INSTITUTE
OF ECONOMICS

Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

Will robot replace workers? Assessing the impact of robots on employment and wages with meta-analysis

Dario Guarascio [1]
Alessandro Piccirillo [1]
Jelena Reljic [1,2]

[1] Sapienza University of Rome, Italy
[2] Global Labour Organization

2024/03 February 2024
ISSN(ONLINE) 2284-0400

Will robot replace workers? Assessing the impact of robots on employment and wages with meta-analysis

Dario Guarascio[1,2], Alessandro Piccirillo[1], Jelena Reljic[1,2*]

[1] Sapienza University of Rome

[2] GLO

[*]Corresponding author: jelena.reljic@uniroma1.it

Abstract

This study conducts a meta-analysis to assess the effects of robotization on employment and wages, compiling data from 33 studies with 644 estimates on employment and a subset of 19 studies with 195 estimates on wages. We identify a publication bias towards negative outcomes, especially concerning wages. After correcting for this bias, the actual impact appears minimal. Thus, concerns about the disruptive effects of robots on employment and the risk of widespread technological unemployment may be exaggerated or not yet empirically supported. While this does not preclude that robots will be capable of gaining greater disruptive potential in the future or that they are not already disruptive in specific contexts, the evidence to date suggests their aggregate effect is negligible.

Keywords: robots, employment, wages, meta-analysis, publication bias
JEL codes: J21, J23, J24, J31, O33

1. INTRODUCTION

As it is common in the field of economics, the most frequently addressed research questions are also those giving rise to inconclusive answers. This is exemplified in the ongoing debate surrounding the robots-workers race (Corrocher et al., 2023). Since the early days of Ricardo and Marx, and later Keynes, the potential role of machines in job displacement and the consequent downward pressure on wages has been a major concern of economists (Calvino ad Virgillito, 2018). More recently, robotization has come back to the fore, with empirical papers mushrooming by the dozen per year (for a recent review, see Montobbio et al., 2023) to examine whether robots are stealing jobs, favouring wage compression and/or increasing income inequalities (among others, Aghion et al., 2020; Acemoglu and Restrepo, 2020).

Several reasons may help explaining such a renewed popularity. First, the advancement of digital technologies (e.g., Artificial Intelligence (AI)) has given new impetus to the capabilities of robots, significantly expanding their scope of application (Fernandez-Macias et al., 2021). Second, the exponential growth in global trade and the emergence of China as 'the world's factory (Baldwin, 2013) have opened the way for a massive diffusion of manufacturing robots, with significant implications in terms of efficiency-enhancing and, consequently, labour-saving effects. Third, the globalization of production and supply chains has favoured the diffusion of automation technologies, including robots, in areas that were relatively less affected by this phenomenon, such as developing countries (Shapiro and Mandelman, 2021), as well as in sectors that are key for the operation of markets, both domestically and internationally, like logistics (Sostero, 2020). Fourth, robotization has been identified as one of the key drivers of the broader process of 'de-industrialization', resulting in the reduction of manufacturing employment in both developing and developed economies, particularly in regions like the US 'rust belt', where a large share of blue collars have been crowded out (Rodrik, 2022). These developments have been linked to the growth of socio-economic inequalities (Acemoglu and Restrepo, 2022) and related phenomena of social and political instability (Anelli et al., 2021).

Another factor inflating the empirical literature on the robots-employment-wage nexus concerns the proliferation of data sources, which provide detailed information on robot adoption at the country, sector and firm-level (Mondolo, 2022).[1] This facilitated more nuanced investigations - considering various sources of heterogeneity such as differences in industrial specialization, institutional set-up (including labour market legislations) and the relative strength of unions - as well as the adoption of state-of-the-art econometric techniques. Rather

---

[1] One of the most important sources of information concerning the adoption of robots is the database provided by the International Federation of Robotics (IFR). The latter reports information about the installations of "multipurpose manipulating industrial robots". Robots are defined as "automatically controlled, reprogrammable, multipurpose, manipulator, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications" (IFR, 2021, p. 28). At the firm-level, country specific surveys as, for example, the German Community Innovation Survey (CIS) or the Italian Rilevazione Imprese Lavoro (RIL), provide detailed information on robot stocks and new installations.

than simplifying things, however, this abundance of empirical evidence has further complicated the picture, making it even more difficult to unambiguously establish whether robots are stealing jobs (and wages) or not (Montobbio et al., 2023).

Heterogeneities and contrasting evidence may be the result of a variety of factors (Reljic et al., 2023). First, the unit of analysis matters. For instance, a positive effect of robotization at the firm level may hide an overall negative effect at the macro level (i.e., country- or industry-level data). This can occur when the apparent positive effect of robot adoption on employment and/or wages detected by looking at firms is actually the consequence of a 'business stealing' effect (Calvino and Virgillito, 2018). Second, structural heterogeneities concerning technological characteristics, positioning within GVCs and labour market institutions can lead to very different results, depending on the country, region or sector under consideration. No less relevant, the available evidence may suffer of peculiar biases related to research methodologies. For example, it is not so unusual for researchers that focus on a specific topic (and the technology-employment relationship is no exception) to replicate the most influential studies[2]. Unsurprisingly, such habit can lead to uniformity regarding the chosen theoretical framework (e.g., general equilibrium models) and/or related empirical set-ups (e.g., relying on the same instrumental variables used in previous studies), thereby raising the risk of systemic biases affecting the entire literature body of literature.

This aspect is far from unimportant, as the adopted theoretical/empirical framework has a fundamental impact on the way phenomena are modelled, empirical analyses are conducted and results are interpreted (Montobbio et al., 2023). According to the neoclassical approach, for example, market (equilibrium) forces may diminish the significance of technological unemployment (and related distributive effects) making it a transitory phenomenon at most. Nonetheless, technology could lead to changes in the composition of labour demand, along with shifts in occupational and income shares. Evolutionary theories, on the other hand, highlight the disruptive nature of technological change, acknowledging the importance of evolving production networks, demand patterns and institutions. According to this perspective, technology can not only negatively affect employment and wages but can do so persistently (Dosi et al., 2021). From a strictly empirical viewpoint, privileging statistically significant results and/or those consistent with prevailing theoretical models, even when contrasting results arise from rigorous analyses with equal or greater scientific and policy relevance, may introduce systematic bias into the available evidence on a specific matter (Doucouliagos and Stanley, 2013; Irsova et al., 2023a).

In this context, a powerful tool emerges as a solution: meta-analysis. The latter can help synthesising the extant literature, enabling the identification of prevailing patterns, unravelling relevant heterogeneities and detecting (and correcting!) publication biases. Although meta-analysis is still relatively underutilized in economics, its popularity is growing (Irsova et al., 2023a). Recently, it has been largely employed to dissect empirical literature focusing on

[2] A seminal work by Acemoglu, D., & Restrepo, R. (2020) is a case in point.

macroeconomic research questions. Relevant examples include the works of Heimberger (2020, 2021, 2022), which explore the connections between globalization and income inequalities, corporate tax and competition and globalization and growth, respectively. Brown et al. (2024) have provided insights into the empirical estimates of loss aversion. Additionally, this methodology has been applied to issues relevant from both a meso and a micro perspective, as seen in the works of Card et al. (2018), Havranek et al. (2018) and Kroupova et al. (2022), among others.

In this work, we meta-analyse a comprehensive set of empirical studies investigating the impact of robots on employment and wages. To the best of our knowledge, this is one of the first attempts to apply meta-analysis to the vast scientific literature studying the relationship between labour and technology (Calvino and Virgillito, 2018). In this respect, meta-analysing the literature focusing on the robot-employment nexus is particularly urgent. In fact, while, on the one hand, the fear that robots are going to jeopardize jobs and wages is steadily growing, on the other, the evidence in this regard is, at best, highly contradictory.

To carry out our meta-analysis, we built a unique dataset comprising 644 estimates stemming from 33 studies investigating the employment impact of robotization and 195 estimates from a subset of 19 studies analysing also the wage effect of robots. The following research questions are addressed. What is the average magnitude of the effect of robotization on employment and wages? Are the reported empirical findings plagued by publication selection bias? What factors contribute to explaining the variation in the effects of robotization on employment and wages? The meta-analysis reveals, on average, negative and statistically significant effects of robotization on both employment and wages, although these effects remain marginal and close to zero. As a result, claims regarding the disruptive impact of robot diffusion and the associated risks of widespread technological unemployment seem to be overstated, or at the very least, lack empirical support so far. However, this does not mean that robots are unlikely to gain greater disruptive potential in the future, or that they are not already disruptive in specific contexts. Instead, the existing evidence suggests that their overall average effect remains minimal.

The article is organised as follows. Section 2 reports a brief review of the literature focusing on the impact of robots on employment and wages, highlighting the recent contributions providing a narrative summary of this literature stream. Section 3 introduces the database and provides a set of descriptive evidence regarding the main characteristics of the studies considered for the analysis. Section 4 deals with the publication bias, while Section 5 reports the main results and the robustness checks. Section 6 concludes by discussing strengths and limitations of the existing literature as well as avenues of future research.

2. LITERATURE REVIEW

As our objective is to meta-analyse the existing evidence on the robot-employment-wages nexus, we avoid to carry out an extensive literature review, advising reference to recent summaries such as Barbieri et al. (2019), Mondolo (2022) and Montobbio et al. (2023).[3] Instead, in what follows we offer a brief account of the relevant literature organising the latter according to the unit of analysis (i.e., country-, region-, sector-, and firm-level) adopted to carry out the empirical investigations. We, firstly, review the literature assessing the impact of robots on employment and, subsequently, that on wages. Notice that most of the contributions included in the following review have also been used to build the database upon which the meta-analysis is based. Therefore, conciseness has been our guiding criterion. A few more details are provided only with reference to the most influential contributions.

2.1 The empirical literature on the robots-employment nexus

As Montobbio et al. (2023) emphasise, the stream of literature focusing on the employment impact of robots is a subset of the larger group of studies that, in the last two decades or so, have attempted to understand the implications of the broader process of digitization-automation. The way has been opened by Frey and Osborne (2017), concentrating their attention on US occupations and relying on a rather generic definition of digitization *cum* automation: the likelihood of substituting human tasks by means of computer-controlled machines. Their predictions have caused much concern. About the 47% of US occupations, mostly middle- and low-skilled professions, turned out to be at high risk of replacement. Since then, empirical studies multiplied, also thanks to the availability of novel datasets and proxies to measure automation, including data on robot adoption. As a result, the worrying scenario foreseen by Frey and Osborne (2017)'s seminal contribution has been significantly scaled down (Arntz et al., 2016; Nedelkoska and Quintini, 2018). Yet, automation is confirmed to be a potential driver of job destruction, although on a much smaller scale (Pouliakas, 2018).

