

New technologies and the rise of wage inequality

Raquel Sebastian

Universidad Complutense de Madrid, ICAE; EQUALITAS

Pedro Salas-Rojo

International Inequalities Institute, LSE; CUNEF; EQUALITAS

Juan César Palomino

Universidad Complutense de Madrid, ICAE; EQUALITAS; INET Oxford

Juan Gabriel Rodríguez

Universidad Complutense de Madrid, ICAE; EQUALITAS; CEDESOG

Raquel Sebastian

Universidad Complutense de Madrid, ICAE;
EQUALITAS

Pedro Salas-Rojo

International Inequalities Institute, LSE;
CUNEF; EQUALITAS

Juan César Palomino

Universidad Complutense de Madrid, ICAE;
EQUALITAS; INET Oxford

Juan Gabriel Rodríguez

Universidad Complutense de Madrid, ICAE;
EQUALITAS; CEDESOG

In addition to our working papers series all these publications are available to download free from our website: www.lse.ac.uk/III

International Inequalities Institute
The London School of Economics and
Political Science, Houghton Street,
London WC2A 2AE

E Inequalities.institute@lse.ac.uk

W www.lse.ac.uk/III

X [@LSEInequalities](https://twitter.com/LSEInequalities)

NEW TECHNOLOGIES AND THE RISE OF WAGE INEQUALITY

Raquel Sebastian 

Universidad Complutense de Madrid, ICAE and EQUALITAS

E-mail: raquel.sebastian@ucm.es

Pedro Salas-Rojo 

London School of Economics and Political Science, CUNEF, and EQUALITAS

E-mail: p.salas-rojo@lse.ac.uk

Juan César Palomino 

Universidad Complutense de Madrid, ICAE, EQUALITAS and INET Oxford

E-mail: jc.palomino@ucm.es

Juan Gabriel Rodríguez¹

Universidad Complutense de Madrid, ICAE, EQUALITAS and CEDESOG

E-mail: juangabr@ucm.es

Abstract

Technological change fuels economic growth, but its impact on wage inequality remains contested. This study presents a unified empirical framework that isolates the effects of new technologies such as automation and AI on the entire wage distribution. We develop a continuous and task-sensitive automation index and propose a distributional counterfactual-based method. Applying the approach to Spanish micro-data for 2000-2019 and instrumenting technology variables, we find automation to be a key driver of inequality: without task displacement the Gini coefficient would be 21.5% lower and significant wage shares would shift from the top 10% towards middle and bottom groups. Automation is found to barely affect the gender gap in the period studied, yet to widen the education premium. Like automation, AI exposure increases inequality, although the mechanisms to impact wages differ: automation tends to negatively impact wages in the middle of the distribution, while AI tends to increase wages at the top. Trade, offshorability, educational attainment, employment rates and mark-ups play secondary, period-specific roles. The results can inform policies on skill formation and inclusive innovation.

Keywords: automation; AI; wage inequality; structural change; job tasks.

JEL Codes: O33; D33; J21; J24; J31.

¹ The author ordering was constructed through American Economic Association's randomization tool (confirmation code: apF2L1YMRkpx). We are grateful for the comments and suggestions from Andrés Barge, José Ignacio Antón, Michele Bavaro, and Guido Neidhöfer. Jorge Batanero provided invaluable research assistance. The authors acknowledge funding from Fundación La Caixa under project 242PR30186. The views expressed are those of the authors not the funders.

1. Introduction

Technological change has been the main driver of economic growth and improvements in real wages and living standards over the last two centuries (Frey, 2000). The importance of new technologies has grown steadily, especially in the last decades. A clear example is robotics, which have added an estimated average of 0.37 percentage points of annual GDP growth (about one-tenth of total growth) between 1993 and 2007 for the United States, Western Europe, South Korea and Australia (Graetz and Michaels, 2018). Still, despite the evidence that technology boosts growth, its effect on employment and wage inequality remains unclear.² The lack of consensus about the disruptive effects of technological advances on the wage distribution can contribute to technological anxiety, as there is little evidence addressing the specific impact of technological change on wage inequality. The few existing analyses approaching the topic focus on estimating changes in the wage structure and job polarization, ignoring a more comprehensive study of the entire distribution.³ In this paper, we propose a methodology for evaluating the impact of new technologies on wage inequality, and discuss its implications for the future of labor.

Not all technologies are made equal. Different technologies may have diverse effects on productivity, and on aggregate labor demand and wages. First, technology developments can complement human labor, allowing for an increasing productivity (*productivity effect*) while contributing to a rising demand for labor in non-automated tasks. Second, they can also destroy jobs because they automate tasks (*displacement effect*); in the worst possible case, ‘so-so technologies’ such as self-checkout machines found at retail stores and automated customer services, displace workers but may not substantially improve productivity, leading to a net decline in labor demand. Third, new tasks and jobs are created either through innovation or because old technologies become so cheap that their demand starts to rise (*reinstatement effect*). In addition, Artificial Intelligence (AI), the newest disrupting technological advancement including machine learning, machine vision, natural language processing, and automated guided vehicles, still poses an unclear

² As Korinek and Stiglitz (2019) said: ‘We economists set ourselves too easy a goal if we just say that technological progress can make everyone better off—we also have to explain how we can make this happen.’

³ Job polarization (growth in employment of both low wage and high wage jobs at the expense of middle-wage jobs) has awoken concerns about whether technology replaces to a greater extent those routine-task jobs held by middle-class workers. If occupations in the center of the wage distribution are more affected than low- and high-skilled occupations, the replacement of labor by machines does not only displace the middle class but also increase wage inequality.

effect on labor demand, wages and inequality.⁴ When considering the net impact of technological progress on wage inequality, the literature has yet to reach a consensus on whether it is positive or negative.

Our proposal distinguishes between automation and early AI introduction when measuring how technological progress affects the distribution of wages. To accurately measure automation, our new “*Unfixed automation index*” improves two important aspects of the main measures proposed in the literature (Autor and Dorn, 2013; Lewandoski et al., 2019). First, our measure is continuous – rather than discrete – so it is not affected by boundary problems. Second, and most importantly, unlike existing automation indices (which use a fixed baseline metric of the tasks related to each occupation), our measure does not only capture changes in relative employment among occupations but also changes in the degree of task routinization of occupations over time. The sensitivity of the unfixed index of automation to changes in both employment and the degree of routinization across occupations is critical to capturing the effects of automation on wage inequality.

To assess the early effects of AI technologies, we leverage on two well-known measures from the literature: the AI Occupational Impact score proposed by Felten et al. (2018) and the occupations’ exposure to AI, robots, and software score proposed by Webb (2020). Both indices focus on technologies that are already present and being innovated upon, and remain agnostic regarding substitution, leaving the door open for augmentation effects.⁵ Nonetheless, both indices measure different aspects of AI technologies and, therefore, may have associated different implications for wage inequality.

Still, the sharp increase in wage inequality worldwide in recent decades (Ahmed, 2023) cannot be fully explained by technological change. Other factors can influence wage and must be considered in our empirical application. For example, the observed general rise in markups.⁶ The interplay between technological transformation and globalization reinforces market concentration, as leading firms capture higher profits. While

⁴ For example, see Brynjolfsson et al. (2025) for the impact of generative AI-based conversational assistant on worker productivity.

⁵ In contrast, other measures of AI evaluate the feasibility of future automation, without guaranteeing that the technology will be adopted and concentrate on AI’s potential to replace human labor (see Frey and Osborne, 2017; Brynjolfsson et al., 2018).

⁶ Empirical results based on large publicly traded companies corroborate that markups growth is widespread on a global scale, with increases ranging between 40 and 60 percent depending on the continent (De Loecker and Eeckhout, 2025).

globalization and declining transportation costs have expanded market size, economic activity has reallocated towards a small number of firms. These firms typically possess technological advantages, operate at the knowledge frontier of their sectors, and rely heavily on – often intangible – capital, enabling them to sustain high markups.

In this context, markups derive from current technology dynamics that favor concentration – even in a context of intense competition – rather than from barriers to market entry. The so-called ‘Superstar Firms’ may have larger markups because they have moderated their costs through scaling, which rises the dispersion of firm efficiency and size. Thus, with equal qualifications, workers in a larger, more efficient, and globalized firm are likely to be more productive and receive a higher pay than those in less efficient, smaller firms. In turn, companies that take advantage of the opportunities offered by new technologies to achieve much higher productivity levels can also attract talent by offering higher salaries, potentially further increasing wage inequalities. Including mark-ups along with technological change in our analysis will allow us to discern clearly the effect of automation and AI separately from the market power impact.

Trade openness has potential mixed effects on the wages of unskilled labor in advanced countries: while it raises the skill premium, it can also increase real wages by lowering import prices. Increased offshorability can also favor concentration in technology-intensive sectors, pushing up demand for – and wages of – high skilled workers. Education also shapes wage inequalities, determining occupational choices, access to jobs, and the overall level of pay. Thus, the effect of increased educational attainment on wage inequality can be either positive or negative depending on the evolution of the skill premium. More efficient labor markets foster economic dynamism by reallocating resources to more productive firms but may propose challenges to workers with low and outdated skills. This can be captured by including employment rates in our analysis, which allows also to control for variations due to the business cycle. Finally, we also consider the distribution of workers across sectors, including a dummy for the manufacturing sector.

Our study of the impact of new technologies on wage inequality proceeds in two steps. In the first step, we group workers aged 25-65 by age brackets, gender, educational levels and regions and regress the change in the average wage (in logarithms) of each group between two time periods on regressors capturing the factors described above. Since the effects of automation on wage inequality can differ when a medium- or a long-term

perspective is adopted, our regression analysis lies on two-time spans. Our main analysis considers the whole 2000-2019 long-term period, but we also consider that period by decades (medium-term), using the 2000-2010-2019 panel with a time dummy for the second decade. For the specific analysis of the early effects of AI on wage inequality, we focus on the 2015-2019 period as the introduction of AI is not relevant before 2015 (Acemoglu and Restrepo, 2022).

We lag all explicative variables and correct for the potential problems of reverse causality and omitted variables by applying instrumental variables. Following Ferri (2022), we construct a shift-share or Bartik instrument for each automation and AI variable. For robustness, and to further address concerns about omitted variable bias affecting our results, we conduct an additional analysis based on the *bounding procedure* proposed by Cinelli and Haznett (2020).

The second step in our approach explores the impact of technological change on inequality by analyzing the distributional consequences of the estimates from the previous step. Given the Bartik instruments results and the various robustness and sensibility analyses, we can interpret the main coefficients for the technological variables as having a causal impact on the change in wages. We then use those estimates to construct a counterfactual wage distribution free of the influence of the technological variable under consideration, which can be compared with the observed wage distribution. In this comparison, we apply an ample set of inequality metrics: the Gini coefficient, the Mean Log Deviation (MLD), and the wage shares of the top 10%, the bottom 50%, the bottom 10%, and of the middle-class group (from percentile 25 to percentile 75). This counterfactual approach also permits us to explore the impact of new technologies on different demographic groups and to assess their effect on the gender wage gap and the educational premium.

Our empirical analysis focuses on Spain, an ideal case to test our proposal. Inequality in Spain has notably risen since the 2000s, with a 6.4 Gini points increase between 2000 and 2016. This process became more intense since the onset of the Great Recession (4.7 Gini points increase in only 8 years, from 2008 to 2016). Indeed, in the period 2000-2019, the Gini index gap between the United States and Spain fell from 8.8 Gini points to only 2.2 Gini points.⁷ Surprisingly, despite this marked trend, a potentially crucial factor driving

⁷ These values were obtained from version 9.8 of the Standardized World Income Inequality Database (SWIID; Solt, 2020). The data corresponds to the Gini coefficient for market incomes.

wage inequality in this period, technological progress, has hardly received any attention in Spain.

Beyond inequality dynamics, the Spanish labor market is characterized by large fluctuations in employment, with intense job destruction in recessions and adjustments falling hard on some segments of the labor force.⁸ Bonhomme and Hospido (2016) found wage inequality in Spain to be largely procyclical, blaming the real estate bubble as the main contributor to the evolution of inequality in Spain during the decade 2000-2010. This institutional setting in the labor market interacts with a remarkable expansion of the belief that technological innovation increases inequality in recent years, especially among workers in the manufacturing sector (Rodriguez and Sebastian, 2022). According to the V Survey on the Perception of Innovation in Spain – conducted by COTEC Foundation and Sigma Dos in 2020/21 – 75% of the Spanish population have a positive opinion on innovation, but the figure has dropped 15 percentage points from 2016 to 2020. Half of Spaniards believe that robots and algorithms will destroy more jobs than they will create, while 70% believe that most current jobs will be automatized over the next 15 years. Shedding light on the effect of technological change on wage inequality in recent times could help understand the economic reality shaping these perceptions and, eventually, serve as a reference to develop a policy agenda to mitigate its effects.⁹

Our main results robustly evidence that automation has been a significant driver of wage inequality in Spain. In our estimation, the negative relationship between wage inequality and automation remains after controlling for market power, trade, offshorability, and education. For the long-run period 2000-2019, and keeping everything else equal, the Gini coefficient would have been – at the end of the period – approximately 21.5% lower in the absence of task displacement linked to our preferred automation measure, the Unfixed index. In this period, automation also contributes to a redistribution of wage shares away from the bottom and middle towards the top of the distribution: while the wage share of the top 10% would have been about 3.9% smaller without technological change, the share of the middle 50% would have been 0.15% higher, and that of the

⁸ A bad design of passive policies, dysfunctions in collective bargaining and ineffectiveness of active employment policies, as well as an excessive use of temporary contracts, have typically caused initial differences in training to be transformed into high permanent wage differentials.

⁹ From a methodology standpoint, the technological indices we use implicitly assume that equivalent occupational tasks in Spain are equally exposed to technology as in the U.S. This bears the advantage that the technological measures calculate for Spain are not endogenous to wage changes, thus enhancing the reliability of our identification strategy.

bottom 10% would have been 2.2% higher in the counterfactual scenario. AI exposure indices (for the period 2015-2019) are also associated with a substantial increase in inequality: using the AI index proposed by Felten et al. (2018) we estimate that the 2019 Gini coefficient would be 9.9% smaller in the absence of AI technologies in the 2015-2019 period. Although both automation and AI are linked to increased wage inequality, the mechanism through which their impact on wages occurs is different. Our results show that AI has the potential to raise wages for higher-skilled workers, while automation tends to lower wages for middle- and low-skilled workers.

While changes in the gender wage gap due to automation are small, it is found to consistently enlarge wage disparities across educational groups. For example, for the long run (2000-2019), the 2019 education premium – the average wage differential between high and low-educated workers – would have been substantially smaller, by about 43%, in the absence of these automation effects in this period. This finding supports the notion that automation technologies complement higher skills while substituting tasks typically performed by lower- and middle-educated workers. AI measures also point towards this direction; both, the Felten and Webb indexes are associated with a significant widening of the education premium.

