



For whom the bell tolls: The firm-level effects of automation on wage and gender inequality

Giacomo Domini ^{a,1}, Marco Grazzi ^{b,1}, Daniele Moschella ^{c,1}, Tania Treibich ^{d,e,c,*,1}

^a Erasmus University College, Erasmus University Rotterdam, The Netherlands

^b Department of Economic Policy, Università Cattolica del Sacro Cuore, Milano, Italy

^c Institute of Economics & Department EMbeDS, Scuola Superiore Sant'Anna, Pisa, Italy

^d School of Business and Economics, University of Maastricht, The Netherlands

^e Science-Po Paris, OFCE-DRIC, Sophia Antipolis, France

ARTICLE INFO

JEL classification:

D25
J16
J31
L25
O33

Keywords:

Automation
AI
Wage inequality
Gender pay gap

ABSTRACT

This paper investigates the impact of investment in automation- and AI-related goods on within-firm wage inequality in the French economy during the 2002–2017 period. We document that most wage inequality in France is accounted for by differences among workers belonging to the same firm rather than by differences between sectors, firms, and occupations. Using an event-study approach on a sample of firms importing automation- and AI-related goods, we find that spike events related to the adoption of automation- and AI-related capital goods are not followed by an increase in within-firm wage inequality or in gender wage inequality. Instead, wages increase by 1% three years after the events at different percentiles of the distribution. Our findings are not linked to the rent-sharing behavior of firms obtaining productivity gains from automation and AI adoption. Instead, if wage gains do not differ across workers along the wage distribution, worker heterogeneity will still be present. Indeed, in agreement with the framework in Abowd et al. (1999b), most of the overall wage increase is due to the hiring of new employees. This adds to previous findings presenting a picture of a ‘labor friendly’ effect of the latest wave of new technologies within adopting firms.

1. Introduction

Since the 1980s, France has experienced an increase in top incomes (both capital and labor incomes), in line with a general trend (Mishel and Bivens, 2021), and a high, though slightly decreasing, gender labor income gap (Garbinti et al., 2018). New evidence has uncovered the role of firms in driving income inequality, both due to expanding differences in wages *between* firms (i.e., wage premia related to size, trade, or productivity), as well as *within* firms and establishments (changes in relative wages between workers at different levels of the wage distribution or changes in worker composition, see Card et al., 2013 and Song et al., 2019).

In this respect, the current advent of new technologies belonging to the so-called ‘Fourth Industrial Revolution’, notably including robots and AI, is expected to produce a significant impact and potentially expand already existing inequalities or create new ones. *First*, on the one hand, such technologies could speed up the process of polarization in the labor market so that workers at the top and at the bottom of

the wage and skill distributions are expected to benefit more from the productivity increase disclosed by the new wave of innovations (see among others Autor et al., 2006; Autor and Dorn, 2013; Goos et al., 2014; Autor, 2015). As put forth in Freeman et al. (2020), such recent changes in the nature of work depended more on changes in work within occupations than on changes due to the shifting distribution of employment among occupations.² As such, the wage gap could also increase within firms and within occupations, depending on the ability of the employee to become familiar with the new technologies or, through a process of hiring, on widening the gap between ‘incumbent’ workers with a long tenure and recently hired employees. *Second*, although most societies are focusing increasing attention on the gender wage gap, such pay differences continue to be very relevant and are particularly large in the upper tail of the wage distribution (Blau and Kahn, 2017; Garbinti et al., 2018). However, the interplay between gender and technology could affect the gender wage gap, as we observe

* Correspondence to: Macroeconomics, International and Labour Economics department, Maastricht University, School of Business and Economics, P.O. Box 616, 6200 MD Maastricht, The Netherlands.

E-mail address: t.treibich@maastrichtuniversity.nl (T. Treibich).

¹ All authors contributed equally to the various stages of the work.

² For a similar concern, see Hunt and Nunn (2019) and van der Velde (2020).

<https://doi.org/10.1016/j.respol.2022.104533>

Received 1 September 2020; Received in revised form 5 January 2022; Accepted 31 March 2022

Available online 4 May 2022

0048-7333/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

a decrease in the share of women in routine tasks (Black and Spitz-Oener, 2010). As a consequence, there exist rising concerns about how new technologies are expected to affect the gender wage gap, even within the same firm, and to date, there exists very little evidence to support policy making.³

In this work, we address such questions by employing matched employer–employee data for France over the 2002–2017 period. We identify relevant investment episodes in AI and automation through purchases of selected categories of imported capital goods, adapting a taxonomy developed by Acemoglu and Restrepo (2022). As shown in Domini et al. (2021), acquisitions of such goods display the typical spiky nature that characterizes investment in capital goods (Nilsen et al., 2009; Grazzi et al., 2016).

We combine such data on automation- and AI-related investment spikes at the firm level with detailed information on firms' employees to investigate the effects of the adoption of AI and automation on wage inequality within and across firms. The descriptive evidence that we provide suggests that most wage inequality occurs within firms, occupations and sectors. Such a finding further corroborates a pattern already shown for Brazil and Sweden (Akerman et al., 2013; Helpman et al., 2017). This suggests that France is no exception and that a thorough analysis of the impact associated with the adoption of automation and AI on wage inequality must focus on the different *within* components.

Employing an event study methodology, we focus on the observed trend in wages and on some measures of wage inequality around a spike of investment in automation and AI. We find that employees at firms adopting these new technologies enjoy a small wage increase and that such a positive effect is detectable at most of the wage distribution percentiles. This effect is mostly driven by the fact that firms pay a higher wage to newly hired workers after an automation/AI spike,⁴ compared to incumbents. Overall, firm wage inequality is substantially unaffected. Focusing on the gender pay difference, we find that investments in automation and AI do not appear to be associated with a change in the gender wage gap. Within our methodological framework, we do check that our results are not driven by pre-spike trends in dependent variables. However, we cannot rule out that other contemporaneous shocks (for example, demand shocks) are endogenous to the decision to automate. For this reason, and following Bessen et al. (2020b), we will interpret the coefficients mostly as describing the evolution of firm outcomes around the spike.

Our work builds upon several streams of literature to which we aim to contribute new empirical evidence. *First*, we contribute to the discussion on wage inequality due to job polarization (see, among many, Autor et al., 2008) and, relatedly, on the effects of automation and AI technologies on labor market outcomes. Among the theoretical frameworks on which this literature builds is the model developed by Acemoglu and Restrepo (2019), which provides a rationale for both differences in wages between and within firms due to automation. In addition to the known displacement effect, according to which automation replaces human tasks, they describe productivity and a deepening effect, according to which automation makes labor and capital more productive and raises the demand for labor. The net impact on the overall wage level from these different forces becomes an empirical question. It may also depend on the specific types of technology, whereby AI and robots might have a more pronounced displacement effect than other automated machines, which require complementary labor to be operated (think, for example, of industrial robots in the car industry versus machines that streamline assembly but

require hand sorting of pieces). Instead, provided automation changes the relative demand for workers performing different tasks, both types of mechanisms exert a positive pressure on wage inequality, on the one hand by displacing some workers more than others, and on the other hand by making some workers more productive than others.

At the firm level, other explanations can also apply. If the productivity effect of automation is large, we can also expect to observe rent sharing, whereby the firms' higher profitability leads to a higher wage for all workers in the firm (Blanchflower et al., 1996). The wage profile in more productive firms can also be driven by a sorting mechanism, according to which they attract high-wage workers (Abowd et al., 1999b). In this framework (labeled AKM in the related literature), both firm characteristics (productivity, size) and individual characteristics (observable, such as seniority and education) explain wage differences across and within firms. Following such a sorting and matching approach, the authors also highlight competition among firms to hire the best employees, as well as the role of wage bargaining in explaining observed outcomes (Cahuc et al., 2006). Against this framework, changes in firm technology, productivity or size might modify the profile of the new hires and, through this channel, the wage distribution within firms.

In recent years, the empirical evaluation of the labor market effects of automation, particularly robots, has attracted much attention. Initially, much effort was exerted to predict the potential loss of employment associated with automation and AI technologies; see, among others, Brynjolfsson and McAfee (2014), and Frey and Osborne (2017). Thus far, the empirical evidence has been quite reassuring in suggesting a complementary, more than replacement, effect of automation. While aggregate-level studies have failed to find a consensus (the effect of automation on aggregate employment is negative according to Acemoglu and Restrepo 2020 and Acemoglu et al. 2020, neutral according to Graetz and Michaels 2018 and Dauth et al. 2018, and positive according to Klenert et al. 2020), firm-level evidence has been more consistent in showing a positive effect on employment in firms that adopt automation (Domini et al., 2021; Koch et al., 2019; Acemoglu et al., 2020; Bonfiglioli et al., 2020; Aghion et al., 2020).⁵ Some studies, together with employment, consider the impact of robot adoption (Koch et al., 2019; Humlum, 2020) or automation intensity (Dinlersoz et al., 2018a; Aghion et al., 2020) on the average firm wage. Humlum (2020) and Dinlersoz et al. (2018a) find a positive impact, while Aghion et al. (2020) and Koch et al. (2019) do not report a significant effect. Finally, Bessen et al. (2020a) focus on individual workers' outcomes in the Netherlands and show that after an automation cost spike, daily wages increase, although days of work decrease.

However, much less investigated is the potential impact of automation and AI on wage inequality within firms. Humlum (2020) uses an event study and a structural model (controlling for selection effects) to measure the impact of the adoption of industrial robots in Danish firms. He identifies that the overall positive effect on wages is driven by the impact on tech workers, while production workers observe a wage loss. In a study of Norwegian firms in the manufacturing sector, Barth et al. (2020) find that robots increase wages for high-skilled workers and managerial occupations, thus positively affecting wage inequality. As explained below, we focus on firm (instead of occupation)-level inequality as identified through the wage distribution; in addition, our measure includes but is not confined to robots; hence, it is much broader. Finally, using survey data from France, Fana and Giangregorio (2021) highlight the role of tasks and institutions in shaping the evolution of wage inequality.

³ In this regard, Pavlenkova et al. (2021) document a slight negative impact of automation on the gender pay gap in Estonian manufacturing firms.

⁴ Here, we use the “automation/AI” expression for conciseness, but a more complete label would be “automation- and/or AI-related” or “embedding automation and/or AI technologies”. We will use these different expressions in an interchangeable way in the text.

⁵ Note that there are some potential caveats to this conclusion. It could indeed be that the effects of automation technologies are not yet fully visible in the data or that a mild increase in employment registered at adopting firms is more than compensated by a decrease in employment in non-adopting competing firms via a spillover effect, as shown by Acemoglu et al. (2020).

Second, while there already exists extensive evidence reporting the ubiquitous presence of a gender wage gap (among the recent reviews we refer to [Blau and Kahn, 2017](#)), much less is known about the impact of the newest technologies on such a wage gap and on the job flows as broken down by gender. Among the existing works, [Brussevich et al. \(2019\)](#) investigate differential gender exposure to automation by referring to the routine task intensity of the occupation. On this basis, since women tend to be more represented in such tasks, they face a higher risk of displacement than men. This is also the conclusion reached by [Sorgner et al. \(2017\)](#), who take a broader perspective into consideration by noting several dimensions of the gender equality issue. Focusing more specifically on the gender pay gap, [Aksoy et al. \(2020\)](#) employ country–industry level data and report that a 10% increase in robot-related investments (data being sourced from the International Federation of Robotics) is associated with a 1.8% increment in the gender wage gap. As a common limitation of many contributions in this stream of literature, the authors cannot directly observe the effect on employment and wage associated with an investment within the firm, as data are available at the country, industry and demographic levels. Still, at the aggregate level, employing data from US commuting zones, [Ge and Zhou \(2020\)](#) report contrasting evidence on the change observed in the gender wage gap following investments in robots versus computers. While the former decreases the wage of male workers more than that of female workers, thus reducing the gap, the latter increases the gap. In our work, the data and the empirical setting enable us to investigate what happens to the gender pay gap both across adopting and non-adopting firms and, more specifically, within adopting firms.

The paper is organized as follows. Section 2 first presents the data sources and the variables that are used in the paper and then illustrates the construction of the different samples used in the analysis. In Section 3, we provide descriptive statistics on wage distribution, including an analysis of variance that decomposes overall wage inequality into different components. We also show trends in wage inequality and introduce our measure of investment in automation- and AI-related goods. Section 4 presents the event study framework and discusses the results. Section 5 concludes.

2. Data and variables

2.1. Sources

Our dataset contains data from all French firms with employees over the 2002–2017 period, obtained by merging different administrative sources, using the unique identification number of French firms (SIREN). The first source is the *Déclaration Annuelle des Données Sociales* (DADS), a confidential database provided by the French national statistical office (INSEE) and based on the mandatory forms that all establishments with employees must submit to the social security authorities. To be more precise, we use the DADS *Postes* dataset, in which the unit of observation is the ‘job’ (*poste*), defined as a worker–establishment pair.⁶ We extract from DADS the following worker-level variables: gross yearly remuneration, number of hours worked, age, gender, and occupation,⁷ as well as the sector of the firm defined

⁶ Note that DADS *Postes* does not allow the tracking of workers over time, since the worker identification number is not constant across years.

⁷ The occupation variable is the *Catégorie Socio-professionnelle*, which reflects the hierarchical structure within firms and the levels of management or ‘production hierarchies’ (see also [Caliendo et al., 2015](#); [Guillou and Treibich, 2019](#)). We also retrieve worker-level variables on the ‘type of job’ from DADS, which allows us to identify apprentices and clean them out, and on the start and end dates of job posts, necessary to identify workers present on a specific date (see Section 2.3).

according to NAF rev. 2 classification (corresponding to the European NACE rev. 2).⁸

The second source is the transaction-level international trade dataset by the French customs office (*Direction Générale des Douanes et des Droits Indirects*, DGDDI), containing detailed information on import and export flows, among which are found trade value, country of origin/destination, and an 8-digit product code, expressed in terms of the European Union’s Combined Nomenclature, an extension of the international Harmonized System (HS) trade classification. From this source, we retrieve firm-level information on the value of yearly imports that are related to automation and AI (see below in this section), as well as on the total value of yearly imports per product category.

In addition to our two main sources, we also use FICUS and FARE, two private datasets provided by INSEE, which are based on the fiscal statements that all French firms must make to the tax authorities and which contain detailed balance sheet and revenue account data. FARE has succeeded FICUS since 2008 and collects data from a larger set of tax regimes than FICUS. We use this source to extract firm-level information on value added, which is then used to construct our labor productivity measure, as valued added over the number of hours worked.⁹

2.2. Variables

Wage-related variables

The outcome variables of our analysis are firm-level wage measures based on worker-level variables extracted from DADS.¹⁰ For each worker, we divide the gross yearly remuneration by the number of worked hours to obtain hourly wage.¹¹ This information is then combined at the firm level as well as at the level of specific categories of workers within the firm. First, we construct each firm’s wage distribution moments, in particular the mean and standard deviation, as well as percentiles (p10, p50, p90). In the regressions, we use the log transformation of the level variables (mean wage and wage percentiles) to obtain comparative measures of the effect of automation at different locations of the wage distribution. As measures of within-firm wage inequality, we consider the standard deviation and the p90/p10 ratio. The p90/p10 ratio is a standard measure of wage inequality used in both macro- and microeconomic literature (see [Cirillo et al., 2017](#); [Mueller et al., 2017](#)); the standard deviation is also chosen because it reflects an overall measure of the dispersion of wages within a firm.

Furthermore, wage information can also be constructed for specific categories of workers within a firm (hence, measures of wage inequality between categories can be constructed). In particular, we are interested in comparing the wages of females *vis-à-vis* males. We calculate a firm’s *gender ratio* (corresponding to the gender pay gap) as the mean hourly wage of female workers divided by the mean hourly wage of male workers. Likewise, we calculate gender ratios at various percentiles, that is, the ratio between a certain percentile of the female hourly wage distribution and the same percentile of the male distribution.

An important note must be made here on our definition of gender wage inequality. Since we normalize the wage by the number of hours

⁸ In fact, the sector code (*Activité Principale Exercée*, APE) is expressed in DADS in terms of the NAF rev. 1 classification until 2007. To ensure consistency over the observed time span, we establish a mapping between 4-digit NAF rev. 1 and NAF rev. 2 codes, as explained in [Domini et al. \(2021, fn. 7\)](#). Furthermore, as a firm’s APE may vary across years, we assign each firm a permanent 2-digit sector based on the most frequent APE occurrence.

⁹ Information from FICUS/FARE is not available for 4.42% of the firms in Sample 2 (see below for a definition of the sample).

¹⁰ Note that while plant-level information is available in DADS, we need to focus on the firm level to match DADS data with firm-level customs data.

¹¹ We deflate wages (as well as imports; see below) using yearly value-added deflators for 2-digit NAF divisions provided by the INSEE.

worked and only consider employed persons, two important sources of inequality in earnings between men and women are removed. In France, females are most affected by part-time work, yielding lower monthly wages: based on the ILOSTAT data, approximately 50% of female work during our period of study is part-time, while only 30% of male work is, ILO (2020). As a consequence, if the gender wage per hour gap in France is estimated at 15.5%, right at the EU-27 average, the overall gender earnings gap is exactly double, at 31% (EUROSTAT, 2015).

Adoption of automation- and AI-related technologies

To date, there is a lack of systematic firm-level information on the adoption of digital and automation technologies at the firm level, which has only recently started to be collected by national statistical offices. Exceptions are the Netherlands, where (Bessen et al., 2020a) use information on automation costs included in the national survey from the Dutch statistical office (CBS), and the U.S., where (Dinlersoz et al., 2018b) obtain a proxy of automation intensity via a technology index from a U.S. Census Bureau survey. Nevertheless, trade flows reported by firms to customs offices offer a useful solution to this, as fine product-level decomposition allows identifying the adoption of specific technologies via the import of related goods. We construct a measure of firm-level adoption of technologies related to automation and AI based on product-firm-level customs data. This approach has been employed by several recent studies on the effect of robotization (Dixon et al., 2019; Bonfiglioli et al., 2020; Acemoglu et al., 2020; Aghion et al., 2020) and automation in general at the firm level (Domini et al., 2021). Note, with regard to the French context, that Aghion et al. (2020) instead choose two broader measures (industrial equipment and machines and change in electric motive power) that can be applied to all manufacturing firms, including domestic firms.

More specifically, we identify imports of goods that embed automation- and AI-related technologies based on their 6-digit Harmonized System (HS) product code. Automation-related imports are identified by using a taxonomy presented by Acemoglu and Restrepo (2022), partitioning all HS codes referring to capital goods (divisions 82, 84, 85, 87, and 90) into several categories of automated and non-automated goods. Imports embedding automation technologies include, among others, industrial robots, dedicated machinery, numerically controlled machines, and a number of other automated capital goods.¹² To the automation-related categories listed by Acemoglu and Restrepo (2022), we add 3-D printers, the HS code of which is identified by Abeliensky et al. (2020). In addition to these automation-related categories, we identify some other categories of imports that are expected to be related to AI, namely, automatic data processing machines and electronic calculating machines.¹³

Considering AI-related imports, in addition to automation-related imports, is important for ensuring our measure is representative of the adoption of new technologies in the whole economy. Indeed, the former tend to be less concentrated than the latter in the manufacturing sector: one-fifth of all AI-related imports are accounted for by manufacturing firms *vis-à-vis* one-half of automation-related imports.¹⁴

Some potential limitations of our import-based measure of adoption of automation- and AI-related technologies must be acknowledged and discussed. First, firms might purchase automation- and AI-related goods

domestically instead of internationally; thus, they may be wrongly labeled as non-adopters in our analysis. With respect to this, notice that France has a comparative disadvantage (cf. Balassa 1965) and a negative trade balance for the goods that compose our measure;¹⁵ hence, imports are likely to be the most important source of automation- and AI-related goods for French firms. Second, the import-based nature of our measure restricts the scope of our analysis to firms involved in international trade: this restriction decreases the probability that we wrongly label firms in our sample as non-adopters; however, we do not consider firms that are only active in the domestic market and that may buy automation- and AI-related technologies from domestic suppliers (though unlikely, as argued above). Moreover, the impact on the wage dynamics of these firms may be different, as they tend to be smaller and less productive on average than firms involved in international trade. Third, there exists the possibility that firms resort to an intermediary rather than import goods themselves (Ahn et al., 2011; Bernard et al., 2010; Blum et al., 2010); however, this is less likely for more complex goods (Bernard et al., 2015) that are highly relation-specific, such as the ones that compose our measure. Finally, firms that import automation- and AI-related goods may re-sell them, either in the domestic market or abroad. We will address the possibility of resellers with two different robustness checks in Section 4.4, showing that the main results are largely unchanged.