A large number of contributions studying the robot-employment nexus uses IFR data and, hence, adopts a sectoral/geographical perspective. Graetz and Michaels (2018) combine IFR and EUKLEMS data to build a 17 countries panel ranging from 1993 to 2007. On average, their findings show no negative employment effects of robots, except the reduction of the low-skilled workers' employment share. A similar approach is adopted by Chiacchio et al. (2018) and Acemoglu and Restrepo (2020). The former use IFR data to study the impact of robots on EU regional labor markets. In this case, a negative association is detected, as the introduction of one more robot per thousand workers contributes to a reduction of the employment/population ratio of about the 0.16-0.20%. These findings are rather consistent with those of Acemoglu and Restrepo (2020), focusing on US local labor markets (LLMs) over

[3] Notice that these contributions do not focus exclusively on robots but address, more broadly, the impact of new technologies on the labour market.

the 1990-2007 period.[4] This approach is adapted to the German case by Dauth et al. (2021). Analyzing German regions between 1994 and 2014, they find no negative impact of robots on employment, although a significant compositional effect is documented (the growth of non-manufacturing seem to compensate the reduction of manufacturing employment).

The presence of different 'robotization regimes' in Europe is documented by Reljic et al. (2023). Focusing on European manufacturing industries and using IFR data to measure robot adoption, they find that, on average, robots have a positive impact on total employment. Nonetheless, such a labour-friendly impact of robotization is detected only in core (Austria, Germany, Denmark, Finland, France, Netherlands, Sweden) and service-oriented (Belgium, Estonia, Ireland, United Kingdom, Slovenia) countries and for those at the top of the occupational structure (i.e., managers and technicians). In turn, peripheral countries and manual workers do not seem to benefit at all from robotization. A different measure of automation based on patents' keywords is provided by Mann and Püttmann (2023) who distinguish between automation and non-automation-related patents to study the impact of such technologies on US LLMs. In line with Dauth et al. (2021) findings, these authors report a positive impact of robots on jobs which, however, is mainly driven by the growth of employment in the service sector.[5]

Another group of contributions have assessed the employment impact of robotization taking advantage of firm-level information on robot adoption. Such studies are important as they rely on fine-grained indicators concerning the actual adoption of robotization technologies by firms. On the other hand, estimations stemming from firm-level studies may fall short of providing a final answer on the net employment impact of robotization as, for example, growing labour demand in robot-intensive firms may very well be matched by a more than proportional job destruction in other firms (i.e., the so-called 'business stealing' effect). Either way, several contributions find that labour demand tends to grow more in companies where robotization is more intense. This is the case of Bessen et al. (2020)'s contribution. These authors focus on Dutch manufacturing and no-manufacturing firms, documenting that 'automating firms' tend to display faster employment and revenue growth than those not adopting automation technologies. Similarly, Koch et al. (2021) rely on Spanish firm-level data showing that robot adoption led to net job creation, together with output gains and a reduction of the labour cost share. Robotization seem to have a positive impact also when the analysis focuses on French companies. Aghion et al. (2021) show that robot adoption is associated with job creation concerning both skilled and unskilled workers.

Contrasting evidence are in order even among this group of contributions, though. A good example is Acemoglu et al. (2020), relying on French data as Aghion et al. (2021) yet reaching opposite conclusions. While a slightly positive effect of robotization is confirmed when the

---

[4] Other studies finding a negative impact of robots on employment include: Blanas et al. (2019) and Bonfiglioli et al. (2021).

[5] Other studies finding positive and non-significant effects of automation technologies/robot on employment include: Jäger et al. (2016), de Vries et al. (2020) and Petit et al. (2023).

analysis is caried out using firm-level information, it turns out that, at the industry-level, a 20% increase in robot instalments has a negative effect on employment of about 3.2%. Instead, a strong heterogeneity across workers with different skill-level is detected by Balsmeier and Woerter (2019), analysing the employment impact of digitization-automation technologies - including robots, 3D printing and IoT - by focusing on Swiss firms. Overall, investments in digital-automation technologies are associated with increasing employment among high-skilled workers while the opposite goes for low-skilled one. Yet, they document a slightly positive net effect. A somewhat different framing is proposed by Benmelech and Zator (2022): robots are modelled as a way to react to labour shortages rather than to increase efficiency and, thus, reduce costs. Using data on German firms, these authors show that robot-intensive firms increase employment but, as documented by Acemoglu et al. (2020), the effect turns negative (although small in magnitude) when the analysis is restricted to automation-exposed industries and regions.

Even if not included in our meta-analysis, another channel through which robotization may have an impact on employment concerns company-level decisions on production offshoring/reshoring (Faber, 2020).[6] Two are the main hypotheses. The first is that robotization, by increasing the efficiency of domestic productions, reduces the incentives for offshoring opening the way, in turn, for a process of reshoring. Symmetrically, the second hypothesis points to the potential negative effect that robotization in the home country may have on employment in the destination country. Krenz et al. (2021) have tested the first hypothesis focusing on 43 countries analysed over the 2000-2014 period. They show that, in manufacturing, an increase by one robot per 1000 workers is associated with a 3.5% increase of reshoring. Nonetheless, it is not so clear whether this channel would ultimately lead to a net positive employment effect or, rather, to the growth in demand for high-skill workers only, which does not necessarily translate into a general increase in employment. Regarding the second hypothesis, Artuc et al., (2019) show that an increase in robotization in the US has a negative impact on employment in Mexico, focusing, in particular, on those areas that are traditionally more exposed to Foreign Direct Investments (FDI). These results are in line with the findings of Stemmler (2023), comparing Brazilian industries that are more/less exposed to foreign automation. More exposed industries display a significant reduction in the share of manufacturing employment, while the opposite occurs in less exposed ones (e.g., mining). Consistent evidence on Colombia is provided by Kugler et al. (2021), showing that increasing robot adoption in the US negatively affects employment in those Colombian LLMs that are more exposed to US automation. Concerning the ultimate impact on employment, the evidence is mixed even in this case.

Overall, this brief literature review confirmed how finding a final answer to such a key research question - Do robots destroy jobs? – is all but an easy task. Many forces are at work, plenty of

---

[6] It is worth noticing that contributions analysing the impact of robotization in one country on the employment dynamics of other countries

heterogeneity sources may mediate the impact of robots (e.g., industry- and firm-level characteristics, skill-level and occupational profiles, institutions) and even the unit of analysis is not neutral concerning the final outcome of the empirical analysis. In what follows, we provide an analogous account of the literature addressing the most consequential research question: Does robotization reduce wages?

2.2 The empirical literature on the robots-wages nexus

Many studies exploring the robot-employment nexus extend their investigation to test whether robotization has also an impact on wages. Theoretically, different channels may shape such a relationship. According to the classical view, i.e., Ricardo and Marx, machines are tools aimed, through the threat of technological unemployment, at 'disciplining workers' (Vivarelli, 1995). Hence, the introduction of labour-saving machines, including robots, is expected to have an overall wage-moderating (if not wage-reducing) effect. Yet, neoclassical theories would predict a rather different outcome. Not wage squeezing, brought about by a generalized technology-driven weakening of workers' bargaining power. Rather, a change in marginal productivities, depending on workers' skill-level and/or on the characteristics of working tasks (as predicted by the Skill and Routine Biased Technological Change (SBTC and RBTC) hypotheses, see Vivarelli, 2014), resulting into the reshuffling of occupation-specific employment and wage shares. If technological change is 'skill biased' (i.e., complementarity between new technologies and advanced competences), high-skilled workers are expected to receive a 'wage premium', while the opposite happens to low-skill ones. If technology is 'routine biased', in turn, those losing ground are workers performing routine tasks (i.e., repetitive/codifiable tasks that machines can more easily perform) while wages are expected to increase for workers, both at the top and at the bottom of the skill distribution, performing tasks including dexterity, creativity, experiential knowledge. Against this background, empirical studies testing whether robots, in addition to stealing jobs, can also take away wages, and whether this varies given workers' skill endowment, start multiplying.

The contributions relying on IFR data provide mixed results. Acemoglu and Restrepo (2020) find that an increase in robot adoption has a negative impact on wages: one more robot per thousand workers leads to a 0.42% reduction of wages at the US LLM-level. A similar analysis is carried out by Borjas and Freeman (2019), focusing on the US and comparing the wage effect of robots and of migrant flows, both modelled as supply shocks having a potentially negative effect on hourly earnings. Their estimations show that robots not only have a negative impact on wages, irrespective gender and education level, but also that such effect is two/three times larger than that of migrants.

An opposite outcome is reported by Graetz and Michaels (2018). Relying on a lager sample, including 17 countries observed between 1993 and 2007, these authors document a positive association between robotization and mean hourly wages. Similarly, Compagnucci et al. (2019) find a positive association between robot adoption and hourly wages looking at 16 OECD countries over the period 2011-2016. Nonetheless, wages tend to increase at a slower rate in sectors where robotization grows at a faster pace. On the other hand, Chiacchio et al. (2018) find no significant association between robot adoption and wages, despite controlling for a number of mediating factors and replicating the analysis across sectors and occupational groups. Analogous evidence can be found in Dauth et al. (2021)'s analysis of robotization in Germany. They find a slightly yet not significant effect of robot adoption on wages. When a significant effect is detected, as in manufacturing, the latter turn out to be more than offset by the increase in wages registered in the service sector.

The evidence is mixed even among firm-level studies. According to the evidence provided by Koch et al. (2021), focusing on a large sample of Spanish firms observed between 1990 and 2016, robot adoption has no significant effect on wages, despite having a positive impact on employment and output. Some heterogeneity across sectors and skill groups is instead found by Bonfiglioli et al. (2021) and Barth et al. (2020). The former show that, while robot adoption has no effect in the manufacturing sector, it contributes to wage growth in services. Analysing the Norwegian case and taking advantage of a rich employer-employee database, Barth et al. (2020) find evidence of a 'skill-biased' dynamics with the high/low-skill wage premium increasing as a consequence of robot adoption. Overall, the evidence points to not significant or slightly positive effects of robotization.

