

For whom the bell tolls: the effects of automation on wage and gender inequality within firms

Giacomo Domini¹, Marco Grazzi², Daniele Moschella³, and Tania Treibich^{*4,3}

¹Erasmus University College, Erasmus University Rotterdam (The Netherlands)

²Department of Economic Policy, Università Cattolica del Sacro Cuore, Milano (Italy)

³Institute of Economics & Department EMbeDS, Scuola Superiore Sant'Anna, Pisa, (Italy)

⁴School of Business and Economics, University of Maastricht (The Netherlands) and Science-Po Paris, OFCE-DRIC, Sophia Antipolices (France)

December 16, 2020

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 period 2002-2017. We document that most of 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 automation/AI spikes are not followed by an increase in within-firm wage and gender inequality. Instead, wages tend to increase at different percentiles of the distribution, revealing a relatively spread allocation of rents from automation/AI within the firm. This adds to previous findings showing picture of a ‘labor friendly’ effect of the latest wave of new technologies.

Keywords: Automation, AI, wage inequality, gender pay gap

JEL classification: D25, J16, J31, L25, O33

*Corresponding author: Tania Treibich. Address: Macroeconomics, International and Labour Economics department, Maastricht University, School of Business and Economics, P.O. Box 616, 6200 MD Maastricht (the Netherlands); Email: t.treibich@maastrichtuniversity.nl.

1 Introduction

Most contemporary societies have witnessed a prolonged widening of the so-called pay ratio, accounting for differences in the compensations of workers at the top and bottom of the wage distribution, both across and within firms. Although recently such gap appears to have somewhat stabilized (ILO, 2016; Pereira and Galego, 2019), over the years it had reached a considerable width. In this respect the trend in the CEO-to-worker pay ratio is quite impressive. From a low of 16-to-1 in 1965, it peaked to 386-to-1 in 2000 (Mishel and Wolfe, 2019), before displaying a slow decrease motivated by the 2009 financial crisis as well as concerns over the effectiveness (Edmans and Gabaix, 2016) and fairness of such compensation schemes. While there is a lot of debate on the CEO-to-worker pay ratio (Adams et al., 2010), and such attention also contributed to the adoption of new policies and disclosure requirements (SEC, 2015), very little is known about other potential sources of inequality that, although not exerting their effects at the two extremes of the pay distribution, will affect most of the workers even within the same occupation and firm.

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 such a significant impact, potentially expanding already existing inequalities or creating new ones. *First*, on the one side, 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 the others Autor and Dorn, 2013; Autor, 2015). Such process could of course expand the wage inequality across occupations, even within the same firm. However, *second*, another process could be at work, too. As put forth in Freeman et al. (2020), who also analyze the impact of AI and robotics on jobs, «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». As such the wage gap could accrue also 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, widening the gap between ‘incumbent’ workers with a long tenure and recently hired employees. *Third*, finally, although most societies are witnessing growing levels of attention to the known gender wage gap, such pay difference continues to be much relevant and it is particularly large in the upper tail of the wage distribution (Blau and Kahn, 2017). As a consequence there exist rising concerns about how new technologies are expected to affect the already existing gender wage gap, even within the same firm; and to date there exists very little evidence to support policy making.

In this work we address such questions by employing matched employer-employee data for France over the period 2002-2017. Adapting to the context the empirical approach developed in Acemoglu and Restrepo (2018), we identify relevant investment episodes in AI and automation through purchases of selected categories of imported capital goods. As shown in Domini et al. (2019), acquisitions of such goods display the typical spiky nature that characterizes investment in capital goods (Grazzi et al., 2016; Nilsen et al., 2009).

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. The descriptive evidence that we provide suggests that most of 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 to the adoption of automation and AI on wage inequality must encompass a focus on the different *within* components.

Employing an event study methodology, we focus on the observed trend in wages and in some measures of wage inequality around an investment spike in automation or AI. We find that employees at firms adopting these new technologies register a small wage increase, and that such positive effect is detectable at most of the percentiles of the wage distribution, with a more pronounced effect in the bottom part. This effect is mostly driven by the fact that firms pay a higher wage to newly hired workers after an automation/AI spike. As a consequence, firm wage inequality first increases and then decreases with an overall negligible and non significant effect after three years. Focusing on the gender pay difference we find that investments in automation and AI do not appear to be associated to a change in the gender wage gap.

Our work builds upon several streams of literature to which we aim to contribute with new empirical evidence. *First*, we contribute to the discussion on wage inequality (see among the many, Autor et al., 2008), and in particular on the effects of automation and AI technologies on wage inequality. To date, a lot of effort has been exerted to predict the potential loss of employment associated to automation and AI technologies, see among the others, Brynjolfsson and McAfee (2014), and Frey and Osborne (2017). So far, the empirical evidence is quite reassuring in suggesting a complementary, more than replacing effect of automation. While aggregate-level studies fail 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 is more consistent in showing a positive effect on the employment of firms that adopt automation (Domini et al., 2019; Koch et al., 2019; Acemoglu et al., 2020; Bonfiglioli et al., 2020; Aghion et al., 2020).¹ However much less investigated is the potential impact of automation and AI on the wage structure across and within firms, a few noteworthy exceptions in this respect being Bessen et al. (2019) and Barth et al. (2020). Bessen et al. (2019) focus on individual workers' outcomes in the Netherlands, finding that after an automation event, incumbent workers are more likely to separate and experience a wage income loss. Barth et al. (2020) is the only study, to our knowledge, which empirically addresses the impact of robots on within-firm wage inequality. In a study of Norwegian firms in the manufacturing sector, the authors find that robots increase wages for high-skilled workers and managerial occupations, thus increasing wage inequality. As explained below, we instead focus on inequality as defined by the wage distribution, not by occupations.

There is also a *second* and complementary element to consider, still pertaining to inequality, that is related to the very nature of this recent waves of technologies. AI and

¹ 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 firms, still within in the same sector, as shown by Acemoglu et al. (2020).

related applications indeed greatly benefit from both almost zero variable costs and network externalities which might easily generate dominant position or quasi-monopoly rents. This is a perspective put forth in 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 more to the top income groups. Although taking a more aggregate perspective and without explicit reference to AI, also De Loecker et al. (2020) detect a generalized increase in market power from 18% above marginal cost in the 1980s up to the current level of 67%. We intend to contribute to this line of research by investigating whether the increasing inequality that is already becoming apparent in the aggregate, is displaying also among workers *within* the same sector, occupation and firm in association with investment in capital goods embedding the newest wave of innovation.