2.3. Data cleaning and sample construction

To construct the dataset employed in our analysis, we perform some cleaning at the worker level; then, we create firm-level variables by aggregating information on workers present in each firm on a specific date of each year (December 31st).¹⁶ We want to make sure that we only include workers that are really attached to a particular firm. In the DADS data, these correspond to workers related to jobs labeled as 'principal' (*non-annexes*) by INSEE, which exceed some duration, working-time, and/or salary thresholds.¹⁷ These can be seen as the 'true' jobs that contribute to the production process (see e.g. INSEE 2010, p. 17), and account for the large majority (three-fourths) of total jobs.¹⁸ We also remove apprentice workers, which represent approximately 3.5% of observations, who are workers with less than

¹² Based on calculations by the authors using COMTRADE data (results are available upon request). This is true on aggregate, as well as for most of the subcomponents of the measures shown in Table A.1 in Appendix A. A notable exception is the category of robots, as well as that of regulating instruments, which, however, represent a minority of the measure.

¹³ Referring to a consistent date across years ensures consistency in the computation of our variables of interest, as a firm's employment varies over the year due to new hires and separations, which may be partly driven by short-term and/or seasonal dynamics. This causes variables related to the within-firm distribution of wages to also change. Furthermore, referring to a specific date is necessary to consistently identify the flows of newly hired and separated workers (and the variables on their wage distribution), as it allows ignoring short-term jobs and temporary fluctuations in employment. Note that this approach is followed in other papers constructing gross worker flows (Domini et al., 2021; Abowd et al., 1999a; Bassanini and Garnero, 2013; Davis et al., 2006; Golan et al., 2007).

¹⁴ See the definition in Section 3.2.1 (pp. 17–18) of the *DADS 2010 Guide méthodologique*. To be classified as *non-annexe*, a job should last more than 30 days and involve more than 120 worked hours, with more than 1.5 h worked per day; or the net salary should be more than three times the monthly minimum salary; otherwise, it is classified as *annexe*.

¹⁵ Non-principal (*annex*) jobs represent 22% of all observations in DADS, and 43% of new hires; 50% of them are full-time (vs 72% of principal jobs), 12% part-time, and 24% small part-time (*faible temps partiel*); 43% have a permanent contract (*Contrat à Durée Indéterminée*); vs 61% of principal jobs, 29% have a fixed-term contract (*Contrat à Durée Déterminée*) vs 24% with a temporary or placement contract (*mission*). After one year, 18% of them become principal, 26% stay annexes, and the rest (56%) leave the firm.

¹² For a full list, including the specific 6-digit HS codes falling under each of the above-mentioned categories, see Table A.1.

¹³ As an additional check that these are in fact relevant categories for our analysis, we use the USPC-to-HS 'Algorithmic Links with Probabilities' (ALP) concordance by Lybbert and Zolas (2014) to see whether their codes match to the US patent classification (USPC) code 706 ('Data processing - Artificial Intelligence'). This is the case for four out of seven of these additional categories.

¹⁴ Based on our calculations using DGDDI data for the year 2017. Detailed figures are available upon request.

Table 1
Samples' composition and relative size, 2002–2017.
Source: Our elaborations on DADS and DGDDI data.

	Nb. obs	Nb. firms	Share in nb. of firms	Share in employment
All firms	20,010,009	3,131,425	1	1
Sample 1	2,703,157	287,901	9.19	54.50
Sample 2	1,111,741	91,593	2.92	51.66
Sample 3	501,667	39,295	1.25	37.24

Notes: Sample 1: All importing firms; Sample 2: Importing firms with at least 10 employees. Sample 3: Firms importing automation- and AI-related goods at least once, with at least 10 employees.

120 h worked in the year,¹⁹ and workers with wages below half of the minimum wage, which represent less than 1% of observations.²⁰ Fig. A.1 in the Appendix shows that this bottom threshold to the wage per hour variable truly eliminates outliers, as the minimum wage in France has a very strong impact on the shape of the wage distribution. Overall, and analogously to what has been done in the related literature (see, for example, Song et al., 2019), these choices exclude workers who are not strongly attached to the firm and/or the labor market.

We consider workers employed in the entire economy, except for the primary sector (NAF/NACE rev. 2 divisions 01 to 09). We also remove firms labeled as 'household employers' (*particuliers employeurs*) and those engaged in public administration (*fonction publique*) between 2009–2017, since they are not available in earlier years. These criteria yield a sample of more than 20 million firm-year observations over the period 2002–2017, or 3 million unique firms (see Table 1, row 1, 'All firms').

However, in our analysis, we need to restrict the sample for the following reasons. First, we can construct our measure of adoption of automation- and AI-related technologies only for importing firms (see Section 2.2), which we label as Sample 1 in Table 1. This sample represents 9% of observations in the overall data, but accounts for more than half of total employment. Second, to ensure that within-firm statistics on the wage distribution are meaningful, we restrict our attention to importing firms with at least 10 employees (Sample 2). This threshold excludes 'micro-firms', according to the Eurostat definition. Note that this further restriction reduces to a large degree the number of firms included in the analysis (which represent 3% of all firms present in the DADS dataset), but it only marginally reduces aggregate employment representativeness (cf. Table 1, row 3). Finally, as the event study carried out in Section 4 will compare the impact of automation- and AI-related investment exploiting the timing of the latter, we will focus on those firms in Sample 2 that import automation- and AI-related goods at least once over 2002–2017 (Sample 3).²¹ This final sample includes only approximately 40 thousand firms; nevertheless, this still

¹⁹ This matches one of the thresholds used for defining non-annexe workers. Note that this also removes workers with zero hours.

²⁰ The existence of a wage below the legal minimum (SMIC in French) is not per se surprising as documented in the literature, see among the others (Gautier, 2017; Delahaie and Vincent, 2021). For instance, in 2016 the share of sectors fulfilling the minimum wage was only 86%, and within specific sectors the share can get as low as 38% (Dares, 2018). According to the literature this is mostly due to delays in keeping up with increases in the minimum wage, and as a result it should be of a limited amount. Even if the "jump" at the level of the minimum wage is very clear (see Fig. A.1 in the Appendix), the number of observations below the minimum is non-negligible and it may be due to misreporting of the number of hours for example. Anyway, since we focus on the 10th percentile of the distribution for one of the two measures of inequality, we believe this issue does not impact our results.

²¹ A potential issue related to the sample construction is due to a change in the reporting threshold over the period of observation. In particular, since 2011, product codes for imports from EU countries have been reported only for firms with more than 460,000 euros of imports in a given year; see

accounts for 7.5 million workers.²² In the following section, which presents descriptive statistics, we will refer to different samples, while in the regression analysis (Section 4), we will only keep firms with a spike (Sample 3).

3. Descriptive statistics

3.1. From the wage distribution of workers to the wage distribution within firms

In what follows, we present some descriptive statistics to direct our approach. We start from an aggregate view and decompose wage inequality among all workers into its *between* (differences across firms, related to sector or structural change dynamics) and *within* components (changes within firms, which is the focus of our empirical analysis). Then, we discuss the characteristics of our measure of adoption of automation- and AI-related technologies. Finally, we dig deeper into the study of firm-level wage distributions and inequality and provide some *prima facie* evidence for the differences between adopting and non-adopting firms.

The wage distribution of workers

Fig. 1 shows the distribution of the deflated wage per hour variable across workers in the entire economy for one year (sample 'All firms'), i.e., around 16 million workers. The wage distribution in France is heavily impacted by the minimum wage of approximately 10 euros per hour and is therefore very positively skewed with high kurtosis. Note that wage inequality among all workers can be driven by differences *across firms* (reflecting their relative productivity, profitability, or aggregate sector and institutional dynamics) or *within firms* (reflecting changes in the labor organization of the firm and remuneration of value across workers). To direct our study of within-firm wage inequality, we perform a decomposition exercise that compares the contribution of both dimensions to the overall wage inequality among workers, as shown in Fig. 1.

Decomposing wage inequality

In this section, we decompose the overall wage inequality among all workers into differences between and within components. More specifically, we leverage worker-level information on hourly wages and their occupation (managers and white-collars; supervisors and technicians; clerks; skilled production workers; unskilled production

also Acemoglu et al. (2020), Bergounhon et al. (2018). We cannot directly measure the bias generated by such a change, but the indirect evidence that we collected is very reassuring. First, as reported below in Table 5, within the subsample of importing firms larger than 10 employees (Sample 2), importers of automation technologies (Sample 3) are much larger and hence are less likely to be affected by the changing threshold. Second, the number of adopters (Sample 3) shows only a very marginal decrease in 2011 (from 72,049 in 2011 to 69,849 in 2012). Finally, within our sample of importing firms, we find that there is no discontinuity in 2011 in the share of firms that import automation- and AI-related imports per our measure and of the related spikes (see Table A.2).

²² It is worth noting that the sample of our analysis is larger than that of other studies on robotization and automation using French data. Acemoglu et al. (2020) use a sample of 55,390 manufacturing firms between 2010 and 2015, of which 598 are robot adopters. Bonfiglioli et al. (2020) use a sample of 103,771 manufacturing firms between 1994 and 2013, of which approximately 800 are robot adopters. Aghion et al. (2020) use a sample of 16,227 manufacturing firms between 1994 and 2015. These figures are compared to the 91,593 manufacturing and service-sector firms (Sample 2) that we observe over 2002–2017, of which 39,295 are importers of automation- and AI-related goods (Sample 3). Such differences are due to including different sectors (firms in manufacturing account for approximately 42% in 2017) and employing a measure of automation that is broader than the implementation of robots alone.

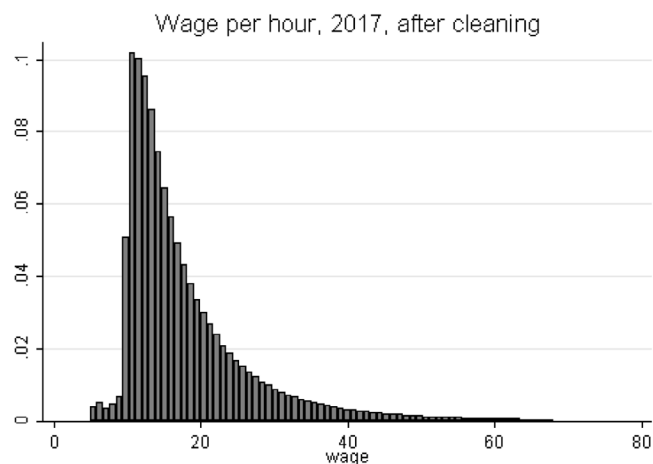


Fig. 1. Distribution of wages per hour among all workers, 2017.
Source: Our elaboration on DADS data.

workers; residual workers),²³ the firm where the worker is employed, and the sector of that firm (defined at the 2-digit level of the NAF classification) to estimate a set of equations as follows:

$$w_i = \delta_j + \varepsilon_i \quad (1)$$

where w is the logarithm of hourly wage, i indexes workers, and δ_j is a set of fixed effects, which represent, depending on the specification, sectors, occupations, sector-occupations, or firms. Using the estimates from Eq. (1), we decompose the overall wage inequality among all workers into a *between* and a *within* component using the following equality:

$$\text{var}(w_i) = \text{var}(\hat{\delta}_j) + \text{var}(\hat{\varepsilon}_i) \quad (2)$$

where $\text{var}(w_i)$ is total variance (T), $\text{var}(\hat{\delta}_j)$ is the between component (B), and $\text{var}(\hat{\varepsilon}_i)$ is the within component (W). Note that the residual term is orthogonal to the other term by construction. In the following tables, we report the share of total variance accounted for by the within component, i.e., W_i/T_i .

Table 2 shows the share of the overall wage inequality accounted for by the within component at different levels of disaggregation (i.e., using different sets of fixed effects), namely, within sectors, within occupations and within sector occupations. Note that the between component (wage inequality due to differences across sectors, occupations, and sector-occupations groups), though not shown, is the mirror image of the values reported in the table. The within sector and within occupation components account for the majority of wage inequality in France in 2017 in all samples, whereas the within-sector-occupation is slightly below 50%. For example, looking at the values for all firms (first row), only 22% of overall wage inequality can be explained by differences in wages between different sectors (e.g., wages in textile manufacturing vs. wages in retail trade) – the remaining 78% being accounted for by differences among workers belonging to the same sector. Furthermore, approximately half of wage inequality occurs among workers belonging to the same occupational category (even within the same sector).

This is consistent among the different samples we defined in the previous section; hence, in the sample that will be used in our regression analysis (Sample 3), the main forces driving wage inequality are the same as those across the entire population of firms. This result confirms

²³ The first three categories are defined at the 1-digit level of the French taxonomy of occupations (*Catégories Socio-professionnelles*) using codes beginning with 3, 4, and 5, respectively, while skilled and unskilled production workers are defined at the 2-digit level using codes beginning with 61–65 and 66–68, respectively.

Table 2

Within-sector, within-occupation, and within-sector-occupation shares of wage inequality, 2017.

Source: Our elaborations on DADS and DGDDI data.

	(%) Within sector	(%) Within occupation	(%) Within sector-occupation
All firms	78	55	46
Sample 1	80	53	46
Sample 2	80	52	45
Sample 3	80	52	45

Notes: Sample 1: All importing firms; Sample 2: Importing firms with at least 10 employees; Sample 3: Firms importing automation- and AI-related goods at least once with at least 10 employees.

Table 3

Within-firm share of wage inequality, 2017.

Source: Our elaborations on DADS and DGDDI data.

	(%) Within firm (sector level)	(%) Within firm (sector-occupation level)
All firms	67	58
Sample 1	75	68
Sample 2	76	70
Sample 3	76	70

Notes: Sample 1: All importing firms; Sample 2: Importing firms with at least 10 employees; Sample 3: Firms importing automation- and AI-related goods at least once with at least 10 employees. The within components are first computed for each sector/sector-occupation separately and then aggregated by taking an employment weighted average.

that within-sector determinants are key to understanding the sources of wage inequality and is in agreement with evidence from other countries (see, for example, Helpman et al., 2017 for Brazil). Finally, it shows that a great amount of wage variance happens not just within sectors but also within occupations. This informs our approach of using measures of inequality based on the whole firm's wage distribution (90/10 ratio and standard deviation) instead of measures based on occupational means (wage of managers vs. wage of production workers).

In Table 3, we report results from a second decomposition exercise in which the within component refers to the share of wage inequality that, within each sector (column 1) and within each sector-occupation (column 2), is accounted for by the within-firm component vs. the between-firm component. In this case, we first solve Eq. (1) for each sector and sector-occupation, where δ_j is a set of firm-level fixed effects, and then leverage equality (2) to compute the within component for each sector and sector-occupation. In Table 3, we report the employment weighted average of these components across the different sectors (and sector-occupations).

Among the population of workers within each sector, on average, 68% of wage inequality is explained by the within-firm component. This means that the wage of a worker in a particular sector is not primarily defined by different characteristics among firms (e.g., firm size). In the other samples that consider importing firms (samples 1–3), this percentage is even greater, approximately 75%: the reason is that within-firm dispersion of wages is larger in large firms. Within-firm dispersion may be driven by the different occupational structures of firms. To account for this, in column 2, we perform the same decomposition for each sector occupation. The within-firm share slightly decreases, but it is still dominant with respect to the between component: in Sample 3, it is as high as 70%.

Overall, this analysis is a further motivation for our focus on within-firm wage inequality. Indeed, the position of the worker within the firm has more impact on his/her wage than the characteristics of the firm, the sector, or the occupation in which he/she is employed.²⁴

²⁴ The within-firm component has also been found to be a sizable factor underlying wage inequality in other studies. See, for instance, Helpman et al. (2017) for Brazil and Song et al. (2019) for the U.S.

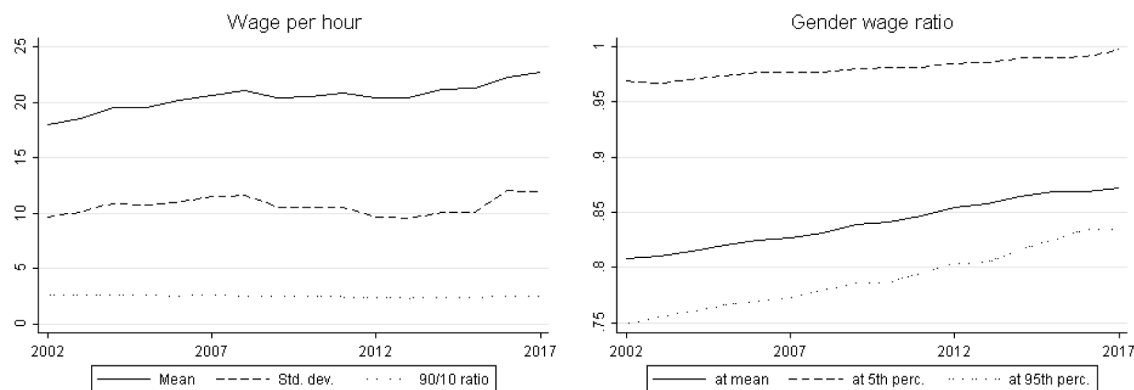


Fig. 2. Evolution of wage characteristics over time, Sample 3, 2002–2017. Note: Sample 3: Firms importing automation- and AI-related goods at least once, with at least 10 employees.

Source: Our elaborations on DADS and DGDDI data.

Trends in wage inequality within firms

Having determined that the within-firm dimension is crucial for understanding overall wage inequality, we now can analyze trends in firm-level wage inequality. Fig. 2 shows the evolution of our most important dependent variables over our period of analysis, namely, the (deflated) wage per hour and the gender wage ratio, for the firms belonging to Sample 3. If the mean wage level increases from 18 euros per hour in 2002 to 23 euros in 2017, the average firm-level wage inequality measures (standard deviation and 90/10 ratio) do not show any trend, except for a bump in the standard deviation variable in the last two years of study. Regarding the gender wage ratio (female/male) at different locations of the wage distribution, we see that in our data, at the bottom of the distribution, it is extremely stable at approximately 1 (no gender wage inequality, which is a positive consequence of the minimum wage). Instead, it starts at 80% at the mean and below 70% at the 95th percentile in 2002 and increases to almost 90% and 80%, respectively, over the period of study. It should be noticed that these impressive dynamics regarding Sample 3 do not reflect the national trend in the mean gender wage per hour gap, which shows no evolution since 2002 (also see EUROSTAT, 2015).²⁵

3.2. Automation and AI imports

We provide here some information to characterize our measure of firm-level adoption of automation- and AI-related technologies, namely, the sectors where it is prevalent and its lumpy statistical properties.

Sectoral distribution of automation investments

We report in Table 4 the list of 2-digit sectors (NAF rev. 2, A88 classification) that are most active in buying automation- and AI-intensive goods in our trade data. This is measured by comparing the share for which a sector accounts with regard to total French automation and AI imports (central column) and the same sector's share in aggregate employment (last column). The electronics (NAF rev. 2 division 26), machinery (28), and automotive sectors (29) are disproportionately represented in automation- and AI-related imports compared to their employment share. The retail sector (46) is a noteworthy case with

²⁵ Part of the explanation has to do with the subset of workers in our sample: in our sample of importers, and even more so in our sample of importers of automation and AI products, wages are higher than in the rest of the economy. At low levels of the wage distribution in our sample (which better mirror the overall wage level in French firms), the gender wage gap is quite stable over time, thus acting similarly to the aggregate dynamics (cf. https://ec.europa.eu/eurostat/databrowser/view/sdg_05_20/default/table?lang=en).

Table 4

Sectors with an automation and AI share larger than their employment share, Sample 3, 2017.

Source: Our elaborations on DADS and DGDDI data.

Sector	A88	Automation and AI share (%)	Employment share (%)
Electronics	26	3.7	2.1
Machinery	28	3.4	2.6
Automotive	29	3.7	3.0
Retail	46	59.5	9.7
IT	62	5.4	3.2

55.1% of those investments, more than six times its share in total employment (9.3%).²⁶

The statistical properties of automation

When looking at the statistical properties of automation- and AI-related imports, it can be observed, as already seen with automation alone in Domini et al. (2021) and Bessen et al. (2020a), that they display the typical *spiky* behavior of an investment variable (Asphjell et al., 2014; Letterie et al., 2004; Grazzi et al., 2016). This means that, first, such imports are rare across firms: approximately 14% of importers import automation- and AI-related goods per year, and fewer than half of them do so at least once over the 2002–2017 period. Second, such imports are rare within firms: among firms that import such capital goods at least once, close to 30% do it only once, and the frequency decreases smoothly with higher values, except for a small group of firms that import AI or automated goods in all years. Finally, the largest yearly event of imports of such goods represents a significantly high share of a firm's total across years: when ranking the shares of each year's imports (out of all years) from largest to smallest, it is apparent that the top-ranked import event displays a predominant share (approximately 70%), while the shares of lower ranks rapidly decrease in value.²⁷ As discussed in Domini et al. (2021), there are two possible explanations for automation adoption being lumpy, and this also applies to AI-related goods. First, the products we select are a subset of capital goods that are automated in nature. As such, they should share similar characteristics as the larger category of physical investment goods (Nilsen and Schiantarelli, 2003). Second, Bessen et al. (2020a) point out that even other dimensions of the adoption

²⁶ Although the relevance of automation technologies in service sectors is largely acknowledged (see among others Sostero, 2020), to account for such an important outlier, we also run the regressions separating manufacturing and services. They are not included in the results due to space constraints but are available from the authors upon request.