Contrasting evidence is also reported concerning the impact of robots on the gender wage gap. A positive relationship is documented by Aksoy et al. (2021), focusing on a panel of 20 European countries and showing that a 10% increase in robotization results into an increase of the gender wage gap of about 1.8%. Contrarily, Ge and Zhou (2020), analysing US LLMs between 1990 and 2015, find that increasing robot adoption is associated with a decrease in the gender wage gap. A more complex picture emerges from Albinowski and Lewandowski (2022)'s study. In this case, robot adoption seems to have a negative impact on young (20-49) male workers employed in routine manual occupations, while positive effect on wages is documented when it comes to young (20-29) women employed in routine cognitive occupations.

Such inconclusive evidence lends further support to meta-analysis as an effective tool to summarize the existing literature. Even in the case of the robot-wage nexus, meta-analysis can help identifying the overall effect, the role of mediating factors in shaping results as well as the impact (if any) of the publication bias. It should be emphasised, however, that we circumscribe the analysis to the subset of papers, which focus on the effect of robotization on employment also including wages into the analysis.

3. DATA COLLECTION

3.1 Search strategy

The selection process for our study is depicted in the PRISMA flowchart in Figure A1. Initially, we conducted a systematic assessment of the literature relying on three databases: Scopus, ResearchGate, and Google Scholar. Articles were retrieved using keywords such as ‘employment’, ‘robot’, and ‘impact’. We limited our search to non-review articles written in English from 2010 to 2022. Over 950 records were screened, based on titles and abstracts. After removing ineligible studies, including nearly 200 duplicates, 60 papers were deemed suitable for a meta-analysis.

Various inclusion and exclusion criteria were applied to determine a study’s relevance. The primary inclusion criterion concerns the specification of the econometric model: we have included articles using employment as a dependent variable and robot adoption as an explanatory one. Therefore, articles providing only descriptive analyses, theoretical models or literature reviews on the robot-employment nexus were excluded. Second, we ruled out studies that did not differentiate between sources of automation (Frey and Osborne, 2017; Schmidpeter & Winter-Ebmer, 2021), those focusing on AI instead of robots (Acemoglu et al., 2022), those using the latter as a dependent variable, those not focusing on employment (Antón et al., 2020; Fiedler and Fidrmuc, 2021) and review papers (Barbieri et al., 2019). Third, out of a more detailed screening, studies that met initial criteria but examined the impact of foreign robot adoption on domestic employment have also been excluded (Gravina and Pappalardo, 2022; Kugler et al., 2020).

We ended up with 33 studies examining the employment impact of robotization, which provided 644 estimates. Among these, 19 studies also offered wage effect estimates, yielding 195 estimates.

3.2 Standardisation

Despite the careful selection of studies for our meta-analysis, satisfying all the established criteria, they differ in various aspects: how they measure employment and robot adoption, sample selection and econometric approaches, among others. To facilitate meaningful comparison across studies, we convert all estimates to a homogenous metric (partial correlation coefficients). This standardisation process follows the methodology outlined in the seminal work by Stanley and Doucouliagos (2012), which has been employed in recent meta-analyses (Cazachevici et al., 2020; Heimberger, 2022). We calculate partial correlation coefficients (r), unitless indicators within the range of -1 to 1, gauging the strength and direction of the relationship between two variables using the following formula:

$$r_{ij} = \frac{t_{ij}}{\sqrt{{t_{ij}}^2 + df_{ij}}} \tag{1}$$

where t stands for the t-statistic of the i^{th} regression coefficient of study j, calculated as the ratio between the point estimate of the coefficient and its standard error, while df represents the degrees of freedom. The standard error of the PCC is then given by: $\sqrt{(1 - r^2)/df}$.

Figure 1 plots an array of partial correlation coefficients derived from different studies exploring the robotization-employment nexus. The average unweighted PCC across all studies is -0.02, marked by the red line, signalling an overall lack of consensus on the direction of this effect. Indeed, the interquartile ranges (IQRs) in Figure 1 reveal significant heterogeneity, with some studies showing a negative impact of robotization on employment (Acemoglu and Restrepo, 2020; Aghion et al., 2019; Brambilla et al., 2022, among others), whilst others report positive or negligible effects (Dixon et al., 2021; Fu et al., 2021; Klenert et al., 2022, among others). Moreover, large IQRs underscore also substantial heterogeneity within studies (i.e., less precision or higher uncertainty in their estimates).

Figure 1: Partial correlation coefficients between robotization and employment

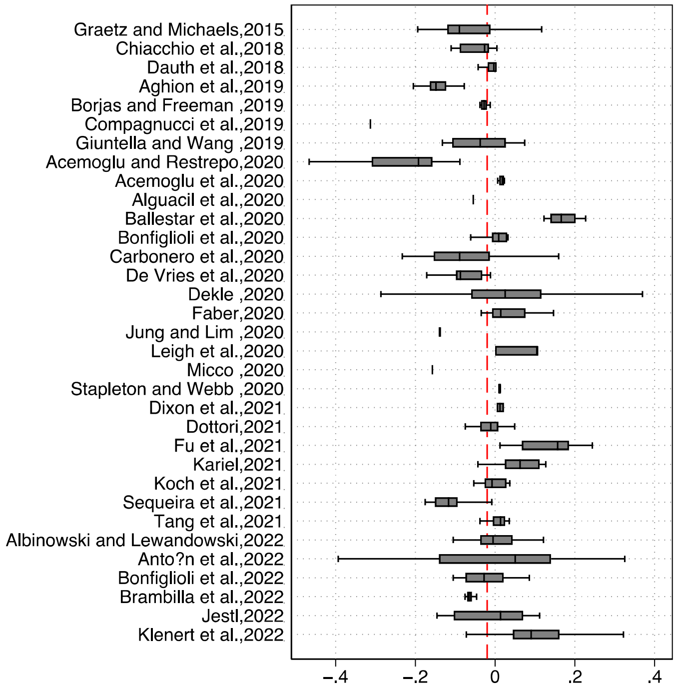

Notes: The figure shows a box plot of partial correlation coefficients of the estimated relationship between robotization and employment across studies; each box represents the interquartile range (P25-P75), while the line inside the box represents the median. Red dashed reference line represents the mean unweighted PCC across all studies. Outliers are excluded from the figure but included in all statistical tests. *Source*: Authors' elaboration

Turning to the relationship between robotization and wages in Figure 2, a similar pattern is observed. The mean unweighted PCC across all studies, again denoted by the red line, is marginally below zero (-0.03), hinting at a slight overall negative correlation. However, the span of IQRs encompassing the zero point emphasises the lack of a predominant effect.

These findings challenge a simplistic narrative that robotization either universally harms or benefits employment and wages. Rather, they underline the heterogeneity of these effects across studies, likely influenced by multiple underlying factors that differ from one study to another. This highlights the context-dependent nature of robotization's impact on the labour market that should and indeed will be taken into account in Section 5.

Figure 2: Partial correlation coefficients between robotization and wages

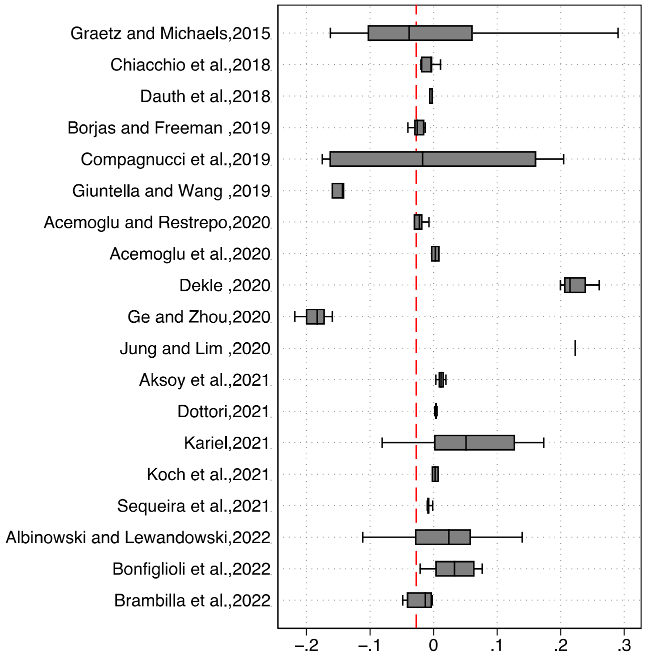

Notes: The figure shows a box plot of partial correlation coefficients of the estimated relationship between robotization and wages across studies; each box represents the interquartile range (P25-P75), while the line inside the box represents the median. Red reference line represents the mean unweighted PCC across all studies. Outliers are excluded from the figure but included in all statistical tests. *Source*: Authors' elaboration

4. PUBLICATION BIAS

A key shortcoming of narrative literature reviews lies in their inability to systematically detect publication selection bias, a frequent issue in economic research. This bias arises from the tendency among researchers and journals to favour and publish findings that are statistically significant or that reinforce prevailing theories, often leading to the exclusion of insignificant or unfavourable outcomes (Stanley and Doucouliagos, 2012). Such selective reporting can lead to a misrepresentation of the available empirical evidence. Meta-analysis provides tools to detect and address these issues, enabling the testing for publication bias using both visual methods, such as funnel plots, and formal tests, like the Funnel-Asymmetry Precision-Effect Test (FAT-PET).