Finally, 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 AI, and more in general the related wave of innovations, on such wage gap and on the job flows as broken down by gender. Among the existing works, Brussevich et al. (2019) investigate the different 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) that take a broader perspective, taking into consideration 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 investments in robots (data are from the International Federation of Robotics) is associated to 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 to an investment within the firm, as data are available at the country, industry and demographic cell. 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 more than that of female workers, thus reducing the gap, the latter increases such difference. In our work, the data and the empirical setting enable to investigate what happens to the gender pay gap both across adopting and non-adopting firms and also, more specifically, within adopting firms.

In the remaining part of this section, taking stock of the theoretical and empirical investigation outlined above, we spell out how the evidence that we provide in the paper is intended to address existing gaps.

We start by noticing that most of the micro literature that has addressed the relationship between automation and wages in general, and automation and wage inequality in particular (Autor et al., 2006; Goos et al., 2014), has done so mostly through the lens of employment polarization, with the partial exception of Bessen et al. (2019). Such perspective, however, is not sufficient to fully highlight the many possible implications that automation can have on wage and wage inequality. Indeed, as shown, among the others, by Hunt and Nunn (2019), Freeman et al. (2020), and van der Velde (2020), occupations cannot fully account for the dynamics of wage, which is relevant also within occupations. It is for this reason that in the empirical

analysis we mostly focus on the individual wage. The fact that a very large share of wage inequality is due to the within component, even inside the same occupation or firm-occupation, confirms that this is a relevant dimension for such analysis.

Focusing on individual wage and through an event study methodology, we then investigate the impact of investment in goods embedding various automation and AI technologies. In this respect, in addition to assessing the impact of adoption of such technologies on the level of wage around such event (as in Bessen et al., 2019; Acemoglu et al., 2020) we aim to understand the implications for inequality, a theme much less investigated so far. To this end, we both study the impact of AI and automation at different percentiles of the wage distribution and on the gender pay ratio.

Finally, a last broad research question concerns the possible sources of the changes in the wage distribution due to automation and AI. While to date there is no wide consensus about the magnitude of the productivity gains brought about by automation and AI, the modest increase that is generally reported² suggests that, at least so far, there is little room for a substantial rent-sharing dynamics, in which the firms pass-through some of the efficiency gains to their workers. In this respect, an hypothesis that we investigate here, also following up the evidence in Domini et al. (2019), is the possibility that part of the change in wage that follows an investment in automation is due to the hiring of new employees as part of the employment expansion that generally follows an event of automation.

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 the wage distribution, including an analysis of variance that decomposes the overall wage inequality in different components. We also show trends in wage inequality and introduce our Automation and AI measure. Section 4 presents the Event Study framework and discusses the results. Section 5 concludes.

2 Data and variables

2.1 Sources

Our dataset contains information on all French firms with employees over the period 2002-2017, obtained by merging two 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 hand in 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.³ We extract from DADS the following worker-level variables: gross yearly remuneration, number of hours worked, age, gen-

² In Graetz and Michaels (2018) and Acemoglu et al. (2020), employing different data and methodologies, they are estimated, respectively, in the order of 0.36% and 4%, in the weighted sample.

³ Notice that DADS *Postes* does not allow tracking workers over time, since the worker identification number is not consistent across years.

der, and occupation;⁴ as well as the sector of the firm defined according to NAF rev. 2 classification (corresponding to the European NACE rev. 2).⁵

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 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 total value of yearly imports per product category.

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.⁶ For each worker, we divide the gross yearly remuneration by the number of worked hours to obtain the hourly wage.⁷ 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 mean and standard deviation, as well as percentiles (p10, p50, p90). In the regressions, we will use the log transformation of the level variables (mean wage and wage percentiles) in order 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 both in the macro and in the micro economic literature (see Cirillo et al., 2017; Mueller et al., 2017); the standard deviation is also chosen as it reflects an overall measure of 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

⁴ 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 from DADS worker-level variables on the ‘type of job’, which allows us identifying apprentices and cleaning them out, and on the start and end dates of job posts, necessary to identify workers present at a specific date (see Subsection 2.3).

⁵ 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. (2019, 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 occurrence.

⁶ Notice that, while plant-level information is available in DADS, we need to focus on the firm level, in order to match DADS data with firm-level customs data.

⁷ We deflate wages (as well as imports; see below) using yearly value-added deflators for 2-digit NAF divisions provided by the INSEE.

male workers. Likewise, we calculate gender ratios at various percentile, e.g. the ratio between the median female hourly wage and the median male hourly wage.

An important note on our definition of gender wage inequality is in order here. Since we normalize the wage by the number of hours worked, and we only consider employed persons, two important sources of income inequality between men and women are removed, as, particularly in France, females are most affected by part-time work, yielding lower monthly wages (based on the ILOSTAT data, around 50% of female work is part-time during our period of study, while only 30% of male work is, ILO, 2020). As a consequence, if the gender pay gap in France is estimated at 15.5%, right at the EU-27 average, the gender overall earnings gap is exactly the double, at 31% EUROSTAT (2015).

Notice also that in our regression exercise (see Section 4), in order to account for the effect of the age and gender of workers on firm wages, we control for the share of female workers and for the average age of workers within a firm. Both variables are based on worker-level indicators (age and sex) from DADS.

Adoption of automation and AI-related technologies

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 robotisation and automation in general at the firm level (Domini et al., 2019; Dixon et al., 2019; Bonfiglioli et al., 2020; Aghion et al., 2020; Acemoglu et al., 2020), as it allows overcoming the lack of systematic firm-level information on the adoption of digital and automation technologies at the firm level, which is only recently starting to be collected by national statistical offices.⁸ Trade flows reported by firms to customs offices offer a handy solution to this, as fine product-level decomposition allows identifying the adoption of specific technologies via the imports of related goods.

In particular, 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 (2018), 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 the others, industrial robots, dedicated machinery, numerically-controlled machines, and a number of other automated capital goods.⁹ To the automation-related categories listed by Acemoglu and Restrepo (2018), we add 3-D printers, the HS code of which is identified by Abeliatsky et al. (2020). Besides 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.¹⁰ Considering

⁸ Notably, the Dutch statistical office (CBS) includes a question on automation costs in their national survey (see Bessen et al., 2019).

⁹ For a full list, including the specific 6-digit HS codes falling under each of the above-mentioned categories, see the Appendix.

¹⁰ One of the criteria for the choice of these categories is that many of the codes belonging to them match to the US patent classification (USPC) code 706 ('Data processing - Artificial Intelligence'), via the USPC-to-HS 'Algorithmic Links with Probabilities' (ALP) concordance by Lybbert and Zolas (2014). We then extend the selection using our own expertise.