²⁷ These statements are based on Fig. A.2 in Appendix A.

of automation technologies, for example, automation costs, share the same characteristics that make investment lumpy: they are irreversible, as they bring about idiosyncratic changes in the production process, and indivisible, as they cannot be carried out in small chunks over time. Because of the very skewed nature of this variable within firms, we define the largest event for each firm as an *automation/AI spike*. As a robustness exercise, we also compare results with alternative definitions of these spikes by first separating automation and AI products,²⁸ and then adopting the spike definition provided in Bessen et al. (2020a,b) with a condition on the value of the imports (see Section 4.4).²⁹

3.3. Firm-level wage inequality and automation

What are the characteristics of the firms that invest in automation and AI-related goods? In Table 5, we compare, within our sample of importing firms above 10 employees (Sample 2), the group of firms that never automate (column ‘No spike’) to that of those who import such goods at least once, and for which we can construct the automation/AI spike variable (column ‘Spike’, corresponding to Sample 3). We also report in the last column the significance level of the mean-difference test comparing those two groups.

In line with previous descriptions from the literature (Koch et al., 2019; Deng et al., 2021; Domini et al., 2021), firms adopting automation and AI are larger, more productive, and pay higher wages than non-adopting firms. Such a difference in the wage level is present at all levels of the wage distribution and is more pronounced at higher levels. We also show that adopting firms have higher within-firm wage inequality according to the two measures used in our exercise (standard deviation of the within-firm wage per hour distribution and 90/10 percentile ratio). Finally, they are more unequal in terms of gender pay, showing a lower female-to-male wage per hour ratio at all locations of the wage distribution.

The static differences highlighted in Table 5 could be due to the impact of automation and AI on wage and employment characteristics, but they might also reflect self-selection into automation and AI adoption. Such a selection effect will be addressed in our empirical strategy by considering only firms that automate (i.e., Sample 3) in our event-study analysis.

The next step is to consider a dynamic approach by evaluating how firm-level wage characteristics evolve around an automation/AI spike. We start with a descriptive exercise in a balanced panel of firms that have an automation/AI spike at time $t = 0$ and that we also observe in the three years before and three years after. Within this subgroup of 17,266 firms, and not controlling for other sectoral, time or firm-level effects (which will be done in the regression analysis), the picture that emerges is that of an increase in wages at all the levels tested here, while the correlation between the spike event and wage inequality is ambiguous (the two measures of wage inequality yield opposite trends). Finally, the gender pay gap seems to slightly decrease, especially at the 90th percentile of the wage distribution (see Table 6).

4. The effect of automation and AI on wages: an event study analysis

4.1. Empirical approach

Automation/AI spikes represent single, major events that we observe for French importing firms during the 2002–2017 period (see Section 3.2). This characteristic makes it suitable to investigate the relationship between automation and wages within an event-study

Table 5

Comparing firms with and without an automation/AI spike, Sample 2, all years (2002–2017).

Source: Our elaborations on DADS and DGDDI data.

	No spike	Spike	T-test
Number of observations	622,506	509,547	
Number of firms	52,298	39,295	
Number of employees	55.63	175.76	***
Value added per hour	75.30	344.95	*
Wage per hour (mean)	18.13	20.42	***
Wage per hour (p10)	11.77	12.59	***
Wage per hour (p50)	15.63	17.45	***
Wage per hour (p90)	28.02	31.96	***
Wage standard deviation	8.62	10.66	***
90-10 wage ratio	2.37	2.52	***
Female-to-male wage ratio (aggregate)	0.88	0.84	***
Female-to-male wage ratio (p10)	1.01	0.98	***
Female-to-male wage ratio (p50)	0.95	0.91	***
Female-to-male wage ratio (p90)	0.83	0.79	***

Notes: ***, significant difference at 1% level; Sample 2: Importing firms above 10 employees.

Table 6

Wage characteristics around an automation/AI spike, balanced panel within Sample 3.

Source: Our elaborations on DADS and DGDDI data.

Years since spike	Wage per hour	Wage standard deviation	90/10 wage ratio
-3	19.573	10.585	2.571
-2	19.679	10.565	2.549
-1	19.885	10.601	2.522
0	20.133	10.646	2.493
1	20.326	10.579	2.477
2	20.598	10.720	2.477
3	21.029	10.838	2.467

Years since spike	Wage per hour (p10)	Wage per hour (p50)	Wage per hour (p90)
-3	11.891	16.524	31.018
-2	11.997	16.624	31.091
-1	12.207	16.813	31.311
0	12.411	17.066	31.452
1	12.605	17.328	31.626
2	12.761	17.576	32.060
3	13.088	18.051	32.466

Years since spike	Gender wage ratio (p10)	Gender wage ratio (p50)	Gender wage ratio (p90)
-3	0.982	0.903	0.773
-2	0.981	0.906	0.782
-1	0.985	0.910	0.787
0	0.984	0.915	0.797
1	0.987	0.918	0.805
2	0.987	0.920	0.809
3	0.988	0.922	0.814

Notes: The sample includes firms belonging to Sample 3 observed for at least three years before and three years after an automation/AI spike, representing a balanced sample of 17,266 firms; Sample 3: Firms importing automation- and AI-related goods at least once, with at least 10 employees.

framework. Such a methodology was used by Bessen et al. (2020b) to study the effect of automation on firm-level outcomes as well as in other contexts to explore differences around a main firm-level event (Balasubramanian and Sivadasan, 2011; Miller, 2017; see also Duggan et al., 2016; Lafortune et al., 2018 for other, non-firm-level, applications).

Given an index t that indicates the difference between the current year and the year in which the automation/AI spike happens for firm i , our main event study specification reads as follows:

$$y_{ijt} = \sum_{k \neq -1; k_{min}}^{k_{max}} \beta_k D_{kit} + \delta_i + \zeta_{jt} + \varepsilon_{it} \quad (3)$$

²⁸ See the distinction in Table A.1.

²⁹ For a more detailed discussion of the statistical properties of automation-related imports, including a comparison to general physical investment, see Domini et al. (2021, Section 3).

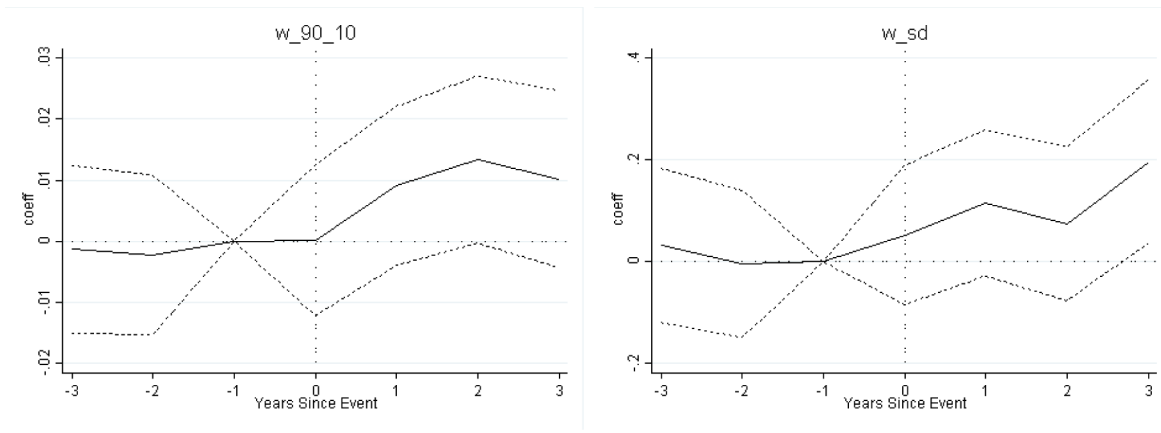


Fig. 3. Automation/AI spikes and within-firm wage inequality. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

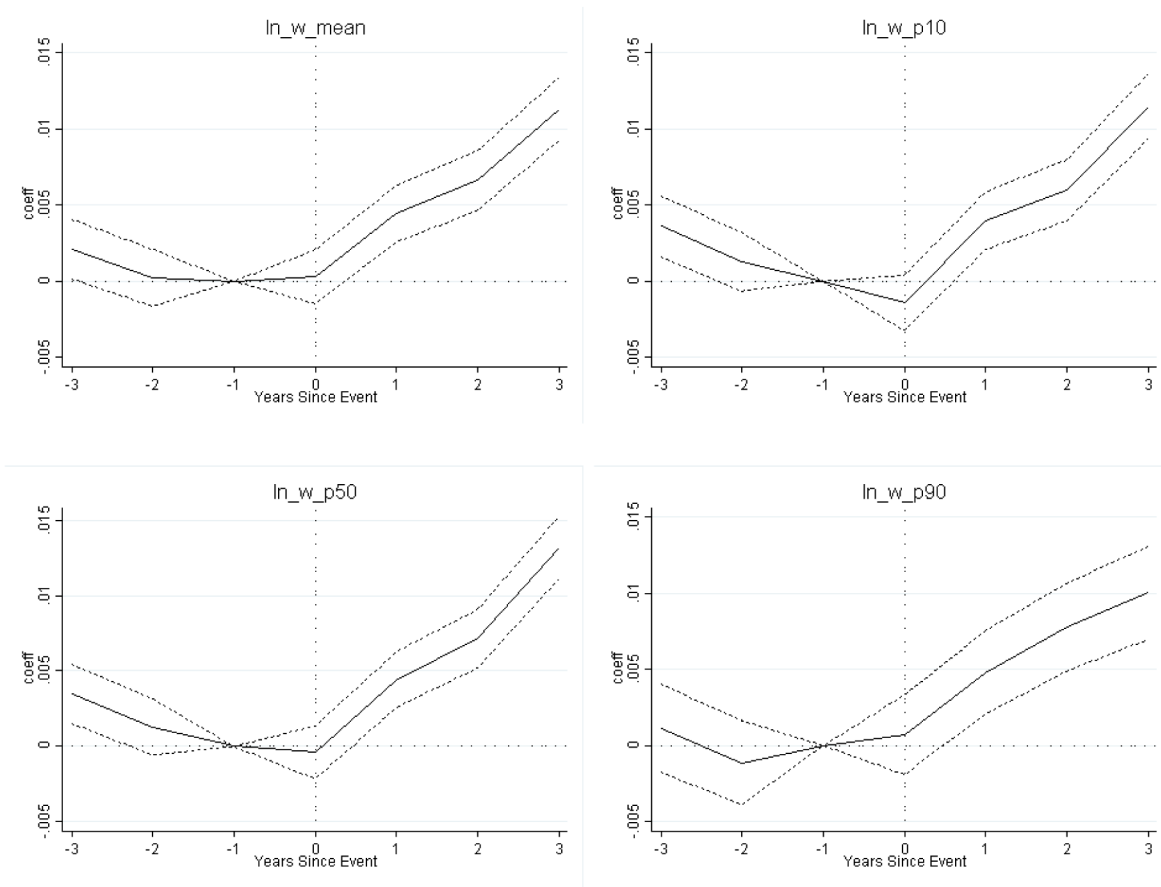


Fig. 4. Automation/AI spikes and the within-firm wage distribution. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

where y_{ijt} is the dependent variable of interest for firm i at time t in sector j ; D_{kit} is a dummy = 1 if index = k for firm i in year t ; δ_i and ζ_{jt} are a set of firm and sector-year fixed effects, respectively, and finally, ϵ_{it} is the error term.

β_k represents the effect of the automation/AI event on outcome y , k years after the event (or before if $k < 0$). These effects are measured relative to a baseline year, in this case, $k = -1$, which is excluded.

The values at which the index is censored (i.e., k_{min} and k_{max}) usually depend on the type of data available. We set $k_{min} = -4$ and $k_{max} = 4$ so that β_{-4} (β_4) represent average outcomes four or more years prior (later) to the event relative to those at $k = -1$. Eq. (3) is thus a flexible tool with which to study the timing of the effects of automation/AI. To focus on short-term effects of automation/AI, which can be more directly attributed to the spike event, we will focus on coefficients

from β_{-3} to β_3 when displaying the results, although other years are controlled for in the regressions.

We perform our main regressions on the sample of spiking firms (Sample 3, see Table 1), including a rich set of firm and sector-year fixed effects. In this way, the coefficients β_k are identified using the variation in the timing of the spike across firms, and they represent the difference between the value of the dependent variable one year before the spike and k years after (or before), net of sector-specific time trends.

It is important to note that in order to provide a causal interpretation of the coefficients, one should assume a counterfactual scenario in which, absent the event, the spiking firm would not experience the observed change. This is similar to the parallel trend assumption of a difference-in-differences regression to which our research design is closely related: in our case, there are only treated firms, but they are treated in different time periods, as in Bessen et al. (2020a,b). Keeping only treated firms makes it more likely that they are following parallel trends, especially given the large differences observed between the groups of firms with and without a spike (see Table 5). On the other hand, a useful characteristic of our event study is that it has built-in placebo tests (Lafortune et al., 2018) that should tell us how far we are from this assumption. In practice, we will check whether the variable of interest shows any specific trend before the spike. Absent that, it is more plausible to assume that the results are not driven by pre-spike differences across firms. In any case, given the non-random nature of an automation/AI spike, one should still be cautious about causal interpretation of our results. In particular, demand or supply shocks that occur the same year of a spike may be endogenous to the decision to automate. For this reason, we will interpret the coefficients mostly as describing the evolution of firm outcomes around the spike, as in Bessen et al. (2020b).

4.2. Results

We will now discuss the results of the estimation of Eq. (3), as displayed in Figs. 3 to 5. In all of these figures, we plot the coefficients β_k from β_{-3} to β_3 , and the dashed lines represent confidence intervals at the 5% significance level. All of the regressions are performed on the number of observations and firms of Sample 3, as reported in Table 1. In this Section 4.2, we report the main results of our analysis, focusing on wage inequality, wages at different locations of the wage distribution, and the gender wage gap. In the next section, we will report findings aimed at explaining those results and uncovering the mechanisms at play (Section 4.3). Then, in Section 4.4, we perform some robustness checks on our findings.

Wage inequality within firms

In Fig. 3, we investigate the impact of automation/AI spikes on within-firm measures of wage inequality, using as proxies of inequality the 90/10 ratio of wages per hour (left) and the standard deviation of hourly wages (right) within a firm. For both investigated measures, the β_k coefficients are not significant, with the exception of a barely significant and positive effect on the 90/10 ratio two years after the spike and on the standard deviation three years after the spike. Note that these coefficients, though significant, are small in size: the increase in the 90/10 two years after a spike is estimated as 0.13, compared to a mean value of 2.52 for firms in Sample 3; and the increase in the standard deviation three years after a spike is 0.20, compared to a mean of 10.66 (see Table 5). Furthermore, these increases are not persistent, as they revert to nonsignificant afterwards in any case, and they are not always found to be robust after our further tests (see Section 4.4).

Our result adds further evidence on the scant literature on technology and within-firm wage inequality.³⁰ A positive correlation between

innovation and within-firm wage inequality is found in Cirillo et al. (2017) in European countries, where, however, a general R&D innovation proxy is considered. Our result conveys a different message. By focusing on adoption rather than innovation, we find that inequality is substantially unaffected by an automation/AI spike. One possible explanation is that adoption did not increase wages to begin with in our case; however, we will see below that this hypothesis is not supported.

Wage increase by percentiles

Having established that within-firm wage inequality does not increase following an automation/AI spike, the question remains whether this is simply the effect of a disconnect between automation and wage dynamics or whether wages increase following an automation spike in a fairly equal way across workers. In that case, differences in wages across firms would be affected by technological adoption. We try to settle this question in Fig. 4. There, we report results where y_{ijt} represents the mean and different percentiles of the within-firm wage distribution. Variables are log-transformed so that coefficients can be interpreted as percentage changes with respect to the value of y_{ijt} one year before the spike.

The first plot of Fig. 4 (top, left) shows the effect of an automation/AI spike on the (log) mean hourly wage of the firm. Following a spike, there is an increase in the mean wage that is, at first, not significant (in the year of the spike) but then reaches significance with an increasing trend. Overall, the effect is precisely estimated to be small: three years after the spike, the mean wage is 1.1% higher than before the spike.

Such an increase in hourly wages seems to be due to positive changes at different percentiles of the distribution. In Fig. 4, the 10th, 50th, and 90th percentiles are respectively 1.1%, 1.3%, and 1.0% higher three years after the spike than before it happened ($t - 1$). Overall, we can conclude that following an automation/AI spike, there is a general increase in workers' wages three years after the event; such an increase is equally distributed across the wage percentiles, reinforcing the message coming from the previous exercises that there is no change in within-firm inequality after such an event.

How do these results compare to other estimates from the literature? Precise estimates of the relation between automation and wages can be found in the case of the adoption of industrial robots (Koch et al., 2019; Barth et al., 2020; Humlum, 2020). Koch et al. (2019) find no significant effect of robot adoption on the average firm wage in Spain; Barth et al. (2020) find a 4% increase in the average log hourly wage in manufacturing firms in Norway, and Humlum (2020) report an 8% increase in the wage bill in the case of Denmark. Finally, a modest relationship between hourly wage and robot adoption is observed in Acemoglu et al. (2020) in a smaller sample including approximately 600 French robot adopters.

For automation, measures of automation intensity (instead of adoption) are used by Dinlersoz et al. (2018b) and Aghion et al. (2020). Dinlersoz et al. (2018b) show a positive relationship between automation intensity (using a technology index) and wages in U.S. plants, while Aghion et al. (2020) do not find any significant effect of changes in electric motive power and average wages at the firm level. Finally, the study by Bessen et al. (2020a) in the case of the Netherlands uses both automation cost spikes and information about automation importers using trade data. Using an event-study methodology, they find an increase in daily wages and the wage bill for importers, while the wage bill decreases among non-importers.

Another element of comparison is the heterogeneity of the wage effect across workers within firms. The study by Webb (2019), using job task descriptions and patents, highlights differences between the predicted labor impact of robots and AI technologies across skill levels. If both Barth et al. (2020) and Humlum (2020) find heterogeneous effects across worker groups, whereby skilled or tech workers benefit from wage gains while unskilled or production workers lose wages after

³⁰ For a review on the link between technology and overall wage inequality, see Acemoglu and Autor (2011).

the adoption of industrial robots, this is not a case supporting automation adoption. Indeed, neither (Aghion et al., 2020) nor (Bessen et al., 2020a) report differences across skill or wage quartile groups. Aghion et al. (2020) conclude that, in the case of France, “the distributional effects of automation in the labor market are subtle”. They attribute this difference to international competition pressure, which is lower in the U.S. (Acemoglu and Restrepo, 2018).

Gender wage gap

While we report an increase in wages per hour at different levels of the wage distribution in adopting firms, an interesting question is whether, within a given percentile of the wage distribution, there is a change in the gender wage gap. Available evidence and theoretical models suggest that intra-firm gender wage gaps may relate to firm-specific characteristics, such as size and bargaining regimes (Oi and Idson, 1999; Heinze and Wolf, 2010; Card et al., 2016), as well as, more generally, the extent to which firms reward job-related characteristics such as temporal flexibility (Goldin, 2014). Only preliminary empirical evidence is available on the direct effect of automation on this gender gap. In a sample of Estonian manufacturing firms, Pavlenkova et al. (2021) show that automation benefits the wages of male workers more than female workers.

We now test this hypothesis by separately estimating Eq. (3) for our gender wage gap measure (the ratio of female-to-male wages) computed at different percentiles of the wage distribution. This takes into account the evidence shown in Table 5, according to which the gender gap does change along the wage distribution, as well as evidence coming from other countries (see, for example, Gardeazabal and Ugidos, 2005 on wage discrimination at quantiles in Spain). The results of this analysis are reported in Fig. 5. We plot the β_k coefficients from Eq. (3), where the dependent variable is the ratio between the female and male wages at the 10th, 50th and 90th percentiles, as well as at the mean. In general, the ratio does not significantly change after a spike, although a larger and more consistent positive increase (though insignificant) emerges for the 90th percentile. This result suggests that the increase in wages following an automation/AI spike is equally distributed not only across the wage percentiles but also within them, across male and female workers.

