# Induced Automation Innovation: Evidence from Firm-Level Patent Data

David Hémous

*University of Zurich and Centre for Economic Policy Research*

Morten Olsen

*University of Copenhagen*

Carlo Zanella

*Swiss Life Asset Managers*

Antoine Dechezleprêtre

*Organization for Economic Cooperation and Development*

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.

This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation program (grant agreement 805007). We

Electronically published April 11, 2025

*Journal of Political Economy*, volume 133, number 6, June 2025.

https://doi.org/10.1086/734778

## I. Introduction

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

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

Our classification aims to identify automation innovations that allow for the replacement of workers with equipment in some tasks. We focus on patents in machinery (e.g., machine tools, handling machines, textile and paper machines.), which account for around 18% of all patents. To classify patents, we follow a two-step procedure. First, we use the text of European patents to identify technology categories in machinery (defined using International Patent Classification [IPC] and Cooperative Patent Classification [CPC] codes) associated with automation. Second, we use these technology categories to classify the universe of machinery patents. This procedure leverages that the combined wording of many patents improves the signal of automation characteristics and allows us to classify patents for which we do not have the text. The resulting classification is transparent, uses highly disaggregated technology categories, and covers a wide range of automation technologies. As a validation exercise, we map our patents to sectors of use and reproduce the cross-sectoral analysis of Autor, Levy, and Murnane (2003). We find that in the United States,

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, Patrick Sproll, and Shi Suo for fantastic research assistance. We thank the editor Rachel Griffith and three anonymous referees for valuable comments and suggestions. Hémous, Olsen, and Zanella are in alphabetical order. We acknowledge the contributions of Dechezleprêtre in the initial stages of this project, which were instrumental in shaping the research direction and framework. The first draft of this paper was written in March 2019.

sectors that use more automation-intensive equipment experience a relative decline in routine tasks and the labor share and a relative increase in the skill ratio.

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

We implement this empirical strategy as follows. We rely on the World Patent Statistical Database (PATSTAT), which contains close to the universe of patents. We link patents to firms and apply our classification of automation and nonautomation patents in machinery. To proxy for firms' international exposure, we use the geographic distribution of their machinery patents before the sample period. We combine these exposure weights with macroeconomic data from 41 countries. Given our focus on international innovation, we restrict attention to patents applied for in at least two countries. Our sample covers the period 1997–2011 and includes 3,236 firms, accounting for 53.5% of global automation innovations in machinery. We run Poisson regressions with a 2-year lag between patent applications and labor costs, and we include firm, industry $\times$ year, and country $\times$ year fixed effects.

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

We supplement this analysis by focusing on two cases where labor cost changes arise directly from policy interventions. First, we build a measure of minimum wages for a subset of countries. We find a positive effect of minimum wages on automation innovations. Second, we focus on a specific labor market shock: the Hartz reforms in Germany. These were a series of labor market reforms implemented in 2003–5. They are credited with increasing labor supply and reducing labor costs notably for low-skill workers (Krause and Uhlig 2012). Therefore, we predict that these reforms reduced automation innovation. In a difference-in-difference exercise, we find that foreign firms that are relatively more exposed to Germany innovated less in automation technologies after the Hartz reforms. Finally, in a triple-difference exercise, we find that the reforms also decreased automation innovations relative to nonautomation innovations.

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

Despite this theoretical literature, existing empirical evidence on the effect of wages on induced automation, adoption, or innovation remains limited.[2] A few papers show that labor market conditions affect adoption of labor-saving technologies in agriculture (Hornbeck and Naidu 2014; Clemens, Lewis, and Postel 2018; Voth, Caprettini, and Trew 2023), health care (Acemoglu and Finkelstein 2008), and manufacturing (Lewis 2011; Acemoglu and Restrepo 2022). Lordan and Neumark (2018) find that minimum wage increases displace workers in automatable jobs, and Fan, Hu, and Tang (2021) find that they induce Chinese firms to adopt industrial robots. None of these papers use firm-level variation in the manufacturing sector, as we do. More importantly, we focus on innovation rather than

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

[2] In contrast, there is an extensive empirical literature on the effects of technology on wages and employment: see, e.g., Autor, Levy, and Murnane (2003), Autor and Dorn (2013), or Gaggl and Wright (2017) for information technology; Doms, Dunne, and Troske (1997) for factory automation; Boustan, Choi and Clingingsmith (2022) for computer numerical control [CNC]; Graetz and Michaels (2018) or Acemoglu and Restrepo (2020) for robots; and Bessen et al. (2023), Mann and Püttmann (2023), and Aghion et al. (2024) for broader measures of automation.

adoption. This is an important distinction: adoption generally refers to firms deciding whether to use an existing technology, while innovation is a continuous process of creating new technologies. While a shock to adoption must die out over time (when all firms have adopted a technology), a shock to innovation can build on itself. As a result, (1) adoption and innovation matter at different time horizons, (2) knowledge spillovers play a different role for innovation than adoption, and (3) the growth literature mentioned above focuses on innovation.

The literature on induced automation innovation is sparser. Acemoglu and Restrepo (2022) find a positive effect of aging on patenting in robotics and numerical control in cross-country regressions, though they focus mainly on adoption. Our paper differs in at least four ways: we build a broader measure of automation innovation in machinery; we are interested in the effect of all wage variations, not only variations arising from demographic trends; we consider policy-induced changes; and, most importantly, we use firm-level panel regressions instead of a country-level cross-sectional analysis. In contemporaneous work, Danzer, Feuerbaum, and Gaessler (2024) exploit an immigrant settlement policy in Germany to show that increases in labor supply discourage automation innovation at the level of local labor markets, and San (2023) shows that a negative shock to agricultural immigration from Mexico induced relatively more innovation related to labor-intensive crops.[3] Neither of these papers exploits firm-level variation, nor do they estimate the effect of labor cost changes on automation innovations.

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

Finally, a recent literature has emerged that classifies patents as automation or not (see, in particular, Webb 2020; Kogan et al. 2022; Mann

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

[4] Other papers have used their methodology, including Noailly and Smeets (2015) on innovation in electricity generation, Coelli, Moxnes, and Ulltveit-Moe (2022) on the effect of trade policy on innovation, and Aghion et al. (2023) on the role of environmental preferences and competition in innovation in the auto industry.

and Püttmann 2023). We compare our approaches in detail. A key difference is that we classify readily available technology codes rather than patents directly so that our classification can be easily applied to other patent data by other researchers without using patent text. This comes at the cost of potentially misclassifying individual patents. Autor et al. (2024) go further and map automation and augmenting patents to specific occupations, highlighting that the same innovations can substitute for some occupations while augmenting others. Our approach does not allow for the same level of granularity, but our results are consistent with theirs insofar as automation patents in machinery appear to substitute for low-skill workers while complementing high-skill workers.

Section II develops our classification of automation technologies. Section III describes the data and our empirical strategy. Section IV presents the results of the main analysis on the effect of wages on automation innovations. Section V discusses policy shocks. Section VI concludes. The appendix (available online) provides an analytical model, additional robustness checks, and details on our methodology.

## II. Classifying Automation Patents

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

### *A. 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. Nonautomation innovations, in contrast, may improve energy efficiency, reduce the cost of producing certain machines, or increase reliability.

We follow a well-established tradition in the empirical literature and use patent data as a measure of innovative activity.[5] We use two databases: the European Patent Office (EPO) full-text database from 2018, which contains the full text of patent applications at the EPO; and PATSTAT from autumn 2018, which contains the bibliographic information (but not the text) of close to the universe of patents. In these datasets, the technological characteristics of patents are recorded in technology codes (notably CPC and IPC codes, hereafter C/IPC codes; see n. 8). Certain

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

types of technologies, such as fossil fuel engines, can be readily identified with existing groupings of C/IPC codes. No such grouping exists for automation in machinery, and we use text analysis to create one.

We employ a dictionary method on patent data and proceed in five steps: (1) we define machinery technology categories based on C/IPC codes; (2) we use the existing literature to identify automation-related keywords; (3) for each machinery technology category, we compute the share of patents at the EPO containing one of our automation keywords; (4) we use this measure to classify machinery technology categories as automation or not based on a cutoff; and (5) we classify worldwide patents as automation if they belong to an automation technology category.

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

### *B. Identifying Machinery Technology Categories*

Technological characteristics of patents are recorded in (generally several) C/IPC codes. The C/IPC codes form a hierarchical classification system.[8]

[6] To give an idea of the increase in the sample size, over the period 1997–2011, there are 549,000 patent families in machinery (defined below), with patent applications in at least two offices (a condition we will impose in our main analysis). Among these, only around 82,000 have an EPO patent with a description in English.

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

[8] The IPC is an international classification, and the CPC is used by the US Patent and Trademark Office (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 three-digit codes (e.g., B25: "hand tools; portable power-driven tools; handles for hand implements; workshop equipment and manipulators"), subclasses have four-digit codes (e.g., B25J: "manipulators; chambers provided with manipulation devices"), and main groups have five- to seven-digit codes (e.g., B25J13: "controls for manipulators"). In the following, we refer to classes, subclasses, and main groups as three-digit, four-digit, and six-digit codes, respectively.

We start by defining the aggregation level at which we classify patents, which we refer to as technology categories. First, we define technology categories as the highly disaggregated six-digit C/IPC codes (e.g., B25J13). The co-occurrence of technology codes can also be informative about the characteristics of a patent. To capture this, we include pairs of four-digit C/IPC codes in the definition of technology categories: for instance, a patent containing both codes B25J (manipulators) and B23K (a type of machine tools) belongs to the technology category {B25J, B23K}. Finally, we include in the definition of technology categories the co-occurrence of four-digit C/IPC codes with the three-digit codes G05 or G06 (e.g., B25J with G05 or G06). The code G05 corresponds to "controlling; regulating" and G06 to "computing; calculating; counting." Aschhoff et al. (2010) use these combinations to identify advanced manufacturing technologies.[9] The six-digit codes alone identify 83% of our automation patents (see app. sec. A.2.3).

Since we are interested in automation in equipment, we restrict attention to C/IPC codes that belong to a group of technology fields, which we call machinery. Out of 34 technology fields (see fig. A1), we focus on machine tools, handling, textile and paper machines, and other special machines with some adjustments.[10] We classify pairs of four-digit C/IPC codes or pairings of four-digit C/IPC codes with G05 or G06 as machinery if at least one four-digit code belongs to that field. This leaves us with 983 six-digit codes, 1,100 pairs of four-digit codes, and 25 groupings of four-digit codes with G05/G06, which form the set of machinery technology categories. We then define a machinery patent as a patent that belongs to one of the machinery technology categories.

### *C. Choosing Automation Keywords*

Next, we choose keywords to identify automation technology categories within the set of machinery technology categories. To tie our hands, we choose most of our keywords from the automation technologies identified in Doms, Dunne, and Troske (1997) and Acemoglu and Restrepo (2022) and supplement them with additional words as described below.[11] Most

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

[10] We make a few minor adjustments, such as excluding weapons and ammunition and adding technology codes referring to program control system. See app. sec. A.2 for details.

[11] 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 Harmonized System categories from Comtrade with automation technology.

of our keywords correspond to the co-occurrence of several words in the same sentence or patent or the repetition of these words a sufficient number of times. Table 1 lists our keywords.

We have eight categories of keywords. Five of them are automation technologies in Doms, Dunne, and Troske (1997) or Acemoglu and Restrepo (2022) ("robot*," "numerical control," "computer-aided design and manufacturing," "flexible manufacturing," and "programmable logic control"). Directly using some of these keywords results in false positives.