AI-related imports, in addition to automation-related ones, is important for our measure to be 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 firm, *vis-à-vis* one-half of automation-related imports.¹¹

Some potential limitations of our import-based measure of adoption of automation- and AI-related technologies should be acknowledged and discussed. First, firms might purchase automation- or AI-related goods domestically, instead of internationally, and thus they may be wrongly labelled as non-adopters in our analysis. With respect to this, notice that France has a revealed comparative disadvantage (cf. Balassa 1965) and a negative trade balance for the goods that compose our measure;¹² 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: firms that are only active in the domestic market may buy automation- and AI-related technologies from domestic suppliers (though unlikely, as argued above); plus, the impact on their wage dynamics may be different, as they tend to be smaller and less productive on average than firms involved in international trade. Finally, 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.

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 at a specific date of each year (31 December). 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.¹³ 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.¹⁴ We also remove apprentice workers, which represent around 3.5% of

¹¹ Based on our calculations on DGDDI data for the year 2017. Detailed figures are available upon request.

¹² Based on calculations by the authors on 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 7 in the Appendix. A notable exception is the category of robots, as well as that of regulating instruments, which however represent a minority of the measure.

¹³ 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 hours worked per day; or the net salary should be more than three times the monthly minimum salary; else, it is classified as *annexe*.

¹⁴ 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*), 24% a temporary or

Table 1: Sample composition and relative size, 2002-2017

	Nb. obs	Nb. firms	Share in nb. of firms	Share in employment
All firms	20,231,242	3,204,497	1	1
Sample 1	2,726,445	291,139	9.08	54.50
Sample 2	1,140,139	96,128	3.02	51.79
Sample 3	506,893	40,087	1.25	37.24

Source: our elaborations on DADS and DGDDI data. Sample 1: all importing firms; Sample 2: Importing firms above 10 employees. Sample 3: Firms importing automation and AI related goods at least once, above 10 employees.

observations, workers with less than 120 hours worked in the year,¹⁵ and workers with wage below half of the minimum wage, which represent less than 1% of observations. Figure 7 in the Appendix shows that this bottom threshold to the wage per hour variable really 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 labelled as ‘household employers’ (*particuliers employeurs*) and the public administration (*fonction publique*) in years 2009-2017, since they are not available in earlier years. This yields a sample of more than 20 million firm-year observations over the period 2002-2017, referring to 3 million unique firms (see Table 1, row 1, ‘All firms’).

However, in our analysis we need to restrict the sample because of a number of reasons. First, we can construct our measure of adoption of automation- and AI-related technologies only for importing firms (see Section 2.2), which represent 9% of observations in the overall data, but account for more than 50% of total employment (see Table 1, sample 1). Second, in order to ensure that within-firm statistics on the wage distribution are meaningful, we restrict the attention to importing firms with at least 10 employees (sample 2). This threshold excludes ‘micro-firms’, according to the Eurostat definition. Notice that this further restriction reduces quite much 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). This final sample includes only around 40 thousand firms, but still 7.5 million workers. In the following section, presenting descriptive statistics, we will refer to different samples; while in the regression analysis

placement contract (*mission*). After one year, 18% of them becomes principal, 26% stay annexes, the rest (56%) leave the firm.

¹⁵ This matches one of the thresholds used for defining non-annexe workers. Note that this also removes workers with zero hours.

(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 - which dimensions matter?

In what follows we present some descriptive statistics to motivate 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). We then 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 on the differences between adopting and non-adopting firms.

The wage distribution of workers

Figure 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 very much impacted by the minimum wage around 10 euros per hour, and therefore very positively skewed and with high kurtosis. Notice 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 organisation of the firm and remuneration of value across workers). In order to motivate our study of within-firm wage inequality, we perform a decomposition exercise which compares the contribution of both dimensions to the overall wage inequality among workers, as shown in Figure 1.

Decomposing wage inequality


We perform two simple variance decomposition exercises. Formally, we decompose the overall wage inequality T_t among all workers into a between B_t and a within component W_t :

$$T_t = W_t + B_t$$

and we then compute the share of total variance accounted for by the within component, i.e. W_t/T_t . In these exercises we exploit worker-level information on hourly wage, their occupation (managers and white-collars; supervisors and technicians; clerks; skilled production workers; unskilled skilled production workers; residual workers),¹⁶ the firm where the worker is employed, and the sector of that firm (defined at the 2-digit level of the NAF classification).

¹⁶The first three categories are defined at the 1-digit level of the French taxonomy of occupations (*Catégories Socio-professionnelles*), respectively as codes starting by 3, 4, and 5; while skilled and unskilled production workers are defined at the 2-digit level, respectively as codes starting by 61-65 and by 66-68).

Figure 1: Distribution of wage per hour among all workers, 2017.



Source: our elaboration on DADS data.

Table 2 shows the share of overall wage inequality accounted for by the within component at different levels of disaggregation, namely within sectors, within occupations and within sector-occupations. Notice 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, within occupation, and within sector-occupation components account for the majority of wage inequality in France in 2017 in all samples. 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, around half of wage inequality occurs among workers belonging to the same occupational category (even within the same sector).

This picture is consistent among the different samples: hence, in the sample that will be used in our regression analysis (sample 3), the main forces driving wage inequality are the same as in the whole population of firms. The result also confirms 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 motivates our approach to use 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 within firm component vs. between firm component. Among the population of workers within

Table 2: Within-sector, within occupations, and within-sector-occupation shares of wage inequality, 2017.

	(%) 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

Source: our elaborations on DADS and DGDDI data. Sample 1: all importing firms; Sample 2: Importing firms above 10 employees. Sample 3: Firms importing automation and AI related goods at least once, above 10 employees.

Table 3: Decomposition of within-sector and within sector-occupation wage inequality, 2017.

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

Source: our elaborations on DADS and DGDDI data. Sample 1: all importing firms; Sample 2: Importing firms above 10 employees. Sample 3: Firms importing automation and AI related goods at least once, above 10 employees.

each sector, on average,¹⁷ 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 mainly defined by market conditions, i.e. the relative positions of firms in the product market, or by different characteristics among firms (e.g. firm size). In the other samples that consider importing firms (samples 1-3), this share is even greater, around 75%: the reason is that within-firm dispersion of wages is larger in large firms. 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%. Instead, labor market conditions (at the sector-occupation level defined here) are not the main driver of wage inequality.

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 where he/she is employed.¹⁸


¹⁷ The numbers in Table 3 are computed for each sector/sector-occupation separately and then aggregated by taking an employment weighted average.

¹⁸ The within-firm component has been found to be a sizable element of wage inequality in other studies too. See, for instance, Helpman et al. (2017) for Brazil and Song et al. (2019) for the U.S.

Trends in wage inequality within firms

Having determined that the within-firm dimension is crucial in understanding overall wage inequality, we analyze here trends in firm-level wage inequality. Figure 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, 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. For what concerns 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 around 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. This is an impressive change which doesn't reflect the national trend in the mean gender wage gap, which shows no evolution since 2002 (also see EUROSTAT, 2015).

Figure 2: Evolution of wage characteristics over time, sample 3, 2002-2017.



Source: our elaborations on DADS and DGDDI data. Sample 3: Firms importing automation and AI related goods at least once, above 10 employees.

3.2 Automation and AI imports

We provide here some information to characterise 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 that are most active in buying automation and AI-intensive goods in our trade data. This is measured by comparing the share a sector accounts for in total French automation and AI imports (central column) and the same sector's share in aggregate employment. The electronics (NAF

rev. 2 division 26), machinery (28), and automotive sectors (29) are disproportionately represented in automation- and AI-related imports, relative to their employment. The retail sector (46) is an incredible case with 55.1% of those investments, more than six times its share in total employment (9.3%).¹⁹

Table 4: Sectors with automation and AI share larger than their employment share, sample 3, 2017.

Sector	A88	Automation and AI share (%)	Employment share (%)
Metal products	25	3.1	2.8
Electronics	26	3.6	1.5
Electric	27	1.5	1.3
Machinery	28	5.1	2.0
Automotive	29	5.6	2.3
Other transport	30	2.1	1.8
Furniture	31	0.5	0.4
Retail	46	55.1	9.3
IT	62	4.2	2.7

Source: our elaborations on DADS and DGDDI data.

The statistical properties of automation

When looking at the statistical properties of automation- and AI-related imports, it can be observed, as already done for automation only in Domini et al. (2019) and Bessen et al. (2019), that they display the typical *spiky* behaviour 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: around 14% of importers import automation- or AI-related goods in each year, and less than half of them do it at least once over the 2002-2017 period. Second, such imports are rare within firms: among firms who import automated 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 who import automated goods in all years. Finally, the largest events of import of such goods represent a significantly high share of a firm’s total across years: when ranking from largest to lowest the shares for which each year’s imports account over a firm’s cumulative imports across years, it is apparent that the top-ranked import event displays a predominant share (around 70%), while the shares of lower ranks rapidly decrease in value.²⁰ Because of the very skewed nature of this variable within firms, we define as an *automation/AI spike* the largest event for each firm.²¹

¹⁹ To account for such important outlier, we also run the regressions separating manufacturing and services, though not included in the results due to space constraints; results are available from the authors upon request.

²⁰ These statements are based on Figure 8 in the Appendix.

²¹ For a more detailed discussion of the statistical properties of automation-related imports, including a comparison to general physical investment, see Domini et al. (2019, Section 3).

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 in the literature (Domini et al., 2019; Koch et al., 2019), firms adopting automation and AI are larger and more productive (as reflected here in the mean wage per hour value) than non-adopting firms. Such difference in the wage level is present at all locations of the wage distribution, and more pronounced at its top. We also show that they 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 as well reflect self-selection into automation and AI adoption. Such selection effect will be tackled in our empirical strategy by considering only firms that automate in our event-study analysis.

Table 5: Comparing firms with and without automation or AI spike, sample 2, all years (2002-2017).

	No spike	Spike	T-test
Number of observations	633,246	506,893	
Number of firms	56,041	40,087	
Number of employees	55.38	177.09	***
Wage per hour (mean)	18.19	20.49	***
Wage standard deviation	8.68	10.73	***
90-10 wage ratio	2.38	2.53	***
Female-to-male wage ratio	0.881	0.840	***
Wage per hour (p1)	10.31	10.45	***
Wage per hour (p10)	11.77	12.60	***
Wage per hour (p50)	15.68	17.49	***
Wage per hour (p90)	28.18	32.11	***
Wage per hour (p99)	46.82	58.86	***
Female-to-male wage ratio (p1)	1.07	1.06	***
Female-to-male wage ratio (p10)	1.01	0.98	***
Female-to-male wage ratio (p50)	0.94	0.91	***
Female-to-male wage ratio (p90)	0.83	0.79	***
Female-to-male wage ratio (p99)	0.73	0.65	***

Source: our elaborations on DADS and DGDDI data. ***: significant difference at 1% level. Sample 2: Importing firms above 10 employees.

The next step is to consider a dynamic approach, 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 event (in time $t = 0$) and that we also observe in the three years before and three years after. Within this subgroup of 16,582 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 wage at all the levels tested here, while the correlation between the spike event and wage inequality is ambiguous (both measures yield opposite trends). Finally, the gender pay gap seems to slightly decrease, especially at the 90th percentile of the wage distribution.

Table 6: Wage characteristics around an automation or AI spike, balanced panel within sample 3.

Years since spike	Wage per hour	Wage standard deviation	90/10 wage ratio
-3	19.344	10.330	2.558
-2	19.558	10.415	2.538
-1	19.822	10.533	2.521
0	20.100	10.623	2.494
1	20.335	10.598	2.477
2	20.615	10.781	2.474
3	21.018	10.755	2.455

Years since spike	Wage per hour (p10)	Wage per hour (p50)	Wage per hour (p90)
-3	11.837	16.365	30.563
-2	11.964	16.549	30.839
-1	12.189	16.765	31.236
0	12.397	17.039	31.433
1	12.594	17.318	31.714
2	12.762	17.559	32.102
3	13.113	18.056	32.410

Years since spike	Gender wage ratio (p10)	Gender wage ratio (p50)	Gender wage ratio (p90)
-3	0.983	0.904	0.773
-2	0.982	0.907	0.782
-1	0.986	0.9011	0.785
0	0.986	0.915	0.797
1	0.986	0.917	0.804
2	0.986	0.909	0.808
3	0.988	0.921	0.812

Source: our elaborations on DADS and DGDDI data. Note: 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 16,582 firms. Sample 3: Firms importing automation and AI related goods at least once, above 10 employees.

4 The effect of automation and AI on wages: Event study analysis

4.1 Empirical approach

Automation/AI spikes represent a single, major event for French importing firms during the 2002-2017 period in which we observe them (see Section 3.2). This characteristic makes it suitable to investigate the relationship between automation and wages within an event study framework. Such a methodology was used by Bessen et al. (2020) 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:

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

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 ; X_{it} is a set of control variables that include the average age of the workforce and the share of female workers for firm i at time t ; δ_i and ζ_{jt} are respectively a set of firm and sector-years fixed effects, and, finally, ε_{it} is the error term.