4.3. Investigating the mechanisms

The exercises above, together with the evidence on a similar dataset in Domini et al. (2021), go against the view that automation/AI adoption affects the relative demand for labor within firms. Indeed, neither the distribution of occupations nor the wage at different percentiles of the distribution appear to change after such an event. Instead, we observe a firm-level effect on wages. Not only is there a between-firm effect (firms that adopt automation- and AI-related capital goods pay higher wages and have higher productivity and profitability than firms that do not; see Table 5) but also, in the sample of firms that have a spike in the period of analysis, we observe higher wages after the event. As discussed in the introduction, several mechanisms may explain the role of automation/AI adoption in wage inequality across firms.

We explore below the different channels according to which automation/AI can lead to higher wages at the firm level: (i) technology adoption has a positive productivity effect (in line with Acemoglu and Restrepo, 2019), and together with such productivity increase, the firm would share its higher profits with its employees in the form of higher wages (according to the rent-seeking behavior in Blanchflower et al., 1996); and (ii) the firm also changes the profile of its newly hired employees through a sorting and matching effect of technological change (Abowd et al., 1999b; Cahuc et al., 2006; Song et al., 2019).

The productivity channel

One simple explanation for the wage increases at all levels of the distribution would be that it reflects a higher productivity and profitability of the firm, then passed through to wages via a rent-sharing process (Blanchflower et al., 1996). On this basis, we would expect a positive impact of an automation/AI spike on productivity, with a higher coefficient than that found for wages. Fig. 6 shows the change in productivity (value added per hour worked) after an automation/AI spike. Contrary to what we expect from our economic intuition, as well as what is predicted from the model of Acemoglu and Restrepo (2019), we find a negative impact on productivity, which is approximately 3% lower three years after the spike than it was before the spike.

A closer look into the literature on the relation between investment and productivity on the one hand (Power, 1998; Grazzi et al., 2016) and on the impact of productivity shocks on wages on the other hand (Harris and Holmstrom, 1982; Carlsson et al., 2016) provides an economic framework through which to interpret these results.

First, the empirical literature on productivity growth after an investment spike shows that the short-term effect is negative (Power, 1998; Huggett and Ospina, 2001; Grazzi et al., 2016). The learning-by-doing mechanism would explain this and suggest that productivity growth should then turn positive after employees adjust to the new technology and obtain returns from it. However, it is very difficult to observe this positive effect of capital investment on productivity within the firm, even when accounting for a long lag in time after the investment (Power, 1998; Grazzi et al., 2016).

Second, what do we know about the response of wages to productivity shocks? In the model by Harris and Holmstrom (1982), the effect depends on the sign of the productivity shock and is asymmetric: only positive shocks are passed on to wages, while negative shocks are not. Such downward wage rigidity is in particular expected in countries such as France with collective wage bargaining and a large emphasis on the minimum wage and permanent contracts (Babecky et al., 2010; Avouyi-Dovi et al., 2013). In addition, according to Carlsson et al. (2016), in the case of Swedish firms, wages respond much more to sector-level changes in productivity than to firm-specific characteristics due to mobility within sectors. Relatedly, Montornès and Sauner-Leroy (2015) show that in the French context, wage changes are mostly explained by new hires. From this exercise, we conclude that the productivity channel does not explain the increase in wages observed after an automation/AI spike.

The employee-matching channel

In the exercises above, we focused on heterogeneous effects across workers at different levels of the wage distribution as a way to form a proxy for changes in labor demand linked to skills or tasks. However, other sources of heterogeneity in the wage dynamics across workers within the firm should be accounted for. In addition to firm characteristics, individual or “personal” unobservable effects matter greatly in explaining wage dynamics in the French context (Abowd et al., 1999b). We explore this channel by decomposing the overall wage effect into new hires and incumbents, and by specifically looking at the wage of workers who leave the firm, i.e., separated workers.

Newly hired workers

The relevant literature highlights how workers with “good” characteristics get matched with “better” firms, i.e., those firms better able to compete in the labor market and attract the best workers (Cahuc et al., 2006). From this, and from the French institutional context described above, we expect that wage dynamics in the firm could be mainly driven by a change in the profile of new hires relative to incumbents.

To test this possibility, we investigate the effects of automation focusing on the ratio between the hourly wage of newly hired workers per year t , defined as those that are not present in the firm on December 31st of year $t - 1$ but are employed on December 31st of year t , with respect to the wage of incumbents, defined as workers present at both dates.

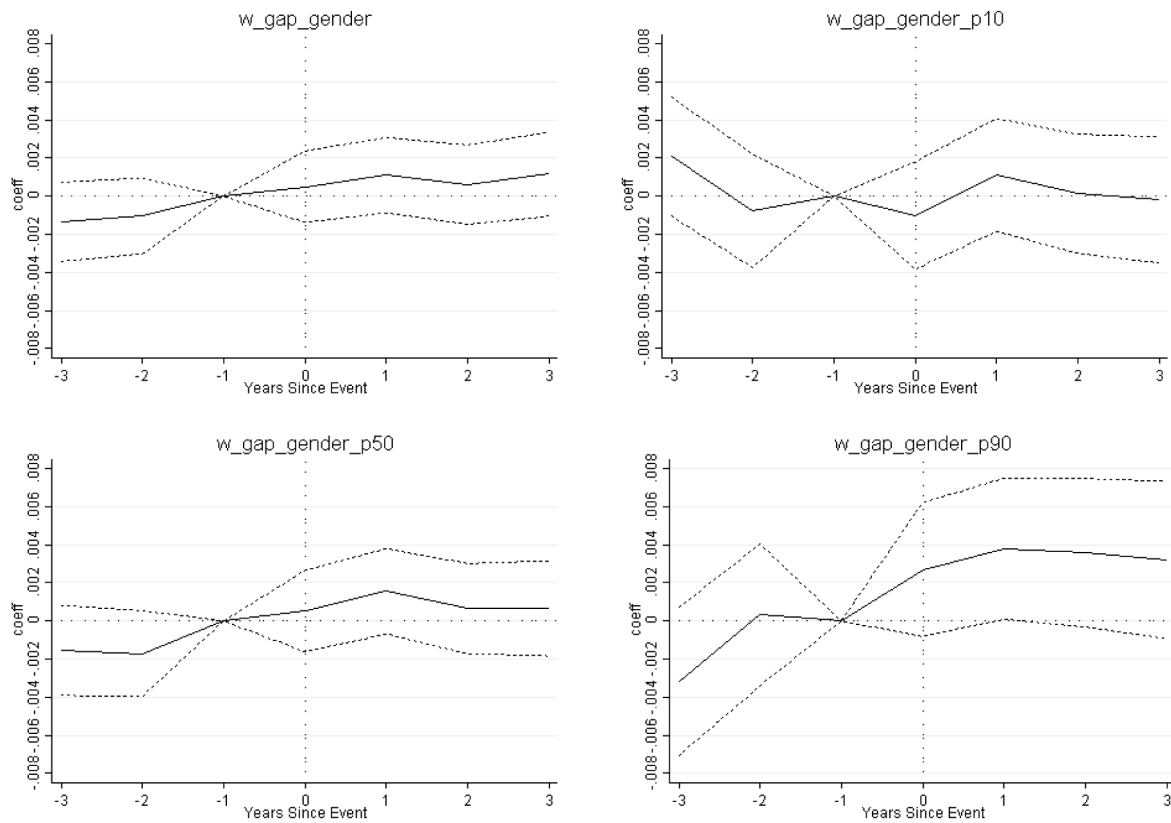


Fig. 5. Automation/AI spikes and the gender wage gap. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

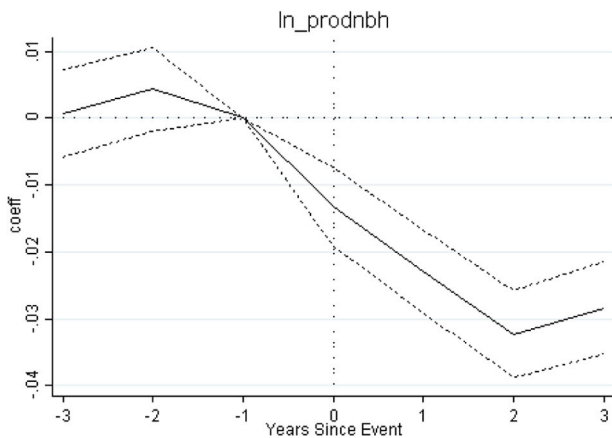


Fig. 6. Automation/AI spikes and productivity. Notes: The plot reports the impact of automation/AI on the log-transformed value of value added per hour worked; The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

Results are reported in Fig. 7. We find that, three years after a spike, the ratio between the mean wage of new hires and that of incumbent workers is one percentage point higher, with respect to one year before the spike. The effect is quite similar at the 50th percentile and slightly larger at the 10th percentile, but it is less prevalent and not significant at the 90th percentile, where the error in the estimation is larger. Note

that, similar to the effect found for all workers (see Fig. 4), the change in the relative wage of new hires and incumbents is slightly delayed, as it starts to be observed at $t + 1$.³¹ Additionally, in this case, there is no evidence that pre-spike trends are significant, suggesting that workers with different wages do not select into the firm before the spike.³² These results suggest that after adopting automation/AI, the profile of new hires changes: one possible explanation, consistent with the employee matching channel, is that automating firms look for workers with more experience and education, including knowledge of the new technology being adopted.

Separated workers

In this final exercise, we compare the wages of separated workers, defined as those that are present in the firm on December 31st of year $t - 1$ but are not present on December 31st of year t , to that of incumbents after the automation/AI investment spike. While the sorting and matching literature focuses on workers' entry into the firm, some empirical works on the employment effects of automation also discuss the characteristics of the workers leaving it. In particular, Bessen et al. (2020a), investigating workers' probability of leaving after an automation spike, find that it does not depend on workers' characteristics such as wage, age or gender.

³¹ Note that not all firms hire new workers each year, so the sample of firms and observations on which the equation for newly hired workers is estimated is slightly smaller, consisting of 473,976 observations (vs. 506,374 of the full sample) and 38,942 firms (vs. 39,580 of the full sample). We have also performed our previous estimations using this smaller sample, and our results remain unchanged.

³² For a similar concern, see Bessen et al. (2020a).

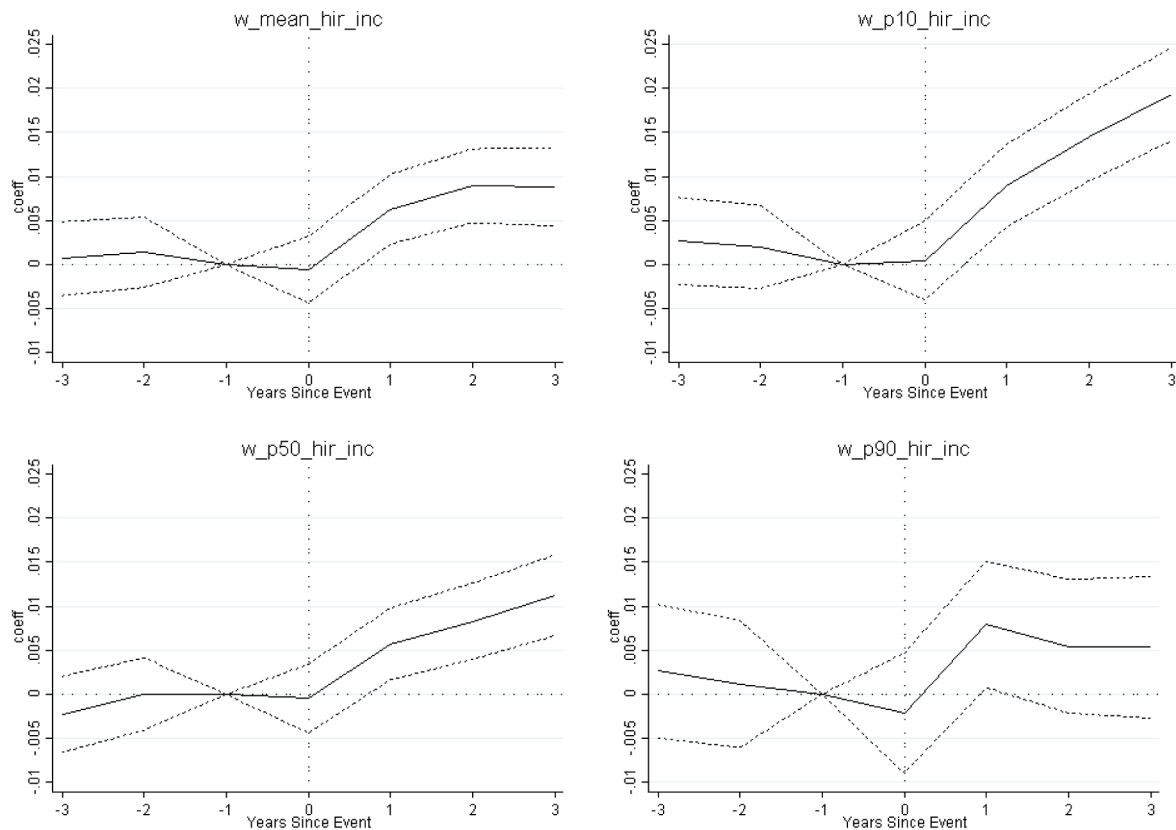


Fig. 7. Automation/AI spikes and newly hired workers' wages relative to incumbent workers' wages. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

In our case, we restrict the comparison to the difference in wages between incumbent and separated workers.³³ Similar to our new hire vs. incumbent wage ratio, we compute the ratio of wages of the separated workers over wages of incumbents at different percentiles of the distribution and track the evolution of this ratio around an automation/AI spike.

We find some heterogeneity across wage percentiles. Indeed, the relative wage of separating workers at the 90th percentile slightly increases after the event and is 1.3 percentage points higher two years later; conversely, we do not find significant differences at the other percentiles of the wage distribution. This result implies that the workers who leave the automating firm have slightly higher wage profiles than workers who stay, but only at the top of the distribution. Those who leave might be those with longer tenures and who therefore do not match as well with the new technology of the firm. Note that at this level of the wage distribution in our sample, almost all workers have a permanent contract (close to 97%), so the decision to separate might be driven by the worker, who may have relatively good re-employment prospects compared to workers at lower levels of the skill distribution (Berson et al., 2020) (see Fig. 8).

4.4. Robustness tests

Here, we discuss a series of robustness exercises as well as the motivation behind performing them. First, we focus on different definitions

³³ The type of data and worker-level information in the paper by Bessen et al. (2020a) allows them to control for more person-specific dimensions, which are not available in our data, such as tenure in the firm as well as information on income and status after separation.

of the spikes, namely: (i) separately identifying automation spikes and AI spikes; and (ii) introducing a size threshold for the identification of an automation/AI spike. Then, we modify the sample of analysis in different ways: we separate manufacturing from services sectors to control for sectoral heterogeneity; we remove firms that import these products but may not change their production processes; and we focus on a balanced sample of firms that are present throughout the years 2002–2017, to assess whether our results are influenced by firm entry and exit. Finally, we address a potential concern regarding a discontinuity in our customs data source, by showing results for the subperiod 2002–2010; and we check whether our interpretations depend on the specific time window we have used for our analysis (from three years before to three years after a spike), by showing results for a larger time window (from five years before to five years after a spike).

In the rest of this section, for the sake of conciseness, we will limit our comments to mentioning whether remarkable differences arise between each robustness test and the baseline results shown so far. Results on our main variables of interests (namely, the 90/10 wage ratio, the gender wage gap, productivity, the mean wage, and the wages of newly hired and separated workers, relative to incumbents) are shown in Appendix B. Full results are available upon request.

Changing the definition of spikes

AI- versus automation-only spikes In our main analysis, we employed spikes based on imports of automation- and AI-related goods. Our spike variable may therefore identify episodes of investment in automation technologies only, AI technologies only, or both at the same time. However, there are reasons to believe that the impact of these two groups of technologies on employment and wages may be different, as

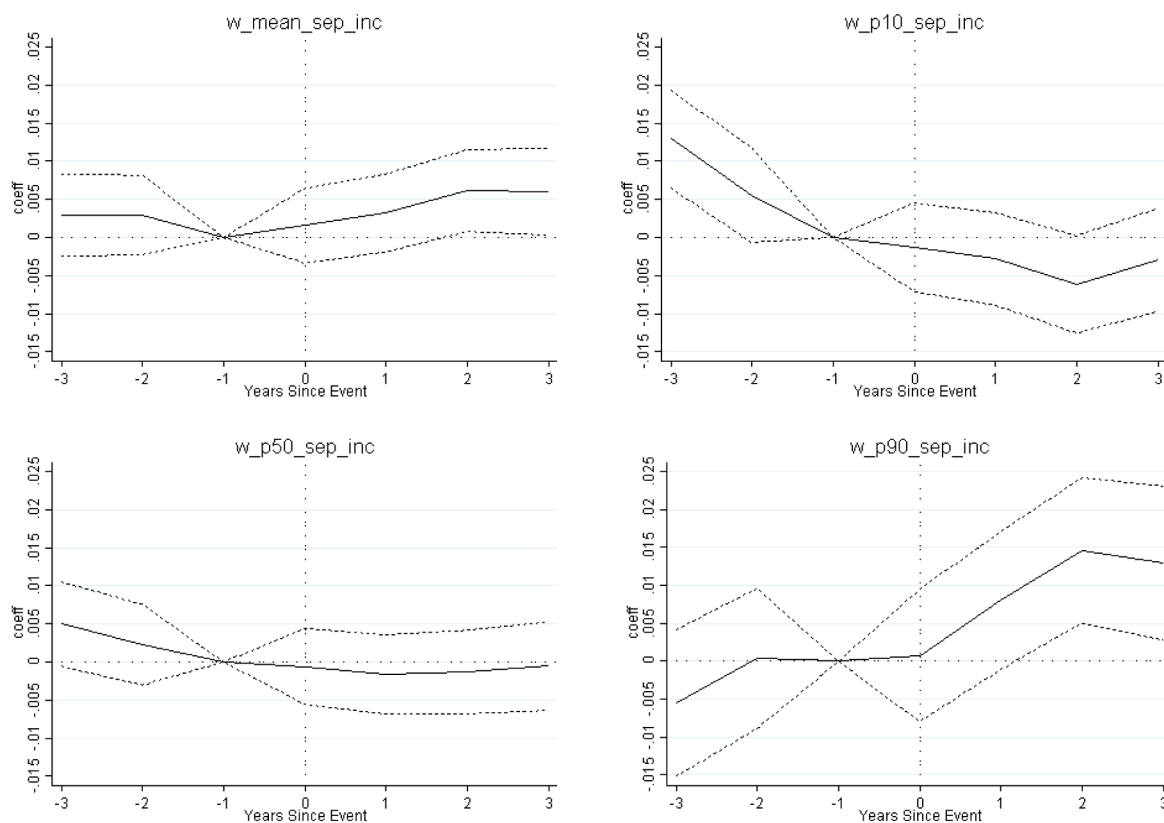


Fig. 8. Automation/AI spikes and separated workers' wages relative to incumbent workers' wages. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

they serve different purposes and replace/complement different types of workers. Webb (2019) observes that while software and robots impact mostly low-skilled workers, AI is directed at high-skill tasks. A further reason for separately analyzing automation and AI is that some previous literature focuses on only one of them. Hence, separate analyzes should enhance the comparability of our findings.

As a consequence, we rerun our analysis using spike variables defined on automation only and on AI only. Table A.1 shows which technologies belong to each group, and Table A.2 shows how many firms import (and have a spike in) goods of either group. We observe that importing automation products is more common than AI products, but the gap is smaller in terms of spikes, especially at the end of the period.

In Figs. B.1 to B.2, we show the results for automation-only and AI-only spikes, respectively. No qualitative differences are detected, except for AI significantly increasing the gender wage gap three years after a spike, although with a negligible magnitude, while automation does not. A general consistency can be observed for other variables (not displayed in Figs. B.1 and B.2) as well.³⁴

More restrictive spike definition We also test the robustness of our findings by setting a size threshold for identifying an automation/AI spike, following the investment spike literature (Nilsen et al., 2009; Grazzi et al., 2016). In our main analysis, we define a spike as the main episode of imports of such technologies, without any restrictions. For this robustness test, we adopt a more stringent definition of spikes as the main episode of imports of automation/AI products that is at least three times larger than the average value of imports by the same firm

in other years (at constant prices). This is similar to the spike definition adopted by Bessen et al. (2020a,b); hence, this robustness test allows use to increase the comparability of our results to theirs.³⁵ Adopting this alternative definition of a spike restricts the sample of spiking firms we use for our regressions (Sample 3) since it causes us to drop spikes that do not meet the relative size threshold mentioned above. The number of observations used in regressions is reduced by approximately one fourth, as a similar (though slightly larger) share of spikes as per our main definition are discarded.

The results from this robustness check, shown in Fig. B.3, convey a substantially unchanged picture, the only noticeable difference being that the post-spike coefficients for the 90th-percentile new hires/incumbents wage ratio (as per Fig. 7, bottom right) and for all percentiles of the separated/incumbents ratios are not significant.

Changing the sample of analysis

Manufacturing vs. Services Our main analysis encompasses the entire French economy, with the exceptions of the primary sector (NAF rev. 2 divisions 01 to 09). However, there are reasons for performing separate analyzes on the manufacturing and service sectors. First, Montobbio et al. (2020) show that labor-saving technologies may challenge different activities in different sectors. Hence, the effect of such technologies on the wage distribution may be different across sectors. Second, restricting the analysis to the manufacturing sector enhances comparability with previous research on the effects of automation,

³⁴ Note that the AI results have larger error bands.

³⁵ They identify automation spikes in year t if automation costs, as a share of a firm's total costs, are at least three times the average firm-level cost share.

which has mainly focused on manufacturing firms.³⁶ Finally, focusing on manufacturing is one possible way of dealing with the issue of resale of imported automation- and AI-related goods, which will be explained below.

We therefore rerun our analysis separately for the manufacturing sector (NAF rev. 2 divisions 10 to 33) and for services (NAF rev. 2 divisions 35 to 96). These subsamples account for 44% and 56% of all observations in Sample 3, respectively. To be more precise, in 2017, manufacturing divisions jointly accounted for 42% of firms, 20% of the value of automation/AI imports, and 33% of employment.

In Figs. B.4 and B.5, we show results for manufacturing only and for services only, respectively. When observing the 90/10 wage ratio, which measures within-firm wage inequality, the results for services appear very close to our baseline results, while the results for manufacturing share the general lack of significance (except for a barely significant and positive coefficient one year after a spike). Likewise, for the gender wage gap, we do not find a significant effect of our automation/AI measure on the mean wage gap within either of the two sectors. The effects of automation/AI spikes on wages are qualitatively similar but of a lower magnitude for manufacturing: three years after a spike, the mean wage is 0.8% higher in manufacturing firms vs 1.3% in services firms.

Some differences appear, however, when looking at the new hires/incumbents wage ratio and the separated/incumbents wage ratio. The former shows stronger and more persistent increases after a spike in manufacturing firms. In contrast, the latter shows nonsignificant dynamics for manufacturing firms but a significant increase after a spike for services. The interpretation of this increase in wages after a spike that we provided above, as due to a sorting mechanism whereby firms adopting automation/AI hire new, better-paid workers, seems to hold particularly well for the manufacturing sector. Conversely, in the services sector, the increase in the new hires/incumbents ratio is less sizeable, while an important role seems to be played by the separation of relatively well-paid workers after a spike.

Excluding potential resellers of automation/AI products A potential drawback of our import-based measure of automation/AI adoption is that some firms that import goods related to these technologies may not use them themselves but instead resell them, either in the domestic market or abroad. When this happens, then our measure identifies firms that in fact are not adopters. This can be expected to be a particularly important issue in industries related to trade; remarkably, Table 4 shows that more than half of the value of automation imports from firms in Sample 3 is accounted for by NAF division 46 (Retail), while this division only accounts for approximately one tenth of the total employment in the same sample. As mentioned above, restricting the sample to the manufacturing sector is one way of dealing with this issue, as such a mismatch between the relevance of the sector in terms of automation/AI imports and in terms of employment and number of firms cannot be detected. However, manufacturing firms are also known to be involved in the (re)export of goods they do not produce, engaging in so-called Carry Along Trade (CAT). The next two robustness checks will address the possibility of re-exporting (by firms in any sector) in two different ways.

First, we exclude from our regressions re-exporting firms, defined as firms that, at least once, import and export automation- and AI-related goods in the same year. This restricts Sample 3 by one fourth.³⁷ We show the results on this smaller sample in Fig. B.6. The results on wage inequality are very similar: we find a positive impact on the 90/10

³⁶ Note that separating manufacturing and non-manufacturing industries allows aligning our results on the gender wage gap with the study by Pavlenkova et al. (2021).

³⁷ Notice that this robustness test is likely to fall short of capturing all resale activities, since we can only observe sales abroad (i.e., exports) but not sales occurring in the domestic market.

wage gap but only (barely) significant after two years and no effect on the gender wage gap. The effect at the mean and at different levels of the wage distribution are unchanged: after three years, there is a 1% increase in wages at all levels of the wage distribution. Finally, this is also the case for the ratio of new hires to incumbent wages. We can conclude that although excluding re-exporters modifies the sample of analysis in a significant way, the results are unchanged.

Another way to deal with the issue of the potential resale of imported automation- and AI-related goods is to exclude firms that import such goods every year, since these firms are more likely to be resellers than firms that import only once. Note that in this robustness check, firms that are identified as resellers can be assumed to resell not only in the export market but also in the domestic market. Again, this does not change the results qualitatively, as shown in Fig. B.7.

Balanced sample In our main analysis, we use an unbalanced panel, therefore including firms present for the whole period (2002–2017) as well as others that enter and exit. In what follows, we check the role of incumbent firms versus entering and exiting firms by running our analysis on a balanced sample. This sample counts 345,216 observations (68% of Sample 3), corresponding to 21,576 unique firms (55%).

Results are shown in Fig. B.8 and are qualitatively in tune with our main findings, although their magnitude is generally smaller. The only difference is that when testing the effect of automation on the new hires/incumbents and separated/incumbents wage ratios, the results are no longer significant. This happens also when rerunning our analysis on the subperiod 2002–2010 (see below): in both cases, we are employing a smaller sample, which is likely to affect statistical significance.

Additional checks

Subperiod 2002–2010 Another potential concern stems from a discontinuity in the import data, which we use to build the firm-level measure of automation/AI adoption. In particular, since 2011, product codes only have to be reported by firms with more than 460,000 euros of imports in a given year within the EU (Bergounhon et al., 2018).³⁸ This exposes our sample to the risk of excluding those smaller importers.

In order to evaluate the importance of this reporting change, we first look at the general trends in our data. In this regard, a first reassurance comes from the fact that we do not observe any discontinuity in 2011 in the share of firms that import automation- and AI-related imports, nor in that of firms characterized by a spike (see Table A.2). To provide further evidence, we rerun our analysis on the subperiod 2002–2010. Notice that this implies re-identifying the spikes, as they must now be based on the maximum value of automation- and AI-related imports within a shorter time period. For example, consider a firm that, when looking at the entire 2002–2017 period, has a spike in 2014: once we restrict the focus to the subperiod 2002–2010, the same firm may either have a spike in a different year (i.e. with a lower import value than the one initially identified in 2014), or no spike at all. The latter case causes our sample to shrink, over and above the simple restriction in the number of years. The resulting number of observations is 231,677, corresponding to 45% of sample 3. Results from this analysis, shown in Fig. B.9, are substantially in line with our main results. The only difference in this smaller sample is that automation and AI events no longer have a significant relation with the ratios between the wages of new hires and incumbents, and between separating workers and incumbents.

Larger time window Finally, we address the concern that the specific time window we use in our main analysis (three years before and after a spike) may miss important long-run dynamics. We therefore rerun our analysis using a larger time window spanning five years before and

³⁸ Acemoglu et al. (2020, Appendix A) estimate that this represents the price of four or five industrial robots.

after a spike. We do this by adding four time dummies ($t - 5$; $t - 4$; $t + 4$; $t + 5$) and changing the long-run dummies ($t < - 5$ and $t > + 5$ instead of $t < - 3$ and $t > + 3$), not shown in the graphs. The regression specification is therefore changed, which may also alter results within the original time window ($t - 3$ to $t + 3$).

Results are shown in Fig. B.10. The only noticeable changes involve the flattening of wage growth, four years after a spike, and the decreases the effects of automation and AI events on the new hires/incumbents and separated/incumbents wage ratios. Overall, these additional insights confirm and further qualify (in temporal terms) the claims made above.

5. Concluding remarks

In this paper, we have shown that within-firm wage inequality is a pervasive phenomenon in the French economy; most wage dispersion in France is accounted for by differences among workers belonging to the same firm rather than by differences between sectors, firms, or occupations. Restricting our attention to a sample of firms importing automation and AI-related goods, we found that major spikes in the imports of such goods are not followed by an increase in wage inequality, but they do tend to increase wages in an equal way at different percentiles and across male and female workers. Indeed, contrary to what has been found in the case of robot adoption (Humlum, 2020; Barth et al., 2020), our study and others focusing on automation (Bessen et al., 2020a; Aghion et al., 2020) do not observe a large distributional impact of automation across workers of different skills/wage percentiles. This hints at a role of the nature of technology: robot adoption displays complementarity with respect to workers at the top of the wage distribution and substitution effects for production/low skilled workers (Webb, 2019), while automation (broadly defined) and AI have a more uniform effect across workers along the wage distribution.

We also note that the magnitude of the effect is smaller than that found in previous studies focusing on robots in Norway and Denmark (respectively, Barth et al., 2016; Humlum, 2020). Adding to the role of the nature of technology highlighted above, the institutional context, especially labor market features, as well as the level of international competition (Aghion et al., 2020), could explain differences across countries. More work, especially across countries, is needed to disentangle the sources of heterogeneity between studies on the topic. These findings should be examined from the perspective and within the institutional context of the French economy, which did not experience any overall significant change in within-firm wage inequality during the examined period. Barth et al. (2020), for example, find that robots increase wage inequality in a sample of Norwegian manufacturing firms. An interesting question is to what extent future results in other countries will lean more towards the 'Norwegian' or the 'French' cases.

Coming to the interpretation of our results, our findings are not linked to the rent-sharing behavior of firms obtaining productivity gains from automation and AI adoption. Instead, we show that if wage gains do not differ across workers along the wage distribution, worker heterogeneity will still be present. Indeed, aligned with the AKM framework putting forward a change in the profile of new hires as a response to changes in firm performance (Abowd et al., 1999b; Cahuc et al., 2006), most of the overall wage increase is due to the hiring of new employees as part of the employment expansion that is generally associated with an event of automation (Domini et al., 2021).

Unfortunately, we do not have data on education or other worker-level characteristics to test whether the higher wages of newly hired workers is due to different skills, experience with similar technology, adaptability, or other individual-specific effects. In particular, our unconditional wage ratio between female and male workers does not take into account systematic differences that can be correlated with both wages and gender. In addition, we cannot follow workers over time, which implies that we cannot control for unobservable personal

characteristics. The lack of this worker-level information, as well as full information on job tenure, is a limitation of our study that we acknowledge. Future work could help identify the relative role of these different factors in explaining the wage impact of automation across workers.

There is also a complementary element to consider, still pertaining to inequality, that is related to the very nature of this recent wave of technologies but is not yet explored in this work. Indeed, both AI and related applications greatly benefit from almost zero variable costs and network externalities, which might easily generate dominant positions or quasi-monopoly rents. This is a perspective put forth by Guellec and Paunov (2020), according to which the growing importance of digital innovation, products and processes based on software and data has increased market rents, with benefits accruing disproportionately for the top income groups. Although taking a more aggregate perspective and not explicitly referencing AI, De Loecker et al. (2020) also detect a generalized increase in market power from 18% above marginal cost in the 1980s up to the current level of 67%. Obviously, the decision concerning the distribution of the returns associated with the adoption of technologies has clear implications for within- and between-firm wage inequality. However, an exhaustive investigation of such a link is beyond the scope of this specific work and is left for future analysis.

Overall, our findings add novel and important evidence to the emerging literature on the firm-level effects of automation. Previous contributions have mostly looked at the employment effects of the adoption of new technologies, usually finding a positive correlation between automation and employment at the firm level (Domini et al., 2021; Koch et al., 2019; Acemoglu et al., 2020). Here, we complement this picture of a 'labor friendly' effect of the latest wave of new technologies for adopting firms by showing that it increases wages without affecting within-firm wage inequality in a significant way. In other words, the increase in wages brought about by the adoption of automation and AI is enjoyed by all workers in the adopting firm, irrespective of their initial wage or gender.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This paper has benefited from the comments of three anonymous referees, Jim Bessen, Valeria Cirillo, Liuchun Deng, and the participants at several conferences: the Workshop on Technology and the Labor Market, IMT Lucca, Dec 2020; European Network on the Economics of the Firm, Jan 2021; Maastricht University, Feb 2021; Technology and Policy Research Initiative, Boston University, Feb 2021; SPRU March 2021; University of Turin, April 2021; PRIN 2017 Workshop, April 2021; SIEPI, University of Parma, June 2021; JRC LET seminars, Sep 2021; LISER-IAB conference Sep 2021; DRUID, Copenhagen, Oct 2021; CAED, Coimbra, Nov 2021; CONCORDi, Nov 2021.

This work has been partly supported by the European Commission under H2020, GROWINPRO, Grant Agreement 822781. Daniele Moschella received financial support from the Italian Ministry of Education and Research under the PRIN 2017 Program, Italy (Project code 201799ZJNS). This work is also supported by a public grant overseen by the French National Research Agency (ANR), France as part of the 'Investissements d'avenir' program (reference: ANR-10-EQPX-17, Centre d'accès sécurisé aux données, CASD). The usual disclaimer applies.

Appendix A

See Figs. A.1, A.2 and Tables A.1, A.2.

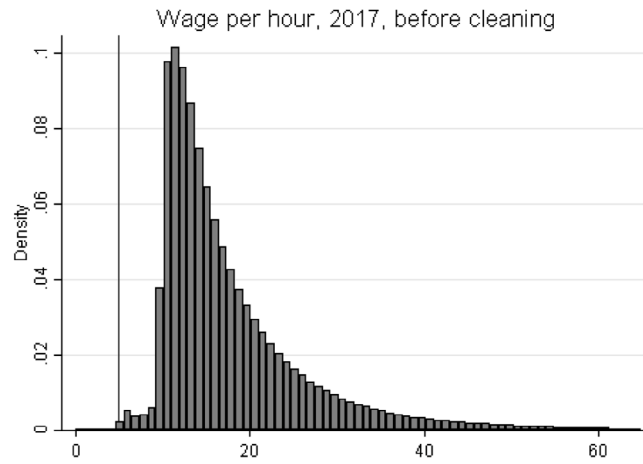


Fig. A.1. Distribution of wages per hour among all workers before cleaning. Note: The vertical line indicates our cleaning threshold (half the minimum wage per hour in 2017). Source: Our elaboration on DADS data.

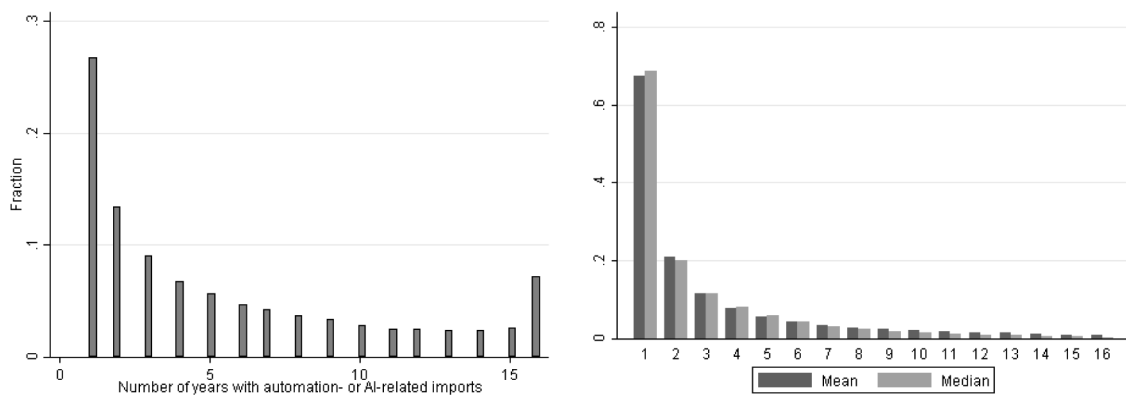


Fig. A.2. Testing the lumpy nature of our spike variable: Number of years with automation/AI imports (left); Automation/AI investment shares by rank (right). Note: Rank 1 is the highest yearly investment share in the firm's timeline. Source: Our elaborations on DGDDI data.

Table A.1
HS-2012 product codes referring to automation- and AI-related technologies.

Label	HS-2012 codes
<i>Automation</i>	
1. Industrial robots	847950
2. Dedicated machinery	847989
3. Automatic machine tools (incl. Numerically controlled machines)	845600–846699, 846820–846899, 851511–851519
4. Automatic welding machines	851521, 851531, 851580, 851590
5. Weaving and knitting machines	844600–844699, 844700–844799
6. Other textile dedicated machinery	844400–844590
7. Automatic conveyors	842831–842839
8. Automatic regulating instruments	903200–903299
9. 3-D printers	847780
<i>AI</i>	
10. Automatic data processing machines	847141–847150, 847321, 847330
11. Electronic calculating machines	847010–847029

Notes: For further details on categories (1)–(8), see [Acemoglu and Restrepo \(2022\)](#) (A-12-A14); on (9), see [Abeliansky et al. \(2020, p. 293\)](#); see also [Domini et al. \(2021\)](#); N.B. Codes for (1)–(8) only refer to automation-related capital goods, while the codes indicated by [Acemoglu and Restrepo \(2022, A-12-A14\)](#) also include non-automation-related capital goods (which are used as a control group in their analysis).

Table A.2
Automation and AI importers and spikes per year, as a share of Sample 2, 2002–2017.
Source: Our elaborations on DGDDI data.

Year	Importers			Spikes		
	Automation only	AI only	Either	Automation only	AI only	Either
2002	11.79	6.67	16.16	3.76	2.51	5.15
2003	11.69	6.36	15.85	2.67	1.78	3.55
2004	12.03	6.90	16.54	2.50	1.88	3.44
2005	12.24	7.09	16.88	2.48	1.87	3.45
2006	12.12	7.34	16.93	2.27	1.93	3.30
2007	12.47	7.03	16.95	2.64	1.64	3.41
2008	12.74	6.95	17.06	2.50	1.59	3.18
2009	12.12	6.44	16.16	1.92	1.23	2.42
2010	12.85	6.75	16.94	2.24	1.38	2.80
2011	12.15	8.61	17.45	1.94	1.95	2.88
2012	12.32	8.36	17.30	1.92	1.65	2.52
2013	13.00	9.60	18.79	1.99	2.08	2.92
2014	13.30	9.98	19.23	2.19	2.33	3.16
2015	13.56	10.52	19.90	2.39	2.83	3.75
2016	14.07	10.78	20.61	2.80	3.12	4.35
2017	14.46	10.71	20.74	3.92	3.68	5.55
Total	12.66	8.08	17.67	2.50	2.07	3.47

Appendix B

See Figs. B.1–B.10.