TABLE 1
CHOICE OF AUTOMATION KEYWORDS

| Keywords | Comments | Source |
|---|---|---|
| "Automat*" | "*Automation*," "*automatization*," or "*automat**" at least five times; also "automat*" or "autonomous" in same sentence with secondary words (see below), "*warehouse*," "*operator*," "arm," "convey*," "handling," "inspect*," "knitting," "manipulat*," "regulat*," "sensor," "storage," "store," "vehicle system," "weaving," or "welding" at least twice | Own; Doms, Dunne, and Troske 1997; or Acemoglu and Restrepo 2022 |
| "Robot*" | Not surgical or medical | Doms, Dunne, and Troske 1997; Acemoglu and Restrepo 2022 |
| "Numerical control" | "CNC," "numeric* control*," or "NC" in same sentence as secondary words (see below) | Doms, Dunne, and Troske 1997; Acemoglu and Restrepo 2022 |
| "Computer-aided design and manufacturing" | "Computer-aided"/"-assisted"/"-supported" in same patent as secondary words (see below); also "CAD" or "CAM" and not "content addressable memory" in same sentence as secondary words (see below) | Doms, Dunne, and Troske 1997 |
| "Flexible manufacturing" | | Doms, Dunne, and Troske 1997 |
| "Programmable logic control" | "Programmable logic control" or ("PLC" and not ("powerline" or "power line")) | Doms, Dunne, and Troske 1997 |
| "*3D printer*" | "*3D print**," "*additive manufacturing*," or "*additive layer manufacturing*" | Own |
| "*Labor*" | Including "*laborious*" | Own |
| | Secondary Words | |
| | "*Machine*," "*manufacturing*," "*equipment*," "*apparatus*," or "machining" | |

NOTE.—The table describe the keywords that we use to identify automation technologies. Keywords include (1) natural adjacent words (i.e., "numerical control" includes "NC," "numerically controlled," and "numeric control"), (2) British/American spelling (i.e., "labour"/"labor"), and (3) hyphenated adjectives (i.e., "computer aided"/"computer-aided design"). "In same sentence as secondary words" refers to at least one secondary word. We added words in italics; the others come from Doms, Dunne, and Troske (1997) or Acemoglu and Restrepo (2022). See appendix for details.

Therefore, we require that our keywords occur in the same patent or in the same sentence as secondary words, such as "machinery" or "equipment," indicating that the text describes a machine. We also add "automation" and "automatization." The stem "automat*" gathers too many false positives, such as "automatic transmission." We resolve this in two ways: we count patents only if the frequency of "automat*" is at least five or "automat*" is combined with a list of words in the same sentence at least twice. This list of words contains our secondary words and additional words that come from Doms, Dunne, and Troske (1997) or Acemoglu and Restrepo (2022) and often refer to tasks (such as "manipulat*," "regulat*," or "inspect*"). Finally, we add "3D printing," which was in its infancy when Doms, Dunne, and Troske (1997) was written, and "labor," which generally indicates that an innovation reduces labor costs.[12] The most important keywords are those associated with "automat*" (app. sec. A.2), and our main results are robust to using only these (app. sec. A.6.2).

### D. *Identifying Automation Technology Categories*

We base our classification on the set of EPO patent applications in machinery from 1978 to 2018 with a description in English (267,997 patent applications), denoted $\Omega_{\text{EPO}}$. We denote $MT_p$ the set of machinery technology categories associated with a patent $p$ and $T_p$ its text.[13] We define the function $k^{\text{any}}(T_p)$, which takes a value of 1 if any of the automation keywords is present in the text and 0 otherwise. For each machinery technology category $t$, we define the prevalence of automation keywords $s(t)$ as the share of patents containing at least one of our keywords:

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

We similarly define the prevalence of each keyword category. 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 section A.2. We manually checked the C/IPC codes extensively and sampled patents from each category to ensure that the process delivered reasonable results.

We define automation technology categories as those with a prevalence measure above a threshold. Figure 1 shows the histogram of the

[12] "Labor" differs from the other keywords in that it does not refer to a technology. We checked that our results are robust to removing that keyword from the whole procedure.

[13] We use all C/IPC codes of the patent family associated with the EPO patent application $p$. See sec. II.E for the definition of the patent family.

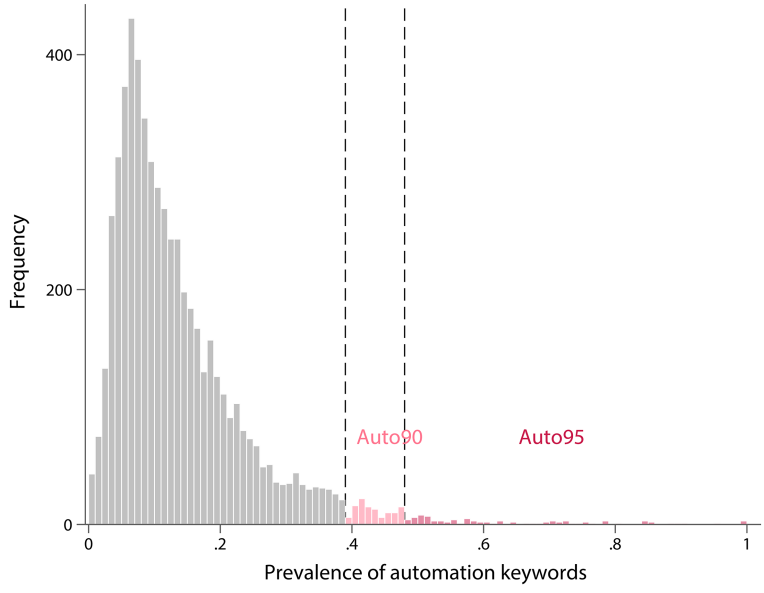


FIG. 1.—Prevalence of automation keywords for C/IPC six-digit codes in machinery.

prevalence of automation keywords for C/IPC six-digit codes in machinery. The distribution is skewed: most C/IPC codes have a low prevalence of automation keywords, but a few codes have a very high value. We choose thresholds at the 90th and 95th percentiles of the distribution of the six-digit code distribution (within machinery), which are given by 0.39 and 0.48, respectively.[14] Therefore, a technology category $t$ belongs to the set of auto90 categories $T^{90}$ if $s(t) > 0.39$ and to the set of auto95 category $T^{95}$ if $s(t) > 0.48$. In appendix section 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.

### *E. Defining Automation Patents*

We now proceed to classify automation patents. To do so, we use PATSTAT, which contains bibliographic information for almost the universe of patents. PATSTAT also allows us to identify patent families, a set of patent applications in different national or international patent offices representing the same innovation. For each patent family, we know the date of the

[14] Choosing different thresholds is easy, and we investigate how robust our results are in sec. IV.E.

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

### *F. Comparison with Other Measures in the Literature*

A recent literature has emerged that uses patent data to identify automation technologies. Our approach—classifying technology categories rather than patents directly and using a dictionary method on the text of patents—is unique, and we compare it here with three alternatives: Mann and Püttmann (2023), Kogan et al. (2022), and Autor et al. (2024).[15]

Mann and Püttmann (2023) manually classify a set of patents as automation or not and then use machine-learning techniques to classify all patents. Relying on keywords instead of a training set presents several advantages. First, manually labeling patents is a difficult task that cannot be easily systematized and delegated. Second, patents are technical descriptions of an innovation and do not always discuss its goal. Only a few words within the text are informative, so the training set needs to be large. Third, our approach is more transparent and more easily replicable and modifiable, and it leaves fewer degrees of freedom, as we choose most of our keywords from the literature.

We compare our classifications in detail in appendix section A.2.6. The measures are positively correlated, though ours is generally more conservative. We classify 9.4% of the common set of machinery patents as auto95, while Mann and Püttmann (2023) classify 29.8% of them as automation. Seventy percent of our auto95 patents are classified as automation patents by Mann and Püttmann (to put this number in context, their algorithm has a 17% false negative error rate on their training set). Focusing on outlier technologies, we find that Mann and Püttmann (2023) classify a number of patents related to elevators and printing machines as automation, which we do not. This is in line with their definition of automation as "a device that carries out a process independently of human

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

intervention," while we seek to identify innovations that replace workers in existing tasks.

In contemporaneous work, Kogan et al. (2022) and Autor et al. (2024) measure the distance between patent texts and the description of tasks in the Dictionary of Occupational Titles database to identify labor-displacing innovations. This approach has the advantages of matching innovations and the affected occupations and of being completely hands-off.[16] However, as Kogan et al. (2022) point out, it captures not only automation innovations where workers are replaced by machines but also innovations where incumbent workers with newly obsolete skills are replaced by new workers.

Compared with these alternatives, our approach has a few drawbacks: we rely on the existence of a set of well-identified automation technologies (e.g., robots, CNC), which is why we apply our method only to machinery patents. Heterogeneity within technology categories will also lead us to misclassify individual patents that are exceptions within their categories (i.e., nonautomation patents in a category with many automation patents, and vice versa). Finally, we do not provide a mapping between innovations and occupations and therefore cannot capture heterogeneous effects between workers and innovation beyond the high skill/low skill dichotomy. At the same time, it is simpler and has the advantage of providing a measure that can be readily used by anyone on any patent data with C/IPC codes.

### *G. Trends in Automation Innovations*

Figure 2 plots the evolution of automation patent families. To restrict attention on innovations of sufficient quality and make the data more comparable across countries, we focus on patent families including patent applications in at least two countries, referred to as biadic patents. Several studies (e.g., de Rassenfosse et al. 2013; Dechezleprêtre, Ménière, and Mohnen 2017) have shown that biadic patents are of higher quality than others.[17] Figure 2*A* shows that machinery patents represent a sizable but

[16] In addition, using another dataset that describes the output of occupations, Autor et al. (2024) also run a parallel mapping between patents and the occupations that they augment.

[17] We count applications and not-granted patents because certain patent offices, notably the Japan Patent Office (JPO), formally grants a patent only if the applicant requests an examination, which they often do only 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.

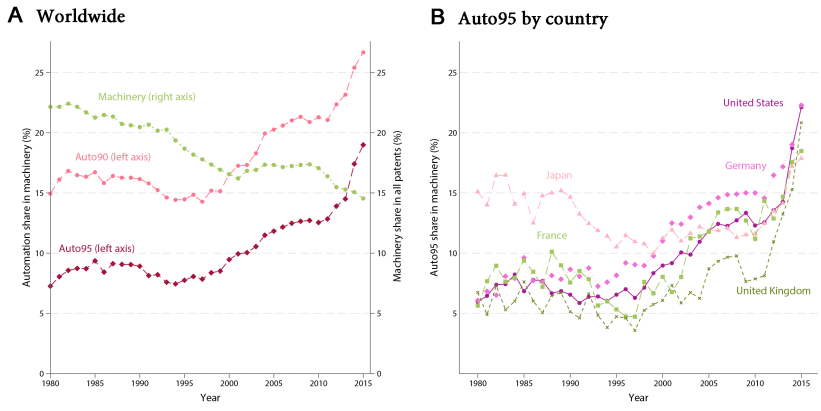


FIG. 2.—Share of automation patents in machinery and share of machinery patents for biadic families. *A*, Share of auto95 and auto90 patents in machinery (left axis) and share of machinery patents in all patents (right axis) worldwide. *B*, Share of auto95 patents in machinery by applicants' nationality. Patents with multiple C/IPC codes are counted only once.

slowly decreasing share of all patents (18.6% on average over this time period), though the absolute number of machinery patents is increasing. Globally, among machinery patents, the share of automation patents declined slightly between the mid-1980s (9.4% in 1985 for auto95) and the mid-1990s (7.5% in 1994 for auto95) before increasing rapidly (reaching 19.1% in 2015 for auto95). Figure A2 (figs. A1–A12 are available online) reports the raw numbers of auto90 and auto95 patents and their share out of total patents. Figure 2*B* shows the trends for auto95 by applicant nationality. In the 1980s, Japan had the highest share of automation patents in machinery, while Germany took that position from the 2000s.

### *H. Automation, Routine Tasks, Skill Composition, and Labor Share*

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

As in Autor, Levy, and Murnane (2003), we run sector-level regressions of the type

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

We focus here on the years 1980–98. $\Delta T_{jk}$ represents the change in tasks of type $k$ in sector $j$. We take this measure directly from Autor, Levy, and Murnane (2003), who define five types of tasks: nonroutine analytic, nonroutine interactive, routine cognitive, routine manual, and nonroutine manual. $\Delta T_{jk}$ is measured as 10 times the annual within-sector change in task input, measured in percentiles of the 1960 task distribution, $\Delta C_j$ is Autor, Levy, and Murnane's (2003) measure of the change in computerization in sector $j$ (available for the period 1984–97), and $\text{aut}_j$ is our patent-based measure of the use of automation technologies in sector $j$. Since patenting is already a measure of knowledge flows, we do not first-difference this measure. We run similar regressions with changes in the skill ratio, the labor share, and employment as alternative outcome variables.

To build $\text{aut}_j$, we allocate patents in machinery to their sector of use, focusing on patents granted by the USPTO. Autor et al. (2020) match USPTO patents with firm-level data from Compustat and provide detailed sectoral information on corporate patents. Using their data, we create a weighted concordance table from C/IPC four-digit codes to four-digit Standard Industrial Classification industries that allows us to map patents to sectors of invention. We then combine this information with the 1997 capital flow table from the US Bureau of Economic Analysis (BEA) to get the sector of use. The capital flow table is akin to an input-output table, but it reports the flows in investment goods instead of intermediate inputs. Sectors that purchase a lot of capital from sectors innovating in machinery generally have high exposure to both automation and nonautomation patents (defined here as pauto90 patents, i.e., machinery patents excluding auto90 patents). The correlation between the log counts of auto95 patents and pauto90 patents across sectors is 0.99, but the correlation between the ratio of auto95 patents over capital purchases and ratio of pauto90 patents over capital purchases drops to 0.76 (see fig. A9).