β_k represents the effect of the automation 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 value at which the index is censored (i.e. k_{min} and k_{max}) usually depends on the kind 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 $k = -1$. Equation 1 is thus a flexible tool to study the timing of the effects of automation/AI. In order 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, though 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 variation in the timing of the spike across firms, and 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 give a causal interpretation to 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. (2019) and Bessen et al. (2020). Keeping only treated firms makes it more likely the assumption that they are on parallel trends,

especially given the large differences observed between the groups of firms with and without a spike (see Table 5 above). On the other hand, a useful characteristic of our event study is that it builds 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 will be more plausible to assume that results are not driven by pre-spike differences across firms. In any case, given the non-random nature of an automation event, one should still be cautious about causal interpretation of our results. For this reason, we will interpret the coefficients mostly as describing the evolution of firm outcomes around the spike, as in Bessen et al. (2020).

4.2 Results

We will now discuss the results of the estimation of Equation 1, as displayed in Figures 3 to 6. In all of them, we plot the coefficients β_k , from β_{-3} to β_3 ; dashed lines represent confidence intervals at 5% and 10%. All the regressions are performed on the number of observations and firms of sample 3, as reported in Table 1.

Wage inequality


In Figure 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 and the standard deviation of hourly wages within a firm. In Figure 3 (left), the β_k coefficients are not significant: the 90/10 ratio stays almost unchanged after the spike. A similar result emerges when considering the standard deviation of hourly wages: the change after an automation/AI spike is not significant at the different time horizons (see Figure 3, right).

Our result adds a further piece of evidence to the emerging, but still scant, literature on automation, AI and within-firm wage inequality. In Barth et al. (2020) robots are found to increase the wage of highly educated workers with respect to low educated workers, and managers and professionals with respect to other occupational categories within Norwegian manufacturing firms.²² Our result conveys a different message. By focusing on a more general measure of wage inequality within firms, we find that inequality is unaffected after an automation event. The difference could be due to the different institutional context of the French labor market. Another possible explanation is that in our case adoption did not increase wages to begin with: we will see below that this is not the case. Finally, notice that we performed the same regressions only on manufacturing (and separately on non-manufacturing) firms, without finding any significant difference in the evolution of within firm inequality measures after the automation/AI spike.²³

²² A positive correlation between innovation and within-firm wage inequality is also found in Cirillo et al. (2017) where, however, a general R&D innovation proxy is considered.

²³ Results are available from the authors upon request.

Figure 3: Automation/AI spikes and within-firm wage inequality.



Note: the solid line reports coefficients β_{-3} to β_3 from the estimation of Equation 1, while the blue and red dotted lines represent confidence intervals at the 5% and 10% significance level, respectively.

Gender wage gap


While we don't find a significant relationship between automation and overall firm-level wage inequality, 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 be related to firm-specific characteristics, like size and bargaining regimes (Oi and Idson, 1999; Heinze and Wolf, 2010; Card et al., 2016) as well as, more in general, to the extent to which firms reward job-related characteristics like temporal flexibility (Goldin, 2014).

To date, no available evidence has yet been produced on the direct effect of automation on such gender gap. We turn now to test this hypothesis by separately estimating Equation 1 for our gender wage gap measure (the ratio of female-to-male wage) computed at different percentiles of the wage distribution. This takes into account the evidence 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 Figure 4. We plot the β_k coefficients from Equation 1 where the dependent variable is the ratio between the female and the male wages within the 10th, the 50th and the 90th percentiles. The ratio seems to stay almost unchanged after the spike; a small change is detectable at the 90th percentile, but it is barely significant.

Wage increase at percentiles

Having established that within-firm wage inequality and the gender wage gap do not increase following an automation event, the question remains open whether this is simply the effect of a disconnect between automation and wage dynamics or whether, on the contrary, wages increase following an automation spike in a fairly equal way across workers. We try to settle this question in Figure 5. There, we report results where y_{ijt} represents the mean and different percentiles of the within-firm wage distribution. Variables are taken in log so that coefficients can be interpreted as percentage changes

Figure 4: Automation/AI spikes and the gender wage gap.



Note: the solid line reports coefficients β_{-3} to β_3 from the estimation of Equation 1, while the blue and red dotted lines represent confidence intervals at the 5% and 10% significance level, respectively.


with respect to the value of y_{ijt} one year before the spike.

The first plot of Figure 5 (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 first not significant (in the year of the spike), then significant with an increasing trend. Overall, the effect is small: three years after the spike the mean wage is around 0.6% higher than before the spike. A modest relationship between hourly wage and robot adoption is observed in Acemoglu et al. (2020) in a smaller sample including around 600 French robot adopters; similarly, Bessen et al. (2020) find that firm-level mean wage increases after an automation spike in a sample of Dutch firms.

Such an increase in hourly wage seems to be due to an increase at different percentiles of the distribution. In Figure 5 (top, right; bottom, left), it is apparent that three years after the spike, hourly wages are around 0.7% - 0.8% higher for the 10th and 50th percentiles. The effect is slightly smaller at the 90th percentile (Figure 5, bottom, right), where hourly wages are around 0.5% higher two years after the spike, but then shows a small decline such that they are only 0.3% higher three years after a spike. Overall, we can conclude that following an Automation/AI spike there is a general increase in workers' wage three years after the event; such an increase is equally distributed across the wage's percentiles, reinforcing the message coming from the previous exercises that there is no change in within-firm inequality after such an

event.

Figure 5: Automation/AI spikes and the within-firm wage distribution.



Note: the solid line reports coefficients β_{-3} to β_3 from the estimation of Equation 1, while the blue and red dotted lines represent confidence intervals at the 5% and 10% significance level, respectively.

Newly hired workers

In all the previous exercises, there was no evidence that coefficients before the spike were significant, suggesting that our results are not driven by differences in pre-spike trends. However, it could be that new workers with different skills and thus with different wages are selected into the firm in anticipation of the automation event. Such case could partly make the interpretation of our coefficients less clear.²⁴

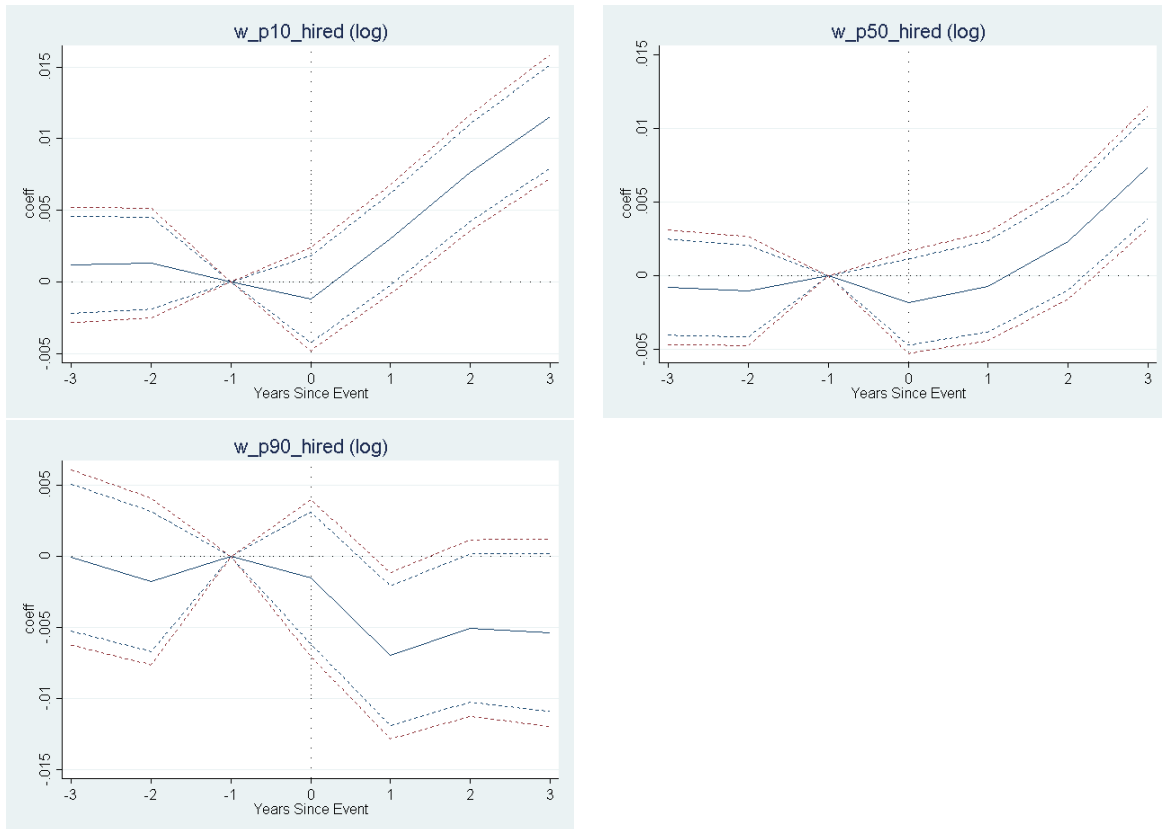
In order to check whether something similar is happening, we investigate the effects of automation focussing on the hourly wage of newly hired workers per each year t , i.e. the effect on the wage of workers that were non present in the firm at year $t - 1$, but turn out to be employed at year t .

Results are reported in Figure 6. Also 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. Also relevant, we find that after three years, at the 10th percentile, firms tend to pay an hourly wage to new workers that is around 1%

²⁴ For a similar concern, see Bessen et al. (2019).

higher with respect to the initial wage paid one year before the spike. The effect is quite similar at the 50th percentile, but it is not present at the 90th percentile, where the error in the estimation is large and there is even a negative, but barely significant decline three years after a spike.²⁵

Figure 6: Automation/AI spikes and newly hired workers' wages.



Note: the solid line reports coefficients β_{-3} to β_3 from the estimation of Equation 1, while the blue and red dotted lines represent confidence intervals at the 5% and 10% significance level, respectively; the outcome variable is the wage of workers hired within the previous year (they were not present in the firm in year $t-1$).

5 Concluding remarks

In this paper we have shown that within-firm wage inequality is a pervasive phenomenon in the French economy: most of wage dispersion in France is accounted for by differences among workers belonging to the same firm, rather than by differences between sectors, firms, and occupations. Restricting the 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

²⁵ Notice that not all firms hire new workers each year, so the sample of firms and observations on which the equation for new hired workers is estimated is slightly smaller, consisting in 473,976 observations (vs.506,374 of the full sample) and 38,942 firms (vs. 39,580 of the full sample).

tend to increase wages in an equal way at different percentiles and across male and female workers.

Coming to the interpretation of our results, our findings on newly hired workers suggest that part of the individual wage increase that follows an investment in automation, and that is spread across all the wage distribution, is due to the hiring of new employees as part of the employment expansion that generally follows an event of automation, as shown in Domini et al. (2019). Unfortunately, we don't have data on education and other worker level characteristics to test whether the higher wage of newly hired workers is due to different skills or are part of a more general rent-sharing dynamics.²⁶ The lack of this worker-level information, as well as full information on job tenure, is a limitation of our study which we acknowledge.

This finding adds a novel and important piece of 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., 2019; 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 by showing that it increases wages as well, without affecting within-firm wage inequality in a significant way. In other words, the increase in wage brought about by the adoption of automation and AI is enjoyed by all workers in the adopting firm, irrespective of their initial wage and gender.

These findings should be read in the perspective of the institutional context of the French economy, which did not experience any overall significant change in within-firm wage inequality during the period. Barth et al. (2020), for example, do find that robots increase wage inequality in a sample of Norwegian manufacturing firms. An interesting question is to what extent future results on other countries will lean more towards the 'Norwegian' or the 'French' cases.

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

²⁶ However, as mentioned in the Introduction, there is a general evidence in the literature on the small productivity effects of automation events.

References

- ABELIANSKY, A. L., I. MARTÍNEZ-ZARZOSO, AND K. PRETTNER (2020): “3D printing, international trade, and FDI,” *Economic Modelling*, 85, 288–306.
- ACEMOGLU, D., C. LELARGE, AND P. RESTREPO (2020): “Competing with Robots: Firm-Level Evidence from France,” *AEA Papers and Proceedings*, 110, 383–88.
- ACEMOGLU, D. AND P. RESTREPO (2018): “Demographics and Automation,” Working Paper 24421, National Bureau of Economic Research.
- (2020): “Robots and jobs: Evidence from US labor markets,” *Journal of Political Economy*, 128, 2188–2244.
- ADAMS, R. B., B. E. HERMALIN, AND M. S. WEISBACH (2010): “The role of boards of directors in corporate governance: A conceptual framework and survey,” *Journal of Economic Literature*, 48, 58–107.
- AGHION, P., C. ANTONIN, S. BUNEL, AND X. JARAVEL (2020): “What Are the Labor and Product Market Effects of Automation? New Evidence from France,” Tech. rep.
- AHN, J., A. K. KHANDELWAL, AND S.-J. WEI (2011): “The role of intermediaries in facilitating trade,” *Journal of International Economics*, 84, 73–85.
- AKERMAN, A., E. HELPMAN, O. ITSKHOKI, M.-A. MUENDLER, AND S. REDDING (2013): “Sources of Wage Inequality,” *American Economic Review*, 103, 214–19.
- AKSOY, C. G., B. OZCAN, AND J. PHILIPP (2020): “Robots and the Gender Pay Gap in Europe,” DP 13482, IZA.
- ASPHJELL, M., W. LETTERIE, P. G.A., AND O. NILSEN (2014): “Sequentiality versus simultaneity: Interrelated factor demand,” *The Review of Economics and Statistics*, 95, 986–998.
- AUTOR, D. H. (2015): “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29, 3–30.
- AUTOR, D. H. AND D. DORN (2013): “The growth of low-skill service jobs and the polarization of the US labor market,” *American Economic Review*, 103, 1553–97.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review*, 96, 189–194.
- (2008): “Trends in U.S. Wage Inequality: Revising the Revisionists,” *The Review of Economics and Statistics*, 90, 300–323.
- BALASSA, B. (1965): “Trade liberalisation and “revealed” comparative advantage,” *The Manchester School*, 33, 99–123.
- BALASUBRAMANIAN, N. AND J. SIVADASAN (2011): “What happens when firms patent? New evidence from US economic census data,” *The Review of Economics and Statistics*, 93, 126–146.
- BARTH, E., M. ROED, P. SCHONE, AND J. UMBLIJS (2020): “How Robots Change Within-Firm Wage Inequality,” IZA Discussion Papers 13605, Institute of Labor Economics (IZA).
- BERNARD, A. B., M. GRAZZI, AND C. TOMASI (2015): “Intermediaries in International Trade: Products and Destinations,” *The Review of Economics and Statistics*, 97, 916–920.
- BERNARD, A. B., J. B. JENSEN, S. J. REDDING, AND P. K. SCHOTT (2010): “Wholesalers and Retailers in US Trade,” *American Economic Review*, 100, 408–13.

- BESSEN, J., M. GOOS, A. SALOMONS, AND W. VAN DEN BERGE (2020): “Firm-Level Automation: Evidence from the Netherlands,” *AEA Papers and Proceedings*, 110, 389–93.
- BESSEN, J. E., M. GOOS, A. SALOMONS, AND W. VAN DEN BERGE (2019): “Automatic Reaction—What Happens to Workers at Firms that Automate?” Law and Economics Research Paper 19-2, Boston University.
- BLAU, F. D. AND L. M. KAHN (2017): “The Gender Wage Gap: Extent, Trends, and Explanations,” *Journal of Economic Literature*, 55, 789–865.
- BLUM, B. S., S. CLARO, AND I. HORSTMANN (2010): “Facts and Figures on Intermediated Trade,” *American Economic Review*, 100, 419–23.
- BONFIGLIOLI, A., R. CRINÒ, H. FADINGER, AND G. GANCIA (2020): “Robot Imports and Firm-Level Outcomes,” Cepr dp14593, CEPR.
- BRUSSEVICH, M., E. DABLA-NORRIS, AND S. KHALID (2019): “Is Technology Widening the Gender Gap? Automation and the Future of Female Employment,” Working Paper WP/19/91, IMF.
- BRYNJOLFSSON, E. AND A. MCAFEE (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, WW Norton & Company.
- CALIENDO, L., F. MONTE, AND E. ROSSI-HANSBERG (2015): “The anatomy of French production hierarchies,” *Journal of Political Economy*, 123, 809–852.
- CARD, D., A. R. CARDOSO, AND P. KLINE (2016): “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women,” *The Quarterly Journal of Economics*, 131, 633–686.
- CIRILLO, V., M. SOSTERO, AND F. TAMAGNI (2017): “Innovation and within-firm wage inequalities: empirical evidence from major European countries,” *Industry and Innovation*, 24, 468–491.
- DAUTH, W., S. FINDEISEN, J. SUEDEKUM, N. WOESSNER, ET AL. (2018): “Adjusting to robots: worker-level evidence,” Institute working paper 13, Federal Reserve Bank of Minneapolis.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): “The Rise of Market Power and the Macroeconomic Implications*,” *The Quarterly Journal of Economics*, 135, 561–644.
- DIXON, J., B. HONG, AND L. WU (2019): “The Robot Revolution: Managerial and Employment Consequences for Firms,” Tech. rep., NYU Stern School of Business.
- DOMINI, G., M. GRAZZI, D. MOSCHELLA, AND T. TREIBICH (2019): “Threats and opportunities in the digital era: automation spikes and employment dynamics,” Tech. rep., LEM Working Paper Series.
- DUGGAN, M., C. GARTHWAITE, AND A. GOYAL (2016): “The market impacts of pharmaceutical product patents in developing countries: Evidence from India,” *American Economic Review*, 106, 99–135.
- EDMANS, A. AND X. GABAIX (2016): “Executive Compensation: A Modern Primer,” *Journal of Economic Literature*, 54, 1232–87.
- EUROSTAT (2015): *Statistics Explained*, Luxembourg: Eurostat, Statistical Office of the European Communities, <https://ec.europa.eu/eurostat/statisticsexplained/>.
- FREEMAN, R. B., I. GANGULI, AND M. J. HANDEL (2020): “Within-Occupation Changes Dominate Changes in What Workers Do: A Shift-Share Decomposition, 2005–2015,” *AEA Papers and Proceedings*, 110, 394–99.

- FREY, C. B. AND M. A. OSBORNE (2017): “The future of employment: How susceptible are jobs to computerisation?” *Technological Forecasting and Social Change*, 114, 254 – 280.
- GARDEAZABAL, J. AND A. UGIDOS (2005): “Gender wage discrimination at quantiles,” *Journal of population economics*, 18, 165–179.
- GE, S. AND Y. ZHOU (2020): “Robots, computers, and the gender wage gap,” *Journal of Economic Behavior and Organization*, 178, 194–222.
- GOLDIN, C. (2014): “A grand gender convergence: Its last chapter,” *American Economic Review*, 104, 1091–1119.
- GOOS, M., A. MANNING, AND A. SALOMONS (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104, 2509–26.
- GRAETZ, G. AND G. MICHAELS (2018): “Robots at work,” *Review of Economics and Statistics*, 100, 753–768.
- GRAZZI, M., N. JACOBY, AND T. TREIBICH (2016): “Dynamics of Investment and Firm Performance: Comparative evidence from manufacturing industries,” *Empirical Economics*, 51, 125–179.
- GUELLEC, D. AND C. PAUNOV (2020): “Digital Innovation and the Distribution of Income,” in *Measuring and Accounting for Innovation in the 21st Century*, ed. by C. Corrado, J. Haskel, J. Miranda, and D. Sichel, Chicago: NBER/ The University of Chicago Press, National Bureau of Economic Research Studies in Income and Wealth.
- GUILLOU, S. AND T. TREIBICH (2019): “Firm export diversification and change in workforce composition,” *Review of World Economics*, 155, 645–676.
- HEINZE, A. AND E. WOLF (2010): “The intra-firm gender wage gap: a new view on wage differentials based on linked employer–employee data,” *Journal of Population Economics*, 23, 851–879.
- HELPMAN, E., O. ITSKHOKI, M.-A. MUENDLER, AND S. J. REDDING (2017): “Trade and Inequality: From Theory to Estimation,” *The Review of Economic Studies*, 84, 357–405.
- HUNT, J. AND R. NUNN (2019): “Is Employment Polarization Informative About Wage Inequality and Is Employment Really Polarizing?” Working Paper 26064, National Bureau of Economic Research.
- ILO (2016): “Global Wage Report 2016/17: Wage inequality in the workplace,” Report, International Labor Organization.
- (2020): “Ilostat–International Labour Organization Database Of Labour Statistics,” .
- INSEE (2010): *DADS Grand Format. Guide méthodologique. Validité 2010*, Dijon: INSEE, Direction Régionale de Bourgogne.
- KLENERT, D., E. FERNANDEZ-MACIAS, AND J.-I. ANTON (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., I. MANUYLOV, AND M. SMOLKA (2019): “Robots and Firms,” Economics Working Papers 2019-05, Department of Economics and Business Economics, Aarhus University.
- LAFORTUNE, J., J. ROTHSTEIN, AND D. W. SCHANZENBACH (2018): “School finance reform and the distribution of student achievement,” *American Economic Journal: Applied Economics*, 10, 1–26.
- LETTERIE, W. A., G. A. PFANN, AND J. POLDER (2004): “Factor adjustment spikes and interrelation: an empirical investigation,” *Economics Letters*, 85, 145 – 150.

- LYBBERT, T. J. AND N. J. ZOLAS (2014): “Getting patents and economic data to speak to each other: An Algorithmic Links with Probabilities approach for joint analyses of patenting and economic activity,” *Research Policy*, 43, 530–542.
- MILLER, C. (2017): “The persistent effect of temporary affirmative action,” *American Economic Journal: Applied Economics*, 9, 152–90.
- MISHEL, L. AND J. WOLFE (2019): “CEO compensation has grown 940% since 1978: Typical worker compensation has risen only 12% during that time,” report, Economic Policy Institute.
- MUELLER, H. M., P. P. OUMET, AND E. SIMINTZI (2017): “Wage inequality and firm growth,” *American Economic Review*, 107, 379–83.
- NILSEN, O. A., A. RAKNERUD, M. RYBALKA, AND T. SKJERPEN (2009): “Lumpy investments, factor adjustments, and labour productivity,” *Oxford Economic Papers*, 61, 104–127.
- OI, W. Y. AND T. L. IDSON (1999): “Firm size and wages,” *Handbook of labor economics*, 3, 2165–2214.
- PEREIRA, J. M. R. AND A. GALEGO (2019): “Diverging trends of wage inequality in Europe,” *Oxford Economic Papers*, 71, 799–823.
- SEC (2015): “SEC Adopts Rule for Pay Ratio Disclosure,” press release, Security Exchange Commission.
- SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): “Firming Up Inequality*,” *The Quarterly Journal of Economics*, 134, 1–50.
- SORGNER, A., E. BODE, C. KRIEGER-BODEN, U. ANEJA, S. COLEMAN, V. MISHRA, AND A. M. ROBB (2017): “The effects of digitalization on gender equality in the G20 economies: Women20 study,” Tech. rep., Kiel.
- VAN DER VELDE, L. (2020): “Within Occupation Wage Dispersion and the Task Content of Jobs,” *Oxford Bulletin of Economics and Statistics*, online first.


Appendix A

Table 7: HS-2012 product codes referring to automation- and AI-related technologies.

Label	HS-2012 codes
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
10. Automatic data processing machines	847141-847150, 847321, 847330
11. Electronic calculating machines	847010-847029

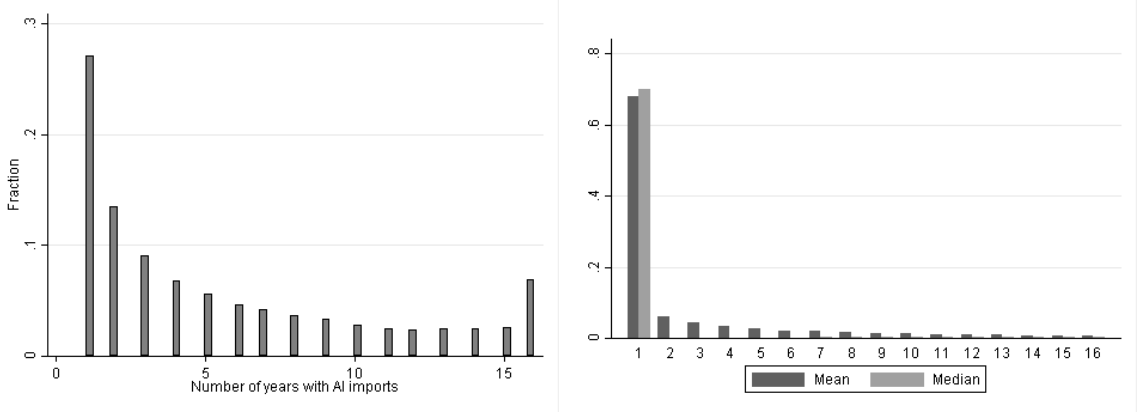
Notes: for further details on the codes automation and 3-D printers, see Acemoglu and Restrepo (2018, A-12-A14) and (? , p. 13), respectively.

Figure 7: Distribution of wage per hour among all workers, before cleaning.



Source: our elaboration on DADS data. Note: the vertical line indicates our cleaning threshold (half the minimum wage per hour in 2017).

Figure 8: Testing the lumpy nature of our spike variable.



Source: our elaborations on DADS data.

Table 8: Automation and AI importers and spikes per year, sample 2, 2002-2017.

Year	Share of automation importers	Share of AI and automation importers	Share of automation spikes	Share of AI and automation spikes
2002	11.58	15.98	3.84	5.31
2003	11.46	15.61	2.68	3.61
2004	11.79	16.31	2.47	3.44
2005	12.04	16.69	2.49	3.46
2006	11.96	16.75	2.28	3.32
2007	12.30	16.79	2.63	3.44
2008	12.53	16.84	2.49	3.20
2009	11.95	15.97	1.9	2.43
2010	12.67	16.76	2.22	2.80
2011	11.98	17.29	1.93	2.92
2012	12.73	17.54	1.99	2.58
2013	13.46	19.05	2.05	2.93
2014	13.21	19.16	2.19	3.23
2015	13.47	19.84	2.39	3.80
2016	13.92	20.53	2.79	4.42
2017	14.46	20.79	3.93	5.61
Total	12.56	17.55	2.51	3.52

Source: our elaborations on DGDDI data.