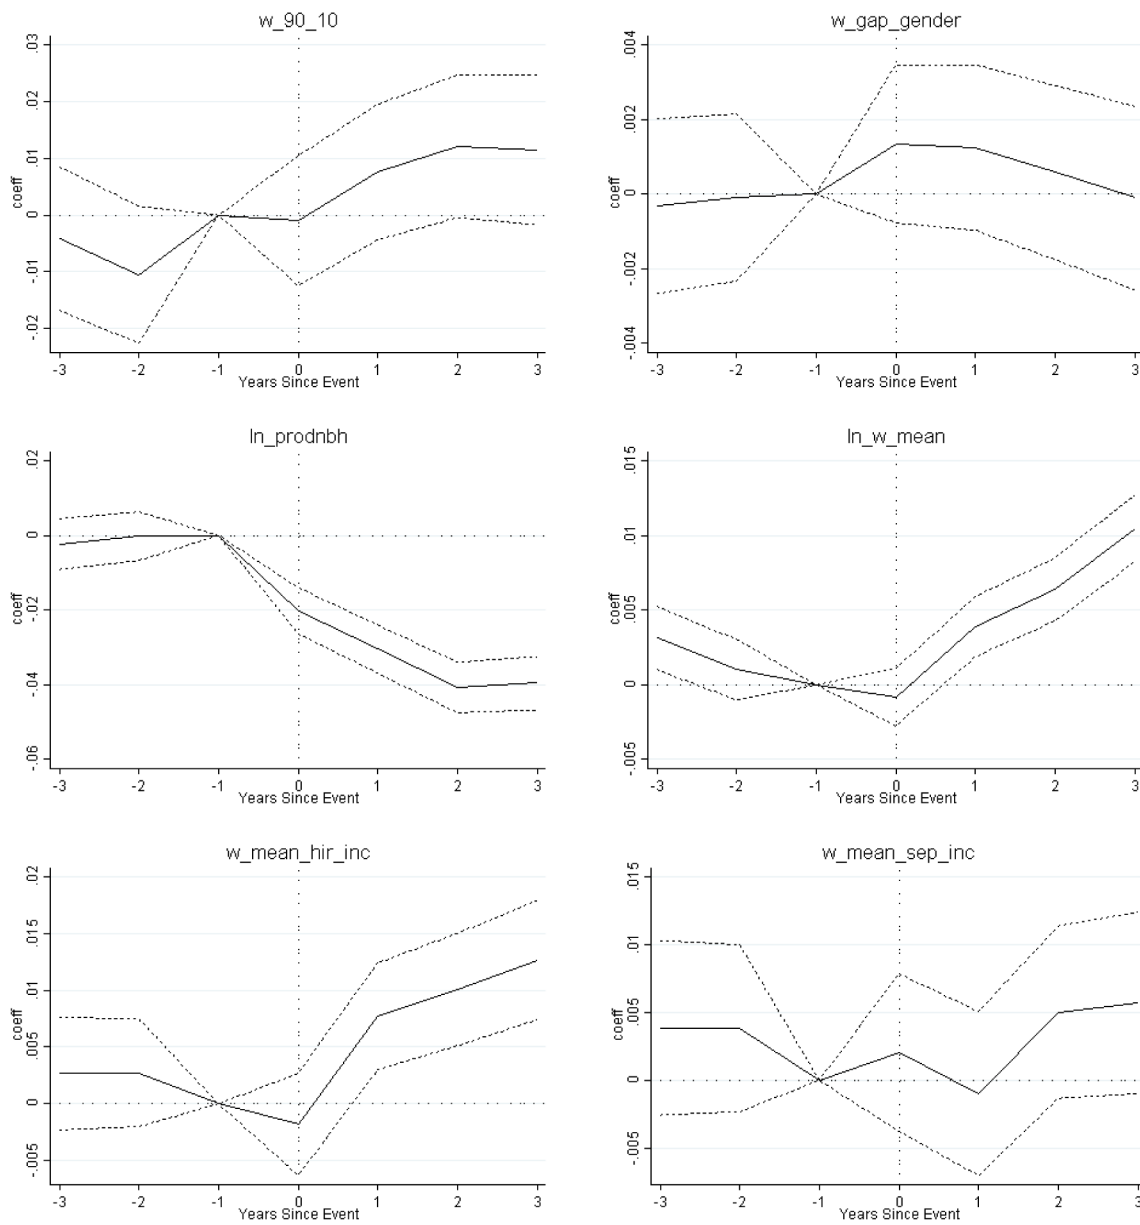


Fig. B.1. Robustness check: Automation-only spikes. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

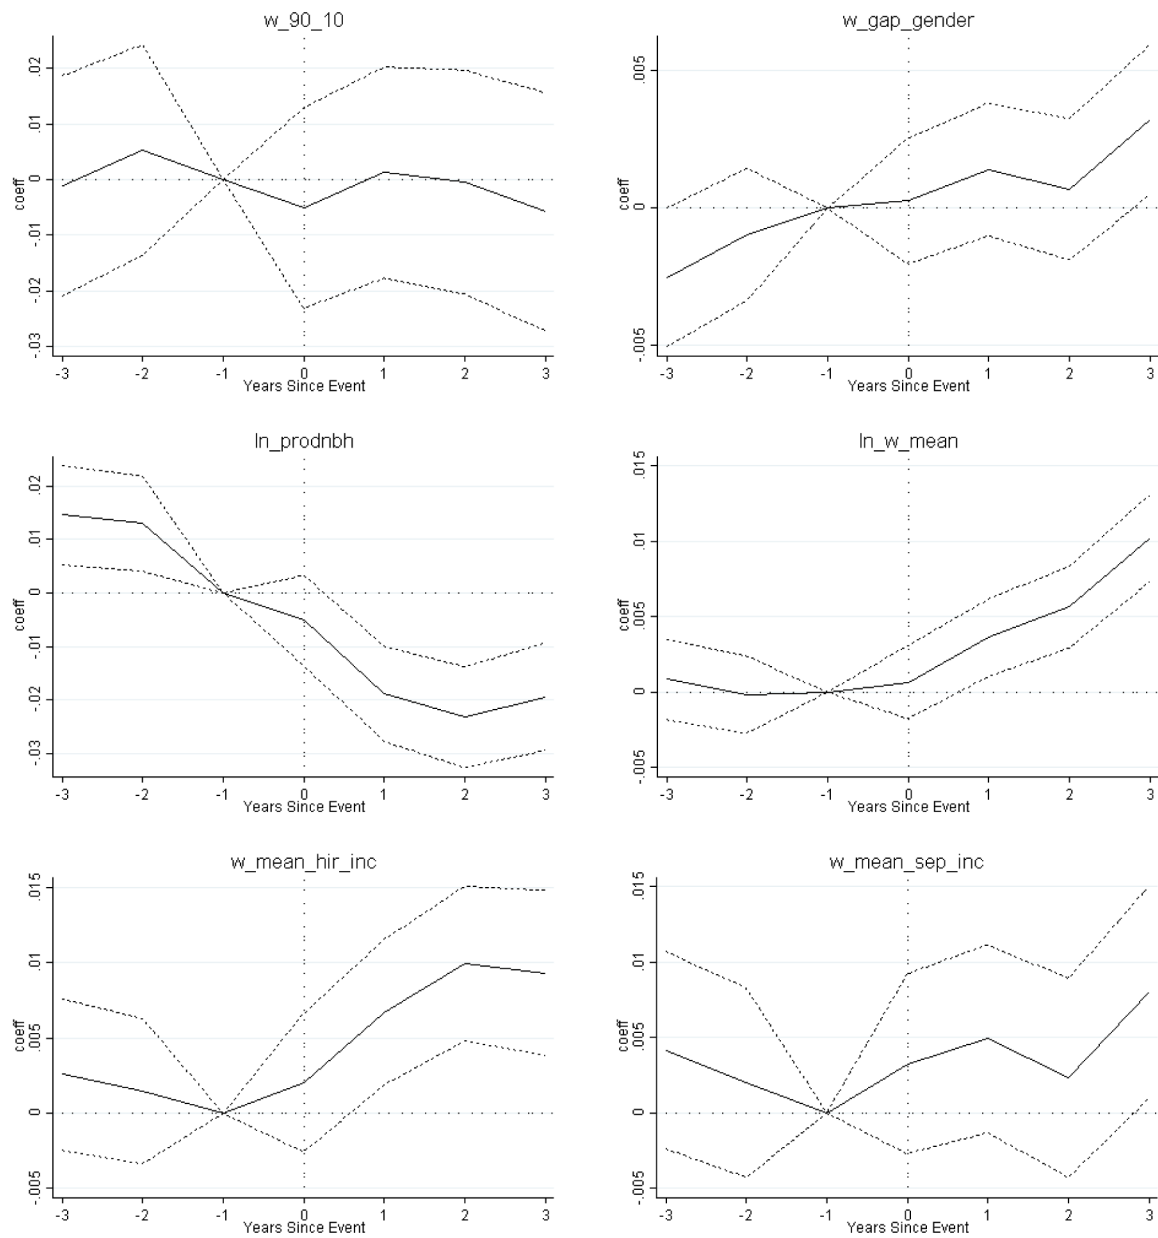


Fig. B.2. Robustness check: AI-only spikes. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

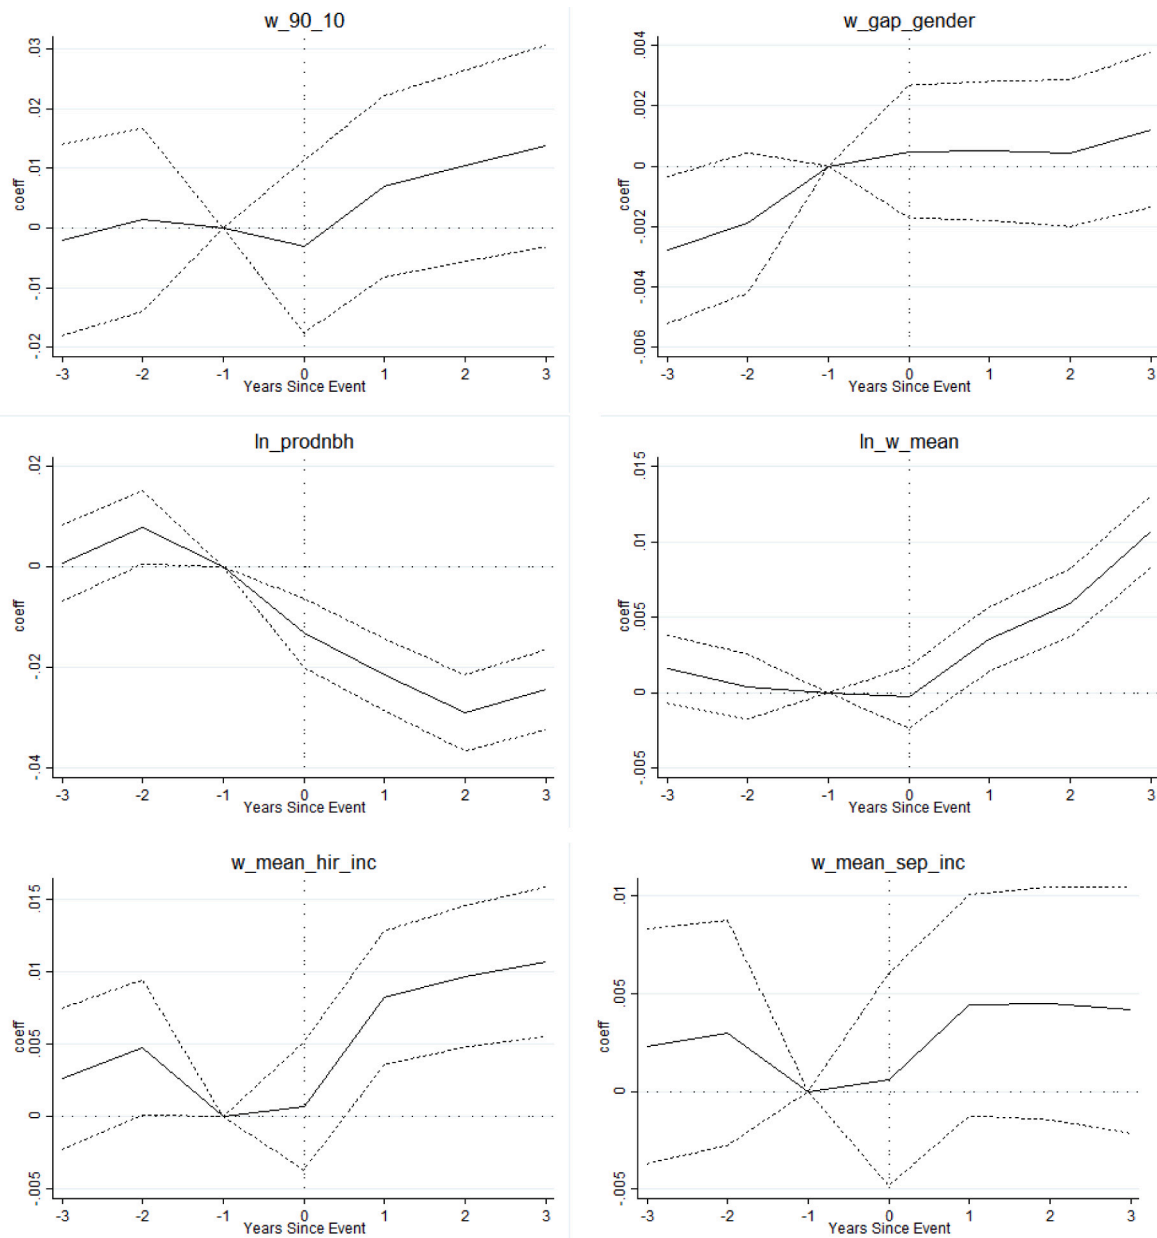


Fig. B.3. Robustness check: Alternative definition of automation/AI spike. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

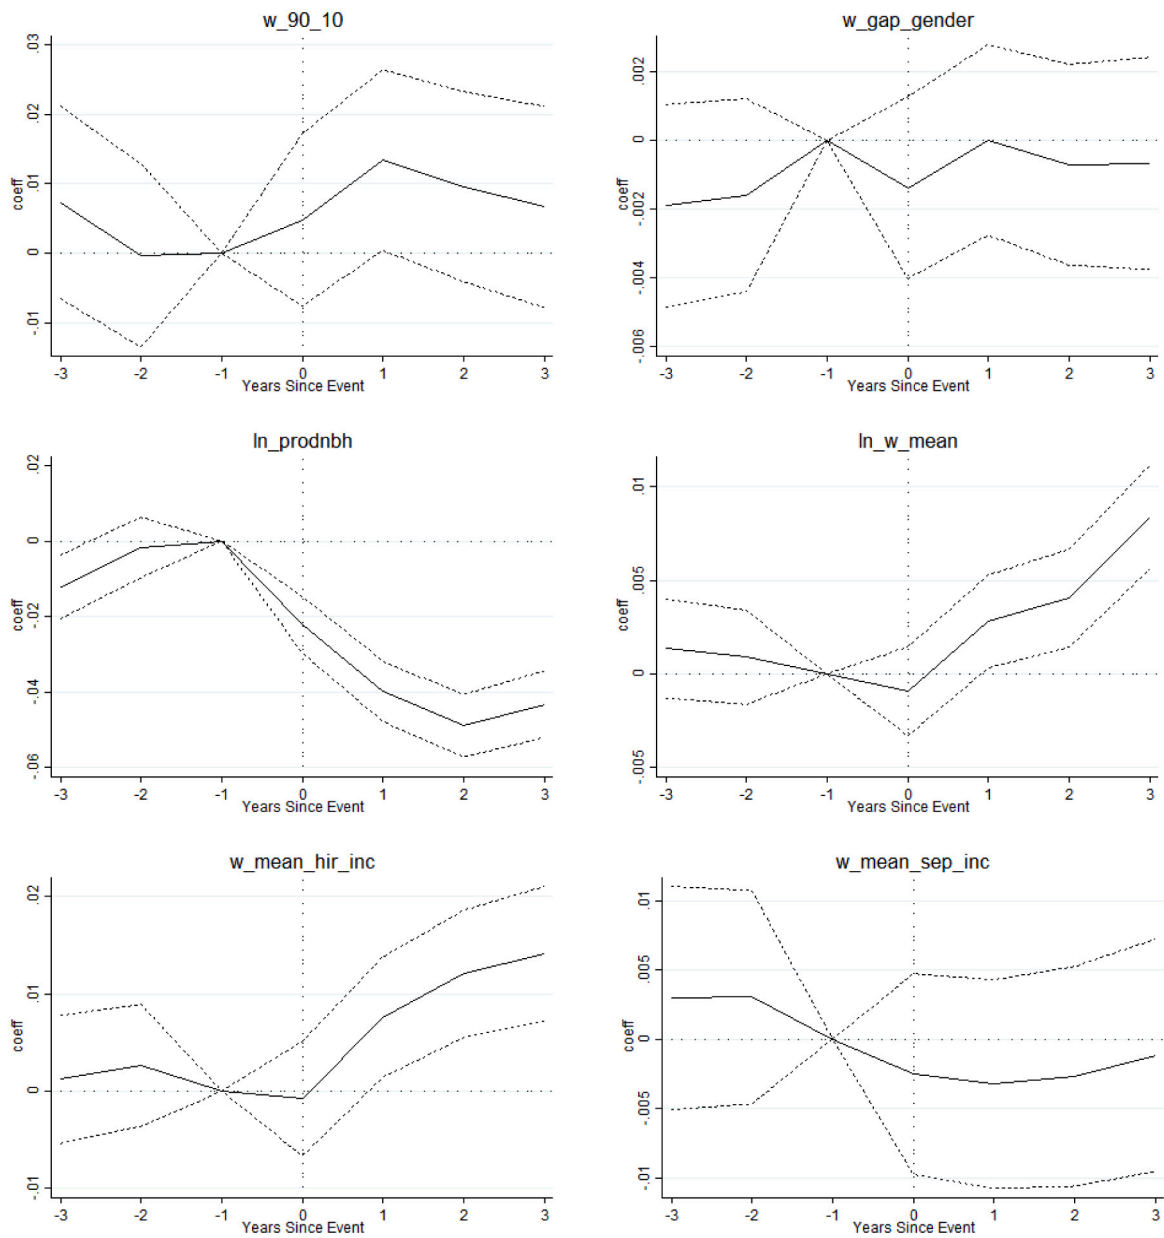


Fig. B.4. Robustness check: Manufacturing only. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

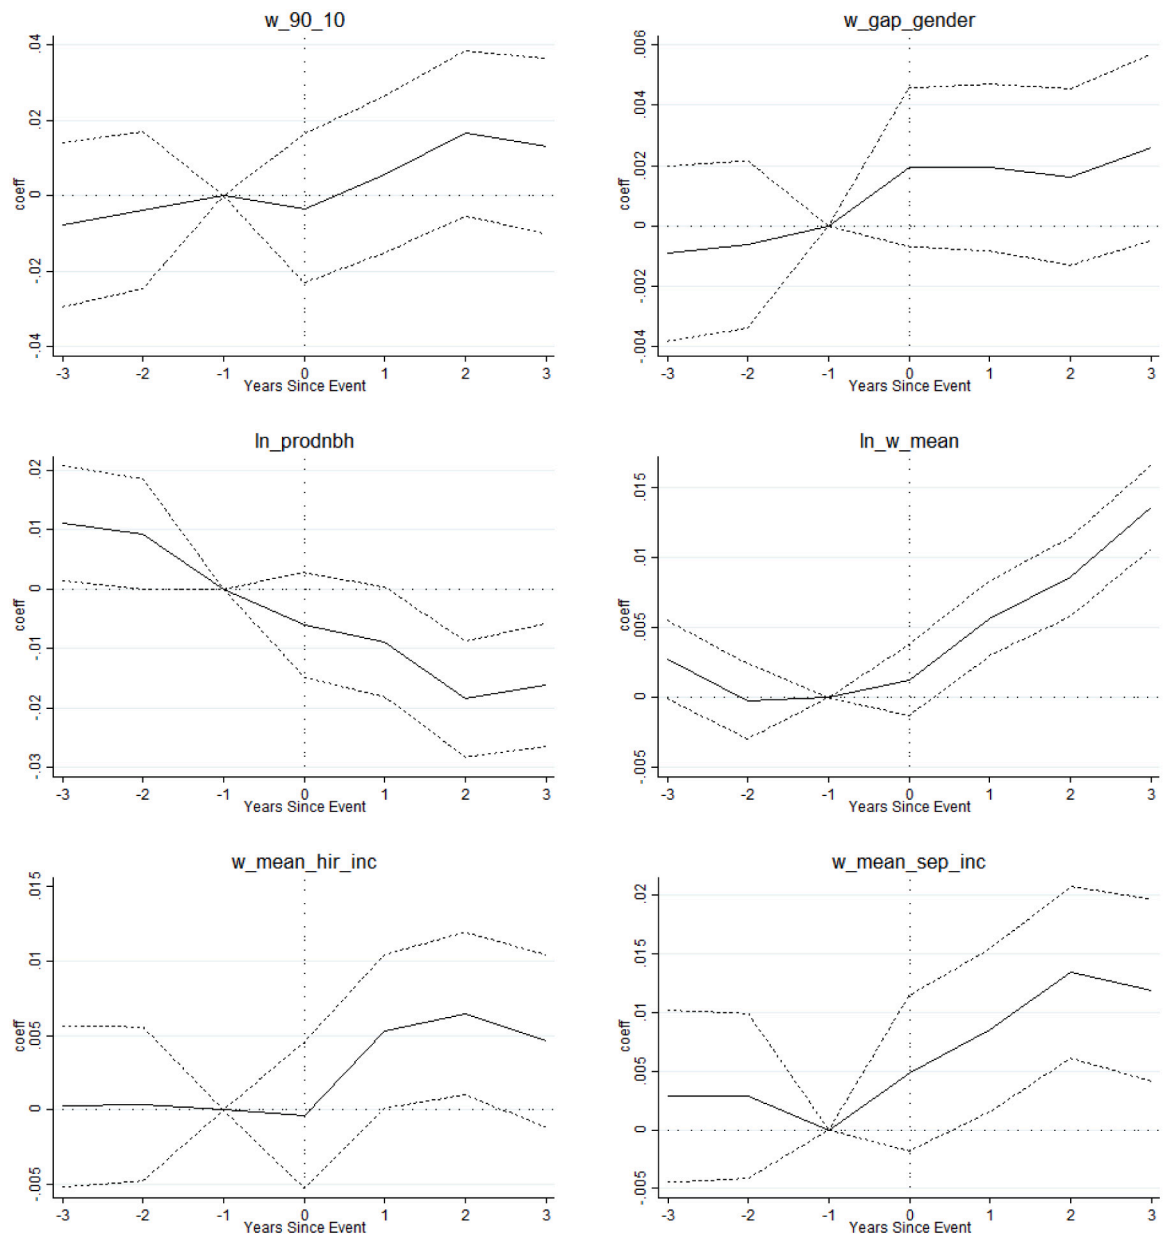


Fig. B.5. Robustness check: Services only. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