For each sector of use $j$, we compute $\text{aut}_j$ as the share of automation patents (auto95 in our baseline) among machinery patents from 1980 to 1998. Our measure captures whether the machinery patents used in a given sector are particularly intensive in automation technologies. We compute this statistic for the 133 sectors with machinery patents (our results are robust to excluding sectors with few machinery patents). Table A1 (tables A1–A43 are available online) reports summary statistics on this measure and our dependent variables: on average, the auto95 share is 7.5%. There is significant variation in the share of automation patents across industries, with a coefficient of variation of 17%. Exposure to automation is on average higher in manufacturing, and we include a manufacturing dummy in our regressions. Interestingly, our automation measure is only weakly correlated with computerization, with a coefficient of 0.08 (and −0.16 when we weigh industries by employment). Table 2

TABLE 2
SECTORS WITH HIGHEST AND LOWEST SHARES OF AUTOMATION PATENTS

| Industry Code | Description | auto95 |
|---|---|---|
| | A. Sectors with Highest Share of Automated Patents | |
| 756 | Automotive services and repair shops | .110 |
| 206 | Household appliances (e.g., radio, television, equipment) | .106 |
| 470 | Water supply and irrigation | .098 |
| 271 | Iron and steal foundries | .096 |
| 130 | Tobacco manufactures | .093 |
| 212 | Miscellaneous plastic products | .093 |
| | B. Sectors with Lowest Share of Automated Patents | |
| 801 | Bowling alleys, billiard and pool parlors | .042 |
| 100 | Meat products | .046 |
| 101 | Dairy products | .046 |
| 102 | Canned and preserved fruits and vegetables | .046 |
| 110 | Grain mill products | .046 |
| 111 | Bakery products | .046 |

NOTE.—The table lists the sectors with the highest and lowest shares of auto95 patents out of all patents in machinery in 1980–98. The industry codes and descriptions of the sectors correspond to the ind6090 industries described in Autor, Levy, and Murnane (2003).

lists the sectors with the highest and lowest shares of automation patents in machinery.[18]

Table 3 reports the results of regression (1), and figure A10 provides scatterplots. Columns 1 and 3 show a clear relationship: a 1 percentage point increase in the automation share is associated with a 1.4 and 1.3 centile decrease in routine cognitive and manual tasks per decade, respectively. This effect is larger than that of computerization (i.e., standardized beta coefficients are larger). At first sight, it may seem surprising that our measure of automation in machinery predicts a decline in routine cognitive tasks. However, Autor, Levy, and Murnane (2003) define routine cognitive tasks as the "adaptability to situations requiring the precise attainment of set limits, tolerances or standards." These correspond to inspection and control tasks that our automation machines may eliminate (fig. A8 gives an example of such a machine). Metalworkers, for instance, are among the occupations with the highest intensity of routine cognitive tasks.

In column 5, we instead use the change in the skill ratio (college graduates over all other workers) as the dependent variable. A 1 percentage point increase in the share of automation patents is associated with an increase in the skill ratio equal to one-third of its average increase over the 1980–98 period. In column 7, we use the change in the labor share for manufacturing industries (using the National Bureau of Economic

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

TABLE 3
EFFECTS OF AUTOMATION ON TASKS, SKILL COMPOSITION, AND LABOR SHARE

| | Δ Routine Cognitive | | Δ Routine Manual | | Δ High-/Low-Skill Workers | | Δ Labor Share (NBER) | | Δ Labor Share (BEA) | |
|---|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| Share automation: | | | | | | | | | | |
| Using industry | −142.48*** | −142.98*** | −128.04*** | −127.30*** | 3.80* | 3.51* | −1.30** | −.85** | −3.64* | −3.58* |
| | (37.37) | (28.34) | (38.85) | (39.81) | (2.16) | (2.01) | (.65) | (.41) | (1.94) | (1.86) |
| Inventing industry | | −22.09** | | −11.00 | | −.21 | | −.46 | | −.21 |
| | | (8.52) | | (8.03) | | (.45) | | (.29) | | (.33) |
| Δ computer use (1984–97) | −18.87*** | −20.51*** | −19.14*** | −21.81*** | .98*** | .92*** | .25* | .27* | .12 | .15 |
| | (6.54) | (5.78) | (7.29) | (7.62) | (.26) | (.27) | (.13) | (.14) | (.31) | (.32) |
| Manufacturing | −1.71* | −2.17** | −.03 | −.25 | .02 | .02 | | | −.05* | −.06* |
| | (.92) | (.86) | (.94) | (.93) | (.03) | (.03) | | | (.03) | (.03) |
| $R^2$ | .29 | .38 | .21 | .23 | .21 | .20 | .18 | .27 | .22 | .23 |
| Industries | 133 | 126 | 133 | 126 | 133 | 126 | 56 | 56 | 60 | 60 |

NOTE.—The table shows the effect of automation technologies on tasks, skill composition, and the labor share. Each column is a separate ordinary least squares regression of the change in an industry outcome between 1980 and 1998 on the share of automation patents in machinery in 1980–98, the annual percentage point change in industry computer use in 1984–97, a dummy variable indicating the manufacturing sector, and a constant. The dependent variable in each column is as follows: change in routine cognitive tasks (cols. 1, 2) and change in routine manual tasks (cols. 3, 4), both measured as 10 times the annual change in industry-level task input in centiles of the 1960 task distribution (see Autor, Levy, and Murnane 2003); change in the ratio of high-skill workers (college graduates) over low-skill workers (others; cols. 5, 6); and change in the labor share in the NBER-CES Manufacturing Industry Database (cols. 7, 8) and in 60 aggregated industries from the BEA (cols. 9, 10). As described in the text, the two automation share measures correspond to a different mapping between C/IPC codes and industries. Using industry allocates patents to their sector of use, while innovating industry (added in cols. 2, 4, 6, 8, and 10) allocates patents to their sector of invention. The regressions are weighted by the mean industry share of total employment in full-time equivalents in 1980 and 1998. Standard errors (in parentheses) are clustered at the level of industry groups that have the same automation share by construction.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

Research [NBER] and Center for Economic Studies [CES] Manufacturing Industry Database). Industries with a high share of automation patents experience a decline in the labor share, consistent with the theoretical literature on automation. A 1 percentage point increase in the automation share is associated with a 1.3 percentage point decline in the labor share. Interestingly, computerization does not have the same effect. We extend this analysis to nonmanufacturing industries using more aggregate data on the labor share for 60 industries from the BEA. We still find a negative effect of automation (col. 9).

In columns 2, 4, 6, 8, and 10, we add a control for the share of automation patents invented by the sector. To allocate patents to the inventing sector, we omit the capital flow table step in the calculation of our automation variable. The coefficients on the automation share in the using sector remain stable. Moreover, the automation share in the using sector has a larger effect than the automation share in the inventing sector.[19]

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

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

## III. Empirical Strategy and Data

We now move to our main empirical exercise and analyze the effect of labor cost shocks on automation innovations. Section III.A presents our empirical strategy, sections III.B and III.C explain how we build our dataset, section III.D describes our estimation equation, and section III.E shows summary statistics.

### *A. Empirical Strategy*

We motivate our empirical strategy with the business structure of the most prominent automation innovators. These are large firms that sell

[19] The standardized coefficients are larger for the using sector than the inventing sector (except in col. 8), as the standard deviation for the share of automation patents in the using and inventing sectors is 1.3% and 6%, respectively.

their automation equipment internationally to downstream firms. Automation equipment enables the replacement of low-skill workers with machines. It can also complement high-skill workers who program, operate, and maintain the machines. Therefore, the incentives for downstream producers to adopt automation technology are determined by labor costs in their local market. Higher low-skill labor costs for potential customers are associated with a larger market for automation machine producers, which in turn would induce innovators to conduct more research on automation technologies.[20] Appendix section A.4 presents a simple model that rationalizes this argument.[21]

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

Ideally, we would measure the expected future labor costs paid by the actual and potential customers of automation innovators. This raises two issues. First, a measure directly based on customers' data would suffer from reverse causality, and we would need an instrument. A natural candidate would be a shift-share instrument. In the absence of direct data on the labor costs paid by innovators' customers, we directly use such a shift-share measure as a proxy. Our regression should therefore be seen as the reduced form of this instrumental approach. Second, we cannot measure expectations, so we use current labor cost shocks as a proxy for shocks on expected future costs—though in section IV.E, we explicitly build a model for expected wages using an AR process.

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

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

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

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

where $w_{\mathrm{L},c,t}$ ($w_{\mathrm{H},c,t}$) are the low-skill (high-skill) labor costs in country $c$ at time $t$ and $\kappa_{i,c}$ is the fixed weight of country $c$ for firm $i$.[22] Similarly, we build controls for several macroeconomic variables—such as labor productivity, gross domestic product (GDP) per capita, or size of the manufacturing sector—that could also affect innovation.

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

### *B. Country-Level Data*

We use data on 41 countries. Most of our data come from the 2013 release of the World Input-Output Tables (Timmer et al. 2015), which contains information on hourly labor costs from 1995 to 2009 across educational attainment groups.[23] We focus on labor costs in manufacturing since our keywords largely come from the SMT. Our results are robust to using labor costs in the entire economy. Although our data cover all labor costs, we refer to them as wages for simplicity. In the data, the low-skill workers are defined as having no high-school diploma or equivalent, whereas the high-skill workers have at least a college degree. Middle- and low-skill wages are highly correlated, and we can interpret our low-skill wage variable as reflecting both.

We calculate labor productivity in manufacturing as value added divided by hours. We gather exchange rate and GDP data from the United Nations Statistics Division and compute the GDP gap to control for business cycles. We deflate all nominal values with the local producer price

[22] More precisely, innovation incentives depend on the costs for automatable tasks, but there are no good international occupational or task-level labor costs data. Since low-skill workers' tasks have been more intensely automated, we use low-skill labor cost as a proxy for the cost of automatable tasks. This proxy will be particularly good if labor markets are flexible across occupations within education groups or if labor shocks affect low-skill workers similarly across occupations. Otherwise, a noisy measure should result in a bias towards zero.

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

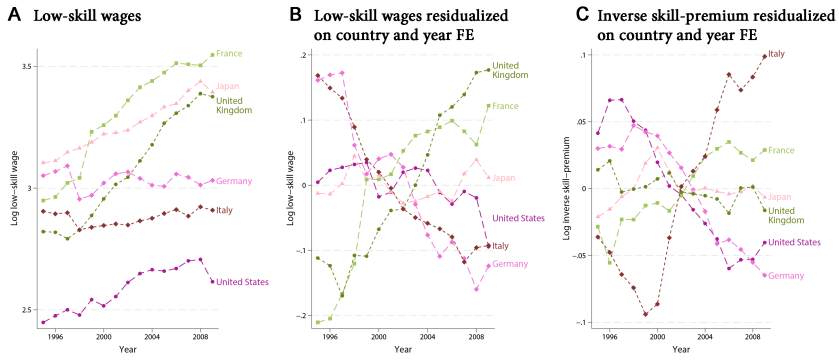


FIG. 3.—Country-level trends in low-skill wages and inverse skill premium. *A*, Raw log low-skill wages. *B*, Log low-skill wages residualized on country and year fixed effects. *C*, Log inverse skill premium residualized on country and year fixed effects.

index for manufacturing (indexed to 1995) and then convert everything into US dollars using the average exchange rate for 1995, the starting year of our regressions. Appendix section A.5.1 provides further details and summary statistics.

Figure 3 uses the six countries with the largest average weights and shows the log low-skill wages, the log low-skill wages residualized on country and year fixed effects, and the residualized log inverse skill premium. Trends vary markedly across countries, providing a significant amount of variation in the data: for instance, the United States sees a relative decrease in the inverse skill premium, while Italy follows a nonmonotonic pattern.[24]

### C. *Firm-Level Data*

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

We use the firm's patent filing history as a proxy for its market exposure to measure the weights $\kappa_{ic}$. This method builds on that of Aghion et al. (2016). Firms differ in their market exposure because of trade barriers, heterogeneous customer tastes, or various historical accidents. A patent

[24] Figure A12 shows more precisely the identifying variation, taking into account the shift-share structure of our measure following Borusyak, Hull, and Jaravel's (2022) methodology.

grants its holder the exclusive right to commercially exploit an innovation in a specific country for a limited period. Inventors must file a patent in each country where they wish to protect their technology. Patenting is costly: a firm must hire lawyers and possibly translators and pay filing costs. Therefore, inventors seek patent protection in a country only if they are confident about the potential market value of the technology (Eaton and Kortum 1996). Indeed, empirical evidence suggests that inventors do not patent widely and indiscriminately, with the average invention being patented in only two countries (Dechezleprêtre et al. 2011).[25]

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 presample period 1971–94 to ensure that they are weakly exogenous.[26] We restrict our attention to patent families with at least one citation (without self-citations) to exclude the lowest quality patents. See appendix section A.5.3 for details notably on how we treat European patents.

The raw patent count indicates whether a firm intends to sell its products in a market but does not capture market size. A larger market attracts more firms, so the market size per firm does not grow one-to-one with country size. To account for this, we weigh each country $c$ by $\text{GDP}_{0,c}^{0.35}$, where $\text{GDP}_{0,c}$ is the 5-year average GDP of country $c$ at the end of the presample period.[27] As a result, the weight of country $c$ for firm $i$ is

$$\kappa_{i,c} = \frac{\tilde{\kappa}_{i,c}\text{GDP}_{0,c}^{0.35}}{\sum_{c'}\tilde{\kappa}_{i,c'}\text{GDP}_{0,c'}^{0.35}}. \tag{3}$$

Following equation (2), we then combine the weights $\kappa_{i,c}$ with the macro variables presented in section III.B to build macro variables, including wages, at the level of firms' customers. We use 1971–94 as a presample period, as PATSTAT's coverage is significantly better from the 1970s onward, and we prefer a long time period for our baseline measure. Importantly,

[25] Aghion et al. (2016) verify that a method similar to ours accounts well for the sales distribution of major auto manufacturers. Coelli, Moxnes, and Ulltveit-Moe (2022) carry out a more systematic exercise and verify that such a method accounts well for aggregate bilateral trade flows and firm exports across eight country groups in a representative panel of 15,000 firms from seven European countries (regressing patent weights on sales weights gives a coefficient of 0.89, with a standard error of 0.008). In app. sec. B.2, we also show that our patent weights correlate well with trade flows.

[26] 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.

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

the weights are stable over time. We show that our results are robust to alternative presample periods and weighing schemes in section IV.E.[28]

### D. Estimation Equation

We now describe how we estimate the effect of an increase in wages on automation innovations. We have a panel of firms with patent data and firm-level wage variables. Since our dependent variable is a count of patents, we use a Poisson specification. We assume that firms' innovation in automation follows[29]

$$\mathrm{E}(\mathrm{PAT}_{\mathrm{Aut},i,t}) = \exp\left(\begin{array}{c}\beta_{w_{\mathrm{L}}} \ln w_{\mathrm{L},i,t-2} + \beta_{w_{\mathrm{H}}} \ln w_{\mathrm{H},i,t-2} + \beta_X X_{i,t-2} \\ +\beta_{Ka} \ln K_{\mathrm{Aut},i,t-2} + \beta_{Ko} \ln K_{\mathrm{other},i,t-2} + \beta_{Sa} \ln \mathrm{SPILL}_{\mathrm{Aut},i,t-2} \\ +\beta_{So} \ln \mathrm{SPILL}_{\mathrm{other},i,t-2} + \delta_i + \delta_{j,t}(+\delta_{c,t})\end{array}\right). \quad (4)$$

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

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

We include controls for knowledge stocks at the firm and country level. $K_{\mathrm{Aut},i,t}$ and $K_{\mathrm{other},i,t}$ denote the stocks of knowledge in automation and in all other technologies of firm $i$ at time $t$. We compute these knowledge stocks using the perpetual inventory method. $\mathrm{SPILL}_{\mathrm{Aut},i,t}$ and $\mathrm{SPILL}_{\mathrm{other},i,t}$ similarly denote the stocks of external knowledge (spillovers) in automation and in other technologies to which firm $i$ has access at time $t$. We compute these spillovers as a weighted average of country-level knowledge stocks, where the weights now reflect the location of the firms' inventors.[30] These controls serve two purposes. First, they ensure that we

[28] We consider two alternative definitions of low-skill wages where weights are based on 1971–89 or 1985–94. For the firms in our baseline regression sample, the correlation between these two wage variables is 0.91.

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

[30] 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) before the sample period in 1971–94. When computing the log of

do not simply capture that some firms or countries are on different automation trends. Second, knowledge spillovers are an important characteristic of innovation processes and can amplify the short-run response of innovation to economic shocks over time. On the empirical side, Popp (2002) finds that including a measure of knowledge stocks is important to obtain correct estimates of energy prices on energy-saving innovations. On the theoretical side, including measures of the stock of knowledge is important for the calibration of macro models of directed technical change (Acemoglu 2023).

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

$\delta_i$ are firm fixed effects such that our variation comes from how changes in wages affect changes in automation innovations.[31] $\delta_{j,t}$ are industry × year fixed effects. A firm's industry $j$ is the manufacturing industry and corresponds to its two-digit industry in Orbis. Table A2 gives the distribution of firms and patents across the main industries in our sample. In some specifications, we include country × year fixed effects $\delta_{c,t}$, where the firm's country (the home country) is defined as the country with the largest weight $\kappa_{i,c}$. We cluster standard errors at the firm level.[32]

### *E. Baseline Sample*

We now describe the sample we rely on to estimate equation (4). Since wages are available for 1995–2009, our baseline datasets rely on firms that applied for at least one biadic automation patent between 1997 and 2011. These firms must also have at least one patent prior to 1995 so that we can compute the patent weights $\kappa_{i,c}$. We also exclude purely domestic firms (i.e., those that patented in only one country before the sample period), though our results are very similar if we include them. Our baseline sample for the auto95 measure corresponds to 3,236 firms.

---

stocks or spillovers, we replace zeros with ones and add a dummy variable to indicate where stocks or spillovers are zero.

[31] We use the Hausman, Hall, and Griliches (1984) 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 eq. (4), as it requires strict exogeneity and hence prevents the lagged dependent variable from appearing on the right-hand side (which it does here to a limited extent through the knowledge stock $K_{\text{Aut},i,t-2}$). Nevertheless, we show in sec. IV.E that our coefficients of interest are not affected by this Nickell bias.

[32] Alternatively, we cluster at the country level in app. sec. A.6, following the approach of Cameron, Gelbach, and Miller (2008), and discuss inference in the shift-share setting in sec. IV.C.

Table A3 shows that our sample of firms covers a considerable share of worldwide automation innovations in machinery. Orbis's coverage is excellent: we can assign 84.9% of all biadic auto95 patent families in 1997–2011 to a firm. Moreover, most heavy patenters had already patented in at least two countries before the sample period: the firms of our sample account for 53.5% of all biadic auto95 patent families.

Table A4 gives descriptive statistics on the number of automation patents per year and lists the sample's 10 biggest automation innovators. The distribution of auto95 patents is strongly skewed: over the period 1997–2011, the median firm in the sample filed two auto95 patent applications, whereas the 99th percentile filed 194. In our empirical analysis, we also look at the effect of wages on nonautomation machinery patents (defined as pauto90 patents, i.e., machinery patents that are not auto90 patents). For that exercise, we restrict the sample further to firms that have at least one pauto90 patent. The average number of citations per patent is slightly higher for auto95 than for pauto90 patents (9.3. vs 7.6).

Table 4 provides descriptive statistics on the country weights for the firms in our sample. The largest country for a given firm has, on average, a weight of 0.47 (for the auto95 sample), and the second largest has a weight of 0.17. For regressions with country × year fixed effects, the latter is more relevant. The three countries with the largest weights on average are the United States, Germany, and Japan. Finally, we compute the

TABLE 4
DESCRIPTIVE STATISTICS OF FIRMS IN OUR BASELINE REGRESSION

| | Total (1) | Foreign (2) |
|---|---|---|
| | A. Average Weights | |
| Largest country | .47 | |
| Second-largest country | .17 | |
| Third-largest country | .10 | |
| Fourth-largest country | .07 | |
| United States | .21 | |
| Germany | .20 | |
| Japan | .17 | |
| France | .09 | |
| | B. Weight Variation | |
| HHI country (%) | 13.3 | 9.0 |
| HHI country × year (%) | .9 | .6 |
| Mean pairwise correlation | .08 | .13 |

NOTE.—The table presents summary statistics for the country weights of firms. Column 1 reports statistics for the total weights. Column 2 presents information on foreign weights (normalized to 1). Panel A reports the average weights of the largest countries. Panel B reports the HHI at the country and country × year level. The mean pairwise correlation is the average pairwise correlation between any two firms (col. 1) or firms within a home country (col. 2).

Herfindahl-Hirschman index (HHI) of the weights. The HHI of country-level weights is 0.13. For regressions with country × year fixed effects, the relevant statistic is the HHI of the foreign weights (i.e., excluding the home country and renormalizing to 1), which is 0.09 at the country level. With 15 years, the HHI of country × year weights is 15 times smaller. The average pairwise correlation of weights is low: 0.08 (and 0.13 for foreign weights of firms from the same country).

Table A5 reports standard deviations and a correlation matrix for the firm-level macroeconomic variables, residualized on firm and industry × year fixed effects. We still find significant variation in the residualized (log) low-skill wages, as the standard deviation is 0.03 (by comparison, the standard deviation is 0.1 when residualizing only on firm fixed effects). Appendix section A.6.1 provides additional statistics computed at the level of the shock of our shift-share variable (see table A30).

## IV. Global Wages and Induced Automation

We present our results in three steps. First, we find a positive effect of low-skill wages on automation innovations. Second, we show that this effect does not exist for nonautomation innovations 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 the magnitude of our effects and provide robustness checks.

### *A. Main Results*

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

the future. In contrast, we find clear evidence of knowledge spillovers: firms exposed to more knowledge in automation technologies tend to undertake more automation innovations. As a result, the overall effect of an increase in low-skill wages on innovation is larger than its short-run effect (see sec. IV.D).

*Country × year fixed effects.*—Unobserved country-level shocks in the innovator's country may impact both wages and innovation by affecting the cost of innovation or the demand for automation equipment through other channels than downstream wages. For instance, a tax reform in Germany could affect both German low-skill wages and directly the incentive to innovate for German firms. Shocks that affect firms mainly through their home country can be captured with (home) country × year fixed effects. As discussed further in section IV.C, our identifying assumption is then that foreign wages are exogenous to the firm's automation innovation given our controls. Columns 4–6 report the results with country × year fixed effects: we still obtain a positive effect of low-skill wages on automation innovations and a negative effect for high-skill wages with similar elasticities. In table A6, we instead define the home country as the country with the largest presample inventor weight or the headquarter country. We find similar effects.[33]

*Foreign wages.*—Building on these results, we decompose the macro variables into their home and foreign components and analyze the effect of foreign wages. We normalize foreign variables such that the coefficients can still be interpreted as elasticities.[34] Again, we find a positive effect of low-skill wages on automation innovation and a negative effect for high-skill wages (cols. 7–9). Neither Aghion et al. (2016) nor other papers using their methodology include country × year fixed effects or focus on foreign variation. As argued below, these fixed effects are generally important for identification in such settings.

The elasticity of automation innovation with respect to low-skill wages ranges from 2.2 to 3.6 when we focus on total wages and is somewhat larger (from 4.2 to 5.3) when we focus on foreign wages. This range does not depend on the inclusion of controls for stocks, spillovers, or the GDP gap

[33] For our sample, the country with the largest patent weight is also the one with the largest inventor weight in 61% of the cases (75% if firms are weighed by their count of auto95 patents).

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

TABLE 5
Baseline Regressions: Effect of Wages on Automation Innovations (auto95)

| | Dependent Variable: auto95 | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Domestic and Foreign | | | | | | Foreign | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| Low-skill wage | 2.97*** | 2.72*** | 3.64*** | 2.24** | 2.61** | 3.64*** | 4.19*** | 5.30*** | 4.43** |
| | (.80) | (.85) | (.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*** |
| | (.73) | (.80) | (.82) | (.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 | .04 | 2.40 | .50 |
| | (2.62) | (2.71) | (2.81) | (6.87) | (6.90) | (7.21) | (4.59) | (4.91) | (5.24) |
| Labor productivity | | .96 | | | −1.77 | | | −2.53 | |
| | | (.92) | | | (1.78) | | | (1.61) | |
| GDP per capita | | | −1.86 | | | −3.45* | | | −.42 |
| | | | (1.32) | | | (1.97) | | | (2.12) |
| Stock automation | −.12*** | −.12*** | −.12*** | −.12*** | −.12*** | −.12*** | −.12*** | −.13*** | −.12*** |
| | (.03) | (.03) | (.03) | (.03) | (.03) | (.03) | (.03) | (.03) | (.03) |
| Stock other | .51*** | .51*** | .51*** | .52*** | .52*** | .52*** | .51*** | .51*** | .51*** |
| | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) |

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Spillovers automation | .61** | .64** | .76** | 1.36*** | 1.34*** | 1.35*** | 1.33*** | 1.29*** | 1.32*** |
| | (.30) | (.30) | (.31) | (.47) | (.47) | (.47) | (.46) | (.46) | (.46) |
| Spillovers other | −.20 | −.25 | −.33 | −.97*** | −.93*** | −.99*** | −.97*** | −.97*** | −.98*** |
| | (.22) | (.22) | (.24) | (.36) | (.36) | (.36) | (.35) | (.35) | (.35) |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Industry × year fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Country × year fixed effects | No | No | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 47,812 | 47,812 | 47,812 | 47,453 | 47,453 | 47,453 | 47,453 | 47,453 | 47,453 |
| Number of firms | 3,236 | 3,236 | 3,236 | 3,233 | 3,233 | 3,233 | 3,233 | 3,233 | 3,233 |

NOTE.—The table presents the results of our baseline regression. The independent variables are lagged by two periods. Coefficients are estimated using conditional Poisson fixed effects regressions (Hausman, Hall, and Griliches 1984). All columns include firm and industry × year fixed effects. Columns 4–9 add country × year fixed effects. In cols. 7–9, the macroeconomic variables are the normalized foreign variables, as defined in the text. Standard errors (in parentheses) are clustered at the firm level.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

(see table A7). To interpret the size of this elasticity, note that our analysis focuses on innovation with a high automation content and reflects the behavior of firms undertaking automation innovations.[35] We discuss the magnitude of our effect further in section IV.D.

*Auto90.*—Table A8 reproduces table 5 but for the auto90 measure of automation. The results are very similar, but the coefficients on low-skill wages tend to be of a smaller magnitude, consistent with auto95 being a stricter measure of automation.

*Skill premium.*—The previous results suggest that the skill premium is a driver of automation innovations since the coefficients on low-skill and high-skill wages are of a similar magnitude but opposite signs. Table A9 directly regresses automation innovation on the log of the inverse of the skill premium. The coefficient on the inverse skill premium is similar to that on low-skill wages in the previous specifications and significant at the 1% level in all specifications.

### B. *Nonautomation Innovation and the Direction of Innovation*

Is the effect of wages on automation innovations specific to automation, or does it affect machinery patents in general? To answer this question, we now look at the nonautomation innovations in machinery undertaken by the sample of firms in our baseline regressions. Specifically, we reproduce the regressions of table 5 but for machinery innovations that are not auto90, denoted pauto90. We recompute the knowledge stocks and spillover variables for pauto90 innovations (own) and for all innovations except pauto90 (other). Table 6 reports the results. The coefficients on low- and high-skill wages are much smaller and significant in only one specification without country × year fixed effects for low-skill wages.[36] Therefore, the same firms react differently following the same shocks in their automation and nonautomation innovations.

The pauto90 patents are a good placebo for the auto95 patents, as they are both machinery patents and tend to target the same sectors (see fig. A9). However, to ensure even greater comparability between the two sets of patents, we focus on a subset of pauto90 patents—refined pauto90—which are technologically closer to auto95 patents. Specifically, we identify the set of four-digit C/IPC machinery codes that contain at least one six-digit auto95 code (i.e., a six-digit machinery code for which the prevalence of automation keywords is in the top 5% percent of the distribution).

[35] By comparison, the elasticities of clean and dirty patents with respect to fuel price in Aghion et al. (2016) are slightly smaller (between 0.5 and 3).

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

Refined pauto90 patents are pauto90 patents that belong to one of these four-digit C/IPC codes.[37] Columns 1–3 in table A10 show that the results are very similar to pauto90. In that table, we also include the full sample of firms with pauto90 innovations (cols. 4–6) and look at all machinery innovations excluding auto95 innovations (pauto95 in cols. 7–9). Again, the coefficients on low- and high-skill wages are much smaller than for auto95 patents and insignificant.

Finally, we show that wages affect the direction and not just the level of innovation in table A11: we regress the count of auto95 innovations controlling for the number of pauto90 patents. We find similar coefficients, so our results are not driven by a general tendency for firms to innovate more.[38]

### C. *Shift-Share Structure and Identification*

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

Since our measure of wages has a shift-share structure, we draw on the recent literature that discusses the identifying assumptions in this type of setup. We interpret our results through the lens of Borusyak, Hull, and Jaravel (2022). In the language of our setting, they show that the random assignment of wage shocks conditional on weights and controls can be sufficient for identification. The estimator is consistent if many country × year pairs are affected by weakly correlated shocks (we argue in app. sec. A.6.1 that these conditions are met).[39]

*Conditionally randomly assigned wage shocks.*—An important feature of our research design is that we include country × year fixed effects and focus on foreign macro variables. As a result, our regression coefficients do not reflect a spurious correlation between low-skill wages and innovation coming from domestic shocks, such as domestic tax policy. Instead, our regression coefficients must reflect the effect of foreign shocks that affect the demand for automation machinery and are correlated with changes in wages. Further, since we control for high-skill wages, these foreign

[37] Refined pauto90 patents tend to be used by the same sectors as auto95 patents: the employment-weighted correlation between the ratio of auto95 patents over capital purchases and refined pauto90 over capital purchases across US sectors rises to 0.90 (vs. 0.76 for all pauto90 patents).

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

[39] As shown in table 4, the Herfindahl index for our foreign weights at the country level is 0.09 and 0.006 at the country × year level. In app. sec. A.6.1, we argue that there is significant variation within countries (see also figs. 3, A12).

TABLE 6
EFFECT OF WAGES ON NONAUTOMATION INNOVATIONS

| | Dependent Variable: Placebo Machinery (pauto90) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Domestic and Foreign | | | | | | Foreign | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| Low-skill wage | .87 | .96 | 1.74* | .34 | .53 | .95 | 1.05 | 1.71 | 1.21 |
| | (.72) | (.78) | (.89) | (.97) | (1.03) | (1.29) | (1.53) | (1.64) | (1.78) |
| High-skill wage | −.47 | −.32 | .30 | −.72 | −.33 | −.35 | −1.51 | −.59 | −1.42 |
| | (.82) | (.79) | (.84) | (1.16) | (1.21) | (1.18) | (1.57) | (1.75) | (1.68) |
| GDP gap | −2.13 | −1.96 | .22 | 3.40 | 3.80 | 4.51 | −.24 | 1.10 | .05 |
| | (1.56) | (1.62) | (1.90) | (4.30) | (4.29) | (4.29) | (2.90) | (3.01) | (2.97) |
| Labor productivity | | −.33 | | | −.86 | | | −1.45 | |
| | | (.74) | | | (1.27) | | | (1.40) | |
| GDP per capita | | | −2.33* | | | −1.42 | | | −.26 |
| | | | (1.32) | | | (1.91) | | | (1.74) |
| Stock own | .36*** | .36*** | .36*** | .35*** | .35*** | .35*** | .35*** | .35*** | .35*** |
| | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) |
| Stock other | .15*** | .15*** | .15*** | .17*** | .17*** | .17*** | .16*** | .16*** | .16*** |
| | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) | (.04) |

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Spillovers own | 2.23*** | 2.24*** | 1.80*** | .88 | .89 | .82 | .92 | .97 | .92 |
| | (.47) | (.47) | (.58) | (.71) | (.71) | (.71) | (.71) | (.71) | (.71) |
| Spillovers other | −2.03*** | −2.03*** | −1.62** | −.85 | −.85 | −.80 | −.91 | −.97 | −.91 |
| | (.56) | (.57) | (.66) | (.75) | (.75) | (.75) | (.74) | (.75) | (.74) |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Industry × year fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Country × year fixed effects | No | No | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 42,538 | 42,538 | 42,538 | 42,405 | 42,405 | 42,405 | 42,405 | 42,405 | 42,405 |
| Number of firms | 2,848 | 2,848 | 2,848 | 2,845 | 2,845 | 2,845 | 2,845 | 2,845 | 2,845 |

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

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

shocks must have asymmetric effects for low- and high-skill workers, a point further supported by the results on the skill premium. Wages themselves are an equilibrium outcome, and we can think of wage shocks as coming from four sources of variation: regulatory changes, labor supply shocks, customer demand shocks, and technology shocks. We discuss these in turn.

First and second, regulatory changes or labor supply shocks present an ideal source of variation. For example, the introduction of a minimum wage or demographic or education shocks is unlikely to affect automation innovations through any channel other than an increase in labor costs. In principle, regulation or labor supply shocks could also affect the production costs of innovating firms and thus innovation. However, as long as production is concentrated in the home country, country × year fixed effects will absorb the effect.[40] In section V, we will focus on minimum wage changes and a specific labor-market shock, the Hartz reforms in Germany.

Third, foreign demand shocks for the customers of innovating firms may affect foreign manufacturing wages and thus automation innovation. But foreign demand shocks may also directly affect the demand for automation equipment and innovation, which could bias our wage coefficients. The asymmetry between low- and high-skill wage coefficients already rules out the possibility that our results are driven by skill-neutral demand shocks (which are also addressed by our controls for GDP gap and GDP per capita). However, sectoral demand shifts toward low-skill intensive sectors could raise low-skill wages (relative to high-skill wages) and may lead to an increase in auto95 innovations, if these innovations tend to target more low-skill intensive sectors (see Autor et al. 2024 for a model with such sectoral shifts). This is why we control for labor productivity in manufacturing in table 5. Further, in columns 1, 4, and 7 of table 7, we control for the size of the manufacturing sector, computed as the other macro variables.[41] Our coefficients on low- and high-skill wages hardly change. Still, skill-biased sectoral demand shifts could also occur within manufacturing. We build a variable to capture these: we weigh time-varying sectoral value added within manufacturing (for 46 sectors) at the country level by the low-skill share of labor costs in each sector in the United States in 1995. We then build this variable at the firm level using a shift-share. Columns 2, 5, and 8 report the results: the coefficients on wages remain

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

[41] 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.

similar. The coefficients on these controls do not suggest a consistent effect for foreign customer demand shocks.[42]

Fourth, foreign technology shocks could also directly affect automation innovations. Skill-neutral technology shocks are already addressed by our control for labor productivity. In contrast, a recent period of higher than usual automation innovation will increase the skill premium and could reduce future automation innovations through another channel than wages, for instance, if it affects the competitive landscape for machinery producers. To address this, we construct a measure of recent innovation analogous to that for low-skill wages: for each country, we compute the number of automation innovations applied for in the past 3 years. Then, we build firm-specific measures using the same patent weights as for wages (while spillover controls use inventor weights). We construct a similar control for other innovations. Columns 3, 6, and 9 of table 7 report the results. Our coefficients on low-skill wages remain similar, and these controls do not show a consistent effect across specifications. Reverse causality would also manifest itself as a technology shock. The effect of a firm on its home market is already captured by country × year fixed effects, but a firm's own innovation may affect foreign wages if the firm is particularly large. In addition, to the above control, we also directly control for the stock of knowledge at the firm level and lag wages. Finally, reverse causality would, if anything, bias the low-skill wage coefficient downward.[43] Therefore, skill-biased technology shocks do not seem to be driving our results.

To summarize, in the presence of country × year fixed effects and a control for high-skill wages, our regression coefficients must reflect the effect of foreign demand shocks for automation equipment producers with asymmetric effect on wages. As we have just seen, these foreign demand shocks are most likely regulatory shocks or labor supply shocks that allow us to identify a causal effect of wages on automation innovations. The stability of our coefficients to various controls can also be seen as a test of the exclusion restriction (Borusyak, Hull, and Jaravel 2022; Aghion et al. 2024). Importantly, our coefficients on low-skill wages should be compared with those from regressions with the placebo innovations (reported in tables 6 and A10). Should our result on the effect of low-skill wages on automation innovations come from a bias, then that bias would have to be absent for other types of machinery innovations undertaken by the same

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

[43] An additional concern might come from low-skill human capital shocks (captured by $\gamma(i)$ in the model of app. sec. A.4), which we cannot directly control for. However, a positive shock to low-skill human capital would be associated with higher wages and less automation innovation and would correspondingly also bias our estimates downward.

TABLE 7
CONTROLS FOR DEMAND EFFECTS AND TECHNOLOGY SHOCKS

| | Dependent Variable: auto95 | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Domestic and Foreign | | | | | | Foreign | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| Low-skill wage | 3.19*** | 2.40*** | 2.59*** | 2.33** | 2.19* | 2.07* | 4.09*** | 5.57*** | 6.94*** |
| | (.81) | (.85) | (.93) | (1.00) | (1.18) | (1.25) | (1.33) | (1.63) | (1.76) |
| High-skill wage | −1.67** | −2.37*** | −2.38*** | −2.07** | −1.97* | −1.14 | −4.01*** | −3.02** | −3.80** |
| | (.75) | (.77) | (.82) | (1.01) | (1.07) | (1.01) | (1.43) | (1.51) | (1.49) |
| GDP gap | .51 | 1.05 | −4.15 | 9.88 | 8.73 | 7.32 | 2.55 | 1.46 | 4.15 |
| | (3.05) | (2.85) | (2.69) | (6.98) | (6.86) | (6.82) | (5.63) | (5.21) | (4.99) |
| Labor productivity | | 1.98** | .73 | | −.68 | −2.36 | | −2.96* | −5.19** |
| | | (.96) | (1.03) | | (1.94) | (1.86) | | (1.79) | (2.11) |
| Manufacturing size | −1.68** | | | −2.74** | | | −1.02 | | |
| | (.67) | | | (1.07) | | | (1.16) | | |
| Manufacturing size (low-skill weighted) | | −1.83*** | | | −1.73* | | | .57 | |
| | | (.61) | | | (.98) | | | (1.02) | |
| Recent auto95 innovation | | | −1.14 | | | −2.51** | | | 1.24 |
| | | | (.76) | | | (1.26) | | | (.93) |
| Recent other innovation | | | .87* | | | 1.56** | | | −.47 |
| | | | (.48) | | | (.78) | | | (.80) |

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Stocks and spillovers | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Industry × year fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Country × year fixed effects | No | No | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 47,812 | 47,812 | 47,812 | 47,453 | 47,453 | 47,453 | 47,453 | 47,286 | 47,453 |
| Number of firms | 3,236 | 3,236 | 3,236 | 3,233 | 3,233 | 3,233 | 3,233 | 3,220 | 3,233 |

NOTE.—The table adds additional control variables. Manufacturing size denotes the log of weighted average of manufacturing value added in the customers' countries. Manufacturing size (low-skill weighted) weighs the value added of each manufacturing sector by the low-skill labor share in labor costs in the United States in 1995. Recent auto95 innovation and recent other innovation denote the log weighted averages of the flow of auto95 and other innovations in the past 3 years in the customers' countries. All columns include firm and industry × year fixed effects. Columns 4–9 add country × year fixed effects. Columns 7–9 use the normalized foreign variables, as defined in the text. The normalized foreign manufacturing sizes and foreign innovations are defined similarly to normalized foreign low-skill wages. Standard errors (in parentheses) are clustered at the firm level.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level

*** Statistically significant at the 1% level.

firms and for the same four-digit C/IPC codes. Finally, section V provides direct evidence from regulatory changes with quantitatively similar results.

*Falsification tests and inference.*—We now turn to inference. Two potential issues arise: residual errors of firms with similar country distributions may be correlated (Adão, Kolesár, and Morales 2019), and our identification variations come from a limited set of country × year observations (Borusyak, Hull, and Jaravel 2022). To address these in our Poisson setting, we implement a Monte Carlo simulation similar to those of Borusyak and Hull (2023). We base our simulation on the regressions of table 5. For each country, we sample with replacement the full path of macro variables (wages, labor productivity, GDP per capita, and GDP gap) from the existing set of countries (e.g., Germany will get the macro variables of Canada). Then, for each firm, we compute the firm-level macro variables as the weighted average of these new country-level variables, keeping the original country weights. We keep the automation activity, the stocks of innovations, and the spillover variables as in the data. We run the regressions, store the coefficients on low-skill and high-skill wages, and repeat 4,000 times. Table 8 reports the $p$-values of the original coefficients on low-skill and high-skill wages based on the simulated distribution of coefficients. The $p$-values are not markedly different from those in table 5. In particular, the low-skill wage coefficients are significant at least at the 10% level (except in col. 4, with $p = .12$) and at the 1% level when we focus on foreign wages. In the language of Adão, Kolesár, and Morale (2019), the set of controls absorbs most of the country-specific shocks affecting the outcome variable, and consequently, no shift-share structure is left in the regression residuals. Figure A3*A*–A3*C* plot the distribution of coefficients for columns 2, 5, and 8. We can also view this permutation exercise as a falsification test, and accordingly, the distributions are centered around zero.

Similarly, we perform an additional falsification test where instead of permuting the macro variables across countries, we permute the weights across firms from the same country. Specifically, for each firm, we keep the automation activity, the stocks of innovations, and the spillover variables as in the data, but we sample (with replacement) their weights from the set of firms from the same country. That is, we may now attribute to Siemens the wages that, in reality, Bosch faces. We repeat the exercise 4,000 times, and table 8 reports the $p$-values of the original coefficients on low-skill and high-skill wages based on the simulated distribution of coefficients. The $p$-values are similar or smaller than those of table 5. Figure A3 plots the distribution of coefficients we obtain for columns 2, 5, and 8. The distributions are centered around zero when we include country × year fixed effects (fig. A3*E*, A3*F*). In figure A3*D*, we use the home country variation for identification, and the mean coefficient is positive

TABLE 8
Falsification Tests including Standard Error Bias from Adão, Kolesár, and Morales (2019)

| | 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*** |
| | [.012] | [.027] | [.002] | [.124] | [.088] | [.020] | [.007] | [.001] | [.008] |
| | {.002} | {.002} | {.000} | {.013} | {.006} | {.002} | {.002} | {.001} | {.008} |
| High-skill wage | −2.23* | −2.64** | −1.56 | −2.81* | −2.04 | −1.87 | −4.47*** | −2.91* | −4.33*** |
| | [.058] | [.026] | [.182] | [.054] | [.150] | [.169] | [.005] | [.058] | [.007] |
| | {.001} | {.003} | {.016} | {.002} | {.064} | {.070} | {.001} | {.060} | {.002} |
| GDP gap | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Labor productivity | No | Yes | No | No | Yes | No | No | Yes | No |
| GDP per capita | No | No | Yes | No | No | Yes | No | No | Yes |
| Stocks and spillovers | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Industry × year fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Country × year fixed effects | No | No | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 47,812 | 47,812 | 47,812 | 47,453 | 47,453 | 47,453 | 47,453 | 47,453 | 47,453 |
| Number of firms | 3,236 | 3,236 | 3,236 | 3,233 | 3,233 | 3,233 | 3,233 | 3,233 | 3,233 |

Note.—The table reproduces the baseline regression and performs two falsification tests. The first test addresses Adão, Kolesár, and Morales's (2019) concern that in shift-share design, observations with similar weights distribution may have correlated errors, leading to an overrejection of the null hypothesis. For each country, we sample with replacement the entire path of macroeconomics variables (wages, labor productivity, GDP per capita, and GDP gap) from the existing set of countries. We then run the same regressions as in our baseline. We repeat 4,000 times and compare the coefficients from the true regression with the distribution of coefficients from simulated regressions. Square brackets report the *p*-values of our original coefficients. In a second exercise, we perform a similar exercise, but instead of redrawing the macro variables, we redraw firms' weights from the distribution of weights for firms in the same country. We run regressions on the simulated data, and curly brackets state the *p*-values. Significance levels use the *p*-values in square brackets.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

but significantly smaller. This exercise highlights that the relevant variation is between firms of the same country and not simply cross-country.

### *D. Magnitude of the Effect*

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

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

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

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

Conceptually, and as emphasized in the introduction, it is important to distinguish the adoption response to a wage shock from the innovation response. In particular, calibrating directed technical change macro models requires estimates of the innovation response (Acemoglu 2023). While a shock to adoption must die out over time (when all firms have adopted a given technology), a shock to innovation can build on itself.

In the short run, technology adoption is dominated by the increased use of existing technologies. In the medium run, however, the innovation response is a key determinant of the overall adoption response, while new adoption of existing technologies plays a smaller role. In the longer run, the innovation response also depends on knowledge spillovers.[44]

*Simulation results.*—In the following, we analyze the economic magnitude of our coefficients and highlight the role played by spillovers. We do so by estimating the long-run effect of a change in the skill premium on task demand and the labor share that runs through endogenous automation innovation. To do so, we run a simulation in which we consider a uniform and permanent decrease in the skill premium by 10% between 1995 and 2009. We run a regression that jointly estimates the effect of automation innovations (auto95) and other machinery innovations (pauto95). We then use our regression results to simulate a Poisson process for each firm, and we recompute the share of automation innovations in machinery over that period. We run alternative simulations that either include or exclude the effects of spillovers and firms' knowledge stocks. Appendix section A.6.4 details our procedure including the exact regression (table A43).

Figure 4 reports the results averaged over 800 simulations.[45] 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 effect is captured by the gap between the data curve and the data + direct effect curve. This gap reflects the elasticity of 2.55 of auto95 innovations with respect to the inverse skill premium (with an elasticity of 0.52 for pauto95). Next, we include the effect of updating firms' own innovation stocks in the data + direct + stock curve, which slightly decreases the effect of low-skill wages, reflecting the negative effect of the automation stock on auto95 innovations and its positive effect on pauto95 innovations.

We then assess the importance of knowledge spillovers by recomputing the spillover variables for the auto95 and pauto95 innovations. This is not a straightforward exercise because we need to predict not only the number of innovations but also their location. The baseline curve corresponds to the median realization for the share of automation patents

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

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

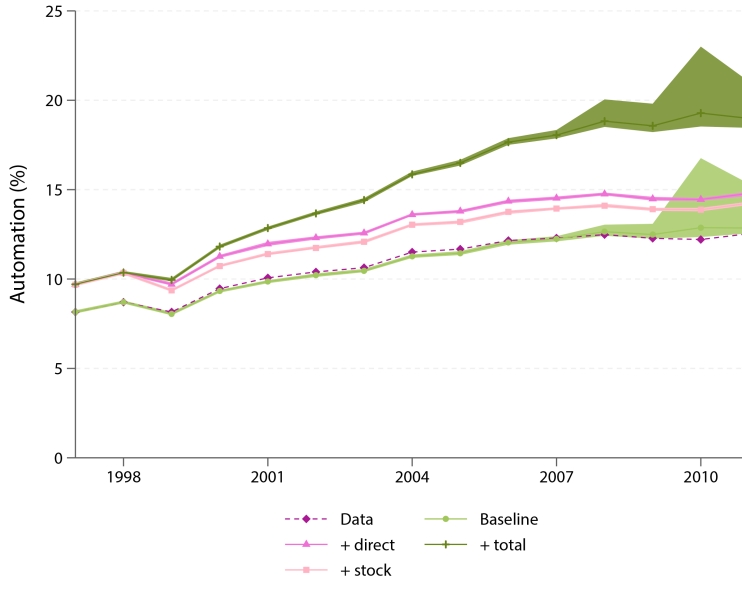


FIG. 4.—Simulation of permanent and global 10% decrease in skill premium on share of automation innovations in machinery. We report the median share of automation patents in machinery across 800 simulations. The shaded areas depict the 33.3%–66.7% range. The effects are computed on the basis of the regression results of table A43. The baseline curve reports the mean realization without any change in the skill premium but when the knowledge spillovers are recomputed every period depending on the simulated number of innovation in the previous periods. The total effect of the skill premium shock is captured by the gap between the total effect curve and the baseline curve.

without a wage change (while the data series is one possible realization). The overall effect of an increase in the inverse skill premium on the share of automation innovation is then captured by the gap between the baseline curve and the baseline + total effect curve. Knowledge spillovers increase the elasticity of the share of automation patents with respect to the inverse skill premium. The average share of automation innovations in machinery between 1997 and 2011 increases by 4.3 percentage points, that is, 2.3 percentage points more than the direct effect (for comparison, the share of automation innovations increased by 4.4 percentage points over the same time period).

Finally, we combine these effects with the results of section II.H on the labor market outcomes of automation use (cols. 1, 3, and 5 of table 3). A 4.3 percentage point increase in the share of automation innovation is associated with a decline in routine cognitive tasks of 6.1 centiles, a decline in routine manual tasks of 5.5 centiles, and a decline in the labor share in manufacturing of 5.6 percentage points per decade (for comparison, routine cognitive tasks, routine manual tasks, and the labor share declined by 1.4 centiles, 1.3 centiles, and 5.2 percentage points, respectively, per

decade in the sectors considered over the period 1980–98). Importantly, we stress that one must not interpret the results of this simulation as predictive, notably because a change in innovation will affect the skill premium. Nevertheless, this partial equilibrium analysis shows that the effect of the inverse skill premium on automation innovations is economically significant.

### *E. Additional Results and Robustness Checks*

This section discusses additional results and robustness checks.

*Shift-share design.*—Borusyak, Hull, and Jaravel (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 section A.6.1, instead of our Poisson regression, we consider a linear setting where such an equivalence result applies: we use the arcsinh of the count of automation patents as the dependent variable and replace our log of the average macro variables with the average of the logs. This linear setting allows us to give summary statistics on our shock variable and unpack the relationship between the inverse skill premium and automation in the data. Figure A11 shows bin scatterplots of the shock-level regressions of residualized automation measures on the inverse skill premium: the relationship appears linear and not driven by outliers. We also report how balanced our shocks are with respect to observables.

Goldsmith-Pinkham, Sorkin, and Swift (2020) show that, alternatively, identification in a shift-share design can be obtained if the weights are exogenous. In our context, this assumption is likely violated because firms' decision to innovate may be affected by other macro shocks in the destination countries, which would affect firms in proportion to the same weights. This is why we rely on Borusyak, Hull, and Jaravel (2022). Nevertheless, we note that our weights appear predetermined and do not reflect firms' expectations of future wage growth: country-level growth rates in low- and high-skill wages between 1995 and 2000 have no predictive power on firm weights in 1995 (see table A12). In addition, we conduct robustness checks on our weights in table A13: we exclude automation patents from the weights, we use a longer lag between the period used to compute the weights and the regression period either by computing weights only up to 1989 or by dropping the first 5 years of the regression, and we drop the earlier years of the presample period when computing the weights. Our results are robust in all cases.

We conduct additional exercises related to our shift-share setting in appendix section A.6.1. First, we show that no single country drives our results by sequentially excluding the six largest countries. Second, Borusyak, Hull, and Jaravel (2022) recommend considering other shock-level variables that may bias the results. We consider the effect of offshoring and the real

interest rate. We also control for R&D costs by building firm-specific wage variables using weights based on the location of inventors instead of patent offices. Finally, we implement the correction suggested by Borusyak and Hull (2023) to address the nonlinearity of logged shift-share measures. Our results are robust in all cases.

*Timing and pretrends.*—We now study different timing assumptions to assess whether there are pretrends and because our choice of a 2-year lag is somewhat arbitrary. In figure A4, we look at alternative lags and leads for the dependent variables. We consider two specifications, both controlling for GDP gap, labor productivity, and country × year fixed effects. In figure A4*A*, we look at total wages, corresponding to column 5 of table 5. In figure A4*B*, we consider only foreign wages, corresponding to column 8.[46] The 2-year lag delivers the highest coefficient for low-skill wages in both cases. This is in line with the empirical literature on induced innovation using patent data, which often finds effects peaking with a 2–3-year lag (see, e.g., Aghion et al. 2016; Popp 2002). A possible interpretation of this fast response is that firms may prioritize existing projects over starting new ones.[47]

Figure A4 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 autocorrelated, and firms may anticipate shocks at short horizons. Importantly, though, we find no significant effect for longer leads, suggesting that there are no pretrends (testing for such pretrends is one of the recommendations of Borusyak, Hull, and Jaravel 2022). Table A14 runs a horse race regression between 2-year lagged macro variables and macro variables with varying leads or lags. The coefficient on lead high-skill wage is significant in only one case, and the effect of 2-year lagged low-skill wages is always positive and significant except in one specification.

Additionally, innovators should care only about current wages insofar as they are predictive of future wages. In table A15, we compute predicted future wages at time $t - 2$ based on an AR(1) process with country-specific trends instead of directly using lagged wages. The results are similar to our baseline.

*Nickell's bias.*—Our regressions include the stock of automation innovations and may suffer from Nickell's bias. Table A16 removes stocks or uses the standard method of Blundell, Griffith, and van Reenen (1999), which proxies for the fixed effect with the firm's presample average of the dependent variable. We obtain similar results.

[46] We keep a lag of two periods for the stock variables; otherwise, the dependent variable would be included in the right-hand side in the lead and contemporaneous cases.

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

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

## V. Labor Market Reforms and Induced Automation

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

### A. *Minimum Wage*

We compute the average minimum wage faced by the downstream customers of a specific firm selling automation equipment. We do this as a shift-share measure of country-level minimum wages exactly as we did for low-skill wages, with the caveat that we have data for only 22 countries instead of 41.[48] We then run panel regressions similar to our baseline where we replace the low-skill wage with the minimum wage. Table 9 reports the results. For regressions on total wages (cols. 1–6), the coefficients on minimum wage are similar to those on low-skill wages in previous tables, and for regressions on foreign wages (cols. 7–9), the coefficients are positive but smaller and, in one case, insignificant.[49] Overall, these results support our hypothesis that higher low-skill labor costs induce automation innovations.

### B. *Event Study: The Hartz Reforms in Germany*

We now examine the effects of the Hartz reforms, a series of labor market reforms in Germany, first drafted in 2002 and implemented between

[48] We use data from the Organization for Economic Cooperation and Development. 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 app. sec. A.5.1.

[49] A smaller coefficient is not surprising: first, we focus on manufacturing, where low-skill wages tend to be above the minimum wage. Second, the minimum wage captures only a portion of the labor costs. Third, the quality of the data is worse, as we lose nearly half of our countries. In table A17, we restrict attention to firms with at least two auto95 patents in the sample period and obtain more precisely estimated coefficients.

TABLE 9
EFFECT OF MINIMUM WAGE

| | auto95 | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Domestic and Foreign | | | | | | Foreign | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| Minimum wage | 2.02*** | 1.80** | 2.22*** | 1.78* | 1.81* | 1.97* | 2.59* | 2.64* | 1.36 |
| | (.73) | (.72) | (.83) | (.98) | (1.01) | (1.12) | (1.34) | (1.42) | (1.56) |
| High-skill wage | −1.53** | −2.33*** | −1.01 | −3.39*** | −3.26** | −2.99** | −3.79** | −3.64* | −5.96*** |
| | (.71) | (.89) | (1.02) | (1.06) | (1.33) | (1.49) | (1.51) | (2.01) | (2.07) |
| GDP gap | −3.34 | −4.01 | −2.49 | 8.82 | 8.91 | 9.50 | 4.38 | 4.55 | −1.27 |
| | (2.67) | (2.69) | (3.00) | (6.53) | (6.58) | (6.98) | (5.53) | (6.12) | (7.20) |
| Labor productivity | | 1.52* | | | −.24 | | | −.19 | |
| | | (.87) | | | (1.65) | | | (1.75) | |
| GDP per capita | | | −.90 | | | −.80 | | | 4.25 |
| | | | (1.48) | | | (2.18) | | | (2.72) |
| Stocks and spillovers | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Firm fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Industry × year fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Country × year fixed effects | No | No | No | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 46,590 | 46,590 | 46,590 | 46,328 | 46,328 | 46,328 | 46,328 | 46,328 | 46,328 |
| Number of firms | 3,154 | 3,154 | 3,154 | 3,151 | 3,151 | 3,151 | 3,151 | 3,151 | 3,151 |

NOTE.—The table replaces the low-skill wage with the minimum wage. All columns include firm and industry × year fixed effects. Columns 4–9 add country × year fixed effects. In cols. 7–9, the macroeconomic variables are the normalized foreign variables, as defined in the text. Standard errors (in parentheses) are clustered at the firm level.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

January 1, 2003, and January 1, 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, attribute an important role to the reforms in the remarkable performance of the German labor market since then, particularly in increasing labor supply and improving matching efficiency.

Such reforms would reduce the incentive to automate low-skill labor by decreasing labor costs both directly and indirectly through an increase in labor supply and a reduction in the expected cost of vacancies. The Hartz reforms are perhaps the most salient labor market reforms in a major country in our time period. This presents an ideal setting: the Hartz reforms are unlikely to have affected the direction of innovation in non-German firms through channels other than the German labor market and were the largest macroeconomic shock in Germany at the time. The reforms had a large and immediate effect. As soon as they were implemented in 2003, low-skill labor costs started stagnating, while high-skill labor costs kept rising, leading to a sharp decline in the inverse skill premium in Germany (see fig. A5). In contrast, there is no such trend for the aggregate rest of the world.

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

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

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

Figure 5*A* reports the results. The coefficient of −2.71 in 2010 means that, on average, a firm with a German weight of 0.1 (the mean value is 0.105) had a 27.1% smaller increase in automation innovations between 2005 and 2010 than a firm with no German exposure. This aligns with our regression results: between 2003 and 2008, the inverse skill premium

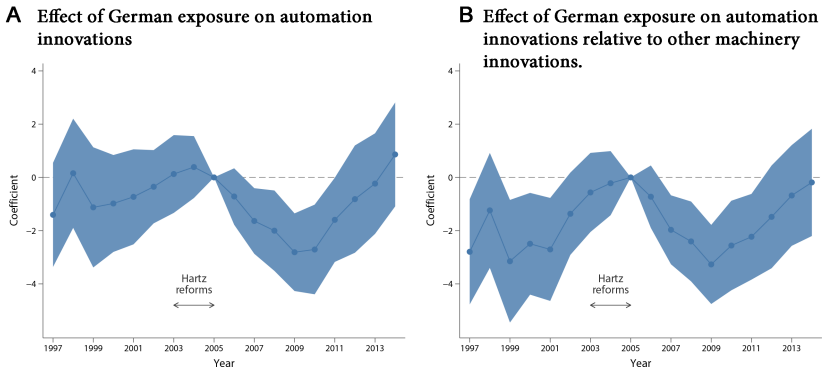


FIG. 5.—Effect of German exposure on automation innovations. *A* reports coefficients on the interaction between the German weight and a set of year fixed effects in a Poisson regression of auto95 innovations controlling for a full set of fixed effects and firm innovation stocks with 2,156 firms. *B* reports coefficients on the triple interaction between the German weight, a dummy for auto95 innovations, and a set of year fixed effects in a regression of auto95 and other machinery innovations controlling for a full set of fixed effects, firm innovation stocks, and the interaction between the German weight and a set of year fixed effects with 6,692 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.

in Germany declined by 12.3% relative to the rest of the world. Using the elasticity of 2.5 in column 4 of table A9, this would correspond to a decline in automation innovations of 30.8% between 2005 and 2010.

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

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

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

where $k$ denotes the type of an innovation that is either auto95 or other machinery innovation (pauto95); $\delta_{k,i}$ represents a set of innovation type

firm fixed effects, $\delta_{k,c,t}$ innovation type country × year fixed effects, and $\delta_{k,j,t}$ innovation type industry year fixed effects; $1_{k=\text{aut}}$ is a dummy for an auto95 innovation; and $K_{\text{Paut},i,t}$ is the stock of other machinery innovations (pauto95) and $K_{\text{other},i,t}$ the stock of nonmachinery innovations. $\beta^{\text{aut}}_{\text{DE},t}$ are the coefficients of interest. For each year, they measure how much exposure to Germany increases the relative propensity to introduce automation innovations compared with other forms of machinery innovations relative to 2005. The coefficients $\beta_{\text{DE},t}$ measure the effect of German exposure that is common to all machinery innovations. Figure 5*B* reports the results: the pattern is, if anything, more pronounced than in figure 5*A*.

To formally test that the Hartz reform created a trend break, we replace the set of year fixed effects $\delta_t$ in $\beta^{\text{aut}}_{\text{DE}} \cdot \delta_t \kappa_{i,\text{DE}} 1_{k=\text{aut}}$ in equation (5) with a time trend $t - 2005$ and a time trend interacted with a post-2005 dummy

TABLE 10
AUTOMATION VERSUS NONAUTOMATION INNOVATION AND EXPOSURE TO GERMANY: TRIPLE-DIFFERENCE EXERCISE

| | auto95 and pauto95 | | | auto95 and pauto90 |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Time trend × auto95 dummy × German exposure × post | −1.05*** | −1.10*** | −1.08*** | −1.12*** |
| | (.33) | (.33) | (.33) | (.33) |
| Time trend × auto95 dummy × German exposure | .49*** | .46*** | .45** | .46*** |
| | (.18) | (.18) | (.18) | (.18) |
| Time trend × German exposure × post | | | .31 | |
| | | | (.22) | |
| Time trend × German exposure | | | −.34** | |
| | | | (.15) | |
| Firm innovation stocks × innovation types | No | Yes | Yes | Yes |
| Year dummy × German exposure | Yes | Yes | No | Yes |
| Industry × year × innovation types fixed effects | Yes | Yes | Yes | Yes |
| Country × year × innovation types fixed effects | Yes | Yes | Yes | Yes |
| Firm × innovation types fixed effects | Yes | Yes | Yes | Yes |
| Observations | 76,037 | 76,037 | 76,037 | 74,102 |
| Number of firms | 5,416 | 5,416 | 5,416 | 5,280 |

NOTE.—The table shows that the effect of German exposure is specific to automation innovations. All regressions control for firm innovation types fixed effects, country × year × innovation types fixed effects, and industry × year × innovation types fixed effects. Innovation types are auto95 and pauto95 in cols. 1–3 and auto95 and pauto90 in col. 4. Columns 2–4 control for innovation stocks lagged by two periods interacted with innovation types dummies. Column 3 controls for a linear time trend times the German exposure instead of yearly dummies times the German exposure. Throughout, German exposure is measured by the German weight. Standard errors (in parentheses) are clustered at the firm level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

$(t - 2005)_{t>2005}$. We focus on the years 2000–2010 to have a panel centered on 2005 and to avoid the effects of the Great Recession on innovation. Table 10 reports the result. Column 2 corresponds exactly to this specification. We find a significant time trend in the effect of German exposure on the relative propensity to innovate in automation between 2000 and 2005. However, the trend sharply reverses in the following 5 years. Column 1 omits the controls for the stock variables. Column 3 replaces the flexible set of year dummies times German exposure, $\delta_{ki,\mathrm{DE}}$, by a time trend times German exposure and a time trend times German exposure after 2005. Finally, instead of looking at auto95 and pauto95 (i.e., all nonauto95 machinery innovations) innovation, column 4 considers auto95 and pauto90 innovations (which we used as the default nonautomation innovations in table 6). In all cases, the trend break on automation innovations remains with a consistent magnitude. Overall, this section shows that, in line with our theory, the Hartz reforms reduced automation innovations by foreign firms highly exposed to Germany both in absolute terms and relative to other types of machinery innovation.

## VI. Conclusion

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

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

Our results highlight that labor market policies can generate technological responses, so the long-term effects of these policies may differ from their short-term effects. Analyzing this feedback loop quantitatively would

require the development of a macroeconomic model, which could be calibrated using our estimates. More generally, our estimates can discipline an emerging literature on automation and economic growth. In addition, future research could adapt our classification method to automation patents beyond machinery and analyze how much the emergence of recent automation technologies for high-skill labor, such as artificial intelligence, results from rising high-skill wages.

## Data Availability

Data and code replicating the results in this paper are available in Hémous et al. (2024) in the Harvard Dataverse, https://doi.org/10.7910/DVN/IMCSW1.

## References

Acemoglu, Daron. 2023. "Distorted Innovation: Does the Market Get the Direction of Technology Right?" *A.E.R. Papers and Proc.* 113:1–28.

Acemoglu, Daron, and Amy Finkelstein. 2008. "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector." *J.P.E.* 116 (5): 837–80.

Acemoglu, Daron, and Joshua Linn. 2004. "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry." *Q.J.E.* 119 (3): 1049–90.

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." *A.E.R.* 108 (6): 1488–542.

———. 2020. "Robots and Jobs: Evidence from US Labor Markets." *J.P.E.* 128 (6): 2188–244.

———. 2022. "Demographics and Automation." *Rev. Econ. Studies* 89 (1): 1–44.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-Share Designs: Theory and Inference." *Q.J.E.* 134 (4): 1949–2010.

Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel. 2024. "Modern Manufacturing Capital, Labor Demand, and Product Market Dynamics: Evidence from France." Working paper.

Aghion, Philippe, Roland Bénabou, Ralf Martin, and Alexandra Roulet. 2023. "Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?" *A.E.R. Insights* 5 (1): 1–20.

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

Andersson, David, Mournir Karadja, and Erik Prawitz. 2022. "Mass Migration and Technological Change." *J. European Econ. Assoc.* 20:1859–96.

Aschhoff, Birgit, Dirk Crass, Katrin Cremers, Christoph Grimpe, Christian Rammer, Felix Brandes, Fernando Diaz-Lopez, Rosalinde Klein Woolthuis, Michael Mayer, and Carlo Montalvo. 2010. "European Competitiveness in Key Enabling Technologies." Mannheim: Centre European Econ. Res. (ZEW).

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. "New Frontiers: The Origins and Content of New Work, 1940–2018." *Q.J.E.* 139 (3): 1399–465.

Autor, David, and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market." *A.E.R.* 103 (5): 1553–97.
Autor, David, David Dorn, Gordon Hanson, Gary Pisano, and Pian Shu. 2020. "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents." *A.E.R. Insights* 2 (3): 357–74.
Autor, David, Frank Levy, and Richard Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Q.J.E.* 118 (4): 1279–333.
Baptista, Rui. 2000. "Do Innovations Diffuse Faster within Geographical Clusters?" *Internat. J. Indus. Org.* 18:515–35.
Békés, Gábor, and Péter Harasztosi. 2020. "Machine Imports, Technology Adoption, and Local Spillovers." *Rev. World Econ.* 156:343–75.
Bena, Jan, and Elena Simintzi. 2019. "Machines Could Not Compete with Chinese Labor: Evidence from U.S. Firms' Innovation." Working paper.
Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge. 2023. "What Happens to Workers at Firms That Automate?" *Rev. Econ. and Statis.* 107 (1): 125–41.
Bessen, James, and Robert Hunt. 2007. "An Empirical Look at Software Patents." *J. Econ. and Management Strategy* 16 (1): 157–89.
Blundell, Richard, Rachel Griffith, and John Van Reenen. 1999. "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms." *Rev. Econ. Studies* 66 (3): 529–54.
Borusyak, Kirill, and Peter Hull. 2023. "Non-Random Exposure to Exogenous Shocks: Theory and Applications." *Econometrica* 91 (6): 2155–85.
Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *Rev. Econ. Studies* 89 (1): 181–213.
Boustan, Leah Platt, Jiwon Choi, and David Clingingsmith. 2022. "Automation after the Assembly Line: Computerized Machine Tools, Employment and Productivity in the United States." Working paper.
Cameron, A. Colin, Jonah Gelbach, and Douglas Miller. 2008. "Boostrap-Based Improvements for Inference with Clustered Errors." *Rev. Econ. and Statis.* 90 (3): 414–27.
Clemens, Michael, Ethan Lewis, and Hannah Postel. 2018. "Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion." *A.E.R.* 108 (6): 1468–87.
Coelli, Federica, Andreas Moxnes, and Karen Ulltveit-Moe. 2022. "Better, Faster, Stronger: Global Innovation and Trade Liberalization." *Rev. Econ. and Statis.* 104 (2): 205–16.
Correia, Sergio, Paulo Guimarães, and Tom Zylkin. 2020. "Fast Poisson Estimation with High-Dimensional Fixed Effects." *Stata J.* 20 (1): 95–115.
Danzer, Alexander, Carsten Feuerbaum, and Fabian Gaessler. 2024. "Labor Supply and Automation Innovation." *J. Public Econ.* 235:105136.
Dechezleprêtre, Antoine, Matthieu Glachant, Ivan Hascic, Nick Johnstone, and Yvan Ménière. 2011. "Invention and Transfer of Climate Change-Mitigation Technologies: A Global Analysis." *Rev. Environmental Econ. and Policy* 5:109–30.
Dechezleprêtre, Antoine, Yann Ménière, and Myra Mohnen. 2017. "International Patent Families: From Application Strategies to Statistical Indicators." *Scientometrics* 111 (2): 793–828.
de Rassenfosse, Gaétan, Hélène Dernis, Dominique Guellec, Lucio Picci, and Bruno van Pottelsberghe de la Potterie. 2013. "The Worldwide Count of Priority Patents: A New Indicator of Inventive Activity." *Res. Policy* 42 (3): 720–37.
Doms, Mark, Timothy Dunne, and Kenneth Troske. 1997. "Workers, Wages, and Technology." *Q.J.E.* 112 (1): 253–90.

Eaton, Jonathan, and Samuel Kortum. 1996. "Trade in Ideas: Patenting and Productivity in the OECD." *J. Internat. Econ.* 40:251–78.
Eaton, Jonathan, Samuel Kortum, and Francis Kramarz. 2011. "An Anatomy of International Trade: Evidence from French Firms." *Econometrica* 79 (5): 1453–98.
Fan, Haichao, Yichuan Hu, and Lixin Tang. 2021. "Labor Costs and the Adoption of Robots in China." *J. Econ. Behavior and Org.* 186:608–31.
Gaggl, Paul, and Greg Wright. 2017. "A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive." *American Econ. J. Appl. Econ.* 9 (3): 262–94.
Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *A.E.R.* 110 (8): 2586–624.
Graetz, Goerg, and Guy Michaels. 2018. "Robots at Work." *Rev. Econ. and Statis.* 100 (5): 753–68.
Hall, Bronwyn, Zvi Griliches, and Jerry Hausman. 1986. "Patents and R&D: Is There a Lag?" *Internat. Econ. Rev.* 27 (2): 265–84.
Hausman, Jerry, Bronwyn Hall, and Zvi Griliches. 1984. "Econometric Models for Count Data with an Application to the Patents-R&D Relationship." *Econometrica* 52 (4): 909–38.
Hémous, David, and Morten Olsen. 2022. "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality." *American Econ. J. Macroeconomics* 14 (1): 179–223.
Hémous, David, Morten Olsen, Carlo Zanella, and Antoine Dechezleprêtre. 2024. "Replication Data for: 'Induced Automation Innovation: Evidence from Firm-Level Patent Data.'" Harvard Dataverse, https://doi.org/10.7910/DVN/IMCSW1.
Hornbeck, Richard, and Suresh Naidu. 2014. "When the Levee Breaks: Black Migration and Economic Development in the American South." *A.E.R.* 104 (3): 963–90.
Humlum, Anders. 2021. "Robot Adoption and Labor Market Dynamics." Working paper.
Kaufer, Erich. 1989. *The Economics of the Patent System.* New York: Rootledge.
Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller. 2022. "Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations." Working paper.
Krause, Michael, and Harald Uhlig. 2012. "Transitions in the German Labor Market: Structure and Crisis." *J. Monetary Econ.* 59 (1): 64–79.
Krusell, Per, Lee Ohanian, José-Víctor Ríos-Rull, and Giovanni Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5): 1029–53.
Lewis, Ethan. 2011. "Immigration, Skill Mix and Capital Skill Complementarity." *Q.J.E.* 126 (2): 1029–69.
Lordan, Grace, and David Neumark. 2018. "People versus Machines: The Impact of Minimum Wages on Automatable Jobs." *Labour Econ.* 52:40–53.
Lybbert, Travis, and Nikolas Zolas. 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." *Res. Policy* 43:530–42.
Mann, Katja, and Lukas Püttmann. 2023. "Benign Effects of Automation: New Evidence from Patent Texts." *Rev. Econ. and Statis.* 105 (3): 562–79.
No, Joung. 2008. "Cities and Growth: Knowledge Spillovers in the Adoption of Advanced Manufacturing Technologies." Canadian Econ. Transition Series, Ottawa.
Noailly, Joëlle, and Roger Smeets. 2015. "Directing Technical Change from Fossil-Fuel to Renewable Energy Innovation: An Application Using Firm-Level Patent Data." *J. Environmental Econ. and Management* 72:15–37.

Popp, David. 2002. "Induced Innovation and Energy Prices." *A.E.R.* 92 (1): 160–80.

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

Siemens. 2018. "Annual Report." Munich: Siemens.

Timmer, Marcel, Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries. 2015. "An Illustrated User Guide to the World Input-Output Database: The Case of Global Automotive Production." *Rev. Internat. Econ.* 23 (3): 575–605.

Voth, Hans-Joachim, Bruno Caprettini, and Alex Trew. 2023. "Fighting for Growth: Labor Scarcity and Technological Progress during the British Industrial Revolution." Working paper.

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

Zeira, Joseph. 1998. "Workers, Machines, and Economic Growth." *Q.J.E.* 113 (4): 1091–117.