

WORKING PAPER

INAPP WP n. 67

Stop worrying and love the robot: An activity-based approach to assess the impact of robotization on employment dynamics

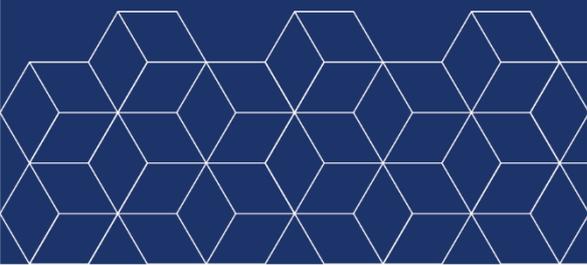
Mauro Caselli

Andrea Fracasso

Sergio Scicchitano

Silvio Traverso

Enrico Tundis



Stop worrying and love the robot: An activity-based approach to assess the impact of robotization on employment dynamics

Mauro Caselli

Scuola di Studi Internazionali, Università di Trento, Trento
Dipartimento di Economia e Management, Università di Trento, Trento
mauro.caselli@unitn.it

Andrea Fracasso

Scuola di Studi Internazionali, Università di Trento, Trento
Dipartimento di Economia e Management, Università di Trento, Trento
andrea.fracasso@unitn.it

Sergio Scicchitano

Istituto nazionale per l'analisi delle politiche pubbliche (INAPP), Roma
Global Labor Organization (GLO), Essen
s.scicchitano@inapp.org

Silvio Traverso

Dipartimento di Sociologia e Ricerca Sociale, Università di Trento, Trento
silvio.traverso@unitn.it

Enrico Tundis

Istituto di Statistica della Provincia di Trento (ISPAT), Trento
enrico.tundis@unitn.it

LUGLIO 2021

Andrea Fracasso and Silvio Traverso gratefully acknowledge the financial support from Fondazione Caritro (Project 2018.0258). The authors are grateful to Antonio Accetturo, Celine Antonin, Claire Lelarge, Matthias Mertens, and the participants in the workshop “Firms and Workers at the Crossroad: Automation and Market Power”, and Ispat for access to their data. All errors are our own.

CONTENTS: 1. Introduction. – 2. A new activity-based measure of local exposure to robot adoption. – 3. Empirical strategy; 3.1 The model; 3.2 Addressing endogeneity concerns. – 4. Data. – 5. Main results; 5.1 Robustness checks; 5.2 Industry-based measure of local exposure to robots. – 6. Employment dynamics across alternative occupational categories. – 7. Closing remarks. – Appendix. – References

INAPP – Istituto nazionale per l'analisi delle politiche pubbliche

Corso d'Italia 33
00198 Roma, Italia

Tel. +39 06854471
Email: urp@inapp.org

www.inapp.org

ABSTRACT

Stop worrying and love the robot: An activity-based approach to assess the impact of robotization on employment dynamics

This work investigates the impact that the change in the exposure to robots had on the Italian local employment dynamics over the period 2011-2018. A novel empirical strategy focusing on a match between occupations' activities and robots' applications at a high level of disaggregation makes it possible to assess the impact of robotization on the shares of workers employed as robot operators and in occupations deemed exposed to robots. In a framework consistently centered on workers' and robots' activities, rather than on their industries of employment, the analysis reveals for the first time reinstatement effects among robot operators and heterogeneous results among exposed occupations.

KEYWORDS: econometric analysis, technological innovation, employment, robot

JEL CODES: J21, J23, J24, O33

1. Introduction

The effects of automation technologies on the economy have been the subject of a long debate, dating centuries back and becoming popular again and again (Mokyr *et al.* 2015). Recently, the discussion has received a boost by the influential contributions by Brynjolfsson and McAfee (2011; 2014), who maintained that the pace of technological innovation and the introduction of ever more sophisticated technologies would have quickly disrupted the labor markets as we know them. Two major strands of research have developed hereafter: the first one focusing on the identification of which occupations are and prospectively will be more exposed to the introduction of robots and other forms of innovative technologies (Frey and Osborne 2017; Arntz *et al.* 2017; McKinsey Global Institute 2017; Dengler and Matthes 2018; Nedelkoska and Quintini 2018); the second one regarding the impact that the introduction of automation in the past has had on workers, industries, local labor markets and, in general, the economy (Acemoglu and Autor 2011; Autor 2015; Dauth *et al.* 2017; Graetz and Michaels 2018; Autor and Salomons 2018; Caselli and Manning 2019; Acemoglu and Restrepo 2019; Blanas *et al.* 2020; Acemoglu and Restrepo 2020; de Vries *et al.* 2020; Klenert *et al.* 2020).

This work contributes to this lively debate by bridging these two strands of the literature and by providing some novel evidence on the impact of robot adoption on occupations and workers, by focusing on the overlap between the specific activities carried out by the different types of robots in the market and those characterising the various occupations. Indeed, robots are not identical among themselves, and each one performs only a narrow set of activities called, in jargon, robot applications¹. This is an important aspect to consider because, following Autor *et al.* (2003)'s path-breaking intuition to appreciate the impact of computers, it is what robots actually do that matters for their interaction with different categories of workers. Accordingly, our analysis introduces three main novelties: first, it matches occupations and robots on the basis of the activities characterising them, thereby improving on previous contributions matching robots and workers on the basis of the industry in which they are employed; second, it focuses on the evolution of local employment in activity-based groups of occupations, rather than in broad occupational groups (e.g., skilled/unskilled); third, it makes it possible to distinguish between *robot operators*, that is workers who are involved in the design, installation, maintenance and operation of forms of automation related to robotization (and whose activities cannot be performed by any robot application), from workers who are deemed *exposed to robots' applications* (or simply exposed workers), in that they perform specific activities that can be matched with specific robot applications (e.g., welders and welding robots).

By developing a novel empirical methodology to classify occupations' exposure to robots, we deliver two innovative results. From the methodological viewpoint, we distinguish at a high (five-digit) level of disaggregation the robot operators from the workers employed in exposed occupations, and we build a new measure of local exposure to robots that is based on workers' specific activities and robots' applications rather than only on industries (details in section 2). From the empirical viewpoint,

¹ See Pratt (2015) for a discussion of the relationship between technological developments and the diversification and applicability of robotics.

by exploring the regional variation across Italian local labor market areas (LLMAs) during the period 2011-2018, we assess the differentiated impact of local robots exposure on the employment dynamics for robot operators and for those categories of exposed workers identified in terms of their specific activities. To account for possible endogeneity issues, we control for several potential confounding factors and we adopt a set of instrumental variables isolating the exogenous changes in the supply of robots (as done by Acemoglu and Restrepo 2020, among others).

To preview our main findings, we show, for the first time, that the shares of local employment of robot operators have grown more where the adoption of robots has been more intense. This finding is consistent with the view that where firms invest more in robots, the number of workers who take care of them and who perform activities complementary to them also grows (Autor 2015), a phenomenon that Acemoglu and Restrepo (2019) define as a reinstatement effect. On the contrary, we fail to detect significant changes either in the employment dynamics of the occupations exposed to robots or in aggregate labor market outcomes (in line with Dottori 2020). We also provide some evidence that the failure to detect an average impact of robot adoption on the dynamics of local employment of exposed occupations can hide heterogeneous effects across professions. These findings help to explain the contrasting outcomes of aggregate studies, that find either negative or no correlation between regional/industrial exposure to robots and employment dynamics, and recent microeconomic studies that show a positive relationship between robot adoption and employment at the firm level (Koch *et al.* 2021; Dixon *et al.* 2020)².

It is worth noticing that, for a number of reasons, Italy represents an ideal case study. To start, it is the second European country, after Germany, for robot use (Dottori 2020). Second, from a methodological perspective, we can take advantage of the 2013 Italian National Institute for Public Policies Analysis (Inapp) Survey of Professions (*Indagine Campionaria sulle Professioni*, Inapp ICP hereafter), that represents the Italian equivalent of the American O*NET and provides information on the generalised tasks and the specific activities and duties of over 800 five-digit level occupations in Italy by combining survey-based worker-level information and post-survey validation by experts' focus groups (more on this in section 4). This ensures that our occupation-related variables rely on a survey that is specifically tailored to the Italian occupational structure, avoiding the problems associated with the use of imperfect crosswalk tables with the American O*NET, which reflects the characteristics of the US labor market³. Third, Italian local labor markets are highly differentiated in terms of (industrial and occupation) employment composition (Garibaldi and Taddei 2013) and quite heterogeneous in terms of labor market dynamics over the last decade. This heterogeneity is particularly valuable given that the empirical strategy relies on cross-regional variation to identify the parameters of interest.

² Other microeconomic studies find negative effects on firm employment once the spurious positive correlation induced by demand effects is netted out (Bessen *et al.* 2020; Bonfiglioli *et al.* 2020), while Acemoglu *et al.* (2020a) obtain contrasting results between adopters and non-adopters but with prevailing negative effects due to the latter set of firms.

³ Data collected through the Inapp ICP survey are post-validated by Inapp experts' focus groups. Although our work is the first to match occupations and robots on the basis of Inapp ICP, this survey has been employed to discuss several other aspects of the Italian labor market, such as work-from-home (Bonacini *et al.* 2021) and exposure to infectious diseases (Barbieri *et al.* 2020).

This work relates to three different strands of the literature. The first one includes those studies exploring the consequences of automation and robot adoption on national or local labor market outcomes. As mentioned above, due to the limited availability of data on firm-level investment in robots, most studies have analysed macro and meso-level effects, thereby exploiting cross-country, cross-region, cross-industry or cross-occupation variation in the empirical analysis. By exploiting the categorization of robots in terms of the industry of adoption, in a seminal paper Acemoglu and Restrepo (2020) create a shift-share measure of the local exposure to robot adoption in the US that combines the national penetration of robots into each industry with the local distribution of employment across industries within each US commuting zone. Acemoglu and Restrepo (2020) show that robots and other computer-assisted technologies negatively affected the US local labor markets during the period 1990-2007. Additional analyses using this industry based approach to match robots and workers were subsequently conducted by Chiacchio *et al.* (2018) on 116 NUTS2 regions of six Western European Union countries, Dauth *et al.* (2017) on German local labor markets, Aghion *et al.* (2019) on French employment zones, and Dottori (2020) and Paba *et al.* (2020) on Italian local labor market areas. These studies fail to ascertain important effects of robot adoption on local labor dynamics, with the only exception of Aghion *et al.* (2019), who however focus on artificial intelligence as well as robots⁴.

In addition to the studies exploiting within-country regional variation of employment dynamics and robots exposure, there exist a few industry-level cross-national studies that are worth mentioning. In particular, Graetz and Michaels (2018) use data from International Federation of Robotics (IFR) and EUKLEMS to estimate national robot density in 14 industries and 17 countries from 1993 to 2007, and they conclude that increased robot adoption had a positive effect on labor, total factor productivity and wages, despite a negative impact on the low-skilled workers' labor share⁵.

With respect to the studies adopting an industry-based match between workers and robots, our analysis innovates in two main directions.

First, by introducing a manual match between occupations' activities and robots' applications, this work identifies specifically those occupations that are exposed to robot applications, namely those in which the main activities performed by the worker overlap with existing robots' applications, and those professions (i.e., robot operators) that carry out activities which are complementary to robot adoption. This improves upon previous works analysing the employment dynamics of workers either

⁴ Chiacchio *et al.* (2018) calculate industry-based shift-share measures of regional exposure to industrial robots and find that its change between 1995 and 2007 produced minor negative effects in labor market conditions. According to Dauth *et al.* (2017), robot adoption in Germany did not cause remarkable job losses in the aggregate, even though it altered the sectoral composition of employment away from manufacturing jobs (especially for young workers) towards the service sector. Aghion *et al.* (2019) conclude that robotization over the 1994-2014 period reduced aggregate employment at the level of employment zones in France. Focusing on Italy, instead, Dottori (2020) does not find robust empirical support for any alleged effect of robots use on local employment dynamics.

⁵ Other works reach divergent conclusions. Using EU Labour Force Survey data instead of EU KLEMS data, Klenert *et al.* (2020) find no relationship between the increased use of robots and changes in total employment levels or in the share of low-skill employment between 1995 and 2015 in 14 countries of the European Union. Negative effects for employment are instead found by Compagnucci *et al.* (2019), who focus on 16 OECD countries over the period 2011-2016, and Blanas *et al.* (2020), who investigate the impact of software and robots on the within-industry share of routine workers in 10 high-income countries over the period 1982-2005.

differentiated in terms of education/skill levels or pooled in broad groups of professions that differ in terms of their relative intensity in routine, manual and cognitive tasks. In doing so, we make a contribution in opening the black box (to use Brandes and Wattenhofer 2016's expression) in the classification of occupational groups with regards to automation/robotization. Second, our methodological innovation makes it possible to build a new measure of local exposure to robots that captures better the composition of the labor force in terms of those activities that robots actually perform. The industry-based approach assumes that every worker in every firm in an industry faces the same level of exposure to robots, regardless of his/her actual activity in the company (and that the distribution of robots within an industry is uniform across regions, conditional on the local employment shares). Our novel approach accounts for workers' heterogeneity within industries. Differently from works focusing on industry-based measures, our approach makes it possible to disregard those workers performing services within industrial companies (as in Autor and Dorn, 2013)⁶.

The second strand of the literature this works relates to includes those studies looking at the extent to which automation and computerization substitute workers in certain occupations. The seminal article by Frey and Osborne (2017) calculates the exposure of US professions to the risk of automation and computerization by building on the SOC/O*NET database that codes several hundred job titles and offers a detailed description of their characteristics, generalized tasks and specific activities. A number of scholars have applied Frey and Osborne (2017)'s methodology on countries different from the US, thereby matching the US O*NET occupations with the professions classified according to the Isco08 system (see, for instance, David 2017). Others (such as Arntz *et al.* 2017; Nedelkoska and Quintini 2018; McKinsey Global Institute 2017) have modified Frey and Osborne (2017)'s methodology so as to take into account the various tasks characterising the occupations, rather than a few specific features. Although these works vary considerably in their quantitative findings⁷, they concur in emphasising the importance of looking at occupations' characteristics to assess their exposure to technological advances⁸. Notably, the studies in this strand of the literature aim to identify what occupations are in principle most susceptible to automation and are not directed to produce measures of the local exposure to robots that account for the actual extent of their adoption. Our work borrows from this literature the idea of looking at workers' activities, rather than at their industry of employment, to measure the local exposure to robots. Yet, we do not stop at calculating the shares of professions potentially exposed to robot applications, but we also build a new shift-share measure of local exposure to robotization that takes into account i) the distribution of local employment in

⁶ As we shall explain, if one had access to data capturing the universe of individual employment conditions, one could go even further and build a measure of local exposure to robots that disregards the industrial structure and is based exclusively on activities and tasks, thereby exploiting workers' heterogeneity within and across firms.