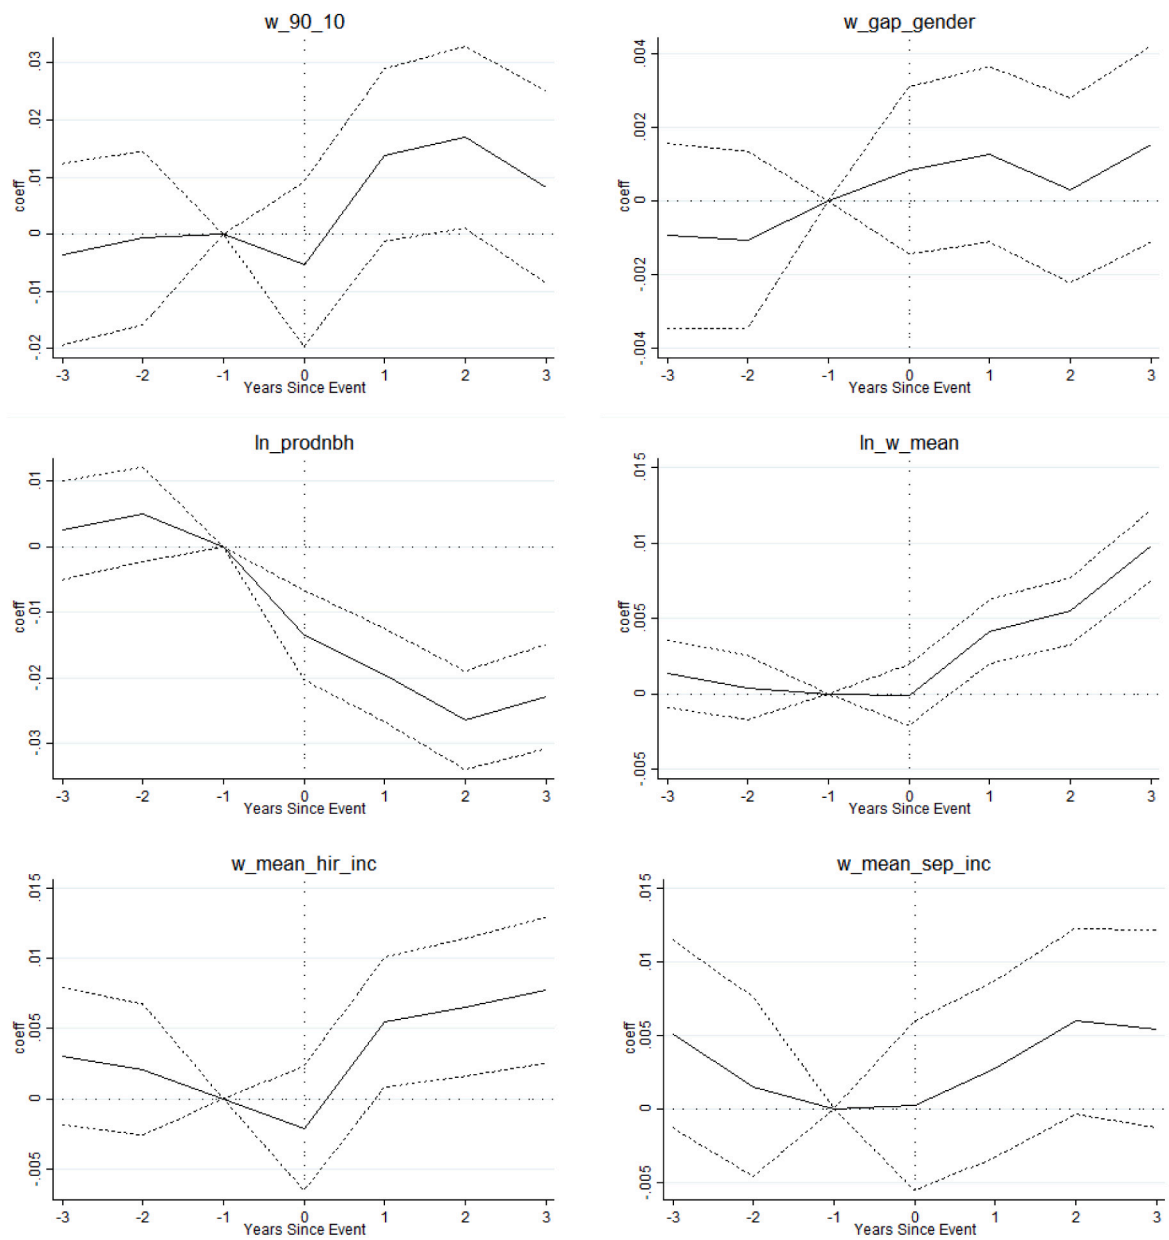


Fig. B.6. Robustness check: Excluding re-exporters. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

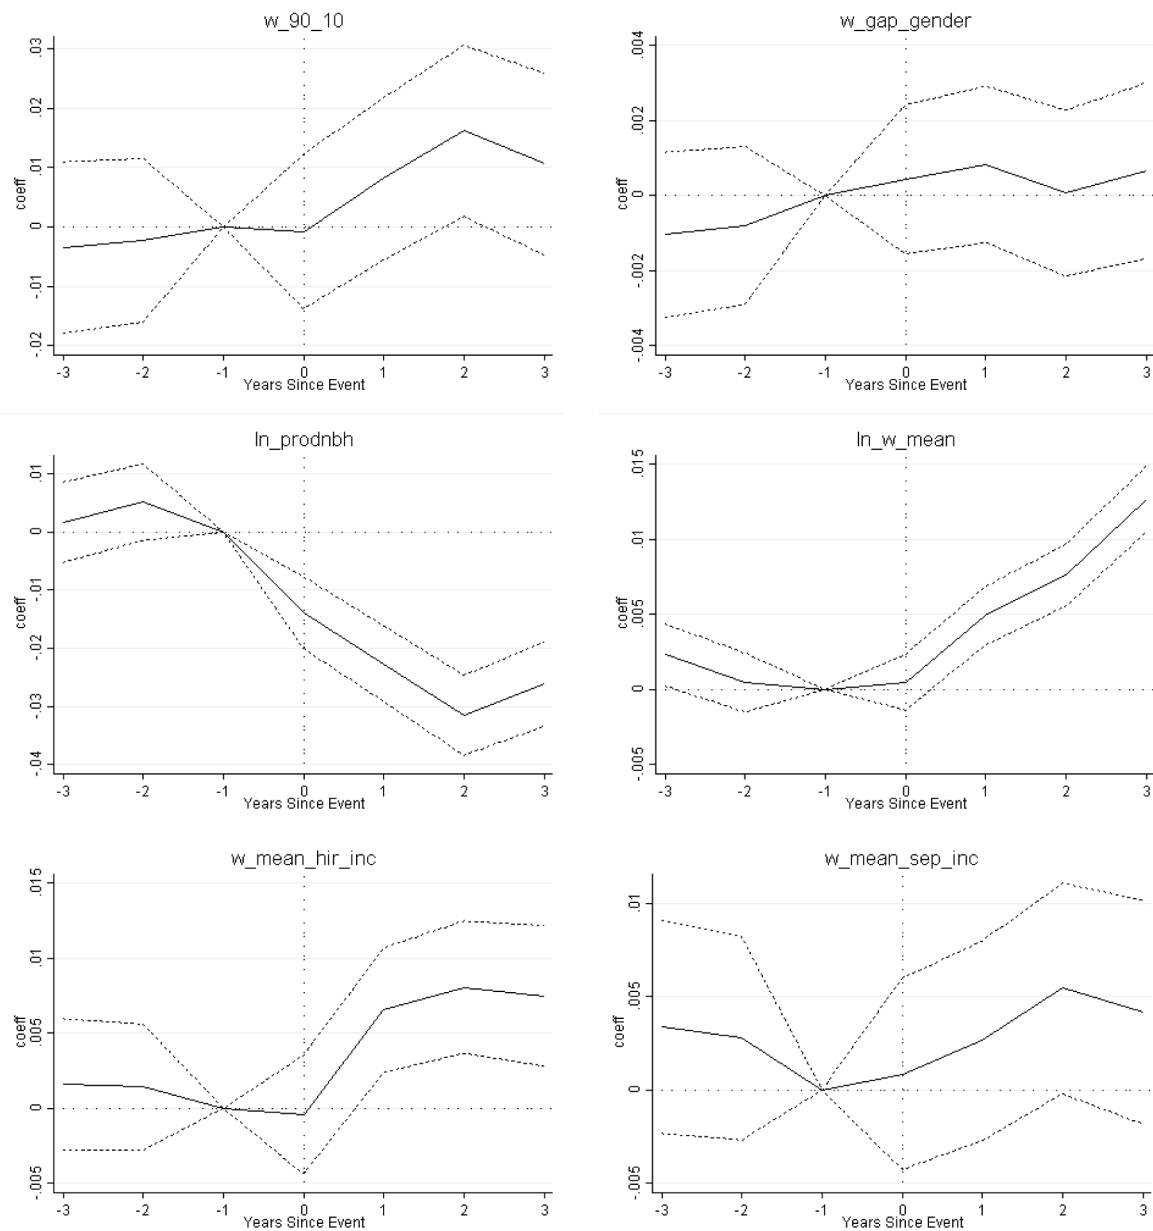


Fig. B.7. Robustness check: Excluding firms that import automation- and AI-related technologies every year. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

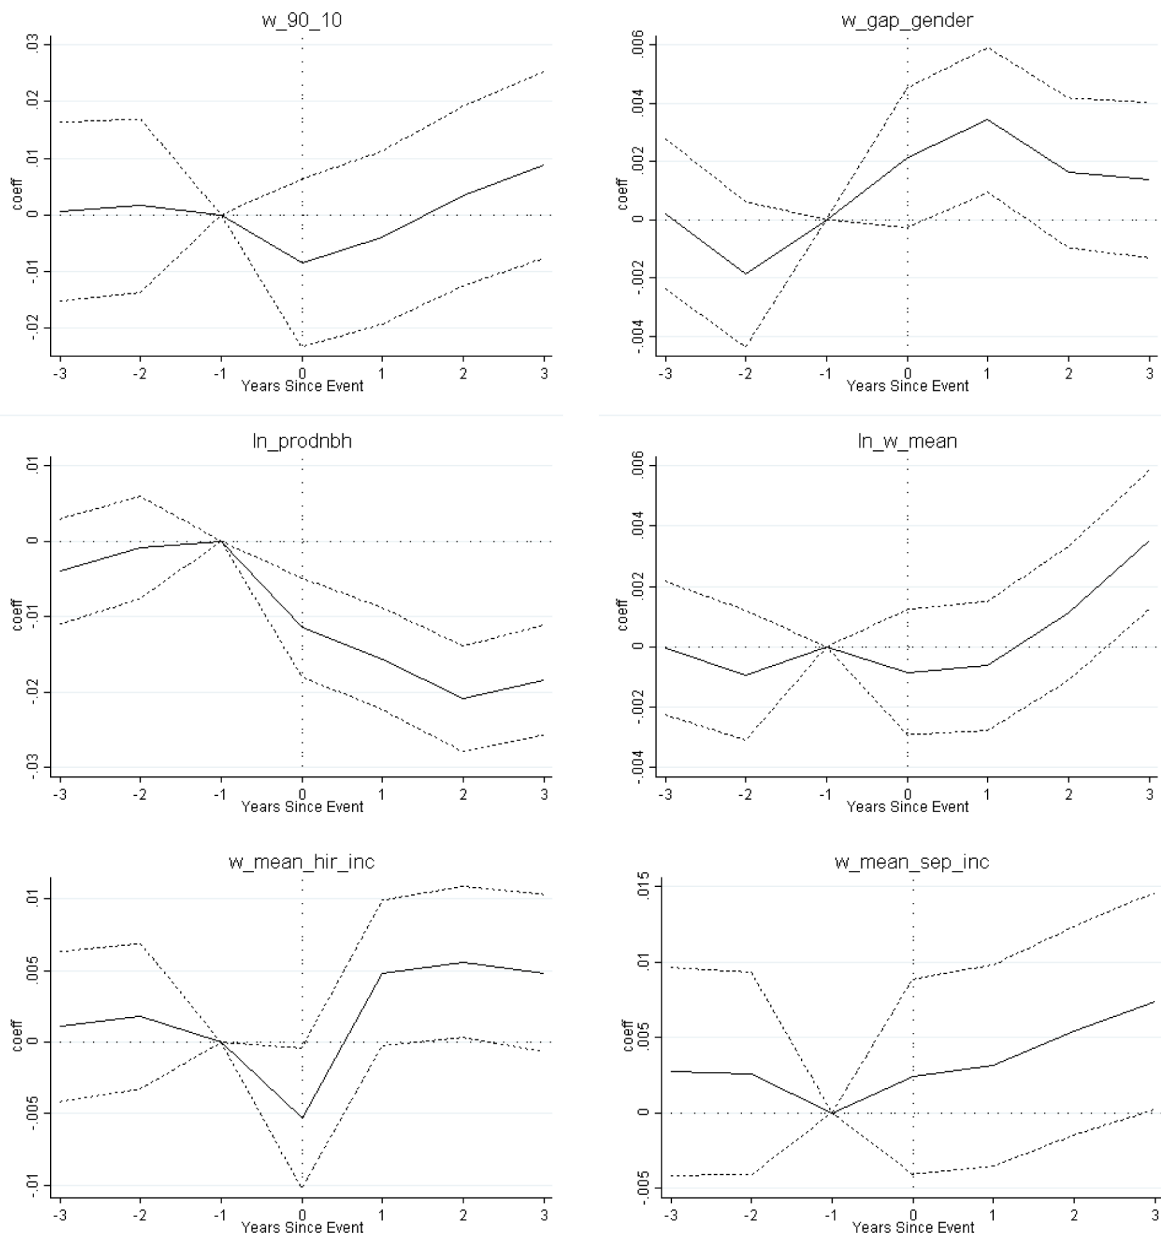


Fig. B.8. Robustness check: Balanced sample. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

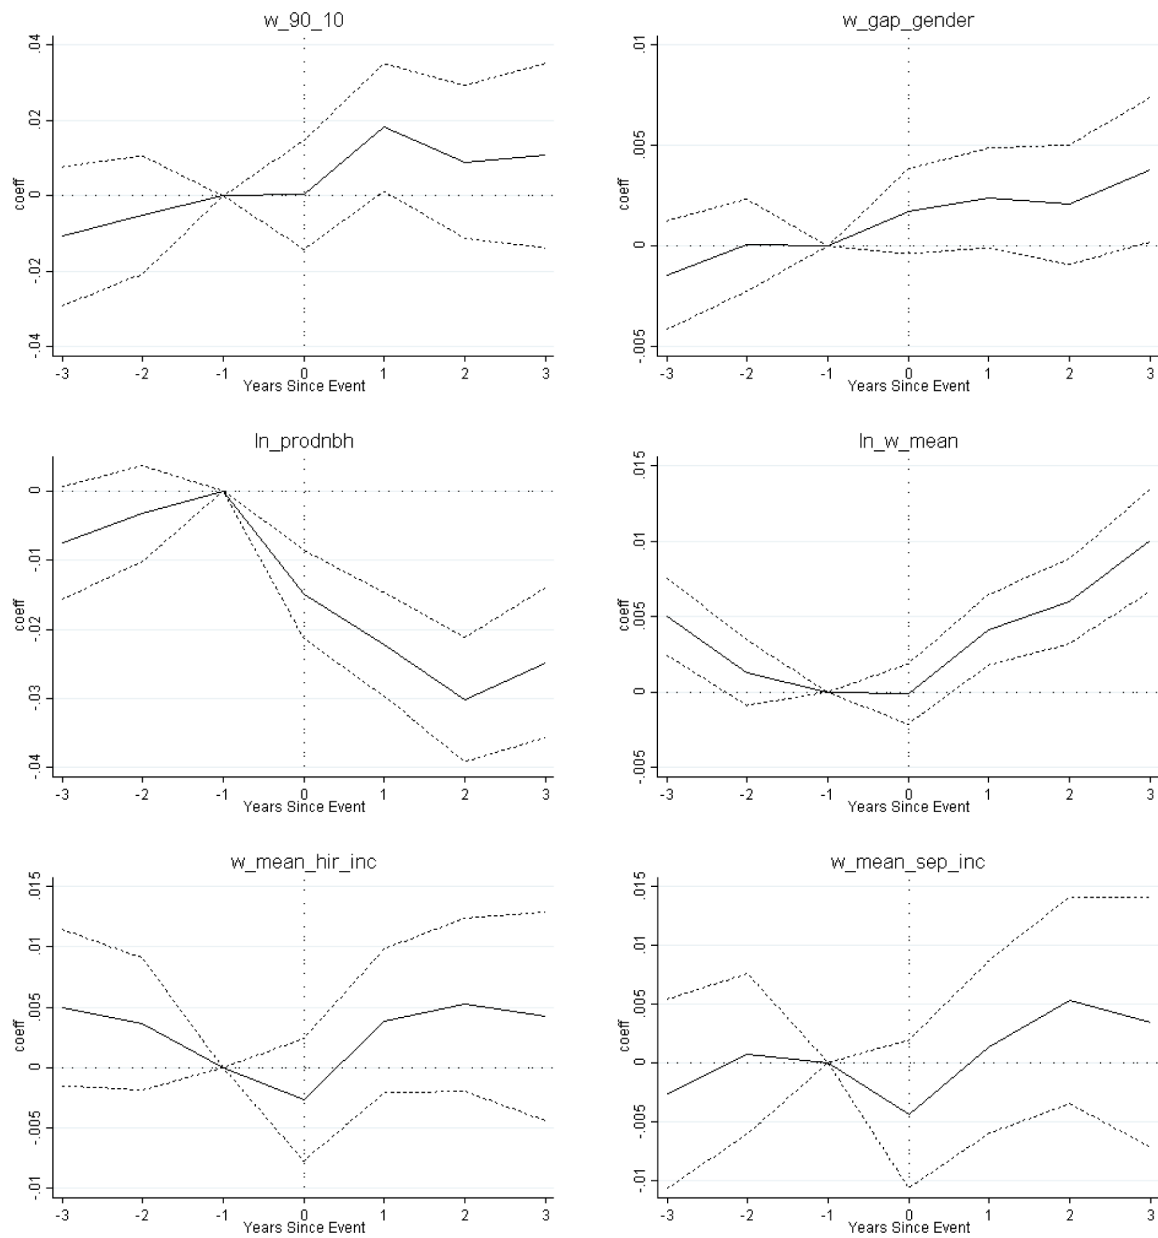


Fig. B.9. Robustness check: Subperiod 2002–2010. Note: The solid line represents coefficients β_{-3} to β_3 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