Asymmetry in the funnel chart serves as a preliminary indication of possible publication bias, and conversely, its absence suggests the lack thereof. Figures 1 and 2 illustrate the partial correlation coefficients plotted against measures of precision (i.e., the inverse of the standard error of the partial correlation coefficients). These figures reveal a significant degree of heterogeneity across the studies under investigation for both employment (Figure 1) and wages (Figure 2). Note that more precise point estimates are placed in the top, while less precise estimates are positioned at the bottom and are more dispersed. There are some outliers in coefficient precision, which are caused not by coding errors but by extremely small standard errors in studies analysing large samples.[7]

Figure 3: Funnel plot of partial correlation coefficients - Employment

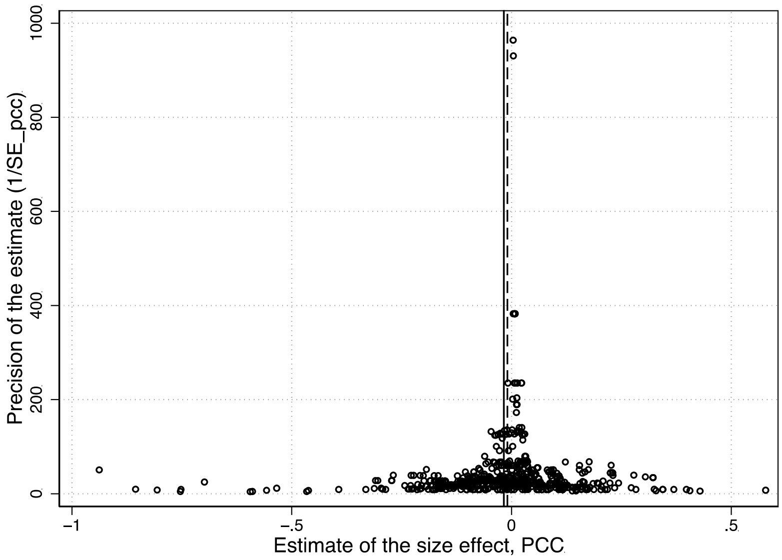

Notes: The figure illustrates the relationship between the partial correlation coefficients of the estimated effects of robotization on employment and a precision indicator, which is calculated as the inverse of standard error. The dashed vertical line in the figure indicates the median, while the solid line represents the unweighted mean of PCCs. Positive PCCs indicate a positive impact of robotization on employment, and vice versa. *Source*: Authors' elaboration

Figure 4: Funnel plot of partial correlation coefficients - Wages

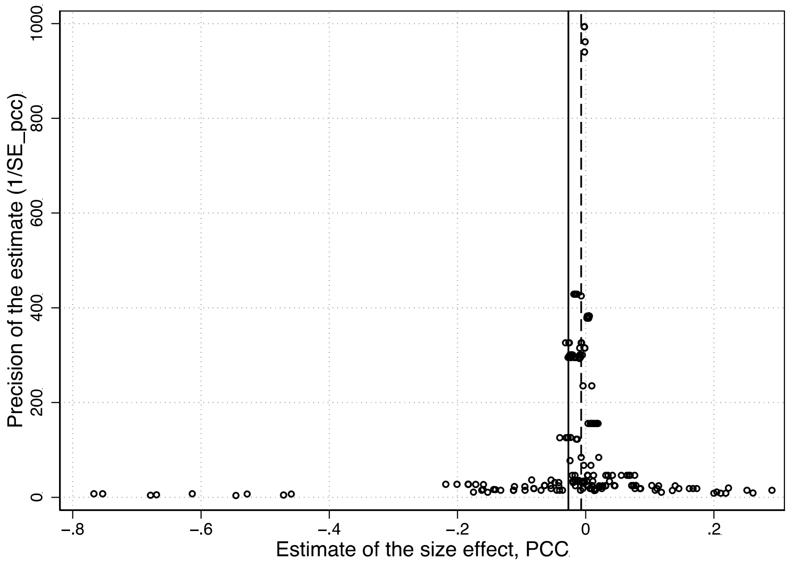

[7] Specifically, Dixon et al. (2021) examined close to 800,000 firms, while Dauth et al. (2021) considered a cohort of nearly 900,000 individuals.

Notes: The figure illustrates the relationship between the partial correlation coefficients of the estimated effects of robotization on wages and a precision indicator, which is calculated as the inverse of standard errors. The dashed vertical line in the figure indicates the median, while the solid line represents the unweighted mean of PCCs. Positive PCCs indicate a positive impact of robotization on wages, and vice versa. Source: Authors' elaboration. *Source*: Authors' elaboration

The funnel charts for both employment and wages exhibit little asymmetry, suggesting a lack of publication bias. This is not surprising, as both positive and negative labour market effects stemming from technology are theoretically possible and, consequently, reported. Nevertheless, reliance on visual inspection of funnel asymmetry may not represent a sufficiently reliable indication of publication bias (Stanley and Doucouliagos, 2012).

We employ FAT-PET to formally assess for publication selection bias by estimating the relationship between PCCs and their standard errors as follows:

$$r_{ij} = \beta_0 + \beta_1 SE_{ij} + \epsilon_{ij} \tag{2}$$

where r_{ij} is the partial correlation coefficient i from study j, SE_{ij} is its standard error and ϵ_{ij} stands for the error term.

The term $\beta_1 SE_{ij}$ serves to control for publication selection bias. The hypothesis H_0: $\beta_1 = 0$ is called the Funnel-Asymmetry Test. If $\beta_1= 0$, it means that there is no publication selection bias. Additionally, we conduct the Precision-Effect Test (PET), which revolves around the hypothesis H_0: $\beta_0= 0$. This test is essential for determining whether remains an empirical effect of robotization on employment and wages after correcting for publication selection bias. To account for possible correlation among multiple estimates from the same primary study, we cluster standard errors at the study level.

Table 2. Linear tests of funnel asymmetry: *employment*

	OLS	BE	FE	Precision	Study	IV
β_1 - Publication	-0.543	-1.047	-0.782	-0.562	-0.824	-0.904
	(0.808)	(0.837)	(1.459)	(0.488)	(0.751)	(1.019)
β_0 - Mean beyond	0.00619	0.00363	0.0177	0.00713	-0.00572	0.0236
	(0.028)	(0.041)	(0.070)	(0.011)	(0.031)	(0.035)
First stage F stat						256.54
N	644	644	644	644	644	644

Notes: Standard errors in parentheses are clustered at the study level. OLS = Ordinary Least Squares, BE = study-level between effects, FE = study-level fixed effects, Precision = Weighed Least Squares with the inverse of the PCC's standard error, Study = Weighed Least Squares with the inverse of the number of estimates reported by study, IV = the standard error is instrumented by the inverse of the square root of the number of observations employed for each estimate. *** p<0.01, ** p<0.05, * p<0.1. *Source*: Authors' elaboration

The first column of Table 3 presents the results from estimating Equation 2 using Ordinary Least Squares (OLS). The coefficient of the standard error is negative but statistically insignificant, suggesting the lack of publication bias (FAT test). The PET test further underscores that the 'genuine' effect of robots on employment is essentially zero. This holds when also when we exploit between-study (BE) and within-study (FE) variation. To address heteroscedasticity concerns, weighted least squares (WLS) is employed. This method assigns greater weight to more precise estimates by using the inverse of standard error as a weighting factor (Stanley and Doucouliagos, 2017). In line with recent studies (Gechert et al., 2022), data are also weighted by the inverse of the number of estimates reported in each study. However, no signs of publication bias seem to emerge. We next turn our attention to endogeneity, acknowledging the potential correlation between the standard error of the PCCs and the error term (for a discussion, see Irsova et al., 2023a). To address this concern, the inverse of the square root of the primary study's sample size is employed as an instrument for the standard error of the PCC. The IV results align with prior findings, affirming the lack of any genuine empirical effect of robotization on employment across all models[8].

Table 3. Nonlinear techniques: *employment*

	Constant (Mean beyond bias)	
Uncorrected mean	-0.0199**	(0.006)
Top 10%	-0.004	(0.004)
WAAP	0.003**	(0.0002)
Endogenous kink	0.007**	(0.002)

Notes: Standard errors in parentheses. Top 10 approach is considering only 10% of the most precise estimates. WAAP = method developed by Ioannidis et al. (2017) focusing on estimates with adequate power (≥ 80%). Kinked = method developed by Bom and Rachinger (2019) that searches for a precision threshold above which publication bias is unlikely. Source: Authors' elaboration

Finally, we reassess the assumption of a linear relationship between publication selection and standard error. Recent discussions (Cazachevici et al., 2020; Gechert and Heimberger, 2022; Irsova et al., 2023b) emphasise that highly precise estimates at or below the 5% level are generally less prone to publication bias. In such cases, employing a linear correction for publication bias may result in an overcompensation, leading to a reverse bias. To mitigate this concern, we employ nonlinear estimation methods of the 'true' effect in Table 3. We first employ the Top 10 method developed in Stanley et al. (2010), who showed that discarding the 90% least precise results can significantly reduce publication bias. The findings are again very much in line with our traditional parametric estimates. Additionally, we implement the Weighted Average of Adequately Powered estimates (WAAP) method proposed in Ioannidis et al. (2017), yielding a positive corrected mean (beyond bias) of 0.003. This figure contrasts with the unweighted average of -0.02 derived from the entire dataset, suggesting the presence of publication bias favouring negative results. WAAP considers only the most informative and potentially least biased estimates. The last row reports the mean beyond bias calculated based on the method proposed in Bom and Rachinger (2019), which seeks endogenously a precision threshold above which publication bias is unlikely. This endogenous kink model yields an estimate of 0.007.

---

[8] Our findings remain robust after winsorizing the partial correlation coefficients (PCCs) and their standard errors at the 1st and 99th percentiles - minimizing the impact of potential outliers.

Our findings, using various linear and nonlinear methods, have consistently pointed to a minimal or non-existent genuine effect of robotization on employment. While there is little evidence of publication bias towards negative results in Table 2, the adjusted mean values are economically negligible (substantially below a rule of thumb 0.3 that denotes a small effect - Doucouliagos, 2011).

Table 4. Linear tests of funnel asymmetry: *wages*

	OLS	BE	FE	Precision	Study	IV
β_1 - Publication	-1.959**	0.939	-3.138***	-0.628	-0.0399	-2.292**
	(0.979)	(0.748)	(0.00582)	(0.780)	(1.106)	(0.952)
β_0 - Mean beyond	0.0436**	-0.0339	0.0861***	-0.00445*	0.00163	0.0555***
	(0.0163)	(0.0356)	(0.000210	(0.00246)	(0.0282)	(0.0173)
First stage F stat						572.096
N	195	195	195	195	195	195

Notes: Standard errors in parentheses are clustered at the study level. OLS = Ordinary Least Squares, BE = study-level between effects, FE = study-level fixed effects, Precision = Weighed Least Squares with the inverse of the PCC's standard error, Study = Weighed Least Squares with the inverse of the number of estimates reported by study, IV = the standard error is instrumented by the inverse of the square root of the number of observations employed for each estimate. *** p<0.01, ** p<0.05, * p<0.1 Source: Authors' elaboration

We now shift our focus to the impact on wages, reporting the results of the FAT-PET test in Table 4. A significant and negative β1 coefficient in both the OLS and IV models indicates the presence of publication bias, implying that researchers tend to publish negative results more often that one would expect. The FE models also shows a highly significant negative coefficient, reinforcing the presence of the bias. The magnitude of this coefficient, exceeding 2, suggests a 'severe' level of publication bias according to most specifications (Doucouliagos and Stanley, 2013). The corrected mean values, switch from negative to positive, though they remain modest, well below the threshold for a moderate correlation (>0.5). Instead, when more weight is given to the more precise estimates with WLS, the corrected mean partial correlation coefficient, while reduced, still retains a negative sign (-0.004).