A closer look into the wage change across demographic profiles shows that older workers and higher- or medium- educated females are often neutral to or positively affected by automation, while those most adversely affected tend to be younger, low-educated individuals. The results also highlight spatial variations in the impact of automation, with regions like Navarra, Aragón and Murcia showing a larger adverse technological impact on wages.

The rest of the article is organized as follows. Section 2 reviews the main existing automation and AI measures and presents our new unfixed automation index. In Section 3, we present an empirical model for the analysis of wage inequality changes due to new technologies. Section 4 presents the databases and describes the construction of our key variables. Sections 5 and 6 develop the descriptive and regression analyses, respectively, while in Section 7 we show the distributional consequences of technological change using a counterfactual scenario. Section 8 concludes.

2. Measures of automation and AI

In this section we go through some of the most popular measures of automation and AI proposed in the literature and argue why our proposed measure overcomes some of their limitations. We also describe how these metrics are used to assign our units of analysis – demographic groups, further explained in Section 3 – with a numeric value that reflects how largely they are affected by new technologies.

2.1. Measures of automation

2.1.1. The ‘classical’ approach

The literature has employed a wide set of measures to approximate automation exposure of workers. Arguably the most influential – which we term the ‘classical’ approach – was proposed by Autor and Dorn (2013) and goes beyond a summary index of routine task activities, measuring the routine, abstract, and manual task content of each occupation.¹⁰ These measures are combined into a relative index of net routine task intensity (*RTI*):

$$RTI_{k,t}^C = \ln(T_{k,t}^R) - \ln(T_{k,t}^A) - \ln(T_{k,t}^M), \quad (1)$$

where $T_{k,t}^R$, $T_{k,t}^A$, and $T_{k,t}^M$ are respectively the routine, abstract, and manual task inputs in each occupation k at time t . Following Acemoglu and Autor (2011) to build this index for our data we use the O*NET database (see Section 4), so the score in each type of task is measured on a one to five scale.¹¹ The *RTI* index is used then to identify the set of routine-intensive occupations, particularly those in the top (demographic-weighted) third of routine task-intensity. Then, the routine employment share (*RSH*) measured for each demographic group j is:

$$RSH_{jt}^C = \sum_{k=1}^K \frac{L_{jkt}}{L_{jt}} \cdot \mathbf{1}_{[RTI_{k,0}^C > RTI_{k,0}^C(P66)]} \quad (2)$$

where L_{jkt} is the employment in occupation k in group j at time t , L_{jt} is the total employment in group j at time t , *P66* refers to the percentile 66, and $\mathbf{1}_{[RTI_{k,0}^C > RTI_{k,0}^C(P66)]}$

¹⁰ Autor and Dorn (2013) collapsed the original five task measures developed in Autor et al. (2003) (routine manual, routine cognitive, nonroutine manual, nonroutine interactive and nonroutine cognitive analytical) into three task aggregates for abstract, routine, and manual tasks.

¹¹ Autor and Dorn (2013) used the DOT database, so tasks’ scores were measured on a zero to ten scale. To avoid non-positive values in the logarithm they used the task score of the 5th percentile for the 5 percent of observations with the lowest task score.

is the indicator function at the initial period of time, equal to one when the occupation is routine intensive according to the previous definition.

2.1.2. The ‘Lewandowski et al. (2019)’ approach

Lewandowski et al (2019) proposed a relative index of routine task intensity, which conceptually differs from the classical approach mainly in not considering manual tasks:

$$RTI_{k,t}^L = \ln(T_{k,t}^{RC}) - \ln\left(\frac{T_{k,t}^{NRA} + T_{k,t}^{NRP}}{2}\right), \quad (3)$$

where $T_{k,t}^{RC}$, $T_{k,t}^{NRA}$, and $T_{k,t}^{NRP}$ are the routine cognitive, non-routine cognitive analytical, and non-routine cognitive personal task inputs in each occupation k at period t , respectively. The routine employment share RSH_{jt}^L in Lewandowski et al. (2019) is calculated in accordance with the proposal by Autor y Dorn (2013), but with $RTI_{k,t}^L$ used instead of $RTI_{k,t}^C$.

2.1.3 The Acemoglu and Restrepo (2020) approach

Due to the clear displacement effect of robots, Acemoglu and Restrepo (2020) propose a robot-focused approach. Other authors, such as Doorley et al. (2023), have recently used this method to model the impact of automation on European economies. However, this approach has the conceptual disadvantage of not capturing automation in a broader sense, as it does not take into account other technologies, such as ICT (information and communication technologies) capital and software. Moreover, the main database on robots that exists – the International Federation of Robotics (IFR) database (see Jurkat et al., 2022) – is a primary source for tracking the adoption of robots, but only industrial ones. This means that, in the case of the Spanish economy, the data show a high concentration of robots in the automotive and electronics sectors, in contrast to other sectors, particularly the service sector and non-manufacturing industries, where robot density is minimal or non-existent. In other words, the Spanish economy, with its large service sector and other non-manufacturing sectors susceptible to automation, is largely excluded from the calculation of automation using the approach of Acemoglu and Restrepo (2020). These limitations make this approach not suitable to analyze automation for whole labor force (i.e., beyond the industrial sector) and we do not use it in our study.

2.1.4 Our proposal: the Unfixed index of automation

Unlike the Acemoglu and Restrepo (2020) approach, the first two methods – the classical and the Lewandowski et al. (2019) approach – have the conceptual advantage of capturing automation in a broader sense, considering robots and other technologies such as ICT capital and software. Still, despite their popularity in the literature, we find three critical flaws in their conceptualization:¹²

i) Manual tasks inputs are used to calculate the index of net routine task intensity. This implies that the distribution of manual task scores by occupations is relevant for the measurement of net routine task intensity. However, while automation is plausibly related directly with routine tasks and inversely related with abstract tasks, the relationship between automation and manual tasks is mixed; automation displaces or replaces some manual tasks, but it is complementary to these tasks in other cases. This makes us argue towards only considering only the distributions of routine and abstract task scores by occupations in the calculation of the *RTI* index.¹³

ii) These indices classify occupations into automatable or not by adopting a discrete criterion, by considering as net routine occupations those in the top third of the occupation’s distribution according to the net routine task score. However, using a continuous criterion, as in our proposed measure, can have measurement advantages: the net degree of automatability of the occupation is considered in its full distribution and the boundary problem is avoided. For instance, according to the ‘classical’ approach, an occupation at the 66th percentile of the occupation distribution according to the net routine task score could look more like the one at the 67th percentile than the one at the 65th percentile, but it would not be considered a routine occupation and, as such, its task content would be ignored. In our metric, instead, we have a continuous range of automatability (net routinization) that assigns each occupation a level according to their task content.

iii) The literature on automation usually fixed the set of routine-intensive occupations at the initial period (t_0). As a result, the variation of the routine employment share of a demographic group only considers how much the group workers move overtime to or

¹² Autor and Dorn (2013) indeed acknowledged that the simplicity of (2) was attractive but there were many plausible ways to construct the measure. These authors addressed this concern by exploring some alternatives, such as replacing the three-factor RTI with a two-factor alternative ($RTI_{k,t} = \ln(T_{k,t}^R) - \ln(T_{k,t}^M)$), measuring the routine share using the top 25 or 40 percent of occupations rather than the top 33 percent, and using the mean RTI as a measure of routine-intensity rather than the routine occupation share.

¹³ Recall that the *RTI* proposed by Lewandowski et al. (2019) also does not consider the distribution of manual task scores by occupations.

from (initially set) routine occupations, measured by the change in the ratio $\frac{L_{jkt}}{L_{jt}}$. However, the change in the degree of routinization of occupations themselves overtime is ignored, and this change may be large. We indeed observe that in our data (see Section 5).

Our measure of automation is thus built to overcome these three important limitations. First, the *RTI* index by occupation is obtained as the routine task input (in logs) minus its abstract task input (in logs) at the period t :

$$RTI_{k,t}^{UF} = \ln(T_{k,t}^R) - \ln(T_{k,t}^A). \quad (4)$$

Then, we calculate the routine employment share for each demographic group j at time t as follows:

$$RSH_{jt}^{UF} = \sum_{k=1}^K \frac{L_{jkt}}{L_{jt}} \cdot RTI_{k,t}^{UF}. \quad (5)$$

Here, the routine employment share at time t is the employment-weighted mean of the routine task intensity, so *RSH* considers the degree of routinization of the occupations and avoids the boundary problem. Since the mean of the routine task intensity is measured at each time t , the index captures change over time in the degree of task routine (it is *unfixed*). Thus, variation in RSH_{jt}^{UF} comes both from changes in the relative employment between occupations and from changes in the degree of routinization of occupations. These advantages of the index RSH_{jt}^{UF} imply that the measure in (5) may be greater than 1. Therefore, for comparability and interpretability, we standardize RSH_{jt}^{UF} between 0 and 1 as follows:

$$RSH_{jt}^{UF} = \frac{RSH_{jt}^{UF} - \min\{RSH_{jt}^{UF}\}}{\max\{RSH_{jt}^{UF}\} - \min\{RSH_{jt}^{UF}\}}. \quad (6)$$

Our empirical application focuses on the above proposed unfixed index of automation, although we perform an extensive robustness analysis. We replicate our entire analysis for the classical approach and propose checks for the Lewandowski et al. (2019) approach, as well as for a variation of RSH_{jt}^{UF} , for which the relative routine task intensity index is solely dependent on routine task input, i.e., $RTI_{k,t} = \ln(T_{k,t}^R)$. We denote this measure by RSH_{jt}^{UF2} .

2.2 Measures of AI

The distinct exposure of occupations to recent advances in AI differs from that of automation, and various indices have recently been proposed to capture this effect. As explained in the introduction, we will consider two of these indices: the AI Occupational Impact score proposed by Felten et al. (2018) and the occupations' exposure to AI, robots, and software score proposed by Webb (2020). Both indices focus on existing technologies and remain agnostic regarding substitution, leaving the door open for augmentation effects.

2.2.1. The AI Occupational Impact (AIOI) score (Felten et al., 2018 and 2019)

Felten et al. (2018 and 2019) link advances in specific AI applications to the skill characteristics required for an occupation. Their scores are based on 2019 O*NET data for occupational descriptions and the Electronic Frontier Foundation (EFF) AI Progress Measurement dataset, which measures progress in AI applications from 2010 to 2015. They find that, on average, occupations impacted by AI experience a small but positive change in wages, although no changes in employment were identified.

The measure is constructed using the growth in EFF's metrics on AI performance in executing different activities (image recognition, visual question answering, image generation, abstract games, real-time video games, reading comprehension, language modeling, translation, and speech recognition). They connect EFF categories to the list of 52 skills that O*NET uses to describe the requirements of each job, using responses from an Amazon's Mechanical Turk (mTurk) survey. The final aggregated score is weighted by the prevalence and importance of abilities within each occupation. After standardization to the 0-1 range, higher values reflect a greater potential effect of AI on the associated occupation.

2.2.2. The occupations' exposure to AI, robots, and software score (Webb, 2020)

Webb (2020) constructs the scores of occupations' exposure to AI, robots, and software by quantifying the textual overlap (verb-noun pairs) of patents to job descriptions. The method for constructing the occupational exposure measure requires two sources of text: patents and job descriptions. For patents it uses Google Patents Public Data, provided by IFI CLAIMS Patent Services, and for job descriptions it uses O*NET.

Webb (2020) finds that low-skilled occupations are more exposed to robots, middle-skilled occupations to software, and high-skilled occupations to AI, with this latter effect being more likely to affect older and highly educated workers than robots and software.

Using the AI exposure scores provided by Felten et al. (2018, 2019) and Webb (2020) as well as the SML index by Brynjolfsson et al. (2018), Acemoglu et al. (2022) conclude that the impact of AI is still too small relative to the scale of the U.S. labor market to have had first-order impacts on employment patterns in the period 2010-18. Here, instead, we focus not on employment but on the potential effect of AI exposure for wage inequality in Spain, and only for the most recent period in our database: 2015-2019. We will use the AI indices described above, and, for greater comparability across results, will standardize them following equation (6).

Albanesi et al. (2023) have recently used the scores provided by Felten et al. (2018, 2019) and Webb (2020) and evidenced a positive association between AI-enabled automation and changes in employment shares for the pooled sample of European countries in the period 2011-2019. The connection with wage changes, however, is less clear in their results (negative and hardly significant for the Felten et al.'s measure and statistically not significant for the Webb's measure).

Note that the Felten and Webb measures of AI exposure were originally constructed using task descriptions for the U.S. economy, so works using them for the European labor market (such as Albanesi et al., 2023) implicitly assume that tasks are equally exposed to technology in the E.U. countries and in the U.S. The same implicit equivalence of tasks within occupations at both sides of the Atlantic occurs with our automation measures, since they also use O*NET data and are also based on the description of tasks for occupations in the US.¹⁴

3. Empirical model

Our ultimate purpose is to estimate how technology and other factors shape the long-run evolution of wage inequality. In principle, one could rely on at least two approaches, both bearing its own limitations. On the one hand, using a macro approach for the estimation considers only one time series of macroeconomic units (e.g. countries or regions) and requires the two periods of analysis to be sufficiently far apart in time. In practice, this approach makes it hardly possible to have enough observations to make appropriate inference. On the other hand, a purely micro approach focusing on individuals faces the

¹⁴ This, can provide an econometric advantage in the estimation, making AI and automation exposure measures in our analysis not endogenous to local wage changes. Nonetheless, to reinforce our conclusions and support a causal interpretation of the results an instrumental variables analysis will also be applied.

same ultimate caveat: scarcity of observations. Even when a panel of workers is available for a sufficiently long period, attrition rises rapidly as we move away in time. To overcome this problem, we follow the literature and use an intermediate approach, estimating the effect of new technologies on wage inequality across demographic groups.¹⁵

3.1. Regression model

Our main sample consists of employees aged between 25 and 65 (see Section 4), from which we define our units of analysis (demographic groups) based on sex, four 10-year age bins (25-35, 36-45, 46-55, 56-65), three educational levels (low, medium and high), and the region of residence (Spain's 17 autonomous communities, excluding the two autonomous cities).¹⁶ We construct 408 groups based on all possible combinations of the grouping variables ($2 \times 4 \times 3 \times 17 = 408$). Then, we estimate a baseline model where we, in the period considered, regress the change in the average wage (in logarithms) of each group of workers on the automation exposure index and the rest of potential explanatory factors:

$$\Delta \ln \mu_{jt} = \alpha + \beta \cdot \ln \mu_{jt-1} + \gamma \cdot RSH_{jt-1} + \lambda^K \cdot X_{jt-1}^K + \varepsilon_{jt}, \quad (7)$$

where μ_{jt} is the real average yearly wage of group j ($j = 1, \dots, J$) in period t , RSH_{jt-1} is the automation exposure index of group j in period $t - 1$, X_{jt-1}^K is the vector of the controls in period $t - 1$, and the error term ε_{jt} represents the aggregate effect of other unobserved factors.¹⁷ Regressions are weighted by the group's demographic size and standard errors are robust to heteroskedasticity.¹⁸

Because the wage distribution is skewed by nature, we log-transform the dependent variable, so its changes are then to be considered in percentages. A common expansion of all wages due, for example, to generalized increases in labor productivity is absorbed

¹⁵ An additional advantage of estimating the effect of technological change across demographic groups is the avoidance of worker selection problem encountered in cross-occupational studies (Böhm et al., 2024).

¹⁶ We also consider full-time public sector workers since they cannot be precisely identified in the initial year of the sample. We could assign task content intensities to unemployed individuals based on the last job they held. Unfortunately, the information about past work is fragmentary and its use could bias the estimations since the lack of this information is likely non-random. Hence, the task contents for unemployed individuals are defined as missing.

¹⁷ All monetary values are Consumer Price Index (CPI) adjusted to 2019 Euros.

¹⁸ We also performed the regressions weighted by the total average hours worked by each group, but the results are similar. They are available upon request.

by the constant term. We include the lag of the initial wage (see, for example, Card et al., 1999, and Fossen and Sorgner, 2022) to account for heterogeneity in starting points. This way, the estimated effects in the period studied are not confounded by preexisting wage differences due to past labor conditions, previous human capital accumulation or preexisting bargaining power. Heckman et al. (2006) have highlighted the role of initial wage in estimating causal effects, while other authors like Autor and Dorn (2013) and Beaudry et al. (2014) have also emphasized the importance of initial conditions for wage dynamics. Including pre-existing wage differences across demographic groups has two additional advantages. First, it addresses potential omitted variable bias or measurement error: prior wage is often correlated with worker ability or overall unobservable productivity so, if one does not control for it, other variables in the regression might wrongly capture those effects. Second, as often studied in the context of labor market dynamics, the parameter on $\ln \mu_{jt-1}$ tests for convergence effects in wages. Demographic groups with lower initial average wages may experience faster wage growth, leading to a negative coefficient on the initial wage.

Our main technological variable is the lagged average automation exposure of the group. As argued above, we focus on our preferred unfixed measure RSH_{jt}^{FU} , but develop a set of robustness analysis for the rest of automation measures. The vector of controls X_{jt-1}^K , includes a broad set of variables. To control for foreign sector exposure, we include trade, which measures external trade intensity, and offshorability, calculated as in Acemoglu and Autor (2011). To control for the distribution of workers across industrial sectors we consider the share of workers in manufacturing jobs, as the substitution of routine jobs by automation is expected to be more intense in the manufacturing sector. Bonhomme and Hospido (2016) show that the real state bubble was the main factor behind the evolution of inequality in Spain during the decade 2000-2010. To account for this, for robustness, we replaced the share of the manufacturing sector with the share of the construction sector. The variable did not turn out to be significant: when the period under consideration is large, the effects of the real state bubble are mitigated.

The average years of education is included to capture the effects of educational attainment on wages. To control for the market power of firms we calculate markup at the industry

level using the method proposed by Akerberg et al. (2015).¹⁹ This method improves the identification of input elasticities. It does not require observing marginal prices or the mode of competition because it only assumes cost minimization and at least one variable input, such as materials, energy, or services. Additionally, it adapts to gross output production, which is ideal for markups involving materials, and provides a flexible process for unobserved productivity.

Finally, we control for the large cyclical variations in employment that recurrently occur in the Spanish labor market (Pijoan-Mas and Sánchez-Marcos, 2010; Bonhomme and Hospido, 2016), using the employment share, denoting the relative size of employed workers. We have also tried using the variables employees and active workers instead of the employment share, and the results remained similar.

The effect of automation changes on wage inequality is only fully appreciated when a medium- and long-term perspective are adopted. We then develop our empirical analysis in two stages. First, we consider the entire ‘long term’ period 2000-2019. Next, we consider that period from a ‘medium term’ perspective (decades) in the form of a 2000-2010-2019 panel (with a time dummy for each decade). Our analysis thus departs from previous studies for Spain, which tend to focus on shorter term effects.

In all cases, we estimate first the base model $\Delta \ln \mu_{jt} = \alpha + \beta \cdot \ln \mu_{jt-1} + \gamma \cdot RSH_{jt-1} + \varepsilon_{jt}$ and then introduce the remaining explicative variables one by one. Finally, we estimate the full model with all explicative variables. For the specific analysis of the effects of AI on wage inequality, we focus, as mentioned, on the period 2015-2019 since the impact of AI before 2015 is still hard to detect in the US (Acemoglu et al, 2022) and Spain is considered a ‘follower’ country in the implementation of technological innovations.

The estimation of wage changes by OLS faces significant endogeneity challenges, even when lagging our automation and AI exposure variables. The adoption of new technologies may be endogenous to labor demand, which also influences wages (*reverse causality*). Furthermore, *omitted variables* correlated with both automation exposure and wage dynamics could bias the estimates. To address these issues and isolate a causal

¹⁹ The markup is calculated at the sector level, so one might think that standard errors should be clustered at this level. However, the level of treatment in our work is the group, and the most recent literature recommends that for this level, clustering is reduced to using robust errors, which is what we do here (see Abadie et al., 2023).

relationship, we apply a shift share (or Bartik) instrumental variable approach to the model in (7).

While this framework is traditionally applied to analyze economic shocks across different geographical regions, we adapt its logic to a different dimension: demographic groups (j) within the same labor market. This novel application allows us to construct a predicted automation exposure for each group that is driven by broad, external trends rather than group-specific, local factors.

Following Ferri (2022), the instrument for the automation exposure RSH_{jt} is constructed as:

$$RSH_{j,t}^{UF*} = \sum_{i=1}^I \frac{L_{i,j,t-1}}{L_{j,t-1}} \cdot RSH_{i,-j,t}^{UF}, \quad (8)$$

where the lagged ‘share’ component $\frac{L_{i,j,t-1}}{L_{j,t-1}}$ is the employment share of industry i ($i = 1, \dots, I$) for demographic group j in period $t - 1$. Crucially, we define $t - 1$ as a 10-year lag before the main analysis period, t . Using these historical employment shares ensures our instrument is based on pre-determined industrial specializations that are not influenced by the contemporaneous technological shocks we seek to analyze. The delocalized “shift” component $RSH_{i,-j,t}^{UF} = \sum_{k=1}^K \frac{L_{i,-j,k,t}}{L_{i,-j,t}} \cdot RTI_{k,t}^{UF}$ is the routine employment share of industry i at time t for all other demographic groups $-j$. This term captures the aggregate, national trend in automation exposure for each industry, which is plausibly exogenous to the specific economic conditions of any single demographic group j . We construct analogous shift-share instruments for the other automation and AI indices used in our analysis.

The Bartik instrument is then utilized within a two-stage least squares (2SLS) regression framework to secure a causal estimate. In the first stage, we regress the actual automation exposure variable on our constructed instrument and the rest of controls. This step isolates the component of automation exposure driven solely by the interaction of historical employment shares and broad, ‘delocalized’ industry trends. In the second stage, these predicted values – now stripped of confounding local shocks and reverse causality – are used in place of the original variable in our main wage equation (7). This 2SLS procedure allows us to interpret the resulting coefficients as the causal impact of automation on wages.

We further check the robustness of our approach with the application of the *bounding procedure* proposed by Cinelli and Haznett (2020), which explores how likely is that omitted variables could affect the results. In particular, this method allows us to assess: *i)* How strong should the unobserved variables be to invalidate the conclusions of our study; *ii)* How robust are our results in a worst-case scenario, where all unobserved variables act together, even after allowing for non-linear interactions across them; *iii)* How strong should unobserved variables need to be in relation to the strength of some observed covariates in order to change our results by a certain amount. We find this analysis reassuring.

3.2. Wage inequality analysis

Once regression (7) has been run and the estimates for the main explanatory factors of the change in the average wage across demographic groups are obtained, the problem lies on how to assess its impact on total wage inequality. The literature has typically measured the effect of factors on wage change and then examined graphically where these effects are most concentrated. Thus, automation is believed to increase wage inequality because it has a negative effect on routine jobs (more exposed to automation) which are concentrated in the middle of the wage distribution (Goos and Manning, 2007; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Acemoglu and Restrepo, 2022). While this approach allows for a precise measurement of the effects on the wage structure, it does not quantitatively assess wage inequality.

Other scholars have relied on measures such as earnings percentile ratios (Juhn et al., 1993; DiNardo et al., 1996; Autor et al., 2006; Lemieux, 2006; Bonhomme and Hospido, 2016) or the variance of logarithms (Levy and Murnane, 1992; Card and Lemieux, 1994; Sala-i-Martin, 1996; Heckman and Honore, 1990) to try to measure the effect of automation on inequality. While these approaches are interesting, their key metrics have an intrinsic limitation: they do not comply with the cornerstone principle of inequality, the Pigou-Dalton principle of transfers (Atkinson, 1970; Foster and Ok, 1999) and, therefore, are not consistent with the Lorenz curve.²⁰

²⁰ This principle requires that any mean-preserving progressive transfer lowers the value of an inequality index (Shorrocks and Foster, 1987).

To address these issues and develop a precise and quantitative analysis of the effect of automation on wage inequality, we propose a direct and effective empirical strategy. Our inequality analysis starts from the regression results corresponding to equation (7), where the use of shift-share instruments and the robustness checks performed allow us to interpret γ , the coefficient on RSH_{jt-1} , as having a causal impact on wage growth. Based on this, we construct a counterfactual wage variable, y_{jt}^c , by subtracting γRSH_{jt-1} from the observed $\ln \mu_{jt}$ as follows:

$$y_{jt}^c = \ln \mu_{jt} - \hat{\gamma} RSH_{jt-1} = \hat{\alpha} + (1 + \hat{\beta}) \ln \mu_{jt-1} + \hat{\lambda}^K X_{jt-1}^K + \hat{\varepsilon}_{jt}. \quad (9)$$

This counterfactual distribution contains all the elements of the predicted fit and the error term, meaning it replicates the original distribution, except for the part of the log wage change during the analyzed period that is linked to the technology component. The comparison between the distribution of counterfactual wages $y_{jt}^c = (y_{1t}^c, y_{2t}^c, \dots, y_{jt}^c)$ and the observed wage distribution $\mu_t = (\ln \mu_{1t}, \ln \mu_{2t}, \dots, \ln \mu_{jt})$ allows us to isolate and size up the direct effect of new technologies on wage inequality. Thus, using an inequality measure $I(\cdot)$, if

$$I(y_t^c) < I(\mu_t) \quad (10)$$

counterfactual wage inequality (in the absence of technological change) is lower than the observed one, suggesting that new technologies increase wage inequality.

We apply this approach to the Gini coefficient, the Mean Log Deviation (MLD), and the shares of different groups of workers in the wage distribution: the top 10%, bottom 50%, bottom 10%, and the ‘middle class’ (percentile 25 to percentile 75).²¹ However, we could use any other index of inequality. We have considered the Gini coefficient and the MLD index because the former is the most widely used and well-known index of inequality, while the latter is widely used in inequality of opportunity measurement.²² For all inequality indices, we report the differences in relative terms, i.e.,

²¹ The Gini coefficient is half the weighted average difference between all pairs of wages divided by the average wage:

$$G(\mu_t) = \frac{1}{2\bar{\mu}_t} \sum_{j=1}^J \sum_{k=1}^J f(\mu_{jt}) f(\mu_{kt}) |\mu_{jt} - \mu_{kt}|,$$

where $\bar{\mu}_t$ is the mean wage of the wage distribution μ_t and $f(\cdot)$ denotes the weight of a given wage. The Mean logarithmic Deviation or Theil 0 index is defined as:

$$MLD(\mu_t) = \sum_{j=1}^J f(\mu_{jt}) \ln \frac{\bar{\mu}_t}{\mu_{jt}}.$$

²² The ranking of income distributions by majority voting and social evaluation functions – the two main alternatives proposed in the literature for aggregating individual preferences – coincide if and only if the inequality index under consideration is the Gini coefficient (Rodriguez and Salas, 2014). The MLD measure

$$RD_t(I) = \frac{I(y_t^c) - I(\mu_t)}{I(\mu_t)} * 100. \quad (11)$$

This approach not only allows us to obtain precise estimates of how much overall wage inequality has changed due to automation and AI in the periods examined, but also to explore wage dynamics across socioeconomic characteristics and regions. Comparing the observed and counterfactual distribution, we can assess the impact of new technologies on the gender gap and the education premium.

The gender gap is calculated as the difference between average wages of men and women, $Gap_S(x) = \bar{x}_{men} - \bar{x}_{women}$, where x is either the observed wage distribution μ or the distribution of counterfactual wages y . The education premium is analogously computed as the difference in average wages of high and low educated workers $Gap_E(x) = \bar{x}_{high} - \bar{x}_{low}$. Then, for each, we calculate the relative difference in the gaps:

$$RD_t(Gap_R) = \frac{Gap_R(y_t^c) - Gap_R(\mu_t)}{Gap_R(\mu_t)} * 100, \quad (12)$$

where R denotes either S or E . The interpretation of $RD_t(Gap_R)$ is straightforward: a negative value would indicate that the gender gap (or education premium) in observed wages is larger than in the counterfactual, suggesting that new technologies increase the gender gap (or education premium). In contrast, a positive value implies that the gap in y is larger than the gap in the observed wages, so new technologies reduce the gap.

We also explore how automation exposure and AI affect each demographic group. We compute the counterfactual wage y_{jt}^c for each group j at period t , and estimate the impact of new technologies as:

$$RD_t(y, \mu) = \frac{y_{jt}^c - \ln(\mu_{jt})}{\ln(\mu_{jt})} * 100. \quad (13)$$

A positive value implies that new technologies reduce the average wage of the group under consideration, with the opposite effect being associated with a negative value. This way, we can study which dimensions (gender, education, age or region) of the demographic groups are most important for understanding the effects of technological change.

is the only inequality index that is additively decomposable into a between-group and a within-group component (Shorrocks, 1980; Cowell, 1980) and has a path-independent decomposition, so the result of the decomposition is independent of which component – the between-group or the within-group – is eliminated first (Foster and Shneyerov, 2000).

4. Data and variables

4.1. Core databases: automation exposure, wages and robots

The O*NET database, developed by the U.S. Department of Labor, serves as the successor to the Dictionary of Occupational Titles (DOT) and constitutes the core of the O*NET program. Within this framework, trained analysts evaluate a wide range of occupational tasks, assigning standardized scores that reflect their relevance and frequency within each occupation. As such, O*NET provides detailed information on occupational characteristics, enabling the construction of indices such as routine task intensity, offshorability, and AI exposure (see Appendix A). The data are collected and updated regularly, with consistent releases since 2003.

To integrate this information into our analysis we map the O*NET-SOC classification to the U.S. Standard Occupational Classification (SOC) system, and subsequently to the International Standard Classification of Occupations (ISCO), as used in the Spanish Labor Force Survey (ES-LFS), at the three-digit level. Appendix Table A2 provides a full description of the crosswalks applied, as well as the corresponding ISCO occupations and NACE industry codes. In this manner, we obtain a score of the manual, routine and abstract components of each occupation, which can be matched exactly to the coding of our other key databases.

In second place, we use data from the Spanish Labor Force Survey (ES-LFS), framed within the European Union Labor Force Survey (EU-LFS). The EU-LFS is compiled by Eurostat based on harmonized national labor force surveys conducted by statistical agencies in EU Member States. It provides comparable information on labor market conditions across the 27 EU countries, as well as Norway, the United Kingdom, Switzerland, Iceland, Turkey, and North Macedonia. For the purposes of this study, we use the Spanish sample (ES-LFS) for the years 1992, 1995, 2000, 2010, 2015, and 2019. Employment is measured according to the International Labor Organization (ILO) definition. Occupations are reported at the three-digit level using the International Standard Classification of Occupations (ISCO). We recode the reported ISCO-88 classifications from 1992 to 2008 into ISCO-08, which is used consistently from 2009 onwards.²³ Similarly, industries at the one-digit level are classified according to the

²³ We exclude occupations with insufficient representation or unavailable task measures in O*NET. Specifically, we drop “Agricultural, Fishery and Related Laborers” (ISCO-88: 92) and “Subsistence Farmers, Fishers, and Gatherers” (ISCO-08: 63) due to limited sample size, as well as “Armed Forces

Nomenclature of Economic Activities in the European Community (NACE) and are recoded from the original NACE Rev.1 (used until 2009) into NACE Rev.2, which is used without changes from 2010 onwards.

While the ES-LFS offers the best source for mapping occupations and assigning automation exposure indices derived from O*NET, it does not include individual-level wage data, essential for analysing earnings inequality. To address this limitation, we use *groups-j* to merge the ES-LFS estimates with wage data drawn from two harmonized sources: the Panel de Hogares (PH) for the year 2000, and the Encuesta de Condiciones de Vida (ECV) for 2005, 2010, 2015, and 2019.²⁴ The PH database – developed by the Instituto Nacional de Estadística (INE) between 1994 and 2001 –, follows the methodology of the European Union Household Panel and includes detailed information on income, covariates, well-being, and living conditions. Since 2004, this survey has been conducted under the name ECV as part of the European Union Statistics on Income and Living Conditions (EU-SILC). Both data series are representative at the national and regional levels, owing to a stratified sampling and a weighting scheme that accounts for variables such as age, sex, nationality, and region of residence. They have been extensively used to explore inequality and poverty dimensions in Spain (see Ayala et al., 2022, for a discussion on data sources on income inequality in Spain). Our sample subset are the full-time individual employees aged 25 to 65 reporting strictly positive net labor income.²⁵ Wages are expressed in constant 2019 euros, deflated using the Consumer Price Index (CPI) provided by the World Bank. The dependent variable is constructed as the weighted average wage across sociodemographic groups, and Table A3 in Appendix A reports the corresponding wage inequality statistics.

Occupations” (ISCO-88: 0; ISCO-08: 0), for which O*NET provides no coverage. We also omit “Unpaid Family Workers,” given their sparse representation. To preserve broad occupational coverage, we merge certain categories across systems: in ISCO-88, we combine “Teaching Associate Professionals” (33) with “Other Associate Professionals” (34); in ISCO-08, we merge “Cleaners and Helpers” (91) with “Food Preparation Assistants” (94).

²⁴ We have verified that the relative position of groups and their wages remain stable across surveys, regardless of the origin of the data. The Spearman rank correlation of group wages between year 2000 (POGHE) and 2005 (ECV) is 0.80, while the correlation between 2005 and 2010 (both ECV) is 0.83. Creating larger groups – for example, defined using only sex, education, and age – yields correlations over 0.9 in all cases. More information and additional checks are available from the authors upon request.

²⁵ We have excluded the self-employed as we do not have appropriate data to disentangle wages from capital in the self-reported labor income. For the employed, the terms ‘wages’ and ‘labor income’ are used interchangeably throughout the text.

4.2. Non-technological variables

While wages are defined at the socio-demographic group level, some explanatory variables are defined at the occupation or industry level. Here, the value for each sociodemographic group is constructed as a weighted average of the industry or occupation values, where the weights used are the share of wages earned by the group workers in the corresponding occupation or industry relative to the total earnings of the group.

The variable '*trade*' reflects commercial openness by regions and is constructed as (export + imports) over GDP at the regional level. Blinder (2009) and Blinder and Krueger (2013) consider that occupations not requiring physical presence can, in principle, be outsourced. Adopting this definition, and following Acemoglu and Autor (2011), we obtain a score of 'offshorability' (see the details of its construction in Appendix A). Since the substitution of routine jobs by automation is expected to be more relevant in the manufacturing sector, we also include the share of workers in the manufacturing sector by region (*sh_manufacturing*) as a control. As discussed above, we have also considered the weight of the construction sector in each region (*sh_construction*). The variable '*avg_years_edu*' denotes the average years of education by autonomous regions.

To control for firms' market power, we consider four different variables of '*markup*', calculated as the average markup at the sectoral level using the method proposed by Akerberg et al. (2015) to control for unobserved heterogeneity. The differences between the four variables are explained by the production function considered (Cobb-Douglas or trans-log production function) and by whether revenue-based weights at the NACE level are used. Output is approximated by value added in all cases. The markup variables bear similar results in all models, but we primarily lean towards the Cobb-Douglas production function with weights because its significance is higher.²⁶

In addition, we consider the evolution and functioning of the labor market throughout the business cycle. We calculated three alternative measures for each region: a relative variable, the employment rate (*employment_rate*) defined as employed workers over active workers, and two absolute variables, namely the number of employed workers (*employees*) and the size of active labor (*active_workers*). However, we opted to

²⁶ All these results are available upon request.

include the first measure in our estimations because it captures best the large employment variations characterizing the Spanish labor market due to the business cycle (Pijoan-Mas and Sánchez-Marcos, 2010; Bonhomme and Hospido, 2016).

Databases are matched using the 408 socioeconomic groups defined above, or the correspondent region, occupation or industry code. As a result, we construct a comprehensive dataset that condenses information from all available sources at the demographic group level.

5. Descriptive analysis

This section provides a descriptive overview of our key variables and offers initial intuitions and context into the raw wage-automatability dynamics, which we will seek to identify more formally in the full results of the econometric models in Section 6 and in the counterfactual analysis of Section 7.

5.1. Comparison of the automation indices

A key argument of this paper is that to best measure automation exposure, we need an index that captures the fact that the tasks involved in a job can change over time. As described in detail in Section 2, our preferred index, RSH^{UF} , is built to precisely account for that. Figure 1 shows the variation in the core component of our index – the Routine Task Intensity (RTI) – for each occupation between 2000 and 2019, plotted against the occupation’s initial wage. The blue line, representing the RTI before standardization for the unfixed index (RSH^{UF}), exhibits substantial fluctuation around zero, both positive and negative, confirming that the degree of exposure to automation in the job tasks is not a static occupational characteristic but one that changes over two decades. Some occupations have seen their RTI vary in the range of -0.41 (Mining, Construction, Manufacturing and Transport) to 0.35 (Chief Executives and Senior Officials) or even 0.61 (Market-Oriented Skilled Agricultural Workers) points, which are considerable variations given that RTI ranges between -1 and 1. In contrast, the green and red markers, representing the RTI for the Classic (RSH^C) and Lewandowski (RSH^L) indices, coincide and show a flat horizontal line at zero. This is a direct visual representation of their core

assumption: the task content of an occupation is fixed at a baseline year, and the index only captures changes in employment shares within each given demographic group.


Figure 1. RTI variation between 2000 and 2019.



Note: Own elaboration. The red marker is not visible in the figure because the RTI for the Classic (RSH^C) and Lewandowski (RSH^L) indices are the same, a flat horizontal line at zero. The UF blue line stands for the Unfixed RTI approach.

In addition, to illustrate the advantage of the unfixed automation index in capturing the variation of automation across the wage distribution over time, we compare in Figure 2 the three automation indices – RSH^{UF} , RSH^C , and RSH^L – by plotting their final standardized scores against occupation-level wages for the years 2000, 2010, and 2019. All three panels illustrate a clear negative gradient, slightly more pronounced for RSH^{UF} , confirming the general observation that occupations with lower wages tend to have higher automation exposure. However, if one focuses on how this wage-automation relation has varied across the periods analyzed, RSH^{UF} displays a noticeably wider range of scores, as its changes include both changes in occupation shares and in the automation component of occupations.

Figure 2. Routine shares (RSH^C , RSH^L and RSH^{UF}) by baseline wage and waves.



Note: Own elaboration.


5.2. Exposure to automation and AI

Who is more exposed to automation? We now assess the automation exposure by key demographic characteristics. Figures 3a and 3b map the baseline automation scores (in the year 2000) across demographic groups, disaggregated by gender, age, and education, for the unfixed and classic indices, respectively. Both reveal a consistent pattern for education across the board: individuals with a low level of education face a significantly higher exposure to automation than their more educated peers. The relationship with age appears to be also negative for most groups in both indices, with a less pronounced negative gradient for less educated workers generally. Only for low-educated female workers in the RSH^C index the relation with age is positive.

A comparison of the two indices uncovers a subtle gender difference. In Figure 3a, which displays our RSH^{UF} index, the scores for female workers consistently lie below those for males within the same education category. This suggests that with this measure, men are systematically more exposed to automation. However, in Figure 3b, for the RSH^C index, this gender gap is far less distinct; the scores for men and women are much closer and


overlap in the mid-education group. Women are more exposed to automation in the low-education group and less in the high-education group.

Figure 3a. Baseline automation for each demographic group in the Unfixed index.



Note: Own elaboration. Automation Index values for each group in the baseline year (2000).


Figure 3b. Baseline Automation for each demographic group in the Classic index.



Note: Own elaboration. Automation Index values for each group in the baseline year (2000).

Turning to the exposure to AI, Figures 3c and 3d present the baseline scores for the Felten and Webb indices, respectively. A clear contrast emerges when looking at the intensity of the rows, which indicate education levels. While automation exposure (Figures 3a and 3b) showed a negative educational gradient, AI exposure displays a clear positive gradient. Across both the Felten and Webb metrics, highly educated individuals face the highest exposure to AI technologies, while those with low education levels show the lowest scores. This aligns with the literature suggesting that AI technologies, unlike earlier waves of automation, are more likely to overlap with tasks performed by high-skilled, and often older, workers. Given the well-known relation between human capital and wages, this different relation with education hints at the different relation that automation and AI have with the wage gradient, which we will explore later.

Figure 3c. Baseline AI exposure for each demographic group in the Felten index.




Note: Own elaboration. AI Index values for each group in the baseline year (2015).

However, the two AI indices diverge notably regarding the gender dimension, anticipating the conflicting results found in our counterfactual analysis for gender. In Figure 3c, the Felten index reveals substantial exposure among high-educated women, often larger than that of their male counterparts. Conversely, Figure 3d, based on the Webb index, depicts a pattern where men, particularly in the high-education bracket,

appear comparatively more exposed. These results suggest that different methodologies for measuring AI may capture distinct sets of tasks that can have opposing implications.

Figure 3d. Baseline AI exposure for each demographic group in the Webb index.




Note: Own elaboration. AI Index values for each group in the baseline year (2015).

5.3. Demographic group wages and levels of automation and AI


Figure 4 shifts the focus to the raw relationship between wages at the demographic group level and the indices of automation and AI for the years 2000, 2015 and 2019. For the automation indices RSH^{UF} and RSH^C a clear, negative correlation exists in all periods. Groups with higher automation exposure tend to have lower average wages. However, in the case of AI, the situation is very different, as the Felten and Webb indices show a clear positive relationship with wages in 2015 and 2019. This evidence suggests that the associations with the wage structure of automation and AI point in different directions. As we will see, in the regression explaining wage changes in the periods analyzed, the coefficients of automation indices and those of AI have different signs (negative and positive, respectively). But – as will be revealed in Section 7 – both automation and AI are found to increase inequality, since the former is associated to depressing wages in the middle of the distribution and the latter to increasing wages at the top. The next subsection provides descriptive evidence for that connection.

Figure 4. Demographic group wages and levels of automation and AI.

4.a. Years 2000 and 2019.



4.b. Years 2015 and 2019.



Note: Own elaboration. The unit of analysis is the demographic group. The size of the bubble represents the weight of the demographic group. A smoother is added to visualize the relation between automation/AI and baseline wage.

5.4. Automation exposure and wage changes between 2000 and 2019


Figure 5 provides a critical link between our descriptive overview and the main regression results. It displays the relationship between a group's initial automation exposure and its subsequent wage change over the 2000-2019 period. For each index, the figure alternates between the raw unconditional relationship, and the same relationship after conditioning on our full set of control variables. In the conditioned plots, a clearer picture emerges: the wage losses, represented by red bubbles, become heavily concentrated in the middle of the baseline wage distribution, limited by the vertical dashed lines that mark the 25th and 75th percentiles of the wage distribution. Losses marked by larger red bubbles tend also to be associated with higher values of automation (Y-axis), especially for our preferred unfixed index, which is consistent with the negative sign of its coefficient in the regression. This visual evidence of a "hollowing out" of the middle is consistent with what the theory of job polarization predicts (Goos and Manning, 2007; Autor and Dorn, 2013; Rodríguez and Sebastian, 2022). This pattern also provides a clear intuition for our later counterfactual finding that automation has increased wage inequality. If the middle of the wage distribution suffers losses and the tails gain or are unaffected, overall inequality increases. The emergence of this pattern in the conditioned plots underlines the importance of our control strategy in isolating the true, polarizing impact of automation.

5.5. AI (and automation) exposure and wage changes between 2015 and 2019

Figure 6 focuses on the more recent 2015-2019 period for which we have data on AI, aiming at capturing the early effects of this technology and comparing them to those of automation. Here the graphs present always conditioned wage changes against the baseline wage for each technology index. The patterns reveal important differences between AI measures and the RSH^{UF} . As we already saw in Figure 4, the AI indices are associated with a slightly positive gradient with the baseline wage, while the RSH^{UF} index shows a negative gradient. Also, we can see that when plotting the AI indices, the larger losses (red bubbles) tend to be in the middle of the wage distribution, a pattern that is clearer in the Felten AI metric. This "hollowing out" effect also helps explain our counterfactual results, where these indices are found to substantially increase overall wage inequality, with the effect being particularly strong and significant for the Felten AI index, and less strong for the Webb AI index. The fact that larger red bubbles are in this


case associated with lower values of AI exposure – particularly when using the Felten index – explains why its coefficient is positive in the regression of wage change on AI.

Figure 5. Log wage change 2000-2019 for the Unfixed and Classic indices. (without controls and conditioning for our set of controls)



Note: Own elaboration. Red bubbles represent negative changes and blue bubbles positive changes. The size of bubbles represents the magnitude of the change.

Figure 6. Log Wage change 2015-2019 for the AI indices and the Unfixed index.
(conditioning for our set of controls)




Note: Own elaboration. Red bubbles represent negative changes and blue bubbles positive changes. The size of bubbles represents the magnitude of the change.

5.6. Control variables by region's average baseline wage

Finally, Figure 7, in panels A through F, plots our main control variables against each region's average baseline wage. The charts make two points clear. First, there is large heterogeneity across Spain's regions in factors like trade openness, industrial structure, offshorability, and firm markups. Second, while there is no uniform pattern, some variables do show a clear raw correlation with regional wage levels. As one would expect, panel C shows that higher-wage regions tend to have a more educated labor force. Interestingly, panel E suggests a similar positive raw correlation for firm markups. This heterogeneity, which does not always follow a simple income gradient, underscores the necessity of including these variables in our models.

Figure 7. Controls value by region ordered by baseline wage.



Source: Own elaboration.

6. Regression results

Here we present the empirical results from our regression analysis, examining the impact of automation and the more recent influence of AI on wage dynamics across diverse

demographic groups in Spain. We primarily focus on our unfixed measure of automation (RSH^{UF}). For comparative purposes, we also refer to results using the RSH^C measure (Autor and Dorn, 2013) which are reported in full in Appendix B. The robustness of these measures, especially within a causal framework, is supported by the performance of their respective Bartik instruments. The results of RSH^{UF} , our preferred index due to the desirable properties discussed earlier, are particularly robust to the shift-share instrumentation.

6.1. Long-term results

We begin by examining the long-term relationship across initial measures of the technological exposure of occupations and the evolution of wages over the extensive 2000-2019 period. Table 1 presents the OLS estimates in Panel A and the corresponding shift-share IV results in Panel B. Regardless of the specific automation measure employed, the OLS estimates displayed in Table 1 indicate a negative and significant association between a higher initial index of automation and wage growth over the two-decade span. For instance, the coefficient for RSH^{UF} is -0.800 and RSH^C is -0.602.

When we turn to the Bartik IV estimates in Table 1, arguably more informative for drawing causal inferences, these negative coefficients generally become larger in magnitude. For RSH^{UF} , the IV coefficient is -0.925, and for RSH^C , it is -1.140. This suggests that linear estimation without instrumentation might underestimate the true impact of automation exposure on wages. As explained above, all models in Table 1 – and throughout the whole empirical exercise – include the lagged logarithm of the group's average wage in the initial period as a control for possible heterogeneity in the effect of automation on wage for different starting wage levels. This consistently shows a statistically significant negative coefficient, indicating a pattern of wage convergence among demographic groups over this long period; that is, groups with initially higher wages tended to experience slower subsequent wage growth.

A central consideration in any instrumental variable regression, including our shift-share Bartik analysis, is the strength of the instrument, typically evaluated using the F-statistic from the first-stage regression. Table 1 shows that the Bartik instrument for our preferred unfixed measure, RSH^{UF} , has an F-statistic of 34.81. Similarly, the instrument for RSH^C also performs well, with an F-statistic of 19.70. In contrast, other alternative automation

measures, RSH^L (based on Lewandowski et al., 2019) and RSH^{UF2} (an alternative specification of RSH^{UF} that considers routine task content only, see Section 2), exhibit F-statistics of 11.63 and 10.05, respectively. While these are near or just above the conventional threshold of 10 (Staiger and Stock, 1997), they suggest a weaker instrumental performance. The superior performance of the Bartik instruments for RSH^{UF} and, to a lesser extent, RSH^C , provides greater credibility to the causal inferences drawn using these specific measures. This supports our primary focus on RSH^{UF} in the subsequent detailed analyses, with RSH^C main results also being reported and available in detail in Appendix B.

Table 1. Automation variables (2000-2019).

	(1) RSH ^C	(2) RSH ^L	(3) RSH ^{UF}	(4) RSH ^{UF2}
<i>Panel A. OLS estimates</i>				
L. Wage Past	-0.587*** (0.0771)	-0.552*** (0.0718)	-0.633*** (0.0682)	-0.538*** (0.0703)
L. RSH ^C	-0.602*** (0.141)			
L. RSH ^L		-0.811*** (0.168)		
L. RSH ^{UF}			-0.800*** (0.114)	
L. RSH ^{UF2}				-0.882*** (0.158)
Constant	5.779*** (0.787)	5.456*** (0.745)	6.500*** (0.717)	5.702*** (0.774)
Observations	408	408	408	408
R-squared	0.291	0.280	0.332	0.283
<i>Panel B. 2SLS estimates</i>				
L. RSH ^C	-1.140*** (0.441)			
L. RSH ^L		-2.937*** (1.087)		
L. RSH ^{UF}			-0.925*** (0.282)	
L. RSH ^{UF2}				-3.143** (1.478)
F-stat	19.70	11.63	34.81	10.05

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 2 (Panel A) presents OLS estimates that incorporate a comprehensive set of control variables, while Table 2 (Panel B) provides the corresponding results that further control for potential endogeneity by using instrumental shift-share analysis. In Table 2, the coefficient on RSH^{UF} remains robustly negative and statistically significant at the 1% level across all specifications. In the most comprehensive model (column 7), which includes all available control variables, the OLS coefficient for RSH^{UF} is -0.886, greater than when no controls were included (Table 1). This implies that a higher baseline of automatable occupational tasks is associated with substantially lower wage growth over the 2000-2019 period.

Several control variables in this full OLS specification (Table 2, column 7) also offer significant insights into long-run wage dynamics. The lagged log wage of the group, maintains its negative and highly significant coefficient, reinforcing the wage convergence hypothesis. Regional trade openness, defined as the sum of exports and imports over regional GDP, shows a positive and significant coefficient, suggesting that groups in regions with greater international trade exposure experienced, on average, faster wage growth. The regional employment rate also shows a positive and significant coefficient, connecting better macroeconomic performance in the labor market with higher wage growth. The markup indicator is also significant in intermediate specifications (column 5), and weakly significant when all controls are considered. Other regional controls in this full specification, such as the weight of the industrial sector in total employment and the average years of education do not show statistically significant coefficients, nor do other controls such as average offshorability of the occupations in the demographic group. The overall model fit, as indicated by the R-squared, is 0.36 in the full specification. For comparison, Table B1 in Appendix B, which utilizes the RSH^C measure, reports a coefficient of -0.697 for RSH^C in its full OLS model for the same period, with an overall R-squared of 0.33.

The Bartik IV estimates presented in Table 2 (Panel B) further strengthen the causal interpretation of the impact of RSH^{FU} . The coefficient in the full model (column 7) is significant and becomes considerably more negative than in the OLS regression, estimated now at -1.014. This estimate is supported by a strong F-statistic of 33.23 for the instrument, reinforcing the likelihood that this reflects a causal depressing effect of initial level of automatable tasks on long-term wage growth in Spain. For RSH^C (see Table B1

in Appendix B) we find a coefficient of -1.530 in the full model and an F-statistic of 17.28.

Table 2. Long-run inequality change (RSH^{UF} , 2000-2019).


	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. OLS estimates</i>							
L. Wage Past	-0.661*** (0.067)	-0.633*** (0.068)	-0.654*** (0.069)	-0.638*** (0.068)	-0.650*** (0.069)	-0.656*** (0.067)	-0.679*** (0.068)
L. RSH^{UF}	-0.853*** (0.114)	-0.800*** (0.114)	-0.832*** (0.112)	-0.774*** (0.117)	-0.832*** (0.118)	-0.842*** (0.113)	-0.886*** (0.122)
L. Trade	0.351*** (0.107)						0.235* (0.13)
L. Manufacturing		1.486 (2.685)					-0.998 (2.869)
L. Education			0.0285 (0.0229)				-0.000 (0.0271)
L. Offshorability				0.155 (0.191)			-0.005 (0.207)
L. Markup					0.354** (0.145)		0.239* (0.134)
L. Employment						0.488*** (0.110)	0.328** (0.132)
Constant	6.661*** (0.699)	6.477*** (0.705)	6.406*** (0.702)	5.997*** (0.959)	6.086*** (0.678)	6.331*** (0.682)	6.246*** (0.934)
Observations	408	408	408	408	408	408	408
R-squared	0.349	0.332	0.338	0.332	0.338	0.349	0.359
<i>Panel B. 2SLS estimates</i>							
L. RSH^{UF}	-0.987*** (0.271)	-0.937*** (0.282)	-0.933*** (0.273)	-0.923*** (0.342)	-0.935*** (0.280)	-0.956*** (0.269)	-1.014*** (0.324)
F-stat	35.92	35.04	36.39	30.62	35.17	35.86	33.23

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

To further address concerns about omitted variable bias affecting our main results, we conduct a sensitivity analysis for the OLS regression of wage growth on automation (RSH^{UF}). Using the methodology proposed by Cinelli and Hazlett (2020), we simulate how much influence unobserved confounders – assumed to be uncorrelated with the observed controls – would need to have to challenge our results. We find a Robustness Value (RV) of 33%, implying that these unobserved factors would need to explain around one third of the remaining variation in both the treatment and the outcome to reduce the coefficient on automation to zero. To make the coefficient statistically insignificant, they would need to explain at least 27% of that residual variation. Figure 8 gives us a sense of scale using one of the observed variables as a reference. It shows that even an omitted

variable with three times the effect of trade exposure ($3/3x L_Trade$) in our regression would barely affect the estimated coefficient of automation (from -0.86 to -0.84).²⁷ Other checks using alternative benchmark variables – available upon request –, consistently support the stability of our results.

Figure 8. Omitted variable analysis benchmark with trade.



Note: Own elaboration following Cinelli and Hazlett (2020). We have used the Stata package produced by the authors: ‘sensemakr’.

6.2. Medium-term results

Although our baseline focus is on the overall long-term effects of inequality in the first two decades of the century, we also analyze wage dynamics over a ten-year period (medium term), using panel data across the sub-periods of 2000–2010 and 2010–2019. We do a preliminary comparison of automation share indices while controlling for baseline group wages in Table B2 (Panel A) in Appendix B. We use pooled data from the two ten-year periods, incorporating a dummy variable for the second decade (2019). Once again, we observe a robust and significant negative association between the initial

²⁷ The unadjusted coefficient of RSH^{UF} (-0.862) is not exactly the one obtained in Table 2a (-0.886) because the package proposed by Cinelli and Hazlett (2020) does not allow the use of weights. We have checked the robustness of the unweighted regression, whose results are available upon request.

automation share (regardless of the index) and wage growth. For instance, the coefficient for RSH^{UF} is -0.662 and RSH^C is -0.462. Similarly, when we turn to the Bartik IV estimates (Table B2, Panel B), the negative coefficients become larger in magnitude. For RSH^{UF} , the coefficient is -0.670, and for RSH^C , it is -0.469.

Now, we focus on the results for our main index of automation and the data pooled from the two ten-year intervals in our full specification with controls. The results for OLS are shown in Table 3 (Panel A), while the results for Bartik IV are represented in Table 3 (Panel B). We first observe that the dummy variable for the second decade is consistently negative and statistically significant across the OLS specifications in Table 3, indicating that, ceteris paribus, wage growth was systematically lower during the 2010-2019 decade compared to the 2000-2010 decade within this pooled framework.

Table 3. Medium-run inequality change (RSH^{UF} , 2000-2010-2019).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. OLS estimates</i>							
Year = 2019	-0.268*** (0.025)	-0.287*** (0.025)	-0.260*** (0.028)	-0.278*** (0.029)	-0.240*** (0.030)	-0.266*** (0.026)	-0.211*** (0.032)
L. Wage Past	-0.579*** (0.054)	-0.555*** (0.054)	-0.579*** (0.055)	-0.552*** (0.054)	-0.568*** (0.055)	-0.570*** (0.054)	-0.607*** (0.055)
L. RSH^{UF}	-0.711*** (0.078)	-0.666*** (0.078)	-0.705*** (0.078)	-0.637*** (0.088)	-0.718*** (0.085)	-0.696*** (0.078)	-0.776*** (0.092)
L. Trade	0.311*** (0.0704)						0.141 (0.098)
L. Manufacturing		3.768*** (1.325)					2.612 (1.590)
L. Education			0.0320*** (0.0114)				0.0158 (0.0143)
L. Offshorability				0.0753 (0.137)			0.001 (0.135)
L. Markup					0.345*** (0.116)		0.326*** (0.116)
L. Employment						0.411*** (0.113)	0.203 (0.124)
Constant	6.042*** (0.561)	5.866*** (0.564)	5.807*** (0.551)	5.603*** (0.719)	5.481*** (0.555)	5.722*** (0.577)	5.45*** (0.691)
Observations	816	816	816	816	815	816	815
R-squared	0.391	0.383	0.384	0.378	0.382	0.385	0.399
<i>Panel B. 2SLS estimates</i>							
L. RSH^{UF}	-0.709*** (0.126)	-0.670*** (0.127)	-0.688*** (0.119)	-0.649*** (0.170)	-0.717*** (0.127)	-0.711*** (0.121)	-0.753*** (0.184)
F-stat	32.16	32.78	32.41	27.36	30.20	31.67	24.78

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

In the pooled medium-term OLS analysis with the whole set of controls (Table 3, column 7), the coefficient for our primary routinization measure, RSH^{UF} , remains robustly negative and highly significant at -0.776, pointing to a substantial wage-decreasing effect associated with higher initial levels of automatable tasks. The group's exposure to higher average industry markups shows a highly significant association in this full pooled model, while other controls are not significant.

The Bartik IV counterpart in Table 3 reports a similar negative coefficient for the RSH^{UF} index of -0.753 in its full specification (column 7). This estimate is supported by a strong F-statistic of 24.78, reinforcing the causal interpretation of its detrimental wage effect in the medium run. Similar patterns for the RSH^C measure are found in Appendix B (see Table B3), with the full Bartik model yielding a coefficient of -0.574 and an F-statistic of 20.68. The Robustness Value of the full specification model, obtained with the Cinelli-Hazlett bounding procedure, is again reassuring, ascending to 31.50%.

6.3. Short-term results and the early irruption of AI

Finally, we explore the short-run period from 2015 to 2019, characterized by the early outbreak of AI technologies. We incorporate explicit measures of AI exposure alongside our RSH^{UF} index. As discussed above, the AI exposure variables considered, based on Felten et al. (2018) and Webb (2020) capture different facets of potential AI impact on occupations, averaged at the demographic group level.

Table 4 presents the OLS estimates for this 2015-2019 period in Panel A. When the AI measures are included individually (with only the lagged wage as a control, columns 2-3), both show positive and highly significant coefficients, suggesting that groups with higher baseline exposure to these AI metrics tend to experience faster wage growth in this recent period. This positive association largely persists when the full set of control variables is included (columns 5-6), and contrasts with the significant negative coefficients associated with RSH^{UF} (columns 1 and 4). When the RSH^{UF} index is introduced alongside each AI measure and all controls (Table 4, columns 7-8), it still displays a negative and statistically significant coefficient; for example, it is -0.488 when included with the Felten AI measure. Importantly, even with the RSH^{UF} measure in the model, the AI exposure index based on Felten et al. (2018) maintains a positive and significant coefficient of 0.256, while the Webb measure remains significant at 0.563.

These OLS results suggest that even in the recent 2015-2019 period, a higher initial routine share is linked to slower wage growth. Furthermore, the Felten and Webb measures of AI exposure appear to capture a dimension of AI's impact that is positively associated with wage growth and is distinct from the traditional routinization effect.

The regional employment share emerges as a particularly relevant factor, displaying positive and significant coefficients in most of the fully controlled specifications (columns 4-8). For example, its coefficient is 1.122 in columns 7 and 1.013 in column 8, when the Felten and Webb measures are considered. This importance of the employment share suggests that in the economic context of 2015-2019, local labor market dynamism played a role in wage growth.

Table 4. Short-run inequality change (AI and RSH^{UF}, 2015-2019).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	RSH ^{UF}	Felten	Webb	All (RSH ^{UF})	All (Felten)	All (Webb)	All (RSH ^{UF} & Felten)	All (RSH ^{UF} & Webb)
<i>Panel A. OLS estimates</i>								
L. Wage Past	-0.503*** (0.095)	-0.452*** (0.094)	-0.423*** (0.0934)	-0.525*** (0.110)	-0.486*** (0.107)	-0.590*** (0.106)	-0.539*** (0.109)	-0.620*** (0.107)
L. Trade				-0.0862 (0.154)	-0.101 (0.154)	-0.226 (0.148)	-0.092 (0.152)	-0.200 (0.146)
L. Manufacturing				3.341 (2.670)	3.043 (2.707)	4.785* (2.480)	3.200 (2.691)	4.564* (2.493)
L. Education				-0.097* (0.056)	-0.072 (0.053)	-0.044 (0.051)	-0.084 (0.053)	-0.054 (0.051)
L. Offshorability				0.099 (0.189)	0.395*** (0.198)	0.651*** (0.184)	0.073 (0.192)	0.338* (0.194)
L. Markup				0.118 (0.357)	0.183 (0.379)	-1.118*** (0.370)	0.216 (0.371)	-0.834** (0.361)
L. Employment				1.202** (0.483)	0.851* (0.454)	0.817* (0.472)	1.122** (0.475)	1.013** (0.485)
L. RSH ^{UF}	-0.637*** (0.118)			-0.626*** (0.146)			-0.488*** (0.144)	-0.343*** (0.125)
L. Felten		0.595*** (0.126)			0.472*** (0.136)		0.256** (0.126)	
L. Webb			0.433*** (0.134)			0.678*** (0.141)		0.563*** (0.133)
Constant	5.160*** (0.970)	3.982*** (0.874)	3.770*** (0.865)	5.144*** (1.180)	3.032*** (1.025)	5.084*** (1.101)	4.883*** (1.162)	6.133*** (1.213)
Observations	408	408	408	406	406	406	406	406
R-squared	0.203	0.185	0.160	0.217	0.203	0.251	0.224	0.262
<i>Panel B. 2SLS estimates</i>								
L. RSH ^{UF}	-0.512*** (0.119)			-0.530* (0.296)			-7.830 (10.06)	-0.682** (0.454)
L. Felten		0.832*** (0.196)			0.887** (0.403)		-7.040 (9.169)	
L. Webb			4.750 (3.980)			2.611 (1.612)		0.938*** (0.304)
F-stat (RSH ^{UF})	12.96			6.55			3.68	3.34
F-stat (AI)		16.07	0.88		8.09	1.39	4.81	12.58

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

The Bartik IV estimates for this short-run period, presented in Table 4, Panel B, highlight the complexities involved in identifying the causal impacts of AI alongside ongoing routinization. When the RSH^{UF} and AI measures are instrumented without controls (columns 1–3), the RSH^{UF} and Felten indices show, respectively, significant negative and positive coefficients, with an F-statistic above the conventional threshold. However, the Webb index is not significant, and its F-statistic is close to zero. When instrumented alone with full controls, both the RSH^{UF} and Felten indices remain significant, although their instrument is below the conventional F-statistic threshold. In this regard, note that the Robustness Value of the full specification model, obtained with the Cinelli-Hazlett bounding procedure, amounts to 21.54% for the RSH^{UF} index and 19.94% for the Felten Index. When attempting to instrument the RSH^{UF} and the Felten indices simultaneously (column 7), the significance of the coefficients disappears, and their instruments are weak. For the Webb index (column 8), the coefficient becomes positive and significant, with an F-statistic above 10. In other words, although the results of the OLS estimates clearly show that AI has had a positive effect on wage growth, the period considered (2015-2019) does not allow for an accurate estimate of the magnitude of the causal effect it has had.

7. Distributional Consequences of new technologies: A Counterfactual Analysis

The previous section established the causal impact of new technologies, proxied by our preferred measure of automation RSH^{UF} , on the average wage growth of different demographic groups in Spain. Building upon those findings, this section studies the distributional consequences of these technological impacts. As described in Section 3.2., we employ a counterfactual approach to assess how automation and AI exposure have shaped wage inequality, measured by the evolution of inequality indices and by specific wage shares. This approach also allows us to assess the impact of automation on the gender wage gap, and the education premium, and to understand which demographic groups have been most affected.

7.1. Impact on overall wage inequality and wage shares

Table 5 presents the core findings of our counterfactual exercise for the shift-share estimates, displaying the percentage difference $RD_t(\cdot)$ in various inequality metrics

between the counterfactual y_t (no-technology effect) and actual μ_t wage distributions: $RD_t(I) = \frac{I(y_t) - I(\mu_t)}{I(\mu_t)} * 100$. Inequality estimates on the original wage distributions can be found in Table A3 (Appendix A). It is crucial to interpret negative values in Table 5 as indicating that technology has effectively *increased* observed inequality; that is, inequality is estimated to be lower in the counterfactual which is absent of that technological effect. The analysis covers standard inequality measures such as the Gini coefficient and the Mean Log Deviation (MLD), as well as wage shares for the top 10%, bottom 50%, bottom 10%, and the middle 50% of the distribution. Results are presented for different periods using our primary automation measures, RSH^{UF} and RSH^C , and the two AI exposure indices.

Table 5. Inequality results (% change in the counterfactual – technology free – distribution of wages relative to the observed distribution).

Series	Variable of technology	Gini	MLD	Wage share Top 10%	Wage share Bottom 50%	Wage share Bottom 10%	Wage share Middle 50%
Long (2000-2019)	RSH^{UF*}	-21.54	-34.82	-3.90	0.83	2.16	0.15
	RSH^C*	-20.64	-35.08	-5.05	0.56	4.06	0.90
Medium (2000-2010-2019)	RSH^{UF*}	-19.33	-30.76	-1.66	0.91	2.84	0.64
	RSH^C*	-11.78	-18.92	-1.99	0.51	-0.06	0.55
Short (2015-2019)	RSH^{UF*}	-14.38	-22.85	-2.72	0.69	0.30	1.02
	Felten*	-9.92	-13.35	-2.15	0.43	0.89	0.57
	Webb	13.62	36.80	0.70	0.02	-1.71	1.27

Note: Own elaboration. MLD stands for Mean Log Deviation. Changes are measured as $RD_t(I) = \frac{I(y_t) - I(\mu_t)}{I(\mu_t)} * 100$, where y_t is the counterfactual distribution and μ_t the observed distribution of log wages. Negative values imply a lower value in the counterfactual distribution (without technology). There is an asterisk * wherever the coefficient is significant in the regression. Results for the 2015-2019 period refer to regressions with AI but without automation variables.

A consistent narrative emerges from Table 5 regarding the impact of automation: both our preferred RSH^{UF} unfixed measure and the classic RSH^C index are found to have increased overall wage inequality. For example, over the long run (2000-2019), the Gini coefficient would have been approximately 21.5% lower in the absence of the RSH^{UF} effect, and the MLD would have been 35% lower. Similarly, for the RSH^C index, the Gini

would have been about 20.6% lower and the MLD nearly 35% lower during the same period. This pattern largely holds for the medium-term periods as well.

For the RSH^{UF} measure in the long run (2000-2019), the analysis indicates that the wage share of the top 10% would have been about 3.9% smaller in the counterfactual scenario without technological change (Table 5). This means that automation, as we measure it, effectively increased the observed wage share of this top group. Conversely, the share of the bottom 50% would have been 0.83% higher, and that of the bottom 10% would have been 2.16% higher in the counterfactual world, suggesting that automation contributed to a redistribution of wage shares away from the bottom and middle towards the top of the distribution, thus increasing inequality in the period 2000-2019 in Spain.²⁸

The results for the AI variables (without including the RSH^{UF} index in the AI regressions) in the short run (2015-2019), are also found to have increased overall wage inequality. First, as a benchmark, automation measured with RSH^{UF} is found to have increased inequality by around 14.4%. In the same vein, the Felten AI exposure index is associated with a very substantial increase in inequality; the Gini coefficient would be over 9.9% lower in its absence. The Webb AI index is not significant in Table 4, so its association to a decrease in inequality should be interpreted with due caution.

As discussed earlier, a key finding that warrants detailed explanation is how AI exposure, which shows mostly positive coefficients in the wage growth regressions (Table 4), can contribute to an increase in overall wage inequality in the same direction than automation (Table 5) which has a negative coefficient for wage growth. Here we must note that an increase in average wages does not automatically translate to a decrease in overall inequality. It all depends on how the variable at stake (AI or automation) is allocated across the wage distribution. If the groups benefiting most from AI (i.e., those with high exposure to AI) are already situated at the upper end of the wage distribution, their accelerated wage growth will further distance them from those at the middle or lower end. This stretching of wage distribution leads to an increase in inequality measures like the Gini coefficient or the Mean Log Deviation (MLD). In contrast, automation (RSH^{UF} and RSH^C) has a negative impact on wage growth but is allocated mostly around the middle of the distribution, as we discussed in Section 5. This tends to harm middle-skill workers, hollowing out the middle of the wage distribution and increasing inequality. Thus, while

²⁸ Note that the magnitude of these effects on shares is somewhat moderated using logarithmic wages in the underlying regressions, which tends to smooth the right tail of the distribution.

the mechanism of wage impact differs (AI potentially boosting mainly higher-skilled groups; automation depressing wages for mostly middle-skilled groups), both ultimately lead to a more unequal overall wage structure.

Finally, AI variables can yield apparently uneven and notably large effects in the counterfactual analysis (Tables 5 and 6). However, as noted above, the two AI indices are constructed using distinct methodologies and aim to capture different facets of AI's potential influence on occupations.²⁹ It is expected that these different dimensions of AI could have somewhat different, and variably intense, effects on the wage structure. However, the index for which the regression coefficient is significant (Felten) is largely aligned qualitatively as inequality increasing in most metrics.

7.2. Consequences for gender and education wage gaps

Beyond overall inequality, our counterfactual approach allows us to study how technology has influenced specific structural wage gaps: the gender wage gap and the education premium. Table 6 presents these findings, showing the relative percentage difference between the gaps in the counterfactual and the observed distribution. Positive values imply a larger gap in the counterfactual distribution (i.e., without the impact of technology). This means that technology has served to reduce the gap, while negative values suggest that technology has widened it.

The RSH^{UF} automation estimate indicates that this form of technological change has had little impact on wage disparities between men and women in the long term (2000-2019). Without the effect of automation as measured by RSH^{UF} , the gender gap would have been approximately 3% smaller, while the RSH^c measure suggests it would have been almost 30% smaller in the long term. This is consistent with a higher exposure to automation for women, especially in the low-education group, when measured with the RSH^c index, as can be observed in Figures 3a and 3b.³⁰ That said, it seems that automation reduced the

²⁹ Recall that the Felten et al. (2018) scores link AI progress to occupational abilities while Webb's scores are based on textual overlap between AI patents and job descriptions.

³⁰ Note that a key feature of RSH^{UF} is that it is agnostic to whether manual tasks are automatable or not, while RSH^c considers them not automatable (they are subtracted from the automation exposure index). Many occupations with high manual task intensity are male-prevalent (construction, transport drivers) which explains the higher automation relative impact for male (recall indices are standardized) in the RSH^c index.

gender gap in the medium term which is also consistent with the result for RSH^{UF} in the short term. Similarly, the impact on the gender gap of the AI index that is significant in the regression analysis (Felten index), suggests that AI contributed to reducing it (the gender gap is almost 32% greater in the counterfactual scenario without AI-impact).

Table 6. Effects on the gender gap and education premium (% change in the counterfactual distribution of wages relative to the observed distribution).

Data Series	Variable of technology	Gender Gap	Education Premium
Long (2000-2019)	RSH^{UF*}	-2,95	-43,03
	RSH^C*	-29,41	-63,89
Medium (2000-2010-2019)	RSH^{UF*}	28,40	-42,33
	RSH^C*	9,18	-25,14
Short (2015-2019)	RSH^{UF*}	22,33	-26,99
	Felten*	32,01	-34,00
	Webb	-187,04	-44,90


Note: Own elaboration The gender gap is computed as $Gap_S(x) = \bar{x}_{men} - \bar{x}_{women}$, where x is either the observed wage distribution μ or the distribution of counterfactual wages y . Following the same nomenclature, the education premium is computed as $Gap_E(x) = \bar{x}_{high} - \bar{x}_{low}$. There is an asterisk * wherever the coefficient is significant in the regression. Results for the 2015-2019 period refer to regressions with AI but without automation variables.

Concerning the education premium, which is the difference in wages between highly and less-educated workers, automation, as measured by both RSH^{UF} and RSH^C , is found to consistently increase this premium for all periods considered. For the long run (2000–2019), the education premium would have been substantially smaller – by about 43% with RSH^{UF} and 64% with RSH^C – in the absence of these automation effects. This finding robustly supports the notion that these automation technologies tend to complement higher skills while substituting for tasks typically performed by low- and middle-educated workers. For the medium run (2000–2010–2019), the effect according to the RSH^{UF} index is similar in size and direction (-42%), while the effect according to the RSH^C index is of the same sign, but much smaller in size (-25%). Furthermore, both AI measures also point in this direction. Indeed, the Felten index is associated with an increase in the education premium of 34%.

7.3. Heterogeneous technological effects across demographic profiles

To gain a more detailed understanding of the impact of automation and AI, we examine how their effects on wages vary across specific demographic groups, defined by combinations of sex, age and level of education. Figure 9 (long-term, 2000-2019) illustrates the percentage change in log wages attributable to the RSH^{UF} automation index, measured as the percentage difference between the counterfactual (automation absent) wage and the actual wage. Thus, positive values on this figure indicate a greater adverse impact of automation on a group's wage growth; that is, their observed wages are lower than they would have been in the absence of this technological effect. The dashed black line represents the average adverse impact across all groups.

Figure 9. The effect of automation (RSH^{UF}) by sex, age and education (2000-2019).



Note: Own elaboration. The dashed black line shows the average difference rate between counterfactual and actual wage across all groups.

Some patterns emerge from these visualizations. Firstly, there is a clear educational gradient: low-educated individuals consistently experience the most substantial negative wage impacts from automation, with their outcome lines typically positioned well above the average adverse effect dashed line. Medium-educated individuals face a smaller negative effect, and mostly only the youngest group. In contrast, highly educated individuals are often situated below the average line, implying that automation has either harmed their wage growth to a lesser extent or, in some instances, may have even

conferred relative benefits. This visual evidence strongly corroborates the earlier finding that automation, as measured by RSH^{UF} , tends to widen the education premium.

Secondly, when comparing by gender within similar education and age categories, the effect is unclear. While low educated females exhibit more adverse impacts (their lines are higher in Figure 9), their male counterparts with high or medium education seem to have been more affected. This observation is consistent with the finding that the RSH^{UF} measure of automation has barely contributed to a reduction in the overall gender wage gap.

Thirdly, there appears to be a general trend for the adverse impact of automation to lessen with age within specific sex and education cohorts, suggesting that older workers might experience less negative wage consequences compared to their younger peers, except for older women in the low education group. These patterns are largely consistent between the long-term and medium-term datasets (see Figure C1 in Appendix C for results in the mid-term 2000-2010-2019 period), and similar qualitative results are obtained when using the RSH^C automation index, differing mainly in the magnitude of the changes.³¹


Figure 10 shows the results of the same analysis for the Felten index of AI. Since the effect of AI on wage growth is positive (see Table 4), note that the average difference rates between the counterfactual and actual wage changes are negative. As before, there is a clear educational gradient: low-educated individuals consistently experience the most substantial negative impacts on wages from AI; medium-educated individuals face a smaller negative effect; and highly educated individuals always fall below the average line. However, when comparing by gender within similar educational and age categories, males experience more adverse impacts (their lines are higher in Figure 10) at all educational levels. Finally, older workers experienced less negative wage consequences compared to their younger peers, except for older women in the high education group.

Figure 11 plots the RSH^{UF} metric (2000–2019 and 2000–2010–2019 periods) and the Felten AI index (2015–2019 period), offering a visualization of the impacts by averaging the technological effect for socioeconomic categories by region. Note that, to make them comparable, the effects of AI have been multiplied by -1, as they are negative, since the counterfactual wage values are lower than the observed ones. The plot shows that the effect on the average wage in absolute terms is greater for long-term automation (2000–

³¹ These results are not shown but they are available from the authors upon request.


2019) than for medium-term automation (2000–2010–2019) and AI (2015-2019). Perhaps more importantly, there is little variation in the three indicators across regions. While there are some regional variations in the impact of automation, with regions such as Navarra and Aragon experiencing larger average adverse technological impacts on wages and regions such as Baleares and Extremadura experiencing smaller average effects than the national mean, these differences are not too noticeable. Some regions, such as Extremadura, Murcia and Asturias, show larger average positive impacts of AI on wages, while others, such as Valencia, Catalonia and La Rioja, show the greatest adverse average effects. However, the differences between regions are again small.

Figure 10. The effect of AI (Felten index) by sex, age and education (2000-2019).



Note: Own elaboration. The dashed black line shows the average difference rate between counterfactual and actual wage across all groups.

Figure 11. The change on wages across regions.
 (RSH^{UF} , 2000-2019 and 2000-2010-2019; Felten, 2015-2019)



Note: Own elaboration Values are estimated as in Equation 13 and aggregated by demographic categories. Values for the Felten index are multiplied by -1 for comparability.

8. Conclusions

Despite the massive economic growth and social development experienced over the past two centuries, reluctance towards rapid technological change remains significant in our societies and is nowadays a major concern for citizens and policymakers. For this reason, a better understanding of the relationship between technological progress and inequality can inform public debate and help implement appropriate economic policies. This issue has only been partially addressed by the literature, and this study aims to fill that gap.

Our findings support the widespread public concern of the negative effects of new technologies on inequality. After developing a dynamic measurement of automation

exposure at different waves, we estimate a new, ‘unfixed’ automation index. Using an ample number of sociodemographic groups as units of analysis, this index captures not only changes in shares of different occupations in the labor force, but also the degree of routinization of occupations over time.

Applying this index in a shift-share instrumented regression framework and constructing a technology-free counterfactual wage distribution enables us to accurately measure the impact of automation and AI on wage inequality. Comparing the counterfactual wage distribution with the observed one reveals that automation was a significant contributor to wage inequality in the period 2000-2019 in Spain, while increasing the education gap. The effect on wages of the AI exposure index during the early introduction of these technologies (2015-2019) also indicates an inequality-increasing effect. Interestingly, our findings reveal that traditional automation and new AI technologies increase inequality through a different mechanism. Exposure to automation – which is negatively related with wage growth – is greater at the central part of the wage distribution, tending to depress wages for workers in the middle of the wage range. Exposure to AI technologies, on the other hand, is found to have a positive relation with wage growth and to be more prevalent at the top of the wage distribution, thus increasing inequality by pulling away the top tail of the salary range. In either case, our results confirm that the public is right to fear the negative effects of automation on inequality and could inform policies aiming to address these imbalances without losing the overall benefits of new technologies. Economists have proposed two main approaches to deal with these disruptive effects. The most common perspective is to address such adverse consequences through a combination of ex-post redistribution to support those left behind by technological progress and investment in education and training to help more people catch up. Others argue that current tax systems that tend to favor capital investment distort economic efficiency, encouraging companies to purchase machinery over hiring more workers driven by tax incentives rather than by efficiency gains. According to this perspective, to level the playing field taxes on labor and capital should be guided by their relative supply elasticities, which are found to be rather similar in practice.

Whatever the appropriate economic policies, a good understanding of the effects that new technologies have on wage distribution is first required. We hope that our contribution will help clarify how new technologies affect wage inequality and how different types of workers are affected. Just as the Industrial Revolution began an unprecedented

transformation that ultimately benefited everyone, AI systems and new technologies have the potential to do the same. However, the distributional effects and potential negative impact on vulnerable groups must be understood and addressed.

References

- Abadie, A., Athey, S., Imbens, G.W. and Wooldridge, J.M. (2023): “When should you adjust standard errors for clustering?”, *Quarterly Journal of Economics*, 138, 1-35.
- Acemoglu, D. and Autor, D. (2011): “Skills, tasks and technologies: implications for employment earnings” in Ashenfelter, O. and Card, D. (eds.), *Handbook of Labor Economics*, vol. 4B, 1043–1171. Elsevier: Amsterdam.
- Acemoglu, D. and Restrepo, P. (2020): “Robots and jobs: evidence from US labor markets”, *Journal of Political Economy*, 128, 2188-2244.
- Acemoglu, D. and Restrepo, P. (2022): “Tasks, automation, and the rise in US wage inequality”, *Econometrica*, 90, 1973-2016.
- Acemoglu, D., Autor, D., Hazell, J. and Restrepo, P. (2022): “Artificial intelligence and jobs: evidence from online vacancies”, *Journal of Labor Economics*, 40, S293-S340.
- Albanesi, S., Dias da Silva A., Jimeno, J.F., Lamo, A. and Wabitsch A. (2024), “New Technologies and Jobs in Europe”, *Economic Policy*, 40, 71–139.
- Ackerberg, D.A., Caves, K. and Frazer, G. (2015), “Identification properties of recent production function estimators”, *Econometrica*, 83, 2411–2451.
- Ahmed, A. (2023): “Global Inequality of Hourly Income, 1980-2020”, *World Inequality Lab Technical Note 2023/03*.
- Atkinson, A.B. (1970): “On the measurement of inequality”, *Journal of Economic Theory*, 2, 244–263.
- Autor, D.H., Levy, F. and Murnane, R.J. (2003): “The skill content of recent technological change: an empirical exploration”, *Quarterly Journal of Economics*, 118, 1279–1333.
- Autor, D.H., Katz, L.F. and Kearney, M.S. (2006): “The polarization of the U.S. labor market”, *American Economic Review*, 96, 189–194.
- Autor, D.H. and Dorn, D. (2013): “The growth of low-skill service jobs and the polarization of the US labor market”, *American Economic Review*, 103, 1553–1597.
- Ayala, L., Pérez, A. and Prieto-Alaiz, M. (2022): “The impact of different data sources on the level and structure of income inequality”, *SERIEs*, 13, 583-611.
- Blinder, A.S. (2009): “How many U.S. jobs might be offshorable?”, *World Economics*, 10, 41–78.
- Blinder, A.S. and Krueger, A.B. (2013): “Measuring offshorability: A survey approach”, *Journal of Labor Economics*, 31(S1), S97–S128.
- Böhm, M.J., Gaudecker, H.-M. and Schran, F. (2024): “Occupation growth, skill prices, and wage inequality”, *Journal of Labor Economics*, 42, 201-243.
- Bonhomme, S. and Hospido, L. (2016), “The cycle of earnings inequality: evidence from Spanish social security data”, *Economic Journal*, 127, 1244–1278.
- Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *American Economic Review Papers and Proceedings*, 108:43–47.

- Brynjolfsson, E., Li, D. and Raymond, L. (2025), “Generative AI at work.” *Quarterly Journal of Economics*, 140, 889–942.
- Card, D. and Lemieux, T. (1994): “Changing wage structure and black-white wage differential”, *American economic Review: Papers and Proceedings*, 84, 29-33.
- Card, D., Kramarz, F. and Lemieux, T. (1999). “Changes in the relative structure of wages and employment: A comparison of the United States, Canada, and France”, *The Canadian Journal of Economics*, 32, 843–877.
- Cinelli, C. and Hazlett, C. (2020): “Making sense of sensitivity: Extending omitted variable bias”, *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82, 39–67.
- Cowell, F.A. (1980): “On the structure of additive inequality measures”, *Review of Economic Studies*, 47, 521–531.
- De Loecker, J. and Eeckhout, J. (2025): “Global Market Power”, *NBER Working Paper* N.º 24768. Revised version, 6th March 2025.
- DiNardo, J., Fortin, N.M. and Lemieux, T. (1996): “Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach”, *Econometrica*, 64, 1001–1044.
- Doorley, K., Gromadzki, J., Lewandowski, P., Tuda, D. and Van Kerm, P. (2023): “Automation and income inequality in Europe”, *IZA Discussion Paper Series* No. 16499.
- Felten, E. W., Raj, M. and Seamans, R. (2018, May), “A method to link advances in artificial intelligence to occupational abilities”, *AEA Papers and Proceedings*, 108, 54–57.
- Felten, E., Raj, M. and Seamans, R. (2019), “The effect of artificial intelligence on human labor: An ability-based approach”, *Academy of Management Proceedings*.
- Ferri, B. (2022): “Novel shift-share instruments and their applications”, *Boston College Department of Economics*, mimeo.
- Fossen, F.M. and Sorgner, A. (2022). “New digital technologies and heterogeneous wage and employment dynamics in the United States: Evidence from individual-level data”, *Technological Forecasting and Social Change*, 175, 121381.
- Foster, J. and Shneyerov, A. (2000): “Path independent inequality measures”, *Journal of Economic Theory*, 91, 199–222.
- Foster, J.E. and Ok, E.A. (1999): “Lorenz dominance and the variance of logarithms”, *Econometrica*, 67, 901–907.
- Frey, C.B. and Osborne, M.A. (2017): “The future of employment: how susceptible are jobs to computerisation?”, *Technological Forecasting and Social Change*, 114, 254-280.
- Frey, C.B. (2019): *The technology trap: Capital, labor, and power in the age of automation*. Princeton University Press, Princeton.
- Goos, M. and Manning, A. (2007): “Lousy and lovely jobs: the rising polarization of work in Britain”, *Review of Economics and Statistics*, 89, 118–133.

- Goos, M., Manning, A. and Salomons, A. (2014): “Explaining job polarization: routine-biased technological change and offshoring”, *American Economic Review*, 104, 2509–2526.
- Graetz, G. and Michaels, G. (2018): “Robots at work”, *Review of Economics and Statistics*, 100, 753–768.
- Heckman, J.J. and Honore, B.E. (1990): “The empirical content of the Roy model”, *Econometrica*, 58, 1121–1149.
- Juhn, C., Murphy, K.M. and Pierce, B. (1993): “Wage inequality and the rise in returns to skill”, *Journal of Political Economy*, 101, 410–442.
- Jurkat, A., Klump, R. and Schneider, F. (2022): “Tracking the Rise of Robots: The IFR Database”. *Jahrbücher für Nationalökonomie und Statistik*, 242, 669–689.
- Korinek, A. and Stiglitz, J.E. (2019): “Artificial intelligence and its implications for income distribution and unemployment” in *The Economics of Artificial Intelligence: An Agenda* (pp.349-390) edited by A. Agrawal, J. Gans and A. Goldfarb, University of Chicago Press.
- Lemieux, T. (2006): “Increasing residual wage inequality: composition effects, noisy data, or rising demand for skill”, *American Economic Review*, 96, 461–498.
- Levy, F. and Murnane, R.J. (1992): “U.S. earnings levels and earnings inequality: A review of recent trends and proposed explanations”, *Journal of Economic Literature*, 30, 1333–1381.
- Lewandowski, P., Park, A., Hardy, W. and Du, Y. (2019): “Technology, skills, and globalization: explaining international differences in routine and nonroutine work using survey data”, *IZA Discussion Paper Series No. 12339*.
- Pijoan-Mas, J. and Sánchez-Marcos, V. (2010), “Spain is different: falling trends of inequality”, *Review of Economic Dynamics*, 13, 154–178.
- Rodríguez, J.G. and Salas, R. (2014), “The Gini coefficient: majority voting and social welfare”, *Journal of Economic Theory*, 152, 214–223.
- Rodríguez, J.G. and Sebastian, R. (2022): *Automatización, polarización laboral e igualdad socioeconómica*. Fundación COTEC.
- Sala-i-Martin, X. (1996): “The classical approach to convergence analysis”, *Economic Journal*, 106, 1019–1036.
- Shorrocks, A.F. (1980): “The class of additively decomposable inequality measures”, *Econometrica*, 48, 613–625.
- Shorrocks, A.F. and Foster, J.E. (1987): “Transfer sensitive inequality measures”, *Review of Economic Studies*, 54, 485–497.
- Solt, F. (2020): “Measuring Income Inequality Across Countries and Over Time: The Standardized World Income Inequality Database”, *Social Science Quarterly* 101, 1183-1199. SWIID Version 9.8, December 2024.
- Staiger, D. and Stock, J.H. (1997): “Instrumental variables regression with weak instruments”, *Econometrica*, 65, 557–586.
- Webb, M. (2020): The impact of artificial intelligence on the labor market. *SSRN Electronic Journal*.

APPENDIX A

In this appendix we describe in more detail the data sources of our measures of employment, the routineness, and offshorability of occupations as well as our measures of artificial intelligence. In addition, we include the statistics of the dependent variable.

A1. Employment

The European Union Labor Force Survey (EU-LFS) gathers data from 29 European countries, collected at the national level using harmonized classifications, definitions, and variables. These data are standardized and centrally processed by Eurostat to ensure comparability across countries. For this study, we concentrate exclusively on Spain, using data from 1992 to 2019. This period provides a consistent and extensive time series that captures significant transformations in the Spanish labor market, including the impacts of technological change, globalization, labor market reforms, and economic crises such as the Great Recession.

Our sample is restricted to individuals classified as employed under the ILO's employment criteria (as indicated by the ELFS variable *ilostat*) aged between 16 and 65. We exclude a minimal number of unpaid family workers, identified through the professional status variable (*stapro*), noting that this exclusion does not materially affect the results. Employment is quantified as the number of employed persons (weighted by ELFS survey weights).

As commented in the main text, occupations are classified using ISCO-88 (three-digit) from 1992 to 2008 and ISCO-08 thereafter; sectors are coded using NACE Rev.1 (one-digit) through 2009 and NACE Rev.2 from 2010 onward. Moreover, to ensure consistency over time, all data are harmonized to include both occupational and sectoral classifications under each system.³²

Occupations representing a negligible share of total employment are excluded from the analysis. Specifically, we drop “Armed Forces” (ISCO-88: 01; ISCO-08: 01), “Agricultural, Fishery and Related Laborers” (ISCO-88: 92), and “Subsistence Farmers, Fishers, and Gatherers” (ISCO-08: 63). To retain broader occupational coverage, we merge certain categories across classification systems. In ISCO-88, we combine

³² The crosswalks can be provided upon request.

“Teaching Associate Professionals” (33) with “Other Associate Professionals” (34); in ISCO-08, we merge “Cleaners and Helpers” (91) with “Food Preparation Assistants” (94).

A2. Routineness

Following the methodology developed by Autor et al. (2003), we construct task-based measures using five indicators derived from O*NET and consistent with subsequent work (e.g., Acemoglu and Autor, 2010; Autor and Dorn, 2013), we aggregate these into three composite task categories (see Table A1). The abstract task index combines non-routine cognitive–analytical tasks (“analyzing data/information,” “thinking creatively,” and “interpreting information for others”) with non-routine cognitive–interpersonal tasks (“establishing and maintaining personal relationships,” “guiding, directing, and motivating subordinates,” and “coaching/developing others”). The routine task index includes both routine cognitive tasks (“importance of repeating the same tasks,” “importance of being exact or accurate,” and “structured vs. unstructured work”) and routine manual tasks (“pace determined by speed of equipment,” “controlling machines and processes,” and “spending time making repetitive motions”). The manual task index reflects non-routine manual activities, including “operating vehicles, mechanized devices, or equipment,” “using hands to handle, control, or feel objects, tools, or controls,” “manual dexterity,” and “spatial orientation.”

The Routine Task Intensity (RTI) index used in this paper is constructed from the O*NET database, where task-level scores are assigned by trained analysts following standardized protocols. We assume that the task content of occupations in the U.S. mirrors that of Spain. To align ONET data with Spanish occupational classifications, we undertake a two-step mapping procedure (see Table A2). First, we link ONET task items – originally coded in ONET-SOC –to the Standard Occupational Classification (SOC) using the crosswalk developed by the O*NET project. Second, we translate SOC codes into the International Standard Classification of Occupations (ISCO-88 or ISCO-08, at the three-digit level) via the official ILO crosswalk, yielding aggregated ISCO categories compatible with our data.

Unlike Autor and Dorn (2013), who compute the Routine Task Intensity (RTI) index for a single base year, we construct RTI measures for multiple time points: 1992, 1995, 2000, 2005, 2010, 2015, and 2019.

Table A1. Abstract, Routine and Manual tasks.

Tasks in Autor and Dorn (2013)	Tasks in Autor, Levy and Murnane (2023)	Question in O*NET
Abstract	Non-routine cognitive: Analytical	Analysing data/information Thinking creatively Interpreting information for others
	Non-routine cognitive: Interpersonal	Establishing and maintaining personal relationships Guiding, directing and motivating subordinates Coaching/developing others
Routine	Routine cognitive	Importance of repeating the same tasks Importance of being exact or accurate Importance of being exact or accurate
	Routine manual	Pace determined by speed of equipment Controlling machines and processes Spend time making repetitive motions
Manual	Non-routine manual: phys. adaptability	Operating vehicles, mechanized devices, or equipment Spend time using hands to handle, control or feel objects, tools or controls Manual dexterity Spatial orientation

Note: Own elaboration.

A3. Offshorability

Blinder (2009) and Blinder and Krueger (2013) argue that any occupation not requiring physical presence – whether its core tasks are abstract, routine, or manual – can be outsourced. Following Acemoglu and Autor (2011), we adopt their offshorability framework and construct an index using data from O*NET. Specifically, we draw on the following task items, each reverse-coded to reflect higher offshorability: “face-to-face discussion”, “assisting and caring for others”, “performing for or working directly with the public”, “inspecting equipment or structures”, “handling and moving objects”, and “maintaining or repairing mechanical and electronic equipment”. To align the task data with our occupational classification, we apply a two-step mapping: first from ONET-SOC to SOC, and then from SOC to ISCO.³³ Departing from Acemoglu and Autor (2011), who construct a single cross-sectional measure, we compute offshorability indices for multiple years – 1992, 1995, 2000, 2005, 2010, 2015, and 2019 – to capture changes over time.

Table A2. Crosswalks from one source to another.

Year	ONET-SOC code	SOC code	ISCO code
2019	ONETSOC2019	SOC 2018	ISCO-08
2015	ONETSOC2010	SOC2010	ISCO-08
2010	ONETSOC2010	SOC2010	ISCO-08
2005	ONETSOC2000	SOC2000	ISCO-88
2000	ONETSOC2000	SOC2000	ISCO-88
1995	ONETSOC2000	SOC2000	ISCO-88
1992	ONETSOC2000	SOC2000	ISCO-88

Note: Own elaboration.

A4. Artificial intelligence

We draw on two established measures of occupational exposure to artificial intelligence: Felten et al., (2018, 2019), and Webb (2020). The original estimates, published by the

³³ The choice of crosswalks depends on the year under analysis, as different occupational coding systems apply across time. See Table A2 for details on the specific crosswalks used.

respective authors, are reported at the occupational level using the U.S. Standard Occupational Classification (SOC) system. To implement these measures in a European context, we first convert SOC codes to the International Standard Classification of Occupations (ISCO), using official crosswalks. We then merge the AI exposure measures with individual-level data from the Spanish Labor Force Survey (ES-LFS), applying survey weights to obtain nationally representative estimates. Finally, we assign the measures to the years 2010, 2015, and 2019, enabling an analysis of the temporal evolution and occupational distribution of AI exposure across the Spanish labor market.

A5. Wages

The PH database (INE, 1994–2001) follows the EU Household Panel methodology and provides detailed different income sources and living-conditions information. Since 2004 it has continued as the ECV within the EU-SILC framework. Both datasets are nationally and regionally representative through stratified sampling and weighting based on age, sex, nationality, and region. Our analysis focuses on full-time employees aged 25–65 with positive net labor income. Wages are in constant 2019 euros. The dependent variable is the weighted average wage for each sociodemographic group, estimated within the PH/ECV and matched with the LFS using the groups defined in Section 3.1. Table A3 shows several inequality measures describing the dependent variable in levels and in logs (as used in the regressions, see Equation 7) in all waves used in the study.

Table A3. Wage Inequality.

Year	Gini	MLD	Wage share Top 10%	Wage share Bottom 50%	Wage share Bottom 10%	Wage share Middle 50%
<i>Panel A: In levels, μ_{jt}</i>						
2000	0,216	0,080	0,204	0,314	0,050	0,461
2010	0,220	0,080	0,166	0,339	0,047	0,483
2015	0,221	0,089	0,204	0,366	0,047	0,481
2019	0,232	0,098	0,205	0,352	0,047	0,482
<i>Panel B: In logs, $\ln(\mu_{jt})$</i>						
2000	0,022	0,001	0,110	0,437	0,097	0,511
2010	0,023	0,001	0,092	0,475	0,093	0,524
2015	0,025	0,001	0,125	0,503	0,103	0,502
2019	0,026	0,001	0,119	0,498	0,103	0,503

Note: Own elaboration. MLD stands for Mean Log Deviation.

APPENDIX B

In this appendix we present the results for the classical measure of automation, RSH^C. The results for the other two indices of automation, RSH^L and RSH^{UF2}, are similar and available from the authors upon request.

Table B1. Long-run inequality change (RSH^C, 2000-2019).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. OLS estimates</i>							
L. Wage Past	-0.634*** (0.0763)	-0.587*** (0.0766)	-0.625*** (0.0747)	-0.602*** (0.0767)	-0.598*** (0.0777)	-0.613*** (0.0771)	-0.655*** (0.0757)
L. RSH ^C	-0.708*** (0.144)	-0.603*** (0.142)	-0.674*** (0.133)	-0.543*** (0.149)	-0.619*** (0.144)	-0.650*** (0.142)	-0.697*** (0.146)
L. Trade	0.436*** (0.109)						0.307** (0.137)
L. Manufacturing		1.878 (2.767)					-0.599 (2.911)
L. Education			0.0398* (0.0210)				0.0101 (0.0248)
L. Offshorability				0.410* (0.220)			0.211 (0.214)
L. Markup					0.259* (0.139)		0.0850 (0.131)
L. Employment						0.493*** (0.111)	0.271* (0.140)
Constant	6.095*** (0.771)	5.756*** (0.775)	5.728*** (0.767)	4.502*** (1.035)	5.454*** (0.771)	5.624*** (0.764)	5.146*** (0.961)
Observations	408	408	408	408	408	408	408
R-squared	0.317	0.292	0.302	0.296	0.295	0.309	0.325
<i>Panel B. 2SLS estimates</i>							
L. RSH ^C	-1.449*** (0.456)	-1.216*** (0.431)	-1.114*** (0.410)	-1.194** (0.059)	-1.136*** (0.433)	-1.204*** (0.430)	-1.530*** (0.582)
F-stat	22.44	18.72	20.88	15.73	19.92	20.23	17.28

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The Robustness Value (RV) of the full specification model, obtained with the Cinelli-Hazlett bounding procedure, ascends to 28.60 % of the remaining variation in both the treatment and the outcome (see main text).

Table B2. Automation variables (2000-2010-2019).

	(1) RSH ^C	(2) RSH ^L	(3) RSH ^{UF}	(4) RSH ^{UF2}
<i>Panel A. OLS estimates</i>				
Year = 2019	-0.340*** (0.0284)	-0.306*** (0.0258)	-0.287*** (0.0251)	-0.389*** (0.0309)
L. Wage Past	-0.471*** (0.0555)	-0.458*** (0.0600)	-0.550*** (0.0542)	-0.438*** (0.0537)
L. RSH ^C	-0.462*** (0.086)			
L. RSH ^L		-0.577*** (0.132)		
L. RSH ^{UF}			-0.662*** (0.0786)	
L. RSH ^{UF2}				-0.533*** (0.0915)
Constant	4.844*** (0.566)	4.727*** (0.624)	5.858*** (0.571)	4.733*** (0.576)
Observations	816	816	816	816
R-squared	0.337	0.331	0.377	0.335
<i>Panel B. 2SLS estimates</i>				
L. RSH ^C	-0.469*** (0.280)			
L. RSH ^L		-1.489*** (0.440)		
L. RSH ^{UF}			-0.670*** (0.120)	
L. RSH ^{UF2}				-5.055 (6.439)
F-stat	19.89	13.47	32.82	0.41

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.


Table B3. Medium-run inequality change (RSH^C, 2000-2010-2019).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. OLS estimates</i>							
Year = 2019	-0.328*** (0.0279)	-0.341*** (0.0283)	-0.317*** (0.0302)	-0.277*** (0.0343)	-0.333*** (0.0313)	-0.324*** (0.0290)	-0.250*** (0.0362)
L. Wage Past	-0.505*** (0.0555)	-0.476*** (0.0552)	-0.505*** (0.0570)	-0.501*** (0.0559)	-0.472*** (0.0558)	-0.491*** (0.0552)	-0.543*** (0.0571)
L. RSH ^C	-0.540*** (0.0870)	-0.469*** (0.0860)	-0.529*** (0.0863)	-0.363*** (0.0931)	-0.467*** (0.0888)	-0.503*** (0.0855)	-0.470*** (0.0947)
L. Trade	0.337*** (0.0722)						0.177* (0.101)
L. Manufacturing		3.824*** (1.367)					2.178 (1.628)
L. Education			0.0352*** (0.0116)				0.0138 (0.0144)
L. Offshorability				0.453*** (0.144)			0.401*** (0.142)
L. Markup					0.0542 (0.112)		0.0942 (0.110)
L. Employment						0.402*** (0.122)	0.168 (0.137)
Constant	5.075*** (0.560)	4.855*** (0.560)	4.810*** (0.550)	3.549*** (0.615)	4.766*** (0.582)	4.711*** (0.576)	3.621*** (0.634)
Observations	816	816	816	816	815	816	815
R-squared	0.352	0.343	0.345	0.348	0.337	0.344	0.364
<i>Panel B. 2SLS estimates</i>							
L. RSH ^C	-0.711*** (0.334)	-0.670*** (0.272)	-0.485*** (0.285)	-0.188** (0.430)	-0.478*** (0.292)	-0.502*** (0.292)	-0.574*** (0.462)
F-stat	22.79	19.17	20.89	18.26	20.22	20.45	20.68

Source: own elaboration. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The RV of the full specification model, obtained with the Cinelli-Hazlett bounding procedure, ascends to 22.29 %.

APPENDIX C

Figure C1. The effect of automation (RSH^{UF}) by sex, age and education categories.
(2000-2010-2019)



Note: Own elaboration. The dashed black line shows the average change rate across all groups.