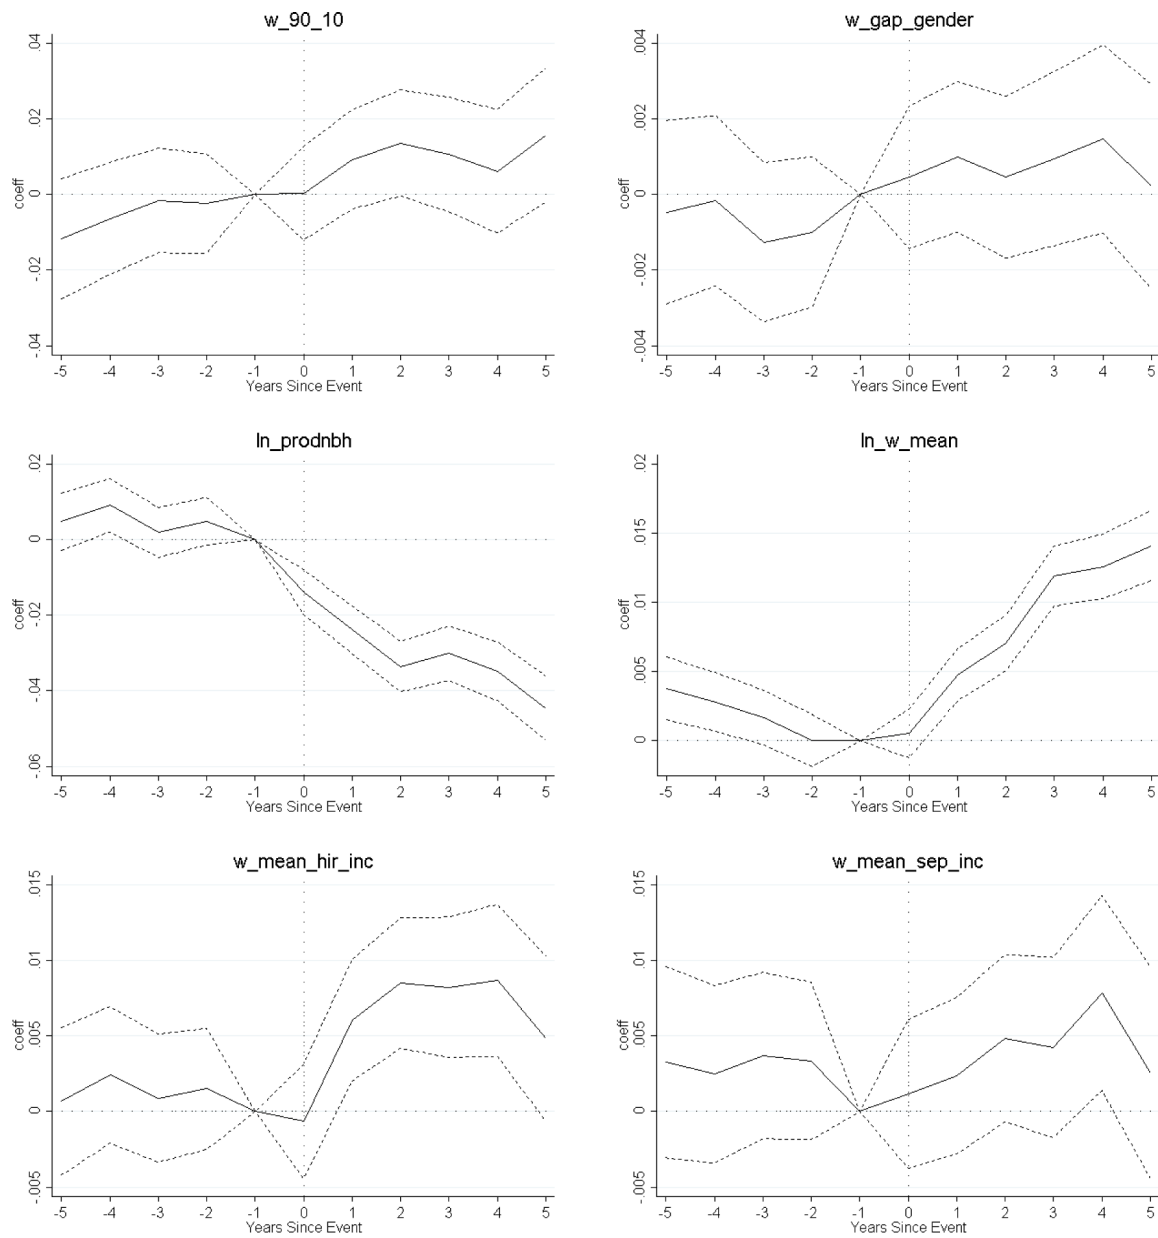


Fig. B.10. Robustness check: Larger time window. Note: The solid line represents coefficients β_{-5} to β_5 from the estimation of Eq. (3), while the dotted line represents the confidence interval at the 5% significance level.

References

- Abeliansky, A.L., Martínez-Zarzoso, I., Prettnner, K., 2020. 3D printing, international trade, and FDI. *Econ. Model.* 85, 288–306.
- Abowd, J.M., Corbel, P., Kramarz, F., 1999a. The entry and exit of workers and the growth of employment: An analysis of french establishments. *Rev. Econ. Stat.* 81 (2), 170–187.
- Abowd, J.M., Kramarz, F., Margolis, D.N., 1999b. High wage workers and high wage firms. *Econometrica* 67 (2), 251–333.
- Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings. In: *Handbook of Labor Economics*, Vol. 4. Elsevier, pp. 1043–1171.
- Acemoglu, D., Lelarge, C., Restrepo, P., 2020. Competing with robots: Firm-level evidence from France. In: *AEA Papers and Proceedings*, Vol. 110. pp. 383–388.
- Acemoglu, D., Restrepo, P., 2018. Demographics and Automation. Working Paper 24421, National Bureau of Economic Research.
- Acemoglu, D., Restrepo, P., 2019. Artificial intelligence, automation, and work. In: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, pp. 197–236.
- Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from US labor markets. *J. Polit. Econ.* 128 (6), 2188–2244.
- Acemoglu, D., Restrepo, P., 2022. Demographics and automation. *Rev. Econom. Stud.* 89, 1–44.
- Aghion, P., Antonin, C., Bunel, S., Jaravel, X., 2020. What Are the Labor and Product Market Effects of Automation?: New Evidence from France. *Sciences Po publications info:hdh:2441/170cd4sul89*, Sciences Po.
- Ahn, J., Khandelwal, A.K., Wei, S.-J., 2011. The role of intermediaries in facilitating trade. *J. Int. Econ.* 84 (1), 73–85.
- Akerman, A., Helpman, E., Itskhoki, O., Muendler, M.-A., Redding, S., 2013. Sources of wage inequality. *Amer. Econ. Rev.* 103 (3), 214–219.
- Aksoy, C.G., Ozcan, B., Philipp, J., 2020. Robots and the Gender Pay Gap in Europe. DP 13482, IZA.
- Asphjell, M., Letterie, W., G.A., P., Nilsen, O., 2014. Sequentiality versus simultaneity: Interrelated factor demand. *Rev. Econ. Stat.* 95 (5), 986–998.
- Autor, D.H., 2015. Why are there still so many jobs? The history and future of workplace automation. *J. Econ. Perspect.* 29 (3), 3–30.
- Autor, D.H., Dorn, D., 2013. The growth of low-skill service jobs and the polarization of the US labor market. *Amer. Econ. Rev.* 103 (5), 1553–1597.
- Autor, D.H., Katz, L.F., Kearney, M.S., 2006. The polarization of the U.S. labor market. *Amer. Econ. Rev.* 96 (2), 189–194.
- Autor, D.H., Katz, L.F., Kearney, M.S., 2008. Trends in U.S. wage inequality: Revising the revisionists. *Rev. Econ. Stat.* 90 (2), 300–323.
- Avouyi-Dovi, S., Fougère, D., Gautier, E., 2013. Wage rigidity, collective bargaining, and the minimum wage: evidence from french agreement data. *Rev. Econ. Stat.* 95 (4), 1337–1351.
- Babečková, J., Du Caju, P., Kosma, T., Lawless, M., Messina, J., Rööm, T., 2010. Downward nominal and real wage rigidity: Survey evidence from European firms. *Scand. J. Econ.* 112 (4), 884–910.
- Balassa, B., 1965. Trade liberalisation and “revealed” comparative advantage. *Manch. Sch.* 33 (2), 99–123.
- Balassubramanian, N., Sivadasan, J., 2011. What happens when firms patent? New evidence from US economic census data. *Rev. Econ. Stat.* 93 (1), 126–146.
- Barth, E., Bryson, A., Davis, J.C., Freeman, R., 2016. It’s where you work: Increases in the dispersion of earnings across establishments and individuals in the United States. *J. Labor Econ.* 34 (S2), S67–S97.
- Barth, E., Roed, M., Schone, P., Umblijs, J., 2020. How Robots Change Within-Firm Wage Inequality. IZA Discussion Papers 13605, Institute of Labor Economics (IZA).
- Bassanini, A., Garnero, A., 2013. Dismissal protection and worker flows in OECD countries: Evidence from cross-country/cross-industry data. *Labour Econ.* 21 (C), 25–41.
- Bergounhon, F., Lenoir, C., Mejean, I., 2018. A Guideline to French Firm-Level Trade Data. Tech. rep.
- Bernard, A.B., Grazi, M., Tomasi, C., 2015. Intermediaries in international trade: Products and destinations. *Rev. Econ. Stat.* 97 (4), 916–920.
- Bernard, A.B., Jensen, J.B., Redding, S.J., Schott, P.K., 2010. Wholesalers and retailers in US trade. *Amer. Econ. Rev.* 100 (2), 408–413.
- Berson, C., Viviano, E., de Philipps, M., 2020. Job-to-Job Flows and Wage Dynamics in France and Italy. Working Paper WP756, Banque de France.
- Bessen, J.E., Goos, M., Salomons, A., Van den Berge, W., 2020a. Automatic Reaction-What Happens to Workers at Firms that Automate? Mimeo.
- Bessen, J.E., Goos, M., Salomons, A., Van den Berge, W., 2020b. Firm-level automation: Evidence from the Netherlands. In: *AEA Papers and Proceedings*, Vol. 110. pp. 389–393.
- Black, S.E., Spitz-Oener, A., 2010. Explaining women’s success: technological change and the skill content of women’s work. *Rev. Econ. Stat.* 92 (1), 187–194.
- Blanchflower, D.G., Oswald, A.J., Sanfey, P., 1996. Wages, profits, and rent-sharing. *Q. J. Econ.* 111 (1), 227–251.
- Blau, F.D., Kahn, L.M., 2017. The gender wage gap: Extent, trends, and explanations. *J. Econ. Lit.* 55 (3), 789–865.
- Blum, B.S., Claro, S., Horstmann, I., 2010. Facts and figures on intermediated trade. *Amer. Econ. Rev.* 100 (2), 419–423.
- Bonfiglioli, A., Crinò, R., Fadinger, H., Gancia, G., 2020. Robot Imports and Firm-Level Outcomes. Cepr dp14593, CEPR.
- Brussevich, M., Dabla-Norris, E., Khalid, S., 2019. Is Technology Widening the Gender Gap? Automation and the Future of Female Employment. Working Paper WP/19/91, IMF.
- Brynjolfsson, E., McAfee, A., 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company.
- Cahuc, P., Postel-Vinay, F., Robin, J.-M., 2006. Wage bargaining with on-the-job search: Theory and evidence. *Econometrica* 74 (2), 323–364.
- Caliendo, L., Monte, F., Rossi-Hansberg, E., 2015. The anatomy of French production hierarchies. *J. Polit. Econ.* 123 (4), 809–852.
- Card, D., Cardoso, A.R., Kline, P., 2016. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Q. J. Econ.* 131 (2), 633–686.
- Card, D., Heining, J., Kline, P., 2013. Workplace heterogeneity and the rise of West German wage inequality. *Q. J. Econ.* 128 (3), 967–1015.
- Carlsson, M., Messina, J., Skans, O.N., 2016. Wage adjustment and productivity shocks. *Econ. J.* 126 (595), 1739–1773.
- Cirillo, V., Sostero, M., Tamagni, F., 2017. Innovation and within-firm wage inequalities: empirical evidence from major European countries. *Ind. Innov.* 24 (5), 468–491.
- Dares, 2018. La Conformité au Smic des Minima de Branches s’est-Elle Améliorée en dix Ans?. Dares Analyses Janvier 2018, n. 005, Ministère du Travail.
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N., et al., 2018. Adjusting to Robots: Worker-Level Evidence. Institute working paper 13, Federal Reserve Bank of Minneapolis.
- Davis, S.J., Faberman, R.J., Haltiwanger, J., 2006. The flow approach to labor markets: New data sources and micro-macro links. *J. Econ. Perspect.* 20 (3), 3–26.
- De Loecker, J., Eckhout, J., Unger, G., 2020. The rise of market power and the macroeconomic implications*. *Q. J. Econ.* 135 (2), 561–644.
- Delahaie, N., Vincent, C., 2021. The SMIC as a driver for collective bargaining. In: *Minimum Wage Regimes: Statutory Regulation, Collective Bargaining and Adequate Levels*. Routledge, p. 162.
- Deng, L., Plümpe, V., Stegmaier, J., 2021. Robot Adoption at German Plants. IWH Discussion Papers 19/2020, Halle Institute for Economic Research (IWH).
- Dinlersoz, E., Kalemli-Ozcan, S., Hyatt, H., Penciakova, V., 2018a. Leverage over the Life Cycle and Implications for Firm Growth and Shock Responsiveness. Working Paper 25226, National Bureau of Economic Research.
- Dinlersoz, E., Wolf, Z., et al., 2018b. Automation, Labor Share, and Productivity: Plant-Level Evidence from US Manufacturing. Ces working paper, Center for Economic Studies.
- Dixon, J., Hong, B., Wu, L., 2019. The Robot Revolution: Managerial and Employment Consequences for Firms. Tech. rep., NYU Stern School of Business.
- Domini, G., Grazi, M., Moschella, D., Treibich, T., 2021. Threats and opportunities in the digital era: Automation spikes and employment dynamics. *Res. Policy* 50 (7), 104137.
- Duggan, M., Garthwaite, C., Goyal, A., 2016. The market impacts of pharmaceutical product patents in developing countries: Evidence from India. *Amer. Econ. Rev.* 106 (1), 99–135.
- EUROSTAT, 2015. *Statistics Explained*. Eurostat, Statistical Office of the European Communities, Luxembourg, <https://ec.europa.eu/eurostat/statisticsexplained/>.
- Fana, M., Giangregorio, L., 2021. Routine-Biased Technical Change Can Fail: Evidence from France. Tech. Rep. 2021/14, JRC Working Papers Series on Labour, Education and Technology.
- Freeman, R.B., Ganguli, I., Handel, M.J., 2020. Within-occupation changes dominate changes in what workers do: A shift-share decomposition, 2005–2015. In: *AEA Papers and Proceedings*, Vol. 110. pp. 394–399.
- Frey, C.B., Osborne, M.A., 2017. The future of employment: How susceptible are jobs to computerisation? *Technol. Forecast. Soc. Change* 114, 254–280.
- Garbinti, B., Goupille-Lebret, J., Piketty, T., 2018. Income inequality in France, 1900–2014: evidence from distributional national accounts (DINA). *J. Public Econ.* 162, 63–77.
- Gardeazabal, J., Ugidos, A., 2005. Gender wage discrimination at quantiles. *J. Popul. Econ.* 18 (1), 165–179.
- Gautier, E., 2017. Les salaires minima de branche en France. *Rev. Fr. Écon.* (1), 94–136.
- Ge, S., Zhou, Y., 2020. Robots, computers, and the gender wage gap. *J. Econ. Behav. Organ.* 178, 194–222.
- Golan, A., Lane, J., McEntarfer, E., 2007. The dynamics of worker reallocation within and across industries. *Economica* 74 (293), 1–20.
- Goldin, C., 2014. A grand gender convergence: Its last chapter. *Amer. Econ. Rev.* 104 (4), 1091–1119.
- Goos, M., Manning, A., Salomons, A., 2014. Explaining job polarization: Routine-biased technological change and offshoring. *Amer. Econ. Rev.* 104 (8), 2509–2526.
- Graetz, G., Michaels, G., 2018. Robots at work. *Rev. Econ. Stat.* 100 (5), 753–768.
- Grazi, M., Jacoby, N., Treibich, T., 2016. Dynamics of investment and firm performance: Comparative evidence from manufacturing industries. *Empir. Econ.* 51, 125–179.

- Guellec, D., Paunov, C., 2020. Digital innovation and the distribution of income. In: Corrado, C., Haskel, J., Miranda, J., Sichel, D. (Eds.), *Measuring and Accounting for Innovation in the 21st Century*. In: National Bureau of Economic Research Studies in Income and Wealth, NBER/ The University of Chicago Press, Chicago.
- Guillou, S., Treibich, T., 2019. Firm export diversification and change in workforce composition. *Rev. World Econ.* 155 (4), 645–676.
- Harris, M., Holmstrom, B., 1982. A theory of wage dynamics. *Rev. Econom. Stud.* 49 (3), 315–333.
- Heinze, A., Wolf, E., 2010. The intra-firm gender wage gap: a new view on wage differentials based on linked employer–employee data. *J. Popul. Econ.* 23 (3), 851–879.
- Helpman, E., Itskhoki, O., Muendler, M.-A., Redding, S.J., 2017. Trade and inequality: From theory to estimation. *Rev. Econom. Stud.* 84 (1), 357–405.
- Huggett, M., Ospina, S., 2001. Does productivity growth fall after the adoption of new technology? *J. Monetary Econ.* 48 (1), 173–195.
- Humlum, A., 2020. *Essays on Automation and Labor Markets*. Princeton University, pp. 1–101, chap. Robot adoption and labor market dynamics.
- Hunt, J., Nunn, R., 2019. *Is Employment Polarization Informative About Wage Inequality and Is Employment Really Polarizing?*. Working Paper 26064, National Bureau of Economic Research.
- ILO, 2020. *Ilostat–International Labour Organization Database Of Labour Statistics*. Tech. rep..
- INSEE, 2010. *DADS Grand Format. Guide Méthodologique. Validité 2010*. INSEE, Direction Régionale de Bourgogne, Dijon.
- Klenert, D., Fernandez-Macias, E., Anton, J.-I., 2020. *Do Robots Really Destroy Jobs? Evidence from Europe*. JRC Working Papers on Labour, Education and Technology 2020–01, Joint Research Centre (Seville site).
- Koch, M., Manuylov, I., Smolka, M., 2019. *Robots and Firms*. Economics Working Papers 2019–05, Department of Economics and Business Economics, Aarhus University.
- Lafortune, J., Rothstein, J., Schanzenbach, D.W., 2018. School finance reform and the distribution of student achievement. *Am. Econ. J.: Appl. Econ.* 10 (2), 1–26.
- Letterie, W.A., Pfann, G.A., Polder, J., 2004. Factor adjustment spikes and interrelation: an empirical investigation. *Econom. Lett.* 85 (2), 145–150.
- Lybbert, T.J., Zolas, N.J., 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 (3), 530–542.
- Miller, C., 2017. The persistent effect of temporary affirmative action. *Am. Econ. J.: Appl. Econ.* 9 (3), 152–190.
- Mishel, L., Bivens, J., 2021. *Identifying the Policy Levers Generating Wage Suppression and Wage Inequality*. Report 215903, Economic Policy Institute.
- Montobbio, F., Staccioli, J., Virgillito, M.E., Vivarelli, M., 2020. *Robots and the Origin of Their Labour-Saving Impact*. IZA Discussion Papers 12967, Institute of Labor Economics (IZA).
- Montornès, J., Sauner-Leroy, J.-B., 2015. Wage-setting behavior in France: Additional evidence from an ad-hoc survey. *J. Bank. Financ. Econ.* 5–23.
- Mueller, H.M., Ouimet, P.P., Simintzi, E., 2017. Wage inequality and firm growth. *Amer. Econ. Rev.* 107 (5), 379–383.
- Nilsen, Ø.A., Raknerud, A., Rybalka, M., Skjerpen, T., 2009. Lumpy investments, factor adjustments, and labour productivity. *Oxf. Econ. Pap.* 61 (1), 104–127.
- Nilsen, Ø.A., Schiantarelli, F., 2003. Zeros and lumps in investment: Empirical evidence on irreversibilities and nonconvexities. *Rev. Econ. Stat.* 85 (4), 1021–1037.
- Oi, W.Y., Idson, T.L., 1999. Firm size and wages. In: *Handbook of Labor Economics*, Vol. 3. Elsevier, pp. 2165–2214.
- Pavlenkova, I., Alfieri, L., Masso, J., 2021. *Effects of Automation on Gender Pay Gap: the case of Estonia*. Working paper, University of Tartu.
- Power, L., 1998. The missing link: Technology, investment, and productivity. *Rev. Econ. Stat.* 80 (2), 300–313.
- Song, J., Price, D.J., Guvenen, F., Bloom, N., von Wachter, T., 2019. Firming up inequality*. *Q. J. Econ.* 134 (1), 1–50.
- Sorgner, A., Bode, E., Krieger-Boden, C., Aneja, U., Coleman, S., Mishra, V., Robb, A.M., 2017. *The Effects of Digitalization on Gender Equality in the G20 Economies: Women20 Study*. Tech. rep., Kiel Institute for the World Economy (IfW), Kiel.
- Sostero, M., 2020. *Automation and Robots in Services: Review of Data and Taxonomy*. JRC Working Papers on Labour, Education and Technology 2020–14, Joint Research Centre (Seville site).
- van der Velde, L., 2020. Within occupation wage dispersion and the task content of jobs. *Oxf. Bull. Econ. Stat.* 82, 1161–1197.
- Webb, M., 2019. *The impact of artificial intelligence on the labor market*. Available at [ssrn](https://ssrn.com).