Table 5. Nonlinear techniques: *wages*

	Constant (Mean beyond bias)	
Uncorrected mean	-0.027*	(0.0109)
Top 10%	-0.0052*	(0.002)
WAAP	-0.004**	(0.002)
Endogenous kink	-0.006***	(0.001)

Notes: Standard errors in parentheses. Top 10 approach is considering only 10% of the most precise estimates. WAAP = method developed by Ioannidis et al. (2017) focusing on estimates with adequate power (≥ 80%). Kinked = method developed by Bom and Rachinger (2019) that searches for a precision threshold above which publication bias is unlikely. Source: Authors' elaboration

The results from nonlinear techniques presented in Table 5, including the top 10%, WAAP with statistical power above 80% and the most precise estimates determined endogenously, are based on the selection of highly precise estimates. These approaches are grounded on the assumption that the latter are less susceptible to publication bias. All three methods corroborate the existence of publication bias, as the mean beyond bias is lower relative to unweighted mean (-0.027) albeit negative. They support the presence of a negative effect of robotization on wages after adjusting for publication bias, although the impact is significantly resized.

5. HETEROGENEITY

This section delves into the reasons behind the heterogeneous estimates stemming from primary studies. Such variability is attributable to multiple sources of heterogeneity, including, among others, level of analysis, country selection, method employed, each of which could influence the relationships at stake. Therefore, alongside fundamental data elements such as effect size, standard error and the sample size, we systematically collect all relevant information that may explain differences in empirical findings. In what follows, we provide an overview of the coded data used in our multivariate meta-regression analysis. The list of variables is given in Table A1.

5.1 List of coded variables

Data characteristics. The articles included in the meta-analysis exhibit differences in the measurement of the dependent variable. To account for this, we first introduce a dummy variable "hours worked" that equals 1 if a study uses hours worked as a proxy for employment and 0 otherwise (i.e., number of employees). Likewise, different proxies for robotization are used, predominantly robot density from the IFR database, alongside patents and investments. We control for these differences by including a corresponding binary variable, IFR. Furthermore, in line with recent studies (Gechert et al., 2022; Heimberger, 2023), we create a variable denoting the average year of the data in the primary studies under examination to test if the selected time span matters, i.e., whether there is a temporal variation in the estimated effect size.

Unit of analysis. Another relevant source of cross-study heterogeneity is the unit of analysis. Micro-level studies often highlight positive effects of technological change, but factors like business stealing, selection dynamics and reallocation effects might yield different outcomes at more aggregated levels (Calvino and Virgillito, 2018). To test this, we construct a series of binary variables that examine systematic differences in estimates arising from diverse empirical settings, including country, local labor market, industry, firm and individual-level studies.

Country selection. The impact of robotization on employment and wages varies significantly across countries, influenced by factors like technological capabilities and institutional contexts (Reljic et al., 2023). To test whether the selection of countries affects results, we classify studies based on the development levels of the countries in their sample, creating dummy variables for developed, emerging economies or a mix of both. We also control if estimates are based on single-country or cross-country studies. In addition, we code a dummy variable equal to 1 for studies using the US data exclusively and another for those based on the EU data.

Manufacturing versus services. The aggregate impact of robotization on employment and wages may hide the underlying dynamics of job displacement in robot adopting sectors and the compensatory reallocation (i.e., creation of new jobs) elsewhere in the economy (Vivarelli, 1995; Dauth et al., 2021). The effects on employment and wages could vary between manufacturing industries - where the majority of robots are employed - and service sectors. To explore this, we create a dummy variable set to 1 if estimates refer to manufacturing sector and 0 otherwise.

Short- vs. long-run analysis. We examine the differences between the short- and long-term effects of robotization. In other words, we test whether market compensation mechanisms operate 'effectively', offsetting the potential short-term displacement effects with the eventual medium- to long-term reallocation dynamics (Vivarelli, 1995).

Qualitative effects. While most primary studies focus on the overall employment/wage impact, some report estimates for worker groups differentiated by skills, gender, age and contract type. We code these by creating a set of binary variables to assess whether the employment/wage impact of robotization varies with demographic and occupational attributes of labour such as education (low, medium and high), age (young, middle-aged and old), gender (female and male), contract (full-time vs. part-time) and level of routine task intensity (routine vs. non-routine).

Method of estimation. We also check whether the use of different estimation techniques affects the magnitude and direction of robotization's impact on employment and wages. A set of binary variables has been created to distinguish among different estimators, namely Instrumental Variables (IV), Ordinary Least Squares (OLS), and others (e.g., difference-in-difference). At the same time, this is critical to understanding whether addressing endogeneity concerns - by instrumenting the robotization variable - yield markedly different results.

Inclusion of relevant controls. An important aspect is whether researchers control for workforce characteristics. We construct a dummy variable equal to 1 if a study considers at least two out of four labour market attributes (e.g., education, occupation, age, gender) and 0 otherwise.

Lastly, the heterogeneity in findings may be influenced by publication characteristics. In line with recent research (Gechert et al., 2022; Heimberger, 2023), we collect and control for the number of citations each primary study has received, as reported in Scopus. It is argued by some that, after adjusting for publication bias, the citation count is also indicative of the study's quality (Cazachevici et al., 2020).

5.2 Estimation and results

To explore heterogeneity, we extend the model (Equation 2) by including a vector of 31 "moderator" variables that aim to capture differences in the estimated impact of robotization on employment and wages across studies:

$$r_{ij} = \beta_0 + \beta_1 SE_{ij} + \gamma X_{ij} + \varepsilon_{ij} \tag{3}$$

Where i^{th} partial correlation coefficient (r) from study j is influenced by a vector of variables (X_{ij}) consisting of study characteristics discussed above.

In line with previous research (Gechert and Heimberger, 2022; Heimberger, 2020; Heimberger, 2023), we adopt a general-to-specific approach for the estimation of the linear regression model (Equation 3). This entails the sequential exclusion of the least statistically significant explanatory variable, proceeding iteratively until all remaining variables exhibit statistical significance at the 10% threshold (Stanley and Doucouliagos, 2012). It is important to note that in our model, one reference category per set of mutually exclusive dummy variables is omitted to prevent multicollinearity. Therefore, the intercept β_0 cannot be interpreted as the genuine effect of robotization on employment and wages, as its value varies with the chosen reference category.

We firstly apply OLS for estimating Equation 3, complemented by a robustness check through winsorising the data at the 1% level. Subsequently, we include all coded variables identified as statistically significant in the previous step and re-estimate with WLS, giving more weight to more precise estimates. Finally, to further ensure robustness and control for any unobserved heterogeneity, we re-estimate Equation 3 using a fixed-effects model. In this process, we will also check the robustness of our findings regarding publication bias (Section 4), because heterogeneity can cause asymmetry in funnel plots, even when such bias is not present (Kroupova et al., 2022).

Table 6 presents the results of a multivariate meta-regression examining the relationship between robotization and employment. Notably, the constant term in all models lacks statistical significance, indicating that - when all other variables in the model are at their reference levels - the employment impact of robotization is not significantly different from zero. This is consistent with findings in Table 2. Likewise, the coefficient for SE is negative, yet not significant, suggesting the absence of publication bias.

A key question, therefore, is: what factors contribute to the heterogeneity observed in empirical studies regarding the robot-employment nexus?

Table 6. Multivariate meta-regression results, *employment*

	(1)	(2)	(3)	(4)	(5)
	Unweighted			Weighted	FE
	benchmark	+SE	Winsorized		
Constant	-0.000580	0.00843	0.0107	-0.0135	0.00658
	(0.0402)	(0.0427)	(0.0388)	(0.0260)	(0.107)
SE		-0.424	-0.318	-0.292	-0.654
		(0.697)	(0.606)	(0.448)	(1.382)
Country ==Developed	0.0725**	0.0820*	0.0756*	0.0519*	0.0121
	(0.0351)	(0.0463)	(0.0412)	(0.0298)	(0.0267)
EU	-0.0823**	-0.0931**	-0.0833**	-0.0304	
	(0.0360)	(0.0441)	(0.0403)	(0.0337)	
US	-0.169***	-0.182***	-0.168***	-0.103***	
	(0.0468)	(0.0580)	(0.0531)	(0.0334)	
Unit==Individual	0.109***	0.0863*	0.0847**	0.0404**	0.0838
	(0.0322)	(0.0433)	(0.0387)	(0.0192)	(0.0684)
Methods==OLS	0.0772***	0.0773***	0.0716***	0.0533***	0.0295**
	(0.0212)	(0.0213)	(0.0191)	(0.0185)	(0.0123)
Controls LM	0.0575*	0.0551*	0.0492*	0.0126	0.0724*
	(0.0320)	(0.0297)	(0.0272)	(0.0177)	(0.0379)
Effect term==Long-run	-0.0744**	-0.0596	-0.0647*	-0.0305	-0.0741***
	(0.0336)	(0.0375)	(0.0352)	(0.0265)	(0.0260)
Task==Non-routine	0.119***	0.109***	0.107***	0.0919***	0.120***
	(0.0394)	(0.0276)	(0.0264)	(0.0329)	(6.17e-05)
Citations	-0.0154*	-0.0147*	-0.0135*	-0.00947	
	(0.00838)	(0.00737)	(0.00727)	(0.00879)	
Observations	644	644	644	644	644
R-squared	0.274	0.281	0.292	0.215	0.068

Notes: Robust standard errors in parentheses clustered at the study level. *** p<0.01, ** p<0.05, * p<0.1 *Source*: Authors' elaboration

Several robust findings stand out: (i) there is a heterogeneous impact of industrial robots in developed versus emerging countries.[9] In the former, robotization appears to complement the workforce, except for the US where a negative employment impact is detected. By the same token, combining the developed countries and the EU coefficients leads to a negative yet minimal effect. Taken together, these findings imply that the effects of robotization cannot be universally generalized across developed economies and are likely influenced by country-specific structural and institutional factors (e.g., labour market regulations); (ii) in line with the 'routinization hypothesis' (Autor et al., 2003), the positive coefficient associated to non-routine jobs suggest that the latter are less susceptible to automation as compared to other occupations.

[9] Note that this result doesn't exclude the option that foreign robot adoption (typically in developed countries) may have negative implications on employment in emerging countries via reshoring channel as some recent studies suggest (Faber, 2020; Kugler et al., 2020).

Indeed, higher robotization seems to create more demand for non-routine jobs in relation to routine ones; (iii) studies failing to address endogeneity concerns tend to report more favourable employment outcomes compared to those employing IV approach and other methods. This suggests that OLS may overestimate the positive effects of robotization or underestimate its disruptive potential. Consequently, the choice of the estimator is confirmed to affect the results; (iv) the empirical setting also 'moderates' the robot-employment relationship. While the overall effect of robots on employment, if any, is minimal, a positive effect is observed at the individual level.

Furthermore, despite these findings should be interpreted cautiously due to a lack of robustness across all specifications, the negative coefficient associated with the long-term effect suggests that the adverse impacts of robotization may only become apparent over a prolonged period. This finding goes against the conventional wisdom, which points to the potential adverse effects in the short run that are counterbalanced by market compensation mechanisms in the long run. This implies that investigating short-term labour market responses to robotization may be insufficient to gauge the potential impact of robotization. Moreover, we find that workers' characteristics can be a confounding factor when exploring the employment effect of robotization. For instance, workers with certain skills or in specific occupations might be more or less affected by automation. Therefore, it's important to control for these worker characteristics to isolate the effect of robotization itself. Finally, although the effect is small, we find that studies with higher number of citations tend to report more adverse effects of robotization on employment.

Summing up, the evidence on robotization and employment is not monolithic. The impact of robotization is 'moderated' by various factors including, but not limited to, the socioeconomic context, methodological rigor and the nature of jobs.

Shifting focus to the effects on wages, we present the multivariate meta-regression findings related to a subsample of our primary studies (Table 7). The constant term, statistically significant in all models, implies a positive but moderate effect of robotization on wages (less than 0.5) when other variables in the model are at their reference levels. This finding stands in contrast to the overall unweighted and uncorrected mean of -0.03 across studies. It indicates robust evidence for a publication bias favouring negative results, even when accounting for cross-study heterogeneity. Notably, the constant is higher compared to the results in Table 4, although it cannot be interpreted as the 'true' effect due to its dependency on the chosen reference categories.

According to our results, which align with the evidence presented above, the impact of robotization depends on the level of economic development in the countries under investigation. Notably, robotization tends to yield greater wage growth in developed than in emerging countries. This does not imply that robotization reduces wages in emerging countries, rather, it suggests that developed countries reap more benefits from it, possibly due to the differences in occupational and skill structure, among other factors. This, however, doesn't hold for the US and EU countries, where estimates are significantly lower, reinforcing the pattern observed in the robot-employment nexus (see Table 6).

Table 7: Multivariate meta-regression analysis: wages

	(1)	(2)	(3)	(4)
	Unweighted		Weighted	FE
	benchmark	Winsorized		
Constant	0.137***	0.143***	0.147***	0.0822***
	(0.0369)	(0.0373)	(0.0413)	(0.0129)
SE	-2.539***	-2.636***	-2.592***	-2.831***
	(0.230)	(0.253)	(0.606)	(0.137)
Country==Developed	0.352***	0.353***	0.241***	
	(0.0644)	(0.0656)	(0.0547)	
US	-0.368***	-0.372***	-0.289***	0.00422
	(0.0515)	(0.0530)	(0.0632)	(0.00515)
EU	-0.351***	-0.355***	-0.281***	-0.103
	(0.0538)	(0.0552)	(0.0626)	(0.0653)
Unit==Country Level	0.382***	0.381***	0.321***	
	(0.0279)	(0.0277)	(0.0161)	
Manufacturing	-0.101**	-0.101**	-0.0799***	
	(0.0355)	(0.0358)	(0.0252)	
Methods==OLS	0.0128**	0.0132**	0.0163**	-0.164***
	(0.00532)	(0.00535)	(0.00588)	(0.0218)
IFR	-0.130***	-0.129***	-0.103***	
	(0.0287)	(0.0288)	(0.0186)	
Controls LM	0.0346**	0.0331**	0.00787	0.0228
	(0.0144)	(0.0132)	(0.00479)	(0.0164)
Skill==Low skill	-0.138*	-0.137*	-0.0549	0.00176***
	(0.0792)	(0.0788)	(0.0528)	(0.000106)
Skill==Medium skill	-0.231***	-0.231***	-0.202***	-0.0365***
	(0.0555)	(0.0575)	(0.0511)	(0.000225)
Age==Young	0.0828***	0.0856***	0.115***	
	(0.0149)	(0.0160)	(0.0240)	
Age==Middle-aged	0.0847***	0.0875***	0.117***	
	(0.0148)	(0.0158)	(0.0235)	
Age==Old	0.0458***	0.0488***	0.0786***	-0.147***
	(0.0152)	(0.0163)	(0.0249)	(3.06e-05)
Contract==FT	-0.267***	-0.268***	-0.178***	
	(0.0624)	(0.0641)	(0.0512)	
Contract==PT	-0.120*	-0.121*	-0.0306	
	(0.0624)	(0.0641)	(0.0512)	
Citations	-0.0110**	-0.0111**	-0.00254***	-0.0788**
	(0.00502)	(0.00524)	(0.000522)	(0.0330)
Observations	195	195	195	195
R-squared	0.782	0.792	0.627	0.762

Notes: Robust standard errors in parentheses clustered at the study level. *** p<0.01, ** p<0.05, * p<0.1 *Source*: Authors' elaboration

The level of analysis also affects the wage impact of robots. Namely, studies conducted at the country level report more positive effects of robotization on wages compared to those conducted at a more micro level (such as the industry or firm level). However, it should be noted that while aggregate analyses can provide insights into the broader economic impact of robotization, they can mask the negative impacts experienced by specific groups of workers or sectors. Indeed, we find that wage premia stemming from robotization are asymmetrically distributed, not only between sectors but also among different skill groups. Medium-skilled workers tend to experience a more pronounced reduction in wages, which partly aligns with the RBTC theory, whereby a reduction in wages is observed for the medium-skilled worker category. However, no robust results emerge for low-skilled workers albeit there is some evidence of wage moderation. Studies focusing on manufacturing sectors only reports more negative effects in relation to non-manufacturing sectors. Interestingly, wages of full-time workers appear to be more susceptible to the increasing adoption of industrial robots. Finally, studies using IFR data generally report a more negative impact compared to other data sources.

Overall, our meta-analysis uncovers a range of interesting patterns. There is a large-to-severe publication bias, leaning towards negative outcomes. Once accounted for this bias, a positive wage effect seems to prevail. Factors such as country, sector and employment type matter: manufacturing sectors, medium-skilled and full-time workers are most susceptible to negative wage effects.

6. CONCLUDING REMARKS

This paper provides one of the first meta-analysis of the literature studying the relationship between robots and employment. Moreover, a subsample of contributions is separately analysed to assess the impact on wages. The analysis builds upon 33 studies, which delve into the employment impact of robotization - yielding 644 estimates - and 19 studies examining wage effects - yielding 195 estimates.

The meta-analysis reveals, on average, negative effects of robotization on both employment and wages, although these effects remain marginal and close to zero. As a result, claims regarding the disruptive impact of robot diffusion and the associated risks of widespread technological unemployment seem to be overstated, or at the very least, lack empirical support. However, this does not mean that robots are unlikely to gain greater disruptive potential in the future, or that they are not already disruptive in specific contexts (e.g., the US and, to a lower extent, the EU, medium-skilled occupations and manufacturing in the case of wages). A publication bias towards negative outcomes is detected concerning the robots-wages relationship. Nonetheless, the size of the effect remains negligible.

Several research questions still need to be addressed. First, there is the issue of measuring robot adoption. Most studies rely on the IFR's industry-level data. While this is perhaps the best measure of robot adoption at the industry level, this indicator has also been imputed to other empirical settings (e.g., LLMs) relying on strong assumptions regarding the distribution of robots across territorial domains. Equally relevant is the fact that, by construction, IFR data ignores technological progress, assuming that a robot installed in 1993 is of the same 'quality' as one installed in 2021 (Klump et al., 2021). Another source of robot data is the UN Comtrade on net imports of industrial robots (Blanas et al., 2019). However, even in this case, the mismeasurement risk of robotization is not negligible, especially for firms, industries and countries that rely on domestically produced robots. Therefore, more research efforts are needed to come up with more sound indicators of robotization.

A second element concerns the complexity of the labour-technology nexus. While technology is an important driver of transformation in labour markets, it is unlikely the only one. In light of this, future studies should explore how robotization interacts with other structural factors concerning supply and demand conditions, countries and industries technological capabilities (Bogliacino and Pianta, 2010; Dosi et al. 2021) and prevailing competitive strategies (i.e., technological vs. cost-competitive strategies). Equally important is the role of institutions that may mitigate or exacerbate robotization effects (Cetrulo et al., 2019; Pianta and Reljic, 2022).

The scope of existing empirical studies predominantly centred on industrial robots. However, a substantial segment of recent advancements in robotics concerns the service sector with increasingly "intelligent" AI-enhanced robots, encompassing various domains of application (e.g., medical robotics and logistics management). Consequently, it is plausible that robots will continue to exert an influence on employment and wages in the foreseeable future, particularly within the service sector. Another promising research avenue would be to explore the effects of service robots.

REFERENCES

Acemoglu, D., & Restrepo, P. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5), 1973-2016.

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

Aghion, P., Antonin, C., Bunel, S., Jaravel, X., 2022. The effects of automation on labor demand. In: *Robots and AI: A New Economic Era*. Routledge, 15–39.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2020). What are the labor and product market effects of automation? New evidence from France. CEPR Discussion Paper, DP14443.

Aksoy, C. G., Ozcan, B., & Philipp, J. (2021). Robots and the gender paygap in Europe. *European Economic Review*, *134*.

Albinowski, M., & Lewandowski, P. (2022). The Impact of ICT and Robots on Labour Market Outcomes of Demographic Groups in Europe. *IBS Working Paper*, *4.*

Anelli, M., Colantone, I., & Stanig, P. (2021). Individual vulnerability to industrial robot adoption increases support for the radical right. *Proceedings of the National Academy of Sciences*, 118(47).

Antón, J.-I., Fernández-Macías, E., Brancati, M. C. U., & Alaveras, G. (2022). The labour market impact of robotisation in Europe. *European Journal of Industrial Relations*, 28(3), 317-339.

Arntz, M., Gregory, T., & Zierahn, U. (2016). The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. *OECD Social, Employment and Migration Working Papers,* 189.

Artuc, E., Christiaensen, L., & Winkler, H. (2019). Does automation in rich countries hurt developing ones?: Evidence from the US and Mexico. Evidence from the US And Mexico (February 14, 2019). *World Bank Policy Research Working Paper*, 8741.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics, 118(4), 1279-1333.*

Baldwin, R. (2013). Trade and industrialization after globalization's second unbundling: How building and joining a supply chain are different and why it matters. *Globalization in an age of crisis: Multilateral economic cooperation in the twenty-first century*, University of Chicago Press, 165-212.

Ballestar, M. T., Díaz-Chao, A., Sainz, J., & Torrent-Sellens, J. (2020). Knowledge, robots and productivity in SMEs: Explaining the second digital wave. *Journal of Business Research*, *108*, 119-131.

Balsmeier, B., & Woerter, M. (2019). Is this time different? How digitalization influences job creation and destruction. *Research policy, 48(8), 103765.*

Barbieri, L., Mussida, C., Piva, M., & Vivarelli, M. (2019). Testing the Employment Impact of Automation, Robots and AI: A Survey and Some Methodological Issues. *IZA – Institute of Labor Economics, 12612*.

Benmelech, E., & Zator, M. (2022). Robots and Firm Investment. *NBER Working Papers,* 29676.

Blanas, S., Gancia, G., & Lee S. Y. T. (2019). Who is afraid of machines?. *Economic Policy*, 34(100), 627-690.

Bogliacino, F., & Pianta, M. (2010). Innovation and Employment: a Reinvestigation using Revised Pavitt classes. *Research Policy*, 39 (6), 799-809.

Bom, P. R. D., & Rachinger H. (2019). A kinked meta-regression model for publication bias correction. *Research Synthesis Methods*, 10, 497–514.

Bonfiglioli, A., Crino, R., Gancia G., & Papadakis, I. (2022). Robots, Offshoring, and Welfare. *Robots and AI,* 40-81.

Borjas, G. J., & Freeman, R. B. (2019). From Immigrants to Robots: The Changing Locus of Substitutes for Workers. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, *5*(5), 22–42.

Brambilla, I., Cesar, A., Falcone, G., & Gasparini, L. (2022). The impact of robots in Latin America: Evidence from local labor markets. *Cedlas Working Papers,* 4451.

Brown, A. L., Imai, T., Vieider, F. M., & Camerer, C. F. (2024). Meta-analysis of Empirical Estimates of Loss Aversion. *Journal of Economic Literature* (forthcoming)**.**

Calvino, F., & Virgillito, M. E. (2018). The innovation-employment nexus: a critical survey of theory and empirics. *Journal of Economic surveys*, 32(1), 83-117.

Carbonero, F., Ekkehard, E., & Weber, E. (2020). Robots Worldwide: The Impact of Automation on Employment and Trade. *IAB-Discussion Paper*, 7.

Card, D., Kluve, J., & Weber, A. (2018). What works? A meta-analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, 16(3), 894-931.

Cazachevici, A., Havranek, T., & Horvath, R. (2020). Remittances and economic growth: A meta-analysis. *World Development*, 134, 105021.

Cetrulo, A., Cirillo, V., & Guarascio, D. (2019). Weaker jobs, weaker innovation. Exploring the effects of temporary employment on new products. *Applied Economics*, *51*(59), 6350-6375.

Chiacchio, F., Petropoulos, G., & Pichler, D. (2018). The impact of industrial robots on EU employment and wages: A local labour market approach. *Bruegel Working Paper, 2*.

Compagnucci, F., Gentili, A., Valentini, E., & Gallegati, M. (2019). Robotization and labour dislocation in the manufacturing sectors of OECD countries: a panel VAR approach. *Applied Economics*, *51*(57), 6127-6138.

Corrocher, N., Moschella, D., Staccioli, J., & Vivarelli, M. (2023). Innovation and the labor market: Theory, evidence and challenges. *IZA Discussion paper*, 16199.

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2018). Adjusting to Robots: Worker-Level Evidence*.* *Opportunity and Inclusive Growth Institute Working Paper*, 13.

de Vries, G. J., Gentile, E., Miroudot, S., & Wacker, K. M. (2020). The rise of robots and the fall of routine jobs. *Labour Economics*, *66*, 101885.

Dekle, R. (2020). Robots and industrial labor: Evidence from Japan. *Journal of the Japanese and International Economies*, *58*, 101108.

Dixon, J., Hong, B., & Wu, L. (2021). The Robot Revolution: Managerial and Employment Consequences for Firms. *Management Science*, *67*(9), 5586-5605.

Dosi, G., Piva, M., Virgillito, M. E., & Vivarelli, M. (2021). Embodied and disembodied technological change: The sectoral patterns of job-creation and job-destruction. *Research Policy*, 50(4).

Dottori, D. (2021). Robots and employment: evidence from Italy. *Economia Politica*, *38*(2), 739-795.

Doucouliagos, C., & Stanley, T. D. (2013): Are all economic facts greatly exaggerated? theory competition and selectivity. *Journal of Economic Surveys*, 27(2), 316–339.

Doucouliagos, H. (2011): How large is large? Preliminary and relative guidelines for interpreting partial correlations in economics. Economics Series Working Paper 5, Deakin University.

Faber, M. (2020). Robots and reshoring: Evidence from Mexican labor markets. *Journal of International Economics*, 127, 103384.

Fernandez-Macias, E., Klenert, D., & Anton, J. I. (2021). Not so disruptive yet? Characteristics, distribution and determinants of robots in Europe. *Structural Change and Economic Dynamics*, 58, 76-89.

Fiedler, P., Fidrmuc, J., & Reck, F. (2021). Automation, Digitalization, and Income Inequality in Europe. *Czech Journal of Economics and Finance,* 71(3), 203-219.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation?. *Technological Forecasting and Social Change*, 114(C), 254-280.

Fu, X. (Maggie), Bao, Q., Xie, H., & Fu, X. (2021). Diffusion of industrial robotics and inclusive growth: Labour market evidence from cross country data. *Journal of Business Research*, *122*, 670-684.

Ge, S., & Zhou, Y. (2020). Robots, computers, and the gender wage gap. *Journal of Economic Behavior & Organization*, 178, 194–222.

Gechert, S., & Heimberger, P. (2022). Do corporate tax cuts boost economic growth?. *European Economic Review*, 147, 104157.

Giuntella, O., & Wang, T. (2019). Is an Army of Robots Marching on Chinese Jobs?. *IZA – Institute of Labor Economics*, *12281*.

Graetz, G. & Michaels, G. (2015). Robots at Work. *IZA Discussion Papers*, *8938*.

Gravina, A. F., & Pappalardo, M. R. (2022). Are robots in rich countries a threat for employment in emerging economies?. *Economics Letters*, 221, 110888.

Havranek, T., Irsova, Z., & Zeynalova, O. (2018). Tuition fees and university enrolment:Ameta-regression analysis. *Oxford Bulletin of Economics and Statistics*, 80(6), 1145–1184.

Heimberger, P. (2020). Does economic globalisation affect income inequality? A meta-analysis. *The World Economy*, 43(11), 2960-2982.

Heimberger, P. (2021). Corporate tax competition: A meta-analysis. *European Journal of Political Economy*, 69, 102002.

Heimberger, P. (2022). Does economic globalisation promote economic growth? A meta-analysis. *The World Economy*, 45(6), 1690-1712.

Heimberger, P. (2023). Do higher public debt levels reduce economic growth?. *Journal of Economic Surveys*, 37(4), 1061-1089.

Ioannidis, J. P., Stanley, T. D., & Doucouliagos, H. (2017). The power of bias in economics research. *The Economic Journal,* 127(605), 236-265.

Irsova, Z., Bom, P. R. D., Havranek, T., & Rachinger, H. (2023b). Spurious precision in meta-analysis. Retrieved from meta-analysis.cz/maive

Irsova, Z., Doucouliagos, H., Havranek, T., & Stanley, T. D. (2023). Meta-analysis of social science research: A practitioner's guide. *Journal of Economic Surveys*, 1-20.

Jäger, A., Moll, C., & Lerch, C. (2016). Analysis of the Impact of robotic systems on employment in the European Union - Update. Publications Office of the European Union, Luxembourg.

Jestl, S. (2022). Industrial Robots, and Information and Communication Technology: The Employment Effects in EU Labour Markets. Wiiw Working Papers, 215.

Jung, J. H., & Lim, D. G. (2020). Industrial robots, employment growth,and labor cost: A simultaneous equation analysis. *Technological Forecasting & Social Change*, *159*.

Kariel, J. (2021). Job Creators or Job Killers? Heterogeneous Effects of Industrial Robots on UK Employment. *CEIS, Fondazione Giacomo Brodolini and John Wiley & Sons Ltd*, 52–78.

Klenert, D., Fernandez-Macias, E., & Anton, J.-I. (2022). Do robots really destroy jobs? Evidence from Europe. *Economic and Industrial Democracy*, *1*(37).

Klump, R., Jurkat, A., & Schneider, F. (2021). Tracking the rise of robots: A survey of the IFR database and its applications. MPRA Working Paper 110390. Munich Personal REPEc Archive, München.

Koch, M., Manuylov, I., & Smolka, M. (2021). Robots and Firms. *The Economic Journal*, *131*, 2553–2584.

Krenz, A., Prettner, K., & Strulik, H. (2021). Robots, reshoring, and the lot of low-skilled workers. *European Economic Review*, 136.

Kroupova, K., Havranek, T., & Irsova, Z. (2022). Student Employment and Education: A Meta-Analysis. *CEPR Discussion Papers* 16550.

Kugler, A. D., Kugler, M., Ripani, L., & Rodrigo, R. (2020). Us robots and their impacts in the tropics: Evidence from colombian labor markets. *National Bureau of Economic Research,* w28034.

Li, J., Zhang, R., Zheng, P., Lu, Y., Bao, J., & Sun, X. (2022). A fusion-based spiking neural network approach for predicting collaboration request in human-robot collaboration. *Robotics and Computer-Integrated Manufacturing*, 78.

Leigh, N. G., Kraft, B., & Lee, H. (2020). Robots, skill demand and manufacturing in US regional labour markets. *Cambridge Journal of Regions, Economy and Society*, *13*, 77–97.

Mann, K., & Püttmann, L. (2023). Benign Effects of Automation: New Evidence from Patent Texts. *The Review of Economics and Statistics*, MIT Press, 105(3), 562-579.

Mondolo, J. (2022). The composite link between technological change and employment: A survey of the literature. *Journal of Economic Surveys*, *36*(4), 1027-1068.

Montobbio, F., Staccioli, J., Virgillito, M. E., & Vivarelli, M. (2023). The empirics of technology, employment and occupations: lessons learned and challenges ahead. *Journal of Economic Surveys*, 12601.

Nedelkoska L., & Quintini, G. (2018). Automation, skills use and training. *OECD Social, Employment and Migration Working Papers, 202.*

Ni, B., & Obashi, A. (2021). Robotics technology and firm-level employment adjustment in Japan. *Japan and the World Economy*, 57, 101054.

Petit, F., Jaccoud, F., & Ciarli, T. (2023). Heterogeneous Adjustments of Labor Markets to Automation Technologies. *CESifo Working Paper, 10237*.

Pianta, M., & Reljic, J. (2022). The good jobs-high innovation virtuous circle. *Economia Politica: Journal of Analytical and Institutional Economics*, 39(3), 783-811.

Pouliakas, K. (2018). Determinants of Automation Risk in the EU Labour Market: A Skills-Needs Approach. *IZA Discussion Papers,* 11829.

Reljic, J., Cirillo, V., & Guarascio, D. (2023). Regimes of robotization in Europe. *Economics Letters*, 232, 111320.

Rodrigo, R. (2022). Robot Adoption, Organizational Capital and the Productivity Paradox. https://rodimirorodrigo.github.io/JMP2021-Rodrigo.pdf.

Rodrik, D. (2022). Prospects for Global Economic Convergence Under New Technologies. *Brooking Reports*.

Schmidpeter, B., & Winter-Ebmer, R. (2021). Automation, unemployment, and the role of labor market training. *European Economic Review*, 137.

Sequeira, T. N., Garrido, S. & Santos, M. (2021). Robots are not always bad for employment and wages. *International Economics*, *167*, 108–119.

Shapiro, A. F., & Mandelman, F. S. (2021). Digital adoption, automation, and labor markets in developing countries. *Journal of Development Economics*, 151, 102656.

Sostero, M. (2020). Automation and Robots in Services: Review of Data and Taxonomy. *JRC Working Papers on Labour, Education and Technology*, 14.

Stanley, T., & Doucouliagos, H. (2017). Neither fixed nor random: Weighted least squares meta-analysis. *Statistics in Medicine*, 34(13), 2116–2127.

Stanley, T. D., Jarrell, S. B., & Doucouliagos, H. (2010). Could it be better to discard 90% of the data? A statistical paradox. *The American Statistician*, 64(1), 70–77

Stapleton, K., & Webb, M. (2020). Automation, trade and multinational activity: Micro evidence from Spain. *CSAE Working Paper Series*, *16*.

Stemmler, H. (2023). Automated Deindustrialization: How Global Robotization Affects Emerging Economies: Evidence from Brazil. *World Development*, 171.

Tang, C., Huang, K., & Liu, Q. (2021). Robots and skill-biased development in employment structure: Evidence from China. *Economics Letters*, *205*, 109960.

Vivarelli, M. (1995). The economics of technology and employment: Theory and empirical evidence. Elgar.

Vivarelli, M. (2014). Innovation, employment, and skills in advanced and developing countries: A survey of economic literature. *Journal of Economic Issues*, 48(1), 123–154.

APPENDIX

Figure A1. PRISMA Flowchart

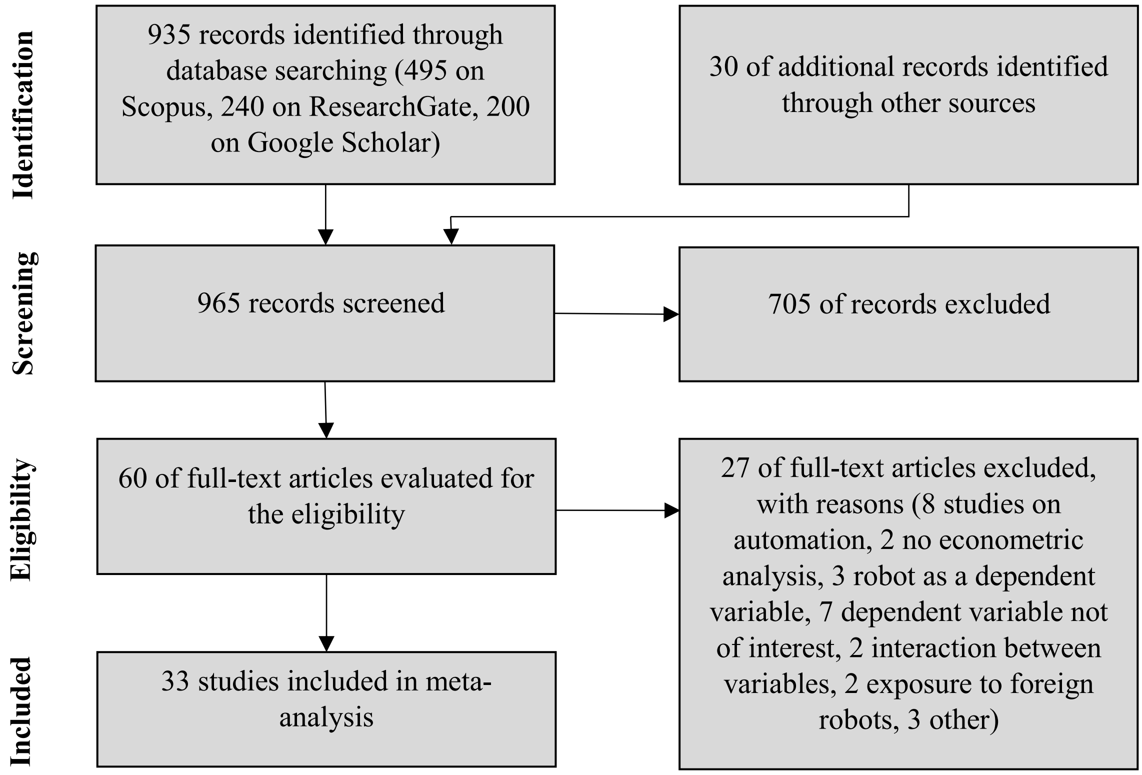

Source: Authors' elaboration

Table A1. Variables used in the meta-regression analysis

	Employment (N=644)		Wages (N=195)	
	Mean	SD	Mean	SD
Partial correlation coefficient (PCC)	-0,020	0,147	-0,027	0,153
Standard error of PCC	0,048	0,034	0,036	0,041
Logarithm of the number of citations	1,501	1,674	1,937	2,433
Unit of analysis				
Country Level == 1	0,043	0,273	0,223	-
Firm Level == 1 [reference]	0,013	0,048	0,003	0,007
Individual Level ==1	0,007	0,002	0,006	0,006
Industry Level == 1	-0,012	0,207	-0,060	0,234
LLM == 1	-0,036	0,102	-0,013	0,070
Sample				
Country Case ==1	-0,032	0,144	-0,047	0,179
Multi-country [reference]	-0,002	0,151	0,006	0,086
US == 1	-0,164	0,199	-0,089	0,199
EU == 1	0,000	0,105	0,005	0,042
Country group				
Developed == 1	-0,022	0,143	-0,027	0,155
Emerging ==1	0,021	0,133	-0,063	0,066
Country mix ==1 [reference]	-0,099	0,205	0,223	n.a.

Manufacturing == 1	-0,001	0,171	-0,146	0,300
Estimation methods				
IV ==1	-0,065	0,142	-0,050	0,169
OLS ==1	0,032	0,123	0,019	0,087
other ==1 [excluded category]	-0,114	0,286	0,038	0,174
Long run == 1	-0,036	0,144	-0,030	0,158
Data characteristics				
Hours worked == 1	-0,021	0,099	-	-
IFR == 1	-0,024	0,157	-0,032	0,159
Midpoint year	2004,262	4,530	2004,231	5,375
Controls == 1	-0,022	0,113	-0,002	0,057
Subsample - Gender				
Female ==1	0,017	0,087	-0,011	0,097
Male ==1	0,031	0,086	-0,023	0,075
Total == 1 [reference]	-0,023	0,150	-0,030	0,167
Subsample - Skills				
High-skilled ==1	0,026	0,099	0,030	0,049
Low-skilled ==1	-0,024	0,207	-0,042	0,070
Med-skilled == 1	-0,031	0,149	-0,096	0,037
Total ==1 [reference]	-0,022	0,144	-0,027	0,160
Subsample - Age				
Young ==1	0,063	0,101	0,029	0,067
Middle-aged ==1	-0,013	0,078	0,033	0,063
Old ==1	-0,021	0,055	-0,012	0,062
Total ==1 [reference]	-0,021	0,150	-0,033	0,161
Subsample - Contract				
Part-time ==1	0,038	0,094	0,126	0,041
Full-time ==1	-	-	-0,020	0,044
Total ==1 [reference]	-0,021	0,148	-0,034	0,156
Subsample - Tasks				
Non-routine ==1	0,032	0,090	-	-
Routine ==1	-0,097	0,038	-	-
Total ==1 [reference]	-0,019	0,149	-	-

Source: Authors' elaboration