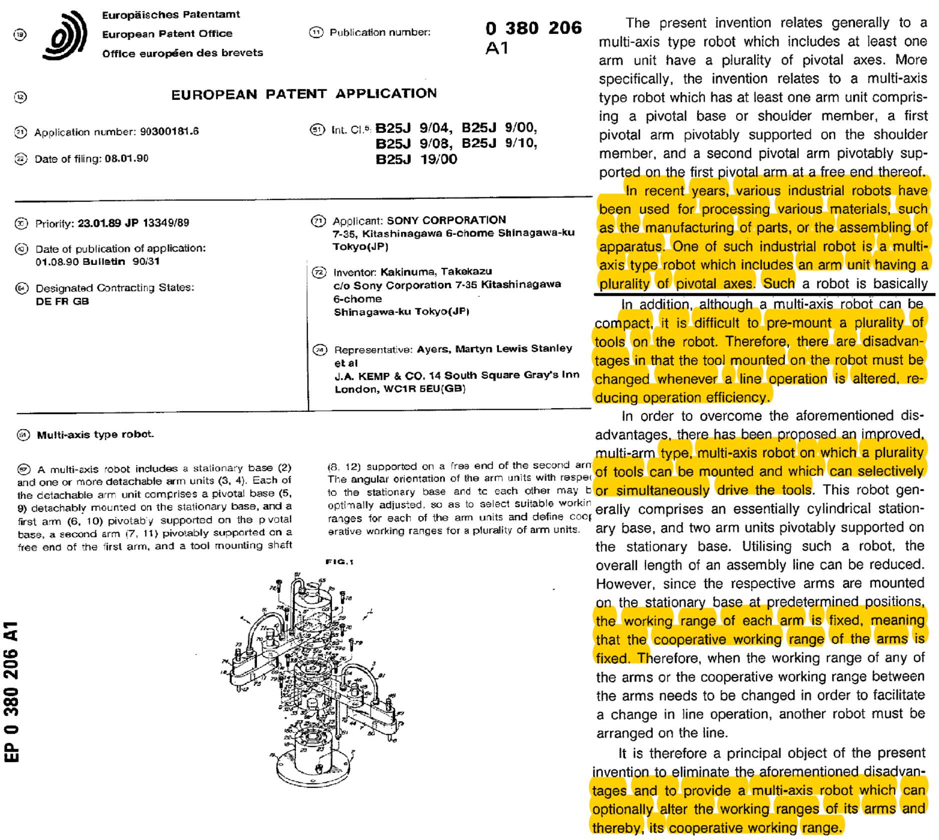

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(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 respe 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 coo 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**

(51) 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**

(84) Designated Contracting States: **DE ES FR GB IT**

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

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

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

(54) **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

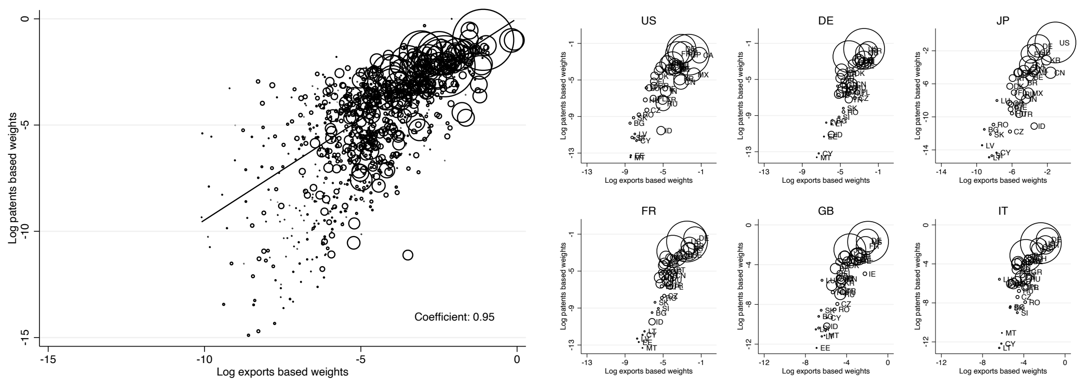

(a) All countries **(b)** Trade from the 6 largest countries

Figure B.8: Bilateral patent flows and trade flows in machinery. 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.[50] This gives us a machinery weight for each industry code and each country. We then

[50]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.

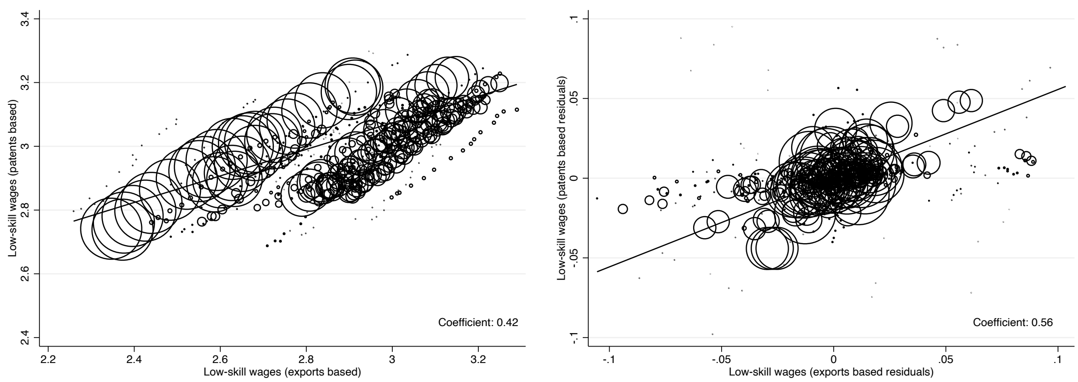

(a) Low-skill wages. **(b)** Residualized low-skill wages

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

multiply sectoral trade flows (after having aggregated the original data to the NACE 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.