⁷ Arntz *et al.* (2017) find that only 9% of US jobs face a risk of automation above 70%, a far lower value (about 38%) than that found by Frey and Osborne (2017). Nedelkoska and Quintini (2018) show that about 14% of jobs in the OECD countries participating in PIAAC have a probability of automation above 70% and about 32 % between 50 and 70%.

⁸ A recent strand of the literature investigates, along this line, the relationship between advancements in Artificial Intelligence and tasks (Felten *et al.* 2018; Fossen and Sorgner 2019).

terms of occupations, ii) a manual match between occupations and robot applications in terms of their activities, and iii) the extent to which robots have been actually adopted in the country.

Our focus on the main activities of individual professions connects this work to a third strand of the literature that looks at skill-biased technical change (SBTC) (Katz and Murphy 1992; Katz and Autor 1999; Acemoglu 2002; Autor *et al.* 2006; Goos and Manning 2007; Goos *et al.* 2009) and routine-biased technical change (RBTC) (Acemoglu and Autor 2011; Autor *et al.* 2003; Autor and Dorn 2013; Autor *et al.* 2015; Goos *et al.* 2014). The works in this strand build on a canonical model that includes different groups of workers performing imperfectly substitutable tasks, and where technological advances affect different workers in different ways⁹. While the SBTC account points to the (mis-)match between technology advances and workers' education, the RBTC approach focuses on the complementarity or substitutability between the occupations' tasks and the features of the technical innovations, in a way that is close in spirit to our work. Autor *et al.* (2015), for instance, distinguish the US occupations in routine and non-routine ones on the basis of their tasks and show that long-term changes in the employment structure of US commuting zones are not significantly associated with the initial specialisation in routine-intensive occupations (suited to computerisation). Despite the common focus on the heterogeneity of tasks across occupations, the studies adopting the RBTC approach differ from ours in a few dimensions. First, these studies typically use the initial employment shares in routine occupations, that are suited to computerization, as the main explanatory variable of employment and wage dynamics, whereas we match occupations and robot applications to build a measure for the extent of robot adoption. Second, RBTC studies emphasize the role of routine and non-routine tasks, whereas we consider the main tasks and activities performed by each occupation, as these latter make it possible to match occupations with specific types of robots. Hence, our analysis mainly focuses on the dynamics of occupations and narrow occupational groups that are directly associated with robots, rather than generally interested by computerization and automation. In sum, our novel matching approach leads to a new measure of local exposure to robots that accounts for robots' and occupations' activities, for the employment composition of the local workforce, and for the actual adoption of robots in the country, thereby bridging all the strands of the literature mentioned above. Finally, two recent articles share some of the features present in our analysis. Adachi *et al.* (2020) exploit a dataset with information on quantities and unit values of robot shipments by industry and application to study the impact of industrial robots on employment in Japan. These features allow them to build an identification strategy leveraging the heterogeneous application of robots across industries (as in the studies mentioned above) and heterogeneous price changes across applications. Their analysis (both at the industrial and commuting zone levels) shows that the decline of robot prices in Japan increased the number of robots as well as employment, suggesting that robots and labor are

⁹ According to the SBTC story, technological change affects negatively (positively) the less (more) skilled workers, thereby exercising a monotonically increasing impact on the occupational structure in terms of skills. The RBTC story, instead, focuses on the occupations' tasks, and hinges on the idea that occupations with routine-intensive tasks are negatively impacted upon by computerisation and digitalization innovations. As routine tasks are typically performed by workers in manufacturing jobs with middle-skill levels, the RBTC hypothesis helps to account for job polarization and between-industry shifts in the structure of employment observed in the US and other advanced economies. For an overview see Fernandez-Macias and Hurley (2016).

gross complementary in production. We share with their work the use of robot applications, albeit with a different identification strategy, and the positive effects of robot adoption on employment. On the other hand, de Vries *et al.* (2020) investigate the relationship between robot adoption and the share of workers performing routine manual jobs across country-industry pairs for 19 industries and 37 countries over the period 2005-2015. They focus on jobs' tasks to distinguish 13 two-digit occupational groups into four task-based broad categories of jobs (i.e., analytic/manual and routine/non-routine), and analyze the differentiated employment outcomes by exploiting the variation across country-industry pairs. A somewhat similar approach has also been adopted by Blanas *et al.* (2020) in their analysis of the impact of various kinds of capital on employment outcomes in European countries. Our work differs from these studies not only for its regional, rather than national, focus. More importantly, instead of focusing on the employment dynamics across broad occupational categories and industries, we determine which occupations are exposed to robots' applications by matching such applications with five-digit occupations on the basis of the main occupation-specific activities and we exploit such information to build a measure of exposure at the level of local labor markets. An advantage of this approach is that it makes it possible to exploit differences in the composition of occupations within each industry, and this reduces the risk of picking up effects associated with unobserved factors affecting industries.

The remainder of the paper is structured as follows. section 2 discusses the construction of our new measure of local exposure to robots based on the functional match between occupations' main activities and robots' applications. Section 3 describes the empirical strategy used to estimate the effects of the local exposure to robot adoption on local labor markets. Section 4 offers details on the construction of our dataset and our data sources. Section 5 presents and discusses the estimation results and introduces, as a robustness check, the estimations that one would have obtained by using a traditional industry-based approach. Section 6 discusses the employment dynamics of alternative groups of occupations that share specific features. Finally, section 7 provides some concluding remarks.

2. A new activity-based measure of local exposure to robot adoption

Due to the scarcity of information about the number of robots installed by individual companies, most applied economic analyses of the impact of automation on labor markets perform cross-regional examinations of the relationship between the local exposure to robots and local employment dynamics.

The local exposure to robots, as explained in the Introduction, is typically calculated on the basis of the combination of the national penetration of robots into each industry, on the one hand, and the local distribution of employment across industries, on the other. The application of this shift-share approach (Bartik 1991) makes it possible for researchers to allocate the national stocks/changes of

robots in an industry across the local areas on the basis of the local concentration of workers employed in that industry¹⁰.

This approach exploits the fact that the IFR classifies robots with respect to the industry in which the companies purchasing them operate. Accordingly, after the seminal work by Acemoglu and Restrepo (2020), the industry-based shift-share approach to build measures of the local exposure to robot adoption has become popular (see, for instance, Chiacchio *et al.* 2018; Dauth *et al.* 2017; Dottori 2020).

While valuable, this industry-based approach presents some potential limitations. First, it prevents from accounting for the high degree of firms' and workers' heterogeneity within an industry, as it implicitly assumes that every worker in an industry is equally exposed to robots, independently from the activities performed. Second, the industry-based approach does not make it possible to differentiate workers who clearly operate on robots from those which are exposed to their adoption. These limitations are at the core of efforts to develop the methodology proposed by Frey and Osborne (2017), who calculate the local share of occupations that are potentially exposed to automation/computerization, in terms of the occupations' title, tasks, activities and characteristics, rather than the industry of employment.

It is worth noticing that the identification of the occupations exposed to robots *à la* (Frey and Osborne 2017) does not suffice to build a measure of the local exposure to robots that accounts for the actual adoption of robots: the local shares of occupations susceptible to robots do not contain information about the robots that have been actually installed over time in the country. Thus, while industry-based shift-share measures of local exposure to robot adoption do not account for robots and workers heterogeneity within industries, the local shares of occupations potentially exposed to robots neglect the actual volume of robots purchased and installed over time. To account for the shortcomings of these alternative approaches, in this work we develop a novel shift-share measure of the local exposure to robot adoption that represents an activity-based alternative to the shift-share measure focusing only on industries. By exploiting the (initial) local composition of the workforce in terms of exposed occupations (whose activities match those performed by the different kinds of robot applications) within industries to allocate the national stocks of robots actually purchased over time, we combine in a single measure both the national penetration of robots with specific applications and the initial local distribution of workers in terms of their actual activities.

Graetz and Michaels (2018) use the IFR robot taxonomy by application to assess whether various groups of professions (at three-digit level) are replaceable by robots (when the professions' titles correspond to at least one of the robot applications) or not. Then, they calculate the national shares of jobs considered as replaceable in each industry and build a variable to instrument the industry variation in robot adoption. While sharing the idea of matching of occupations' activities and robots' applications, our empirical strategy differs remarkably from that proposed by Graetz and Michaels (2018) along several dimensions. First, while they aim to identify the country-industry shares of professions potentially replaceable by robots, we intend to build a shift-share measure of local

¹⁰ This approach mimics the methodology adopted by Card (2001) to redistribute national immigration flows on the basis of the local composition of migrants in terms of the nationality of origin. In the economic literature, proxies of the local exposure to trade competition in each industry have been built along the same line (see Autor *et al.* 2013).

exposure to robot adoption that is able to account for the actual evolution of the national stocks of robots. Moreover, our manual matching procedure to associate occupations' activities with robots' applications is performed at the five-digit level of the occupational taxonomy and is more sophisticated than that used by Graetz and Michaels (2018), as they match robots and jobs at the three-digit level and only on the basis of the jobs' and applications' titles. In appendix C we illustrate a qualitative and quantitative comparison between the classification of occupations based on the two approaches.

In a recent work, Webb (2020) develops a new method to construct a measure of exposure of job tasks to automation based on the overlap between the text of job task descriptions and the text of patents descriptions (i.e., the most frequent verbs and illustrative nouns extracted from robot patents). He calculates an exposure score for each occupation that captures the intensity of patents directed towards the tasks in that occupation. This method recalls both the approach used by Graetz and Michaels (2018) (based on jobs' and robots' titles) and our activity-based methodology to match occupations and robot applications. While both Webb (2020) and we focus on the activities performed by humans and by robots, our approach differs along two dimensions. First, we perform a manual match between occupations and robots according to their functional analogies, that is on the basis of the overlap between workers' activities and robots' applications, whereas he explores the database on US patents and the O*NET database via a language algorithm focusing on jobs' tasks. Second, Webb (2020) calculates an exposure score for each occupation, thereby aggregating the scores from the individual tasks and averaging out both frequent/important and rare/unimportant tasks; instead, we exploit the one-to-one association between occupations and robot applications that descend from our focus on the three main activities characterising each job. Third, Webb (2020) investigates the empirical relationship between the average local exposure score and the employment dynamics within occupation-industry cells, while we study the relationship between the employment dynamics of exposed workers and a measure of the local exposure to robot adoption accounting for the actual evolution in the national stocks of robots.

A few caveats are in order. Our matching approach between occupations' activities and robot applications requires that each occupation is associated only with one application, whereas the same type of robot can be associated with more occupations. As a drawback, this leads to neglect all the minor activities that might be exposed to other robot applications. The assumption underpinning this matching is that firms refrain from purchasing robots for unimportant and infrequently repeated activities in a profession. While this is a demanding assumption, it is worth recalling that, as explained above, any occupational taxonomy and any occupation-robot matching procedure does involve an implicit simplification. The second caveat to consider is that, notwithstanding the wealth of information in the Inapp ICP Survey, a non-negligible dose of judgement is used in matching occupations' activities and robot applications. To the extent that possible errors are not systematic, this should not represent a problem in our estimations that exploit cross-regional differences in the local exposure and in employment dynamics. The adoption of an instrumental variable also contributes to limit the relevance of such potential problem. Being this as it may, in appendix C we provide some comforting evidence showing that the occupations we identify as exposed to robots accord well with those that Graetz and Michaels (2018) and Josten and Lordan (2019), and to a lower extent also Frey and Osborne (2017), find as susceptible of automation.

Before closing this section, it is worth noticing that the novel elements of our measure find an optimal application to the Italian case for two main reasons. The first one regards the availability and quality of the data. The frequency and importance of the various activities in each occupation are assessed on the basis of the Inapp ICP Survey that accounts for the actual structure of the Italian labor markets, the adopted technology and workers' skills, thereby reducing any methodological concerns associated with the use of crosswalk tables linking the classification of Italian occupations with that of US occupations in O*NET. As in Italy there is the only European example of a Dictionary of Occupations that is fully comparable to the US O*NET database at the five-digit Isco08 level (more on this in section 4), this represents an advantage stemming from focusing on Italy¹¹. The second advantage of our dataset is that the description of the occupations in the Inapp ICP survey is not influenced by the most recent developments in the realm of robotics, being it developed at the beginning of the period of interest. Hence, neither the recent changes in the activities associated with the occupations, nor the recent modifications in the array of professions for which new vacancies are posted influence our measure of exposure. Accordingly, as recommended by Graetz and Michaels (2018), the adopted classification is not endogenous with respect to the social and technical innovations introduced in the last decade or so¹².

3. Empirical strategy

3.1 The model

In order to estimate the effects of the local exposure to robot adoption on the Italian local employment dynamics, we estimate different variants of the following empirical model:

$$\Delta y_i = \beta \Delta RX_i + \mathbf{x}'_{i,t0} \boldsymbol{\alpha} + \mathbf{z}'_{i,t0} \boldsymbol{\delta} + \mathbf{p}'_{i,t0} \boldsymbol{\gamma} + \mathbf{r}'_m \boldsymbol{\phi} + \varepsilon_i, \quad (1)$$

where $\Delta y_i = y_{i,t} - y_{i,t0}$ denotes the change in the outcome variable y for local labor market i between $t0$ and t (i.e., from 2011 to 2018), and ΔRX_i indicates the measure of the change in the local exposure to robots, which is based on our matching between workers' and robots' activities, over the period of interest. Besides local employment and unemployment rates, we study the dynamics of local employment dynamics (measured as changes in the share of employment over the working-age population) in either the occupations exposed to robot adoption (because characterised by specific activities that can be performed by specific robot applications) or robot operator jobs. Given the relatively long time span of the analysis, β should be interpreted as the long-term structural effect of robot exposure on labor market outcome y , as in Acemoglu and Restrepo (2020).

¹¹ The German BIBB/IAB dataset also includes job-related information (e.g., job tasks, job skill requirements) (see Spitz-Oener 2006), but only the Inapp ICP survey can be considered equivalent to O*NET.

¹² This is an important feature, as shown by Graetz and Michaels (2018) who use the old (year 2000) Census occupational classifications with a view of avoiding matching occupation titles and robots' applications on the basis of a contemporaneous classification of jobs.

Although the specification in first differences eliminates most of the effects associated with the unobserved time-invariant heterogeneity across local labor markets, to control for further possible confounding factors we include in the estimations a set of lagged demographic controls ($x_{i,t0}$), a set of lagged controls to capture industry-related factors ($z_{i,t0}$), a set of lagged profession-related controls ($p_{i,t0}$), all calculated at the beginning of the period (2011), and four macroarea (North West, North East, Center, South) fixed effects (r_m). The term ε_i indicates the idiosyncratic error term. All the model estimates are weighted on the basis of local labor markets' resident population at the beginning of the period (2011). Weighted regressions are typically adopted in this literature to ensure that sampling issues in small geographical areas do not drive the results. As LLMA are unplanned domains in the design, we shall conduct some robustness checks to show that the results are consistent across samples where very large and very small areas are excluded. We will also show the results of unweighted regressions.

The choice of the economic indicators to include in the specification to deal with possible confounding factors is informed by previous works in the literature. Among the lagged demographic controls, we encompass the logarithm of the resident population, the share of residents older than 65, the share of residents with a tertiary education, the share of female residents and the share of foreign-born residents. Since the size and the age structure of the population are likely to influence (along with other characteristics of the local demography) both the labor market dynamics and the likelihood of robot adoption, the inclusion of control variables of this kind is a common practice in the literature. The lagged industry-related controls include the local exposure to trade competition from China (measured as in Autor *et al.* 2013), the share of employees in the manufacturing sectors, the share of employees in the transportation and logistics industry, and the share of employees that use a personal computer (available only at the regional level). The lagged profession-related controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations calculated at the beginning of the period¹³.

We also conduct some balancing tests to see whether other regional economic indicators are correlated with predicted robots exposure, as in this case they should also be controlled for. We consider the local employment and unemployment rates, calculated at the beginning of the period, as well as the local share of workers employed in the private sector (to control for structural differences in the private and public sector during the austerity period), and the share of local units in non-light manufacturing sectors with at least 50 employees (since Acemoglu *et al.* 2020, find that firms adopting robots tend to be large). The results of the tests, available in table 1, reveal that, conditional on the included controls, the local employment and unemployment rates are not significantly correlated with the predicted change in robots exposure, and need not to be included in the model.

¹³ Autor and Dorn (2013) and Autor *et al.* (2015) distinguish routine and non-routine occupations on the basis of the presence of routine, manual and abstract tasks in each occupation in 1980, and classify as routine occupations those in the top-third of the employment-weighted distribution of the routine task intensity. We follow this approach as well. Autor *et al.* (2015) further elaborate the Autor and Dorn (2013)'s classification of occupations that appear as automatable due to the intensive execution of routine and codifiable tasks. More recently, Josten and Lordan (2019) build on Autor and Dorn (2013)'s classification of automatable works and, by drawing on patent data, identify additional jobs that could be automatable in the next decade.

On the contrary, the local share of workers employed in the private sector and the share of local units in non-light manufacturing sectors with at least 50 employees are significantly correlated with the adoption of robots. Accordingly, we add these variables among the controls in the main specification. Although studying the association of these controls with local labor market outcomes may be of interest *per se*, these terms are included exclusively to mitigate the risks of biased estimates, and therefore we shall not discuss the estimated coefficients.

Table 1. Balancing tests for effects of local robots exposure on employment dynamics

	Employment rate (1)	Unemployment rate (2)	Non-light manufacturing (3)	Private sector (4)
ΔRX	0540 (0.419)	0.105 (0.280)	2.486*** (0.571)	0.147*** (0.066)
Macroarea fixed effects	C	C	C	C
Lagged demographic controls	C	C	C	C
Lagged industry controls	C	C	C	C
Lagged profession controls	C	C	C	C
Observations	377	377	377	377

Notes: The dependent variables are: the local employment rate in 2011 in column (1); the local unemployment rate in 2011 in column (2); the share of non-light manufacturing firms in 2011 in column (3); the share of workers in the private sector in 2011 in column (4). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, and the share of employees that use a personal computer. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

The inclusion of the macroarea fixed effects allows us to control for the large-scale labor market dynamics that characterize the Italian macro regions (i.e., North-West, North-East, Centre, and South and Islands), because of long-lasting structural differences within the country. It is worth noticing that the inclusion of macroarea fixed effects is demanding, as the identification relies only on the variability within such relatively homogeneous areas, but it is fairly common in the literature. Acemoglu and Restrepo (2020) adopt State level fixed effects for the US and Dauth *et al.* (2017) include three macro regions for Germany. In a robustness check we shall report also the results obtained with regional (NUTS2) fixed effects.

As anticipated in section 2, the new measure of the local exposure to robots is the explanatory variable of interest in the empirical model 1 and it can be formally defined as follows:

$$\Delta RX_i = \sum_a \mu_{a,i} \log \left(\frac{R_{a,t}}{L_a} \right) - \sum_a \mu_{a,i} \log \left(\frac{R_{a,t0}}{L_a} \right), \quad (2)$$

where $R_{a,t}$ and $R_{a,t0}$ indicate the national stock of industrial robots in application a respectively at the end and at the beginning of the period, $\mu_{a,i}$ stands for the local share of workers performing

activities mapped to application a in LLMA i in 2011, and L_a represents the total number of workers (at the national level) employed in occupations mapped to robot application a in 2011. This activity-based measure combines the local distribution of occupations and their prevalent activities with the variation in the national stocks of robots performing the various applications.

Census data on the employment conditions of the residents, available for the years 2001 and 2011, do not contain information about their occupations but only about the industry of employment. However, the data on the workforce structure can be retrieved from the continuous labor force survey (*Rilevazione continua sulle forze di lavoro* – RCFL) that is carried out each quarter by the Italian National Institute of Statistics (Istat) at a high level of disaggregation (four-digit level between 2004 and 2009, five-digit level thereafter). Thus, we proceed in three steps. First, from the 2011 Census data we obtain information about the size of the local population aged between 15 and 74, and the number of workers employed in each industry at three digit level. Second, using the national data from the RFCL we calculate the occupational structure of each industry in terms of the various occupations. Third, we calculate the share of workers employed in each industry in those occupations that we match with the different robot applications. Accordingly, the parameter $\mu_{a,i}$ is the local share of workers employed in all the occupations that are matched with robot application a . It is calculated as:

$$\mu_{a,i} = \sum_s \sum_o I_{oa} \theta_{os} \mu_{s,i}, \quad (3)$$

where I_{oa} is an indicator function taking value 1 if occupation o is exposed to robot application a (and zero otherwise), θ_{os} is the share of workers employed in occupation o in the industry s (calculated at the national level based on RFCL data), and $\mu_{s,i}$ is the ratio of workers employed in industry s over the residents (aged between 15 and 74) in local labour market area i (calculated based on 2011 Census data)¹⁴.

3.2 Addressing endogeneity concerns

In the presence of unobservable and unevenly distributed shocks that simultaneously influence local labor market dynamics and the local exposure to robots, OLS estimates of the coefficient β would be biased and inconsistent, as they would not identify the component of robot exposure driven by changes in technology. For example, firms may adopt robots in response to changes in labor demand at the LLMA level, such as a recession in an area with a large exposure to robots, or in response to other changes that also impact on labor demand, such as an increase in workers' bargaining power (Acemoglu and Restrepo 2020).

Accordingly, to capture the impact of the exogenous component of robot adoption on local employment dynamics we resort to an instrumental variable (IV) approach. More precisely, we build a shift-share IV *à la* Bartik that combines historical (i.e., predetermined) local employment shares and the evolution of robot adoption abroad to proxy for the increase of robots in Italy. This strategy, which

¹⁴ If Census data had contained information on the local occupational structure, the parameter $\mu_{a,i}$ could have been calculated directly from the micro-level data. As this is not possible, we combined the Census data on the sectoral structure of local employment with the occupational composition of the industries that the RFCL survey allows us to derive.

has become standard in the literature (e.g., Dauth *et al.* 2017; Caselli *et al.* 2020; Acemoglu and Restrepo 2020), is an effective way to capture only the part of robots exposure that is exogenously determined by advancements in robot technology. As the evolution of the stocks of robots in the European countries may be determined by common demand factors and by regional value chains in the manufacturing sectors, we consider the process of robot adoption in Japan, South Korea and the United States (JPKRUS, hereafter). These are all developed economies whose GDP is only partially correlated with the Italian one and that started adopting robots much earlier than 2011, and are not involved into a rapid catching-up process as the one observed in China. We consider all extra-European countries in a robustness check.

In building this instrument, we allocate the changes in the foreign stocks of robots by application on the basis of the local employment shares in occupations exposed to such applications. We mimic the methodology adopted to construct the explanatory variable ΔRX_i , thus our IV reads as follows:

$$\Delta \widetilde{RX}_i = \sum_a \tilde{\mu}_{a,i} \log \left(\frac{\tilde{R}_{a,t}^{JPKRUS}}{L_a^{2001}} \right) - \sum_a \tilde{\mu}_{a,i} \log \left(\frac{\tilde{R}_{a,t0}^{JPKRUS}}{L_a^{2001}} \right), \quad (4)$$

where L_a^{2001} denotes the aggregate number of workers performing activities matched to robot application a in 2001, $\tilde{\mu}_{a,i}$ is the share of workers in occupations matched to robot application a in local labor market i in 2001, $\tilde{R}_{a,t}^{JPKRUS}$ and $\tilde{R}_{a,t0}^{JPKRUS}$ are the stocks of industrial robots performing application a in JPKRUS, respectively at time t (2018) and $t0$ (2011). To prevent endogeneity problems, the employment shares are lagged ten years prior to the beginning of the period of interest, and the data from the 2001 Census are used to derive the sectoral structure of employment at the local level¹⁵.

This shift-share IV has never been adopted in previous works because it is on our new approach to perform a match between occupations and robots on the basis of the functional overlap between workers' activities and robots' applications. Conversely, previous studies resorted to shift-share industry-based IV exploiting changes in the foreign stocks of robots by industry, together with the local shares of workers employed in every industry. We believe that our set up is preferable as it offers internal consistency: we aim to explain the employment dynamics in the professions exposed to robot applications by exploiting the cross-regional variation in the local exposure to robots, which is calculated on the basis of a match between occupations and robots based on their activities, and by instrumenting such variable with a shift-share IV using lagged shares of occupations and the foreign stocks of robots by application.

¹⁵ To convert these values into local shares of occupations, information on the occupational structure of the different industries is necessary. As representative data about occupations at the four-digit level of aggregation became available only in 2004, we combine the 2001 Census data on sectoral employment shares with the 2004 data on the occupational structure of the industries from the RCFL, using the same methodology employed for equation (3).

4. Data

Our study is based on a panel that, for the years 2011 and 2018, integrates detailed information on the workforce occupational structure and on the demographic composition of 377 Italian local labor markets¹⁶, whose boundaries were identified by the Italian National Institute of Statistics in 2011 on the basis of workers' commuting patterns¹⁷. As anticipated, local exposure to robot adoption is calculated by combining information about labour markets together with data on shipments and operational stocks of industrial robots provided by the IFR. In particular, our measure of local exposure to robot, ΔRX_i , as well as all the variables capturing the evolution of employment at the LLMA level, are constructed by using data on the workforce structure retrieved from the Census of Population and the Census of Industry and Services, organized every ten years by Istat, and from the continuous labor force survey, RCFL, carried out each quarter by Istat.

Other demographic and industry-related variables included as controls in our empirical specifications are retrieved from the above-mentioned censuses by Istat. A more detailed description of these data and of the mapping between occupations' activities and robot applications is provided in what follows. The data on industrial robots, defined by the International Organization for Standardization as "automatically controlled, reprogrammable, multipurpose manipulator[s] programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications", are derived from the IFR database, which contains the annual shipments and stocks of industrial robots for more than 50 countries. Robots are classified according to either their application at the three digit level (e.g., welding, painting, palletizing) or the economic activity of the purchasing companies at the two-digit level (based on ISIC rev. 4).

Statistics on the LLMA occupational structure have been computed on the basis of the data retrieved from the Census of Population and the Census of Industry and Services and from the Istat's RCFL, a quarterly survey which represents the most important source of official information on Italian labor markets. Covering about 1,250 Italian municipalities, the survey involves approximately 77,000 households and 185,000 individuals each quarter¹⁸. The survey gathers information on the employment status of the respondents, as well as on other socio-demographic traits, such as age, gender, nationality, and municipality of residence. For the respondents who are employed, the survey records extremely detailed information on their occupation, sector of employment, type and duration of the contract, number of hours worked and salary.

¹⁶ Local labor markets represent functional geographic areas that go beyond administrative boundaries and correspond to economically integrated spatial units, suitable to investigate various aspects of geographical heterogeneity (see, for instance, de Blasio and Poy, 2017).

¹⁷ The sample of LLMA covered by the RCFL survey, which we use to derive the local shares of robot operators and workers exposed to robots, counts 412 areas, but we trim the sample so as to remove the outliers (top and bottom one percent of the distribution) in each one of the dependent and independent variables. Since the sample is correctly stratified, the absence of some LLMA does not cause problems in terms of its representativeness, although some minor LLMA are disregarded.

¹⁸ As the RCFL is run on a quarterly basis, we have computed annual data by merging the quarters. For the year 2018, however, we had to combine the last two quarters of that year with the first two of the following year. This is because the data relative to the first and second quarters of 2018 only provide information on occupations at a four-digit level and not at the five-digit level as found in the subsequent waves.

Occupations in RFCL and in this work are codified according to the 2011 Inapp - Istat Classification of Occupations (Classificazione delle Professioni – CP2011 hereafter), which is the Italian version of the current ILO International Standard Classification of Occupations (Isco08). At its finest level of aggregation (i.e., five digit level), the CP2011 classification contemplates about 800 occupational units (OU, hereafter) – referred to as “unità professionali” –, that account for as many as 6000 job titles¹⁹. The advantage of resorting to the RCFL survey is twofold. First, it allows to match individual workers’ occupations with detailed information on occupations’ characteristics, as provided by the 2013 Inapp Survey on Workers and Occupations (Indagine Campionaria sulle Professioni, Inapp ICP hereafter). In both surveys, the occupations are codified according to the highly disaggregated five-digit level based on the CP2011 database. The second advantage is that the RCFL provides sample weights, which we use to extend the results of the analysis to the overall population, as done in the literature.

As discussed in section 2, our new measure of local exposure to robot adoption is based on a match between robots’ applications and occupations’ specific activities performed by using the information contained in the Inapp ICP survey. By following the design of the American O*NET, but tailoring it to the Italian context, the survey documents the generalized tasks and the specific activities and duties of the workers belonging to each different OU. In particular, it reports quantitative measures (varying on a scale from 1 to 5) of the frequency and the importance of various specific activities associated with each OU at five-digit level of the taxonomy²⁰. Thanks to all this information, we manually match each IFR robot application with one or more occupations on the basis of the occupation description, its three primary occupation-specific activities according to the Inapp ICP scores and, where necessary, additional information provided in the survey. By relating robots and occupations on the basis of functional analogies and complementarities, we determine whether an occupation is “exposed to a robot application”, that is if its main tasks and activities can be associated with a specific industrial robot application, or not. In table B1 in appendix B we report detailed examples of matches, if any, between occupations and robot applications.

Using information about occupations’ description and specific activities, we also identify those jobs where the workers can be considered as “robot operators” because they are involved in the design, installation, maintenance and operation of forms of automation related to robotization. None of such activities can be carried out by a robot, and thus it is not possible to associate occupations and robot applications on the basis of the activities performed by robots and workers. These occupations deserve to be singled out as they are expected to grow unequivocally with the increase in robot adoption.

¹⁹ The CP2011 OUs include for instance “Welders and flame cutters” (code 6.2.1.2.0), and this group contains several job titles (namely, Lead Welder; Oxyacetylene Welder; Plumber Welder; Autogenous Cutter; Oxyacetylene Cutter; Flame Fret-Sawing Worker; Fusion Welder; Water Gas Welder; Braze Welding Welder; Tin Welder; Atomic Hydrogen Welder; Aluminium-Thermal Welder; Autogenous Welder; Special Metal and Metal Alloy Welder).

²⁰ For example, consider the OU “Shoe manufacturers” (code 6.5.4.2.3). According to the frequency and importance indices, it includes both rare and unimportant specific activities (such as performing springing of the uppers, with a value of 1.2 for importance and 0 for frequency) and specific activities that are frequent and important (such as packaging or wrapping finished products, with a value of 3.4 for importance and 2.9 for frequency).

Overall, we identify 123 OUs that we categorize as exposed to different robot applications, 8 as robot operators, and 669 OUs neither exposed to robot applications nor operating on robots.

Table 2 shows some further summary statistics for these variables along with the unemployment rate and robots exposure for the set of LLMA in our sample. These descriptive statistics reveal a great deal of heterogeneity in the evolution of local employment across LLMA, as anticipated in the Introduction. We point out that, on average, the share of robot operators over the working-age population grows relative more than the share of exposed workers, but both are positive. While these rates of change may appear as small, it is worth noticing that the local relative weight of robot operators in the working-age population increases on average by 50% in only 7 years.

Table 2. Descriptive statistics

	Obs.	Mean	Std. Dev.	Min	Max
Δ Unemployment rate	377	0.019	0.040	-0.088	0.167
Δ Employment rate	377	0.017	0.033	-0.120	0.173
Δ Share of robot operators over population	377	0.004	0.007	-0.023	0.055
Δ Share of exposed workers over population	377	0.001	0.016	-0.066	0.079
Δ Robots exposure (applications)	377	0.016	0.009	-0.025	0.113

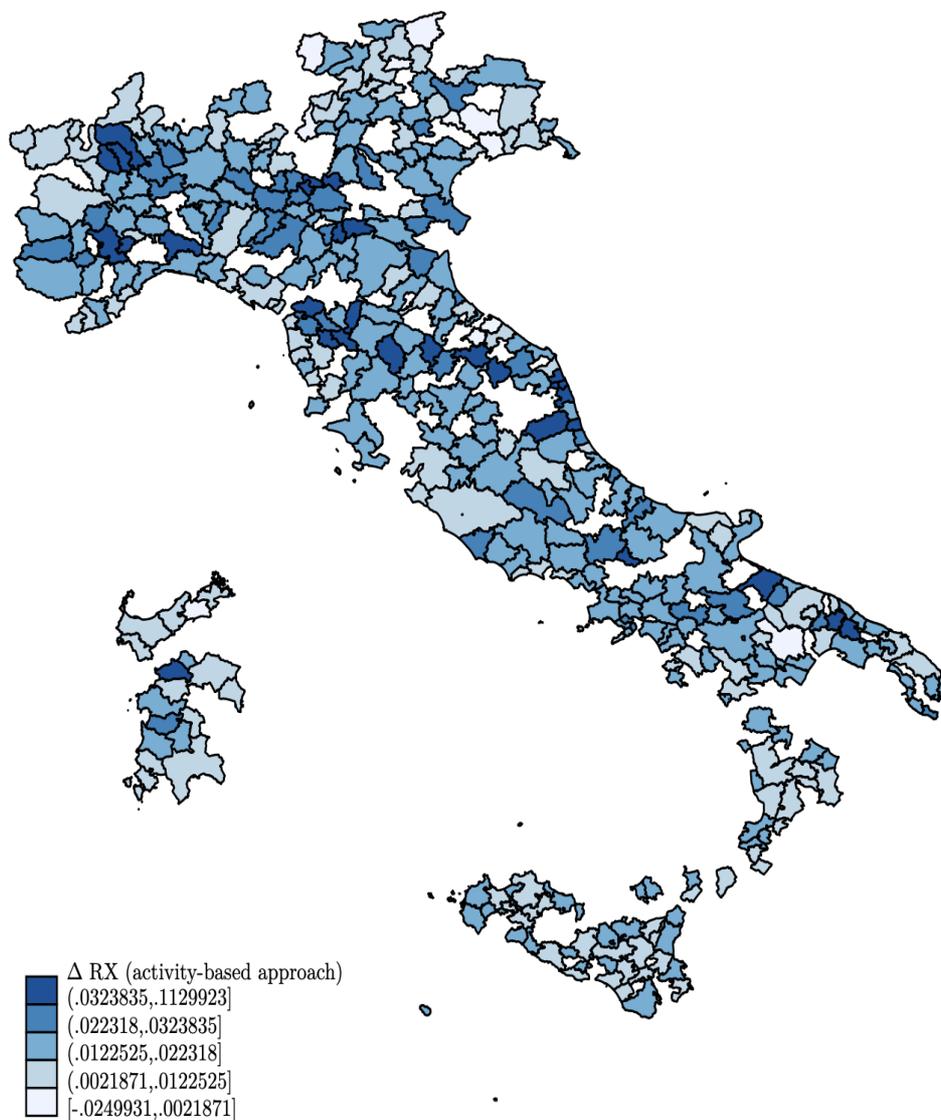
Notes: All the variable are measured in differences between 2018 and 2011 at the LLMA level and weighted on the basis of the resident population in 2011. The employment to population ratios for operators and exposed workers are computed with respect to working-age residents (15-74 years old).

Figure 1 shows a map of Italian LLMA and changes in robots exposure across them. Figure 2 show that there seems to be a positive relation between changes in robot adoption by application and changes in employment in jobs exposed to robotization. This seems to suggest that indeed we should be careful in interpreting robot adoption as job displacing, even for those jobs that include activities that in principle could be fully replaced by robot applications, and that its effects may be highly differentiated across occupations and workers. Finally, figure 3 shows that our IV based on robot adoption in Japan, South Korea and the US seems to be informative for the part related to the change in the stocks of robots in Italy.

Furthermore, thanks to the detailed information provided in the Inapp ICP Survey, we can identify other groups of occupations in terms of their characteristics, rather than using our manual categorization into robot operators, exposed occupations, and non-exposed jobs. For instance, following the literature, we can study whether the change in local exposure to robot adoption is associated with a change in the share of routine intensive occupations. Differently from other works, we test for, rather than assume, a positive relationship between higher levels of routinization at work and exposure to robotization. Accordingly, we empirically assess whether changes in the local shares of routine-intensive occupations are associated or not with the evolution of local robots exposure calculated using our manual match between robots' applications and occupations' activities. Notably, we do not contend that a number of routine jobs can be exposed to robots and, as discussed at length in section 6, our taxonomy of exposed occupations sensibly correlates with that based on routine-task intensity. However, as explained by Autor *et al.* (2015), routine-task intensity has a dual source in blue-

collar production occupations and in white-collar office and clerical activities: while both are exposed to automation, only the former are potentially exposed to robots²¹.

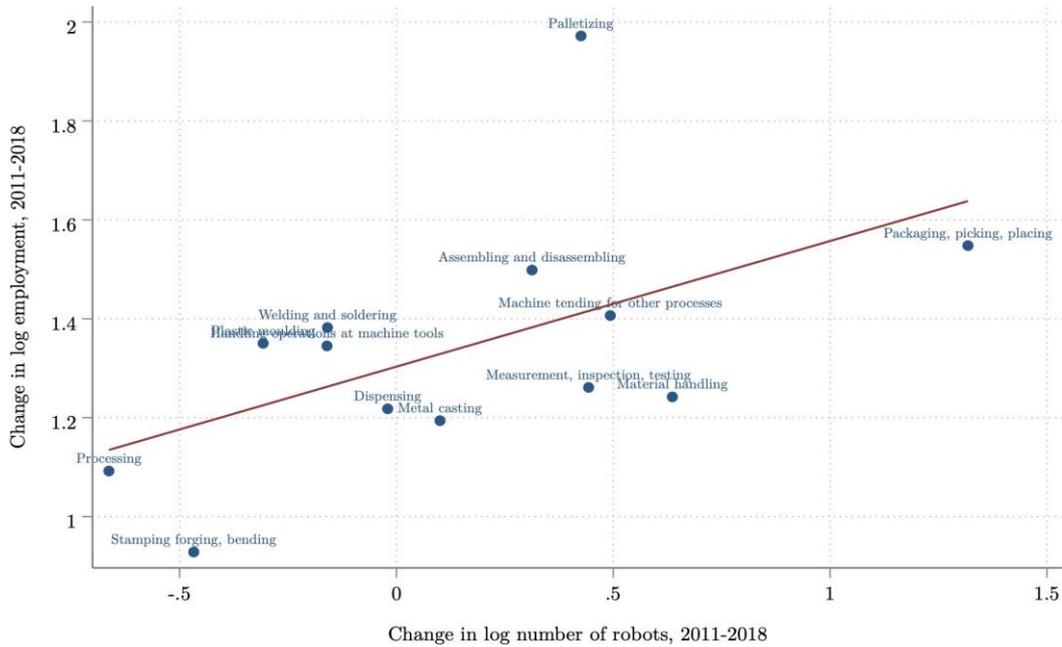
Figure 1. Changes in local robots exposure across Italian LLMA, 2011-2018



Notes: The figure reports the change in local robots exposure in 377 Italian LLMA between 2011 and 2018. Robots exposure is measured following the activity-based approach to match occupations and robot applications.

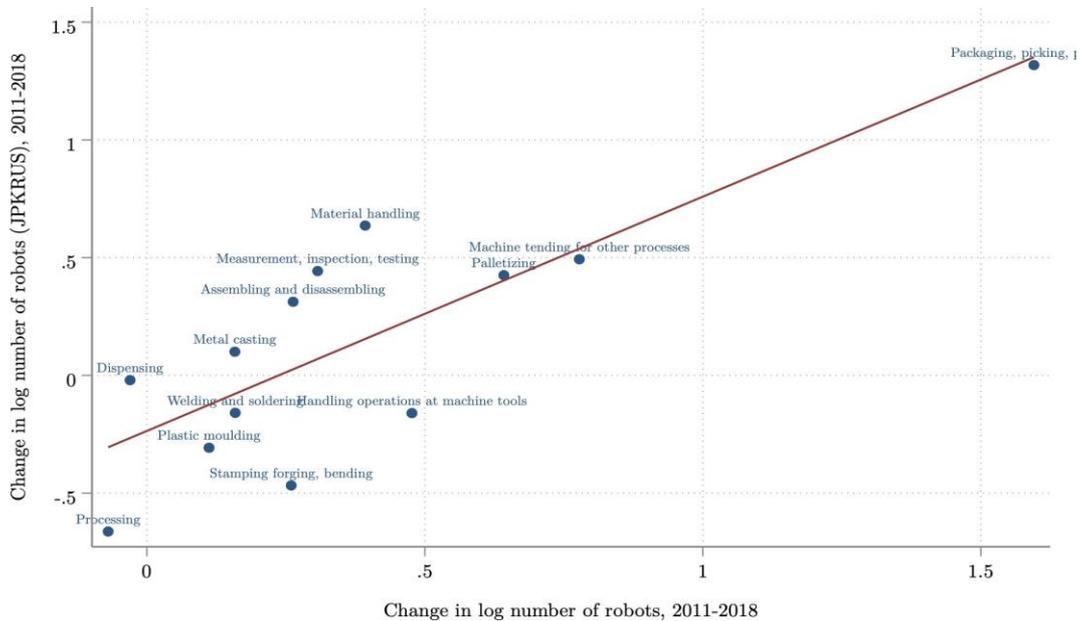
²¹ In all estimations we control for the initial local shares of routine occupations. Together with the other aspects of our empirical strategy, this control variable strengthens the identification as it reduces the chances of picking up unobserved occupation-specific factors associated with the local average diffusion of routine-intensive occupations, rather than robot adoption. Notably, the intensity of routine tasks is assessed at the beginning of our time period, being the rank of tasks stable in the short time (Akcomak *et al.* 2016).

Figure 2. Changes in employment vs changes in robot adoption by application, 2011-2018



Notes: The figure reports changes in log employment in 13 robot applications in Italy between 2011 and 2018 on the vertical axis and changes in the log number of robots in Italy between 2011 and 2018.

Figure 3. Changes in robot adoption in Italy vs IV countries (JPKRUS) by application, 2011-2018



Notes: The figure reports changes in the log number of robots in Italy and in the set of countries used in the main IV (Japan, South Korea and the US) between 2011 and 2018.

In addition to testing groupings of occupations based on the intensity of routine tasks, we also explore the questions in the Inapp ICP survey to study further occupational features and build alternative groups of occupations. This will be discussed in section 6.

5. Main results

To assess the relationship between the change in local exposure to robot adoption, calculated on the basis of our novel approach, and local employment dynamics, we estimate equation 1 by focusing on the evolution of the local employment rate, the local unemployment rate, the local share of robot operators over the working-age population, and the local share of workers employed in exposed occupations over the working-age population. As can be seen in table 3, all the OLS estimates for the variable of interest ($\Delta RX_{i,t}$) are not statistically different from zero. As discussed, a number of possible endogeneity problems plagues the OLS estimation. To address them, we follow the literature and instrument the explanatory variable with a relevant and valid instrument with a view to exploiting only the exogenous component of the growth in adopted robots in Italy, driven by technological advances common to all countries, and not by local demand factors. In this way, we reduce the likelihood that demand shocks and other idiosyncratic (area-specific) factors may bias the estimates. For internal consistency, as explained in section 3.2, we resort to the shift-share IV $\Delta \widetilde{RX}$ that combines the variation in the stocks of robots in Japan, South Korea and the US, application-by-application, between 2011 and 2018, and the 2001 occupational shares of workers exposed to robots in each LLMA.

Table 3. Effects of local robots exposure on employment dynamics, OLS

	Δ employment (1)	Δ unemployment (2)	Δ operators (3)	Δ exposed (4)
ΔRX	-0.125 (0.226)	0.201 (0.238)	0.074 (0.049)	0.029 (0.136)
Macroarea fixed effects	C	C	C	C
Lagged demographic controls	C	C	C	C
Lagged industry controls	C	C	C	C
Lagged profession controls	C	C	C	C
Observations	377	377	377	377
R-squared	0.219	0.273	0.197	0.162

Notes: The dependent variables are the difference between 2018 and 2011 in the employment rate in column (1), in the unemployment rate in column (2), in the share of robot operators over working-age population in column (3), and in the share of exposed workers over working-age population in column (4). All regressions are estimated via OLS. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results of our IV estimates are reported in table 4 and table 5. Columns (1)-(4) in table 4 report the estimates for the variation in the local employment and unemployment rates, while columns (1)-(4) in table 5 reproduce the estimated coefficients of $\Delta RX_{i,t}$ for the shares of local employment over working-age population for robot operators and exposed workers. For each dependent variable of interest, we show both a minimal specification (i.e., we only include among the controls the

explanatory variable of interest, $\Delta RX_{i,t}$, and the macro-area dummy variables) and the specification with all the controls described in section 3.1. Independently from the specification, the results from the first-stage estimations (see column (1) of table A1 in appendix A for the specification with the full set of controls) are satisfactory: the Kleibergen-Paap F statistics are larger than the critical values.

Table 4. Effects of local robots exposure on employment dynamics II

	Δ employment		Δ unemployment	
	(1)	(2)	(3)	(4)
ΔRX	-0.203 (0.312)	0.165 (0.402)	-0.950*** (0.348)	-0.292 (0.398)
Macroarea fixed effects	C	C	C	C
Lagged demographic controls		C		C
Lagged industry controls		C		C
Lagged profession controls		C		C
Observations	377	377	377	377
Kleibergen-Paap F stat	149	130.1	149	130.1

Notes: The dependent variable is the difference in employment rates between 2018 and 2011 in columns (1) and (2), and the difference in unemployment rates between 2018 and 2011 in columns (3) and (4). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

As in other studies, the estimated parameters for the aggregate employment conditions are not significant: this implies that the local exposure to robot adoption is not correlated with variations in the local employment and unemployment rates. This result is in line with the consensus view that, in most countries, robot adoption has so far not been associated with important net effects at the aggregate level.

Regardless of whether the controls are included in the regressions or not, the estimated coefficient of $\Delta RX_{i,t}$ is positive and significant for what concerns the change in robot operators. This finding is in line with expectations, and it is a new result in this strand of the literature. Previous studies have shown a positive effect of robot adoption on skilled workers employed in jobs intensive in cognitive tasks (which can be indirectly associated with automation and computerization), yet our categorization is different. We specifically focus on those jobs that are explicitly associated with activities complementary to robot installation, maintenance and use. Hence, this represents novel empirical evidence in favour of the reinstatement effect discussed by Acemoglu and Restrepo (2019). According to the estimates, an increase by 1% in robot adoption leads to an increase by 0.27 percentage points in the local share of robot operators. Thus, the overall increase in robot operators in Italy between 2011 and 2018 can be entirely explained by the increase in robot adoption that occurred during the same period.

Table 5. Effects of local robots exposure on employment dynamics I

	Δ operators		Δ exposed workers	
	(1)	(2)	(3)	(4)
Δ <i>RX</i>	0.175*** (0.062)	0.293*** (0.095)	0.205 (0.146)	0.301 (0.194)
Macroarea fixed effects	C	C	C	C
Lagged demographic controls		C		C
Lagged industry controls		C		C
Lagged profession controls		C		C
Observations	377	377	377	377
Kleibergen-Paap F stat	149.0	130.1	149.0	130.1

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in columns (1) and (2), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in columns (3) and (4). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

No statistically significant effect of robots exposure on the share of workers employed in exposed occupations is found (columns (3) and (4) in table 5). This result can be explained in various ways. The first reason is that the displacement effect of robots on workers has been very limited in Italy during this period, as also suggested by Dottori (2020) based on an earlier period and an industry-based specification of robots exposure. An alternative explanation refers to the fact that each occupation consists of very many activities and only few of them can be performed by robots: if a robot application can substitute a worker only in a limited range of his/her activities, then that worker is not displaced, but simply forced to redirect his/her efforts on the remaining activities. A third reason why the estimated parameters may be insignificant is that they reflect the fact that contrasting effects across occupations may average out. Besides robots displacing workers, there might be occupations that benefit from the introduction of robots as these professions are and remain relatively intensive in activities where humans have a comparative advantage (Acemoglu and Restrepo 2019; Caselli and Manning 2019). Clearly, it is also possible that our taxonomy of exposed occupations is imperfect, and the estimates suffer from the inclusion and exclusion of certain jobs; while possible, we show in the appendix that our taxonomy accords well with other classifications of professions exposed to computerization, automation and robotization. Finally, it is possible that the displacement effect is counteracted by an increase in the scale of production of firms adopting robots. This would be in line with previous studies on the firm-level determinants of robot adoption, as they concur in discarding labor cost-saving objectives and in emphasizing firms' attempts at standardizing quality, increasing control over the production process, introducing new sophisticated products and reach new destination markets (Aghion *et al.* 2020; Backer *et al.* 2018). This is likely to be true in Italy, where firms cannot undertake collective dismissals to save labor unless they face financial difficulties, a

condition that is at odds with large investments in robots. However, our IV strategy is aimed at minimizing this last possibility.

As anticipated in the Introduction, thus, our findings help to interpret the inconsistent results reached by the studies conducted with industrial/regional data and the analyses using firm-level data: the former find either negative or nil effects of robots exposure on local employment dynamics, whereas the latter suggest that firms investing in robots also expand the number of employees. All in all, opposite effects may indeed coexist, and their net effect may be limited. The taxonomy of occupations that one adopts to estimate the impact of robots exposure on employment influences the results, as it represent a net average effect of a highly differentiated impact. In turn, this suggests to explore alternative occupational categories in the attempt at uncovering the underlying mechanisms associating robots exposure with labour market dynamics, as done in section 6.

5.1 Robustness checks

In this section we carry out a battery of auxiliary regressions to show that our results are robust across specifications and sample compositions. The first set of robustness checks concern the sample composition. To show that the results are driven neither by the few very large LLMA, nor by the multitude of small ones, we build two subsamples: the first one excludes the four cities with more than 1 million residents (i.e., Milan, Rome, Naples and Turin), whereas the second drops the LLMA in the bottom 20% of the distribution in terms of population (namely, those with less than 33363 residents). As can be seen in table 6, the main results remain valid with these subsamples, and the first stage of the 2SLS estimation is not affected negatively.

Table 6. Effects of local robots exposure on employment dynamics, robustness I: Sub-samples

	Δ operators		Δ exposed workers	
	no big (1)	no small (2)	no big (3)	no small (4)
ΔRX	0.302*** (0.093)	0.266*** (0.096)	0.299 (0.202)	0.236 (0.242)
Macroarea fixed effects	C	C	C	C
Lagged demographic controls		C		C
Lagged industry controls		C		C
Lagged profession controls		C		C
Observations	373	301	373	301
Kleibergen-Paap F stat	130.9	134.5	130.9	134.5

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in columns (1) and (2), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in columns (3) and (4). The shares are computed with respect to the working-age resident population (15-74 years old). The subsample “no big” excludes the four LLMA with more than 1 million residents, columns (1) and (3); the subsample “no small” excludes the 20% smallest LLMA, columns (2) and (4). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

As a second robustness check, we show the results obtained by running unweighted regressions. The estimates in table 7 reveal that our main results are confirmed, and thus are not driven by the largest LLMA. In fact, these unweighted estimations would even suggest a positive and significant impact of robots exposure on exposed professions as well. Although stimulating, we do not dwell much on these findings as unweighted regressions are not warranted given our use of data obtained from a representative survey.

Table 7. Effects of local robots exposure on employment dynamics, robustness II: Unweighted

	Δ operators (1)	Δ exposed workers (2)
ΔRX	0.380*** (0.134)	0.555** (0.253)
Macroarea fixed effects	C	C
Regional fixed effects	C	C
Lagged demographic controls	C	C
Lagged industry controls	C	C
Lagged profession controls	C	C
Observations	377	377
Kleibergen-Paap F stat	57.09	57.09

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in column (1), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in column (2). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

The third set of robustness checks is concerned with our assumption about the error term. In the first check, we assume that the error term is correlated across observations within provinces and, thus, we cluster standard errors at the (NUTS3) provincial level. Second, we substitute macroarea fixed effects with (NUTS2) regional fixed effects. This specification is particularly demanding as the identification of the parameter of interest hinges on the variation in labour dynamics and robots exposure across few LLMA within the same region. Table 8 shows that the results remain qualitative and quantitatively similar to the main results. The first-stage regressions of the 2SLS approach show that the IV remains informative.

The fourth robustness check concerns the IV variable. So far, we have focused on robot developments in Japan, South Korea and the United States to construct our informative and valid IV. This group of countries should in principle minimize the risk of using an invalid IV that co-varies with fluctuations in aggregate demand in Italy. It is worth noticing that, differently from those studies using industry-based shift-share measures of robots exposure, our measure is based on the change in the stocks of robot applications: this reduces even further the possibility of inadvertently picking-up simultaneous trends in employment and robot adoption across industries. Yet, as a robustness check, we repeat the main estimations with an alternative IV that considers the variation in the stocks of robot applications in all extra-European countries.

The results, reported in table 9, confirm qualitatively our main findings. In this case, however, also the estimated parameter for exposed workers appears positive and significantly different from zero. It should be noted, however, that the Kleibergen-Paap F statistic has a much lower value, close to the Stock-Yogo critical values, and this suggests to take the results based on this alternative IV with caution.

Table 8. Effects of local robots exposure on employment dynamics, robustness III: Error terms

	Δ operators		Δ exposed workers	
	(1)	(2)	(3)	(4)
ΔRX	0.292*** (0.093)	0.244*** (0.096)	0.301 (0.202)	0.243 (0.242)
Macroarea fixed effects	C		C	
Regional fixed effects		C		C
Lagged demographic controls	C	C	C	C
Lagged industry controls	C	C	C	C
Lagged profession controls	C	C	C	C
Observations	377	377	377	377
Kleibergen-Paap F stat	113.1	141.1	113.1	141.1

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in columns (1) and (2), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in columns (3) and (4). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Standard errors in parentheses are clustered at the province level in columns (1) and (3), while robust standard errors are provided in columns (2) and (4). *** p < 0.01, ** p < 0.05, * p < 0.1.

The IFR accounting methodology to calculate the stock of robots assumes a full active service life of 12 years for each robot and its immediate withdrawal from service afterwards. In previous studies, some scholars (see, among others, Graetz and Michaels 2018) recalculated the evolution of the stocks of robots by means of a perpetual inventory method. This approach avoids large discontinuities in the robot stock series and is compatible with a smoother usage of robots over the large business cycle fluctuations observed in Italy in the period of interest. Hence, we reconstruct the national stocks of robots for Italy, Japan, South Korea and the US by using the perpetual inventory method with an annual depreciation rate of 10%, annual investment in robots reported by IFR, and the initial value of the stock based on 2006²². Hence, our fifth robustness check uses this alternative method for calculating the stocks of robots, in terms of robots applications, both for the measure of local exposure and the IV. Table 10 reports the estimates obtained using these alternative measures of robot stocks. The results are, once again, qualitatively and

²² We choose 2006 for the initial value for two reasons. First, before this year, a substantial portion of robots were assigned by IFR as “unspecified” applications or sectors. Second, in an effort to improve the harmonization of cross-national data, between 2004-2005 the IFR reviewed considerably the way in which information was collected.

quantitatively similar to those obtained in our main specification. Notably, the first stage appears to be very robust, with the Kleibergen-Paap F statistic approximately equal to 300.

Table 9. Effects of local robots exposure on employment dynamics, robustness IV: Alternative IV

	Δ operators (1)	Δ exposed workers (2)
ΔRX	0.547*** (0.192)	0.664** (0.334)
Macroarea fixed effects	C	C
Lagged demographic controls	C	C
Lagged industry controls	C	C
Lagged profession controls	C	C
Observations	377	377
Kleibergen-Paap F stat	21.01	21.01

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in column (1), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in column (2). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in extra-European countries. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 10. Effects of local robots exposure on employment dynamics, robustness V: PIM

	Δ operators (1)	Δ exposed workers (2)
ΔRX	0.291*** (0.098)	0.304 (0.195)
Macroarea fixed effects	C	C
Lagged demographic controls	C	C
Lagged industry controls	C	C
Lagged profession controls	C	C
Observations	377	377
Kleibergen-Paap F stat	293.6	293.6

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in column (1), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in column (2). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and the growth of robots in Japan, South Korea and the US. Stocks of robot are calculated with the perpetual inventory method and a 10% depreciation rate. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The next robustness check concerns the existence of pre-trends that could affect the results and their interpretation. If the adoption of robots in the period of interest (2011-2018) is statistically correlated, conditional on the controls variables used in the specification, with the employment dynamics in the previous period, one cannot exclude that our main regressions pick up pre-existing local trends rather than contemporaneous effects. As can be shown in table 11, this is not the case. Conditional on the control variables included in our specification, future robots exposure does not seem to be correlated with past growth in the employment shares of the occupation categories of interest.

Table 11. Effects of local robots exposure on employment dynamics, pre-trends

	Lagged Δ operators (1)	Lagged Δ exposed workers (2)
ΔRX	-0.110 (0.106)	-0.315 (0.192)
Macroarea fixed effects	C	C
Lagged demographic controls	C	C
Lagged industry controls	C	C
Lagged profession controls	C	C
Observations	360	360
Kleibergen-Paap F stat	124.8	124.8

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2011 and 2004 in column (1), and the difference in the share of exposed workers over working-age population between 2011 and 2004 in column (2). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Our shift-share IV can be seen as the inner product of the (initial) local shares of workers in occupations matched to a given robot application (e.g., $\tilde{\mu}_{a,i}$ for application a) and the exogenous shocks due to changes in the stocks of each robot application in Japan, South Korea and the US (e.g., $g_a = \tilde{R}_{a,t}^{JPKRUS}$). This IV is consistent with a research design based on pooling the differentiated local exposure to the exogenous shocks across the thirteen robot applications. As shown by Goldsmith-Pinkham *et al.* (2020), our two-stage least squares estimation is equivalent to a generalized method of moments estimator using the local shares of workers exposed to robot applications as instruments, and a specific weight matrix constructed using the robot application shocks. A shift-share IV estimator can then be seen as a weighted combination of just-identified estimations, each using the local shares of workers exposed to a given robot application as a separate instrument. Goldsmith-Pinkham *et al.* (2020) also show that the shift-share IV estimator can be decomposed so as to derive the underlying Rotemberg weights, which can be interpreted as the sensitivity-to-misspecification elasticities associated with each instrument. All this implies that a sufficient condition for the validity of the shift-share IV estimator is the validity of the local shares $\tilde{\mu}$ as instruments, conditional on the controls in

the specification, in particular the shares associated to those applications with the largest Rotemberg weights²³.

Using this framework developed by Goldsmith-Pinkham *et al.* (2020), we can identify which robot applications have the largest weights ($\hat{\alpha}_a$) in the Rotemberg decomposition. As reported in Panel C of table 12, the largest-weight robot application (0.85) is “Handling for machine tending” (Application 115), which includes, for instance, the robots developed for handling operations in glass, ceramics or food production. One would expect these kinds of robots to accompany the workers, who remain engaged in various complementary activities. Accordingly, to be in line with the theoretical mechanism inspiring the analysis, one would expect the estimated coefficient for the largest-weight robot application in the regression for robot operators to be positive and similar to the average effect reported in table 5. This is indeed the case, as the estimated coefficient ($\beta_{a=\hat{115}}$) is equal to 0.275, which is close to the value in the main regression. Among the other results from the decomposition, we would like to mention the small share accruing to negative weights (Panel A). The validity of our identification strategy is also supported by the large correlations (Panel B) between the local shares of workers exposed to a given robot application (α_a), on the one hand, and the aggregate shocks g_a , the F-statistic (F_a) from the first stage of the just-identified estimations using the local share of workers associated to robot application a as an instrument, and the dispersion of these shares across locations ($\text{Var}(\mu_a)$), on the other.

Table 12. Rotemberg weights

Panel A: Negative and positive weights					
	Sum	Mean	Share		
Negative	-0.145	-0.018	0.112		
Positive	1.145	0.229	0.888		
Panel B: Correlations of Application Aggregates					
	α_a	g_a	$\hat{\beta}_a$	F_a	$\text{Var}(\mu_a)$
α_a	1				
g_a	0.513	1			
$\hat{\beta}_a$	0.173	0.094	1		
F_a	0.887	0.662	0.252	1	
$\text{Var}(\mu_a)$	0.845	0.155	-0.110	0.598	1
Panel C: Top 3 Rotemberg weight applications					
	$\hat{\alpha}_a$	g_a	$\hat{\beta}_a$	App Share	
Machine tending (115)	0.850	0.778	0.275	3.221	
Packaging, picking, placing (118)	0.267	1.595	-0.163	0.822	
Palletizing (117)	0.016	0.642	0.163	0.500	

Notes: This table reports statistics about the Rotemberg weights based on the regression for robot operators. Panel A reports the sum, mean and share of negative and positive Rotemberg weights. Panel B reports the correlations between the weights (α_a), the national component of growth (g_a), the just identified coefficients from the regression for robot operators $\hat{\beta}_a$, the first-stage F-statistic of our local shares of workers associated to robot application a (F_a), and the variation in robot application shares across locations ($\text{Var}(\mu_a)$). Panel C reports the top three robot applications according to the Rotemberg weights.

²³ As we have a large sample of locations and a fixed number of robot applications, the consistency of our shift-share IV estimator requires the exogeneity of the shares as the number of locations goes to infinity.

Finally, we also check the effect of robotization on two new sets of related dependent variables, i.e., the change in the share of weekly hours worked by robot operators and workers exposed to robots over the total number of weekly hours worked, and the growth rate in the hourly pay for all workers, robot operators and workers exposed to robots. All variables are measured over the same period, 2011-2018. The results in table A2 and table A3 in appendix A show that an increase in exposure to robots leads to an increase in the number of hours worked by robot operators relative to other workers, which implies that our main result is not driven by an increase in part-time workers and a corresponding substitution of hours per worker, but rather by an actual increase in full-time equivalent employees. On the other hand, we find no effect on hourly pay for any category of workers, despite the increase in demand for at least one category of these workers, i.e., robot operators. In part, this could be related to the generally weak labour market in Italy over this time period.

5.2 Industry-based measure of local exposure to robots

Next, we compare the results obtained using our new method to calculate local robots exposure and those using the traditional industry-based shift-share approach without the match between occupations and robot applications. As done in the literature, we calculate for the latter approach the variable $\Delta RX_i^{industry}$ and the corresponding industry-based shift-share IV, $\Delta \widetilde{RX}_i^{industry}$. Following the notation above, we have:

$$\Delta RX_i^{industry} = \sum_s \mu_{s,i} \log \left(\frac{R_{s,t}}{L_s} \right) - \sum_s \mu_{s,i} \log \left(\frac{R_{s,t0}}{L_s} \right), \quad (5)$$

where L_s denotes the aggregate number of workers in sector s in year 2011, $\mu_{s,i} = L_{s,i} / L_i$ is the 2011 share of workers in sector s in local labor market i , and $R_{s,t}$ and $R_{s,t0}$ are the Italian stocks of industrial robots in sector s at time t (2018) and $t0$ (2011). Similarly, the shift-share industry-based instrumental variable $\Delta \widetilde{RX}_i^{industry}$ can be defined as:

$$\Delta \widetilde{RX}_i^{industry} = \sum_s \tilde{\mu}_{s,i} \log \left(\frac{\widetilde{R}_{s,t}^{JPKRUS}}{L_s^{2001}} \right) - \sum_s \tilde{\mu}_{s,i} \log \left(\frac{\widetilde{R}_{s,t0}^{JPKRUS}}{L_s^{2001}} \right), \quad (6)$$

where L_s^{2001} denotes the aggregate number of workers in sector s in year 2001, $\tilde{\mu}_{s,i} = L_{s,i}^{2001} / L_i^{2001}$ is the 2001 share of workers in sector s in local labor market i , and $\widetilde{R}_{s,t}^{JPKRUS}$ and $\widetilde{R}_{s,t0}^{JPKRUS}$ are the stocks of industrial robots in Japan, South Korea and the US in sector s at time t (2018) and $t0$ (2011).

The results of this alternative empirical strategy are shown in table 13. The results are qualitatively similar, but quantitatively different. The estimated impacts are smaller and less significant, in line with an attenuation bias caused by the approximation related to the treatment of all workers in an industry as equally exposed to robotization. Indeed, the IV barely passes the tests for the validity of the instrument: the Kleibergen-Paap F statistic $\Delta RX_i^{industry}$ for is approximately equal to 23, while the Stock-Yogo critical value at the 10% significance level is equal to 16.4²⁴. While obtaining qualitatively

²⁴ The first stage is reported in column 2 of table A1 in appendix A.

similar results is to a certain extent reassuring, it is important to keep in mind that the identification of robot operators and exposed robots (i.e., the dependent variables) still requires our novel procedure to match occupations and robot applications. Hence, internal consistency calls for building a measure of local robots exposure that incorporates the observation that not all occupations within an industry are equally exposed to the robots installed in that industry (as the industry-based approach implicitly entails).

Table 13. Effects of local robots exposure: activity-based vs industry-based approach

	Δ operators		Δ exposed workers	
	Activity-based (1)	Industry-based (2)	Activity-based (3)	Industry-based (4)
ΔRX	0.293*** (0.095)		0.301 (0.194)	
$\Delta RX^{industry}$		0.076** (0.038)		0.102 (0.069)
Macroarea fixed effects	C	C	C	C
Lagged demographic controls	C	C	C	C
Lagged industry controls	C	C	C	C
Lagged profession controls	C	C	C	C
Observations	377	377	377	377
Kleibergen-Paap F stat	130.1	22.98	130.1	22.98

Notes: The dependent variable is the difference in the share of robot operators over working-age population between 2018 and 2011 in columns (1) and (2), and the difference in the share of exposed workers over working-age population between 2018 and 2011 in columns (3) and (4). The shares are computed with respect to the working-age resident population (15-74 years old). The local robots exposure is measured following an activity-based approach in columns (1) and (3) and an industry-based approach in columns (2) and (4). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Ideally, had we access to census data with detailed information on the sector and on the occupation of each worker, as well as information on both the sector and the application of each robot, we could build different and more sophisticated measures of local exposure to robots. In this work, due to the limited information on IFR data and the desire to exploit 2001 census data, we have adjusted the sectoral composition of the local workforce by the occupational structure of each industry.

6. Employment dynamics across alternative occupational categories

As discussed in section 5, it is worth investigating the impact of local robots exposure on the employment dynamics of alternative (narrower) occupational categories. This exercise might contribute to improve our understanding of the underlying mechanisms at work, and provide some novel evidence on the relationship between (our new measure of) local robots exposure and other meaningful groups of occupations that share specific features (such as routine intensity, exposure to automation, and the like). It is worth stressing that this analysis does not aim at covering every possible decomposition of the labor force, but it explores those occupation-specific features that may contribute to determine the sign and size of the impact of robot adoption on employment.

It is worth noticing that our measure of local robots exposure does not change and it is still based on the identification of professions exposed to different robot applications, on the basis of which the shift-share variable is built. In practice, thus, we exploit our novel activity-based matching procedure of occupations and robot applications to investigate the employment dynamics for categories of workers whose identification does not depend on our activity-based taxonomy.

The first analysis we conduct investigates whether the RBTC hypothesis finds supportive evidence in our setting or not. The RBTC literature maintains that intensively codifiable activities are well suited to computerization and, accordingly, exposed to automation and digitalization. As illustrated in the Introduction, Autor and Dorn (2013) and Autor *et al.* (2015) distinguish routine and non-routine occupations on the basis of the relative importance of routine, manual and abstract tasks. For instance, they classify as routine occupations those that are in the top-third of the distribution of the routine-task intensity²⁵. It is plausible to think that also robots (as well as computers and other forms of automation) can affect workers in relation to the manual and routine contents of the tasks characterising every profession. Webb (2020), for instance, shows that robot patents refer significantly more often to manual than routine tasks. Accordingly, rather than grouping workers in terms of the activities characterising the professions, as done so far in our main analysis, we study the local employment dynamics of groups of workers that perform activities that are similar in terms of the routine-task intensity.

To conduct these estimations, we first identify the occupations that are traditionally more routine intensive than others in Italy. We adopt the operationalization proposed by Autor *et al.* (2003), and implemented by Esposito and Scicchitano (2020) and Cassandro *et al.* (2020) on Italian data, using the Inapp ICP Survey. We also exploit the richness of the Inapp ICP Survey, which contains information on different types of generalized tasks, activities, risks, competences and skills associated to occupations at the five-digit level, mimicking what was done by Autor *et al.* (2003) with the information contained in O*NET. Following Autor *et al.* (2003) and considering occupations as routinary if they are in the top third of the distribution of the Routine Task Index, we distinguish the following groups: routine intensive occupations (RTI); routine cognitive (RC); routine manual (RM); non-routine cognitive: analytical (NRCA); non-routine cognitive: interpersonal (NRCI); non-routine manual (NRM); non-routine manual: interpersonal adaptability (NRMIA).

²⁵ See Fernandez-Macias and Hurley (2016) for a discussion of the alternative operationalizations of the concept of routine in the RBTC literature.

As can be seen in the table 14, the estimates suggest that there are no statistically significant effects on local employment dynamics on the basis of occupations’ routine intensity. However, our findings suggest that those who are employed in occupations that are intensive in routine cognitive tasks or non-routine tasks associated with interpersonal relations increase their relative shares. These findings seem to confirm that the relation ship between the change in the local exposure to robots and in the local employment conditions is not characterized by a net substitution of workers in routine manual jobs. As our results for the category of robot operators reveal, workers capable of developing, supervising, maintaining, fixing and operating robots do seem to increase in absolute and relative size. As Pfeiffer (2016) explains, “what workers [interacting with robots] need to do is anything but mere routine work”. This result is in line with recent studies assessing the relevance of the RBTC hypothesis in Italy, as they find that the phenomenon is much less significant than in other developed economies (Basso 2020; Cassandro *et al.* 2020; Marcolin *et al.* 2018), also because of the country’s specialization in low-tech and low-skill intensive sectors.

Table 14. Effects of local robots exposure on employment dynamics, routine

	Δ employment shares						
	RTI (1)	RC (2)	RM (3)	NRM (4)	NRCA (5)	NRMI (6)	NRMIA (7)
ΔRX	0.248 (0.283)	0.742*** (0.365)	0.446 (0.354)	0.182 (0.426)	0.149 (0.299)	0.462* (0.271)	0.122 (0.387)
Macroarea fixed effects	C	C	C	C	C	C	C
Lagged demographic controls	C	C	C	C	C	C	C
Lagged industry controls	C	C	C	C	C	C	C
Lagged profession controls	C	C	C	C	C	C	C
Observations	377	377	377	377	377	377	377
Kleibergen-Paap F stat	130.1	130.1	130.1	130.1	130.1	130.1	130.1

Notes: The dependent variable is the difference in the share of workers in the top-third of the distribution of different routine indices between 2018 and 2011: routine-intensive occupations (RTI); routine cognitive (RC); routine manual (RM); non-routine cognitive: analytical (NRCA); non-routine cognitive: interpersonal (NRMI); non-routine manual (NRM); non-routine manual: interpersonal adaptability (NR-MIA). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

If the intensity of routine tasks does not seem to determine jobs’ susceptibility to robotization, other characteristics of occupations may determine the nature of the interaction between workers and robots. To this end, we explore the questions in the Inapp ICP Survey, so as to draw information on a number of features characterizing each occupation at the five-digit level. We consider, in particular, the answers given to the following six questions: 1) “To what level is it necessary in your profession to be able to use the abdominal and lumbar muscles to support the body in a continuous and repeated manner over time without fatigue or giving in to the effort?” (Torso strength, question D35); 2) “To

what level is it necessary in your profession to perform physical activities that require moving the entire body or a substantial use of the arms and legs, such as climbing, climbing stairs, balancing, walking, bending, and manipulating materials?” (Performing general physical activities, question G16); 3) “To what level is it necessary in your profession to use hands and arms to handle, install, position, and move materials or to manipulate objects?” (Handle and move objects, question G17); 4) “How much of your time do you use your hands to manipulate, check, or feel objects, tools, or control systems?” (question H40); 5) “How automated is your work?” (question H49); 6) “How important is it to keep track of sequences of machinery and equipment in your work?” (question H55). Respondents can answer using a scale from 1 to 7 for the first three questions and a scale from 1 to 5 for the last three questions.

Similarly to the RTI, each of these questions can be used to create an index varying over the interval 0-100. For each of these questions, we identify those (five-digit level) occupations that are in the top third of the distribution, and we calculate the changes in the share of employment over the working-age population at the LLMA level for each one of them. Accordingly, we obtain six groups: torso-intensive occupations (Tol); physical movement-intensive occupations (Pml); manual-intensive occupations (Manl); object-intensive occupations (Obl); automation-intensive occupations (Aul); machinery-intensive occupations (Macl).

Table 15. Effects of local robots exposure on employment dynamics, occupations’ features

	<i>Δ employment shares</i>					
	Tol (1)	Pml (2)	Obl (3)	Manl (4)	Aul (5)	Macl (6)
<i>Δ RX</i>	-.0634** (0.307)	-0.229 (0.317)	-0.030 (0.334)	-0.105 (0.355)	0.546* (0.314)	0.623* (0.335)
Macroarea fixed effects	C	C	C	C	C	C
Lagged demographic controls	C	C	C	C	C	C
Lagged industry controls	C	C	C	C	C	C
Lagged profession controls	C	C	C	C	C	C
Observations	377	377	377	377	377	377
Kleibergen-Paap F stat	130.1	130.1	130.1	130.1	130.1	130.1

Notes: The dependent variable is the difference between 2018 and 2011 in the share of workers in the top-third of the distribution of various occupations’ characteristics: torso-intensive occupations (Tol); physical movement-intensive occupations (Pml); manual-intensive occupations (Manl); object-intensive occupations (Obl); automation-intensive occupations (Aul); machinery-intensive occupations (Macl). The shares are computed with respect to the working-age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 15 reports the estimates for the employment dynamics of the above-mentioned groups of occupations. This analysis reveals contrasting, but not contradictory, results and it confirms both the heterogeneous impact of robot adoption on workers and the implications of using certain grouping of

occupations instead of others. In particular, the estimates reveal a negative effect of changes in the local exposure to robots on the dynamics of the employment share of occupations characterized by activities and tasks requiring continuous and repeated movements of the body: this is in accordance with a substitution of workers in those activities that can be best performed by robots' arms.

This result is an example of a displacement effect that can be detected only by performing an analysis of employment dynamics at a high level of occupational disaggregation, as done in this paper. Conversely, the results confirm a positive and (mildly) significant effect of exposure to robots on occupations characterized by a high level of automation and intensive in machinery. These findings corroborate our conclusions on the impact on robot operators: the integration of robots in firms is accompanied by the expansion of occupations connected with them. It should not be surprising, thus, that the local share of workers who answer that their job is highly automated grows more in those areas where robots exposure expands.

7. Closing remarks

In this work we empirically investigate the effects that the change in local exposure to robots had on the Italian local employment dynamics over the period 2011-2018. Our empirical strategy extends previous research in this field as it focuses on occupations' activities and robots' applications, rather than on the industries in which robots and workers are employed. Moreover, the analysis focuses on novel groups of occupations identified on the basis of a manual match between occupations' activities and robots' applications that we conduct at a high level of disaggregation (five-digit level) for the occupations.

This approach makes it possible to identify robot operators, that is workers employed in occupations clearly associated with robot installation, maintenance and use. Economic theory predicts that these jobs tend to grow together with the volume of adopted robots, as their activities are complementary to those carried out by specific robot applications. With the same method, we also single out those professions that, because of the main activities they consist of, are most exposed to the robot applications available in the market. Workers employed in these occupations are considered as exposed to robots and, according to economic theory, their number may either grow or fall according to the net impacts of displacement effects and reinstatement effects. Finally, we use this information also to build a new shift-share measure of local exposure to robots that does not refer only to the local industry composition, but it also takes into account the occupational composition of each industry and the different exposure to robots of each occupation. Besides being more in line with an interpretation of exposure to robots in terms of tasks and activities, this measure is also internally consistent with the groups of workers whose local employment dynamics we study in this work.

From a theoretical viewpoint, as mentioned, our approach is more consistent with the theoretical task-based production frameworks, whereby tasks are allocated to labor and capital according to the available technology, whose progress alters labor demand across professions according to their inherent tasks (Autor *et al.* 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2019). From an empirical viewpoint, our methodology makes it possible to build consistent measures of local employment dynamics and of local exposure to robots, thereby reducing the possible mismatch between industry-based measures of local robots exposure and the employment dynamics of groups

of workers formed on the basis of occupations, skills and tasks. Besides offering greater internal consistency, this new approach also reduces possible interference in the estimation stemming from the presence of industry-related confounding factors.

Our empirical results suggest a limited aggregated impact of robot adoption in Italy in the most recent period of time. Although this finding is at odds with the seminal studies conducted on the US (Acemoglu and Restrepo 2020), it is in line with the conclusions reached with industry-based approaches by Dauth *et al.* (2017) for Germany and by Dottori (2020) for Italy. As pointed out by Gentili *et al.* (2020), the process of robot adoption and its implications on local labor markets are time-sensitive and depend on the local socio-economic system. Hence, it should not be controversial that the negative impact of robots exposure found in the US does not carry over to Germany and Italy.

More interestingly, we show for the first time that the lack of significant effects of robot adoption on aggregated labor market dynamics hides more complex employment dynamics for specific groups of workers, when properly identified. Our novel approach to look at the occupations (at five-digit level) allows us to discern clear signs of reinstatement effects for robot operators, that is workers employed in occupations complementary to the installation and use of robots. On the other hand, our analysis fails to detect any significant impact of local exposure to robots on the evolution of the shares of exposed workers. While this empirical finding may stem from the absence of replacement effects, further empirical work on other sub-groups of occupations suggest that different forces may be at work. Certain occupations, such as those performing activities that require an intensive use of the torso, are significantly reduced by robot adoption; others, such as routine cognitive and non-routine manual occupations, seem to grow together with robot adoption. This suggests that the lack of significant effects of local robots exposure on local employment dynamics is probably due both to the heterogeneous impact of robot adoption on different jobs, and to inappropriate pooling of professions. In turn, our findings appear consistent with recent microeconomic evidence showing that the firms that install robots also expand employment. Thus, our findings suggest that future work should focus on the effects of robotization on new highly disaggregated groups of workers and occupations that can provide a link between microeconomic and macroeconomic studies.

Our findings suggest the importance of supporting education and training policies that focus on the skills associated with the activities carried out by robot operators. Indeed, investment in human capital is fundamental to generating the skills that complement new technologies (Autor 2015). In Italy, whose economy is based on low-skill and low-tech productions, this is particularly relevant. The case will be even stronger in the aftermath of the Covid-19 pandemic, as this latter has pushed for a more intense use of robots and other innovative technologies that can help reduce the risk of contagion (Hantrais *et al.* 2021; Caselli *et al.* 2020).

Appendix

A. Additional tables and figures

Table A1. Effects of local robots exposure, first stage

	ΔRX (1)	$\Delta RX^{industry}$ (2)
$\Delta \tilde{RX}$	0.353*** (0.028)	
$\Delta \tilde{RX}^{industry}$		0.319***
Macroarea fixed effects	C	C
Lagged demographic controls	C	C
Lagged industry controls	C	C
Lagged profession controls	C	C
Observations	377	377
Kleibergen-Paap F stat	130.1	22.98

Notes: First-stage results for the effects of robots exposure on employment dynamics. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2. Effects of local robots exposure on employment dynamics, weekly hours

	$\Delta operators$ (1)	$\Delta exposed workers$ (2)
ΔRX	0.470** (0.197)	0.394 (0.386)
Macroarea fixed effects	C	C
Lagged demographic controls	C	C
Lagged industry controls	C	C
Lagged profession controls	C	C
Observations	377	377
Kleibergen-Paap F stat	130.1	130.1

Notes: The dependent variable is the difference in the share of weekly hours worked by robot operators between 2018 and 2011 in column (1), and the difference in the share of weekly hours worked by exposed workers between 2018 and 2011 in column (2). The shares are computed with respect to the working age resident population (15-74 years old). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A3. Effects of local robots exposure on employment dynamics, hourly pay

	Δ HP - all (1)	Δ HP - operators (2)	Δ HP - exposed workers (3)
Δ RX	-0.153 (0.732)	-0.301 (2.402)	0.593 (1.111)
Macroarea fixed effects	C	C	C
Lagged demographic controls	C	C	C
Lagged industry controls	C	C	C
Lagged profession controls	C	C	C
Observations	377	248	342
Kleibergen-Paap F stat	130.1	132.5	125.8

Notes: The dependent variable is the difference in the average hourly pay between 2018 and 2011 for all workers in column (1), for robot operators in column (2), and for exposed workers in column (3). All regressions are estimated with the 2SLS estimator. The instrument for robots exposure is measured following a shift-share approach with lagged weights and growth of robots in Japan, South Korea and the US. Four macroareas are included: North-West, North-East, Centre, South and Islands. Lagged control variables are measured in 2011. Lagged demographic controls include the logarithm of the resident population, the share of residents older than 65, the share of female residents, the share of residents with tertiary education, and the share of foreign-born residents. Lagged industry-related controls include the local exposure to trade competition from China, the share of employees in the manufacturing sector, the share of employees in the transportation and logistics industry, the share of employees that use a personal computer, the share of local units in non-light manufacturing sectors with at least 50 employees, and the local share of workers employed in the private sector. Lagged profession controls include the local share of workers occupied as craftsmen, plant operators and unskilled professionals (categories 6, 7 and 8 of ISCO 08), and the local diffusion of routine occupations. The estimations are weighted on the basis of the 2011 resident population. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B. A new matching approach

Table B1. A new matching approach: Exposed and not exposed

5 digit	Occupation	Act 1	imp freq	Act 2	imp freq	Act 3	imp freq	IFR code	Type
6.1.1.2.0	Stone-cutters	Cutting stones	4.5 4.4	Sanding stones	4.3 3.9	Polishing stones	4.1 3.8	190	Exposed
6.2.2.1.1	Blacksmiths	Cutting metals	4.1 4.2	Assembling metal pieces	4 4.2	Welding and repairing metal pieces	3.9 3.9	190	Exposed
6.2.1.2.0	Welders and flame cutters	Using tools	4.7 4.6	Welding pieces	4.6 4.6	Supervising machines	4.2 4.2	160	Exposed
6.2.3.6.0	Inspector mechanics	Identifying malfunctions or anomalies	4 3.5	Performing tests on engines and machines	3.9 3.4	Reading technical designs	3.8 3.7		Not exposed
6.3.4.5.0	Bookbinders and after print finishers	Binding books	3.9 3.7	Gluing materials	3.8 3.7	Folding paper	3.8 3.7		Not exposed
6.5.3.2.1	Weavers	Working on looms	3.9 3.3	Replacing threads	3.8 3.4	Maintaining machinery and equipment	3.6 2.1		Not exposed

Notes: IFR robot application: 190 -processing, enduring changing; the robot leads the workpiece or the tool; material removal. 160 - Welding and soldering. 110 -Handling operations/ Machine tending. Act is the abbreviation of activity. Imp is the abbreviation for importance. Freq is the abbreviation for frequency.

Table B2. A new matching approach: Operators

5 digit	Occupation	Description	Act 1	Act 2	Act 3
2.2.1.3.0	Electrical and factory automation engineers	Design, engineer, control, manufacture and maintain systems, motors, apparatus and equipment (electrical energy)	Verifying safety conditions	Performing tests	Inspecting and producing estimates
3.1.4.1.5	Automated assembly line management and control technicians	Apply procedures and technology to control through automated systems, and manage the operation and safety of automated assembly lines	Carrying out maintenance on equipment or installations	Signaling anomalies	Supervising application of safety norms
7.1.7.1.0	Automated assembly line operators	Lead and control the activity of automated assembly lines and industrial robots	Controlling quality	Supervising process	Choosing tools

C. Comparison to existing measures of job replaceability

The idea of exploiting the classification of robots in terms of their applications was first exploited by Graetz and Michaels (2018) to build an instrumental variable suitable to assess the relationship between the country-industry exposure to robots and the country-level evolution of productivity and employment. More precisely, they consider each profession at the three-digit level (as in the 2000 US Census occupational classification) and attribute a replaceability value of 1 if the profession’s title corresponds to at least one of the IFR application categories, and a value of 0 otherwise. Then, Graetz and Michaels (2018) calculate the share of replaceable professions within each industry and country and use this variable to instrument the measure of robot adoption across country-industry pairs²⁶.

While sharing the same spirit, as well as the manual approach to the classification of the occupations, our strategy differs considerably from that in Graetz and Michaels (2018). The first difference is in the goal. Graetz and Michaels (2018) match US occupations’ titles and robots’ applications to calculate the national shares of replaceable professions, but they do not use these shares to build local measures of robots exposure. As explained in the main text, these measures do not take into consideration the actual evolution of the stocks of robots. The second difference with their approach is in the matching method: the association of robot applications with occupations performed by Graetz and Michaels (2018) is based on the denomination of the professions at a high level of aggregation (three-digit), whereas we examine the frequency and importance of the activities characterizing each occupation at the five-digit level. Third, our activity-based matching approach makes it possible to single out robot operators, predicted to be complementary to robots, whereas Graetz and Michaels (2018)’s approach is directed exclusively at identifying the professions at risk of being replaced.

Although our empirical analysis explores the geographical variation in robots exposure within the country, rather than the heterogeneous distribution of the local share of replaceable occupations, it is interesting to compare our taxonomy of exposed occupations with those identified by Graetz and Michaels (2018) as susceptible of being replaced by robots. For the sake of completeness, we also

²⁶ Their identification strategy is based on the positive correlation between the share of replaceable occupations at the beginning of the period and the subsequent adoption of robots by country-industry pairs.

compare our classification of exposed professions with the automatable jobs proposed by Josten and Lordan (2019), who build on and refine Autor and Dorn (2013)'s classification of routine occupations, with Frey and Osborne (2017)'s classification of occupations at risk of automation, and with Webb (2020)'s classification of occupations exposed to robots based on their tasks and matched patents related to robotization.

We start by recalling that we consider an occupation (at the five-digit level) as exposed to robots any time that a robot application matches with one (or more) of its most relevant and frequent activities. Graetz and Michaels (2018), instead, look at the titles of three-digit level occupations and assign a replaceability value of 1 if the occupation name contains at least one IFR application. It follows that our approach is more refined as we can distinguish different types of jobs within the same occupation at the three-digit level, and this produces a more accurate assessment of the susceptibility of individual occupations to robotization. To compare the two taxonomies, we need to perform some non-trivial adjustments to associate Graetz and Michaels (2018)'s three-digit level SOC1990 classification with our five-digit level CP2011 classification. We match 795 of our occupational units (OU) with theirs and we assume that a CP2011 OU is replaceable according to Graetz and Michaels (2018)'s definition if at least half of the SOC1990 occupations included in that OU are categorized as replaceable. We find that 644 OUs can be identified as not susceptible to robotization in any of the two classifications. On the other hand, 89 OUs are assessed as exposed to robotization using both methods. Only 28 OUs are considered as replaceable according to our method, but not according to Graetz and Michaels (2018)'s approach, whereas we assess as non-exposed 34 OUs that Graetz and Michaels (2018) consider replaceable to robots. To calculate the correlation between our measures and the dummy variable associated with Graetz and Michaels (2018)'s approach, we build two variables based on percentages: one is the percentage of SOC1990 occupations included in a CP2011 OU that Graetz and Michaels (2018) deem as replaceable; the second one is the percentage of workers in a CP2011 OU employed in SOC1990 occupations replaceable according to Graetz and Michaels (2018). The correlation of our variable with the workers-weighted and unweighted measures is 0.71 and 0.74, respectively²⁷. These are very high correlation rates if one considers the noise created by the crosswalk from SOC1990 to CP2011. This provides some indirect evidence on the plausibility of the matching procedure we develop to build the shift-share measure of local robots exposure.

Another relevant comparison to make regards the classification of automatable jobs by Josten and Lordan (2019). As before, we first match the OCC90 occupations at the three-digit level with the CP2011 classification at the five-digit level. Again, OCC90 categories may include several CP2011 OUs, and viceversa. We match 797 of our 800 OUs with theirs. Assuming that a CP2011 OU is automatable according to Josten and Lordan (2019) if at least 50% of the OCC90 occupations included in that OU are automatable, we find that 395 OUs can be identified as not susceptible to automation in any of the two classifications. Only 3 OUs are considered as exposed to robots according to our method, but not susceptible of automation according to Josten and Lordan (2019)'s approach. We assess as non-replaceable by robots 279 OUs that, instead, Josten and Lordan (2019) consider as automatable. Finally, 120 OUs are assessed as likely susceptible with both methods. On this basis, we build two

²⁷ These two variables can also be transformed into dummy variables using the 50% threshold. The correlations become 0.68 and 0.70, respectively.

variables to capture Josten and Lordan (2019)'s classification of automatable jobs in terms of CP2011 occupations. One variable is the percentage of OCC90 occupations included in a CP2011 OU considered as automatable by Josten and Lordan (2019); the second one is the percentage of workers in a CP2011 OU who are employed in occupations automatable according to Josten and Lordan (2019). The correlations between these workers-weighted and unweighted measures and our dummy variable (taking value 1 for any OU that we consider as exposed to robots) are, respectively, 0.42 and 0.44²⁸. Considering that exposure to automation is a broader and more inclusive concept than exposure to robotization, these high correlation rates support the reasonableness of our approach based on workers' activities.

We also compare our classification with that produced by Frey and Osborne (2017), although it differs remarkably from ours along several dimensions. First, Frey and Osborne (2017)'s classification does not adopt an activity-based approach to establish which occupations are at risk of automation, but it rather exploits the presence (or lack thereof) of technical features that prevent exposure to full automation²⁹. Second, they consider the broader concept of automation rather than robotization. Having clarified these substantial differences, Frey and Osborne (2017) end up attaching a risk of automation (ranging from 0 to 1) to 702 (SOC10) occupations depending on the presence of some key technical features. To compare our activity-based classification with theirs, we first match the Isco08 four-digit level classification to theirs, and then match the latter to the five-digit level CP2011 classification. We then identify 731 occupations that are present in both studies. We generate two variables: one is the average automation (from 0 to 1) for the SOC10 professions included in each CP2011 OU, and the other is a dummy variable taking value 1 if the risk is above 80%. The correlation between the two measures and our dummy variable for OUs exposed to robots is 0.37 and 0.41 respectively. Considering that Frey and Osborne (2017) focus on automation rather than robotization and that they look at occupations' features rather than activities, these results provide additional evidence of the reasonableness of our approach.

Finally, we conduct a similar comparison between our match between workers' occupations and robots' applications and Webb (2020)'s measure of the exposure of tasks to robotization based on the overlap between the text of job task descriptions and the text of patents. As for the previous comparisons, we first match SOC2010 occupations with the CP2011 classification at the five-digit level. In this case, we obtain a match for 776 CP2011 OUs. Assuming that a CP2011 OU is exposed to automation by industrial robots according to Webb (2020) if it falls in the top third of the exposure distribution, we find that 412 OUs can be identified as not susceptible to robotization in any of the two classifications, while 103 OUs are susceptible with both methods. On the other hand, only 16 OUs are considered as exposed to robots according to Webb (2020)'s method, but not exposed according to ours, while we assess as exposed by robots 245 OUs that, instead, Webb (2020) does not consider

²⁸ We also build the corresponding dummy variables using the 50% threshold. The weighted and unweighted correlation indices are equal to 0.36 and 0.41, respectively.

²⁹ Frey and Osborne (2017) consult with machine-learning experts and manually classify 70 (out of 703) occupations in O*NET as either automatable or not. Then, they estimate a probability model on these 70 observations using as explanatory variables nine O*NET variables associated with the main technological limitations to implementing total automation, and finally based on the estimated parameters they predict the probability of automation for all the remaining occupations.

as exposed. On this basis, we build two variables to capture Webb (2020)'s classification of exposed jobs in terms of CP2011 occupations. One variable is the average exposure of SOC2010 occupations included in a CP2011 OU; the second one is a dummy equal to 1 if the occupation falls in the top third of the exposure distribution. The correlations between these measures and our dummy variable (taking value 1 for any OU that we consider as exposed to robots) are, respectively, 0.23 and 0.36. In addition, Webb (2020) constructs similar measures of exposure to automation due to software and artificial intelligence. In these cases, the correlation rates are much lower, 0.19 and 0.21 for software, and 0.01 and -0.01 for artificial intelligence. Thus, once again, we can conclude that our matching procedure produces reasonable results.

References

- Acemoglu D. (2002), Technical change, inequality, and the labor market, *Journal of Economic Literature*, 40, n.1, pp.7-72
- Acemoglu D., Autor D. (2011), Skills, tasks and technologies. Implications for employment and earnings, in Card D., Ashenfelter O. (eds.), *Handbook of labor economics Vol. 4A*, New York, Elsevier, pp. 1043-1171
- Acemoglu D., Lelarge C., Restrepo P. (2020), Competing with robots. Firm-level evidence from France, *AEA Papers and Proceedings*, 110, pp.383-388
- Acemoglu D., Restrepo P. (2020), Robots and jobs. Evidence from US labor markets, *Journal of Political Economy*, 128, n.6, pp.2188-2244
- Acemoglu D., Restrepo P. (2019), Automation and New Tasks. How Technology Displaces and Reinstates Labor, *Journal of Economic Perspectives*, 33, n.2, pp.3-30
- Adachi D., Kawaguchi D., Saito Y.U. (2020), *Robots and Employment. Evidence from Japan 1978-2017*, RIETI Discussion Paper n. 20-E-051, Tokyo, Research Institute of Economy, Trade and Industry
- Aghion P., Antonin C., Bunel S. (2019), Artificial Intelligence, Growth and Employment. The Role of Policy, *Economie et Statistique / Economics and Statistics*, n.510-512, special issue 50th anniversary, pp.149-164
- 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 n.14443, London, CEPR
- Akcomak S., Kok S., Rojas-Romagosa H. (2016), Technology, offshoring and the task content of occupations in the United Kingdom, *International Labour Review*, 155, n.2, pp.201-230
- Arntz M., Gregory T., Zierahn U. (2017), Revisiting the risk of automation, *Economics Letters*, 159, issue C, pp.157-160
- Autor D.H. (2015), Why are there still so many jobs? The history and future of workplace automation, *Journal of Economic Perspectives*, 29, n.3, pp.3-30
- Autor D.H., Salomons A. (2018), *Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share*, NBER Working Paper n.24871, Cambridge MA, NBER <<https://bit.ly/36rVh2G>>
- Autor D.H., Dorn D. (2013), The growth of low-skill service jobs and the polarization of the us labor market, *American Economic Review*, 103, n.5, pp.1553-1597
- Autor D.H., Dorn D., Hanson G. H. (2015), Untangling trade and technology. Evidence from local labour markets, *The Economic Journal*, 125, n.584, pp.621-646
- Autor D.H., Dorn D., Hanson G.H. (2013), The China Syndrome. Local Labor Market Effects of Import Competition in the United States, *American Economic Review*, 103, n.6, pp.2121-2168
- Autor D.H., Katz L.F., Kearney M.S. (2006), The polarization of the US labor market, *The American Economic Review*, 96, n.2, pp.189-194
- 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, n.4, pp.1279-1333
- Backer K.D., DeStefano T., Menon C., Suh J.R. (2018), *Industrial robotics and the global organisation of production*, OECD Science, Technology and Industry Working Papers n.3, Paris, OECD <<https://bit.ly/3wo17wE>>

- Barbieri T., Basso G., Scicchitano S. (2020), *Italian Workers at Risk During the Covid-19 Epidemic*, GLO Discussion Paper n.513, Essen, Global Labor Organization
- Bartik T.J. (1991), *Who Benefits from State and Local Economic Development Policies?*, Kalamazoo MI, W.E. Upjohn Institute for Employment Research <<https://bit.ly/3yEOjUi>>
- Basso G. (2020), The Evolution of the Occupational Structure in Italy, 2007-2017, *Social Indicators Research*, 152, n.2, pp.673-704
- Bessen J., Goos M., Salomons A., van den Berge W. (2020), Firm-level automation. Evidence from the Netherlands, *AEA Papers and Proceedings*, 110, pp.389-393
- Blanas S., Gancia G., Lee S.Y.T. (2020), Who is afraid of machines?, *Economic Policy*, 34, n.100, pp.627-690
- Bonacini L., Gallo G., Scicchitano S. (2021), Working from home and income inequality. Risks of a 'new normal' with Covid-19, *Journal of Population Economics*, 34, pp.303-360
- Bonfiglioli A., Crinò R., Fadinger H., Gancia G. (2020), *Robot Imports and Firm-Level Outcomes*, CEPR Discussion Paper n.14593, London, CEPR
- Brandes P., Wattenhofer R. (2016), *Opening the Frey/Oborne black box. Which tasks of a job are susceptible to computerization?*, arXiv, 1604.08823v2 <<https://bit.ly/2VuOvsN>>
- Brynjolfsson E., McAfee A. (2014), *The second machine age. Work, progress, and prosperity in a time of brilliant technologies*, New York, W.W. Norton & Company
- Brynjolfsson E., McAfee A. (2011), *Race against the machine. How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*, Digital Frontier Press
- Card D. (2001), Immigrant inflows, native outflows, and the local labor market impacts of higher immigration, *Journal of Labor Economics*, 19, n.1, pp.22-64
- Caselli F., Manning A. (2019), Robot arithmetic. New technology and wages, *American Economic Review: Insights*, 1, n.1, pp.1-12
- Caselli M., Fracasso A., Traverso S. (2021), Globalization, robotization, and electoral outcomes. Evidence from spatial regressions for Italy, *Journal of Regional Science*, 61, n.1, pp.86-111
- Caselli M., Fracasso A., Traverso S. (2020), Mitigation of risks of Covid-19 contagion and robotization. Evidence from Italy, in CEPR, *Covid Economics. Vetted and Real-Time Papers*, issue 17, CEPR Press, pp.174-188 <<https://bit.ly/2TRP51m>>
- Cassandro N., Centra M., Esposito P., Guarascio D. (2020), *What drives employment-unemployment transitions? Evidence from Italian task-based data*, GLO Discussion Paper n.563, Essen, Global Labor Organization
- 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 n.2, Brussels, Bruegel
- 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, n.57, pp.6127-6138
- Dauth W., Findeisen S., Südekum J., Woessner N. (2017), *German Robots. The Impact of Industrial Robots on Workers*, CEPR Discussion Paper n.12306, London, CEPR
- David B. (2017), Computer technology and probable job destructions in Japan. An evaluation, *Journal of the Japanese and International Economies*, 43, issue C, pp.77-87

- de Blasio G, Poy S. (2017), The impact of local wage regulation on employment. A border analysis from Italy in the 1950s, *Journal of Regional Science*, 57, n.1, pp.48-74
- 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, article 101885
- Dengler K., Matthes B. (2018), The impacts of digital transformation on the labour market. Substitution potentials of occupations in Germany, *Technological Forecasting and Social Change*, 137, pp.304-316
- Dixon J., Hong B., Wu L. (2020), *The Robot Revolution. Managerial and Employment Consequences for Firms*, New York, NYU Stern School of Business
- Dottori D. (2020), *Robots and employment. Evidence from Italy*, Questioni di economia e finanza n.572, Roma, Banca d'Italia <<https://bit.ly/3xu5BmV>>
- Esposito P., Scicchitano S. (2020), *Educational mismatches, technological change and unemployment. Evidence from secondary and tertiary educated workers*, GLO Discussion Paper n.465, Essen, Global Labor Organization
- Felten E.W., Raj M., Seamans R. (2018), A Method to Link Advances in Artificial Intelligence to Occupational Abilities, *AEA Papers and Proceedings*, 108, pp.54-57
- Fernandez-Macias E., Hurley J. (2016), Routine-biased technical change and job polarization in Europe, *Socio-Economic Review*, 15, n.3, pp.563-585
- Fossen F.M., Sorgner A. (2019), Digitalization of work and entry into entrepreneurship, *Journal of Business Research*, 125, special issue, pp.548-563
- Frey C.B., Osborne M.A. (2017), The future of employment. How susceptible are jobs to computerisation?, *Technological Forecasting and Social Change*, 114, issue C, pp.254-280
- Garibaldi P., Taddei F. (2013) *Italy: a dual labour market in transition. Country case studies on labour market segmentation*, Employment Working Paper n.144, Geneva, ILO <<https://bit.ly/2S1uZk2>>
- Gentili A., Compagnucci F., Gallegati M., Valentini E. (2020), Are machines stealing our jobs?, *Cambridge Journal of Regions, Economy and Society*, 13, n.1, pp.153-173
- Goldsmith-Pinkham P., Sorkin I., Swift H. (2020), Bartik instruments. What, when, why, and how, *American Economic Review*, 110, n.8, pp.2586-2624
- Goos M., Manning A. (2007), Lousy and lovely jobs. The rising polarization of work in Britain, *The Review of Economics and Statistics*, 89, n.1, pp.118-133
- Goos M., Manning A., Salomons A. (2014), Explaining Job Polarization. Routine-Biased Technological Change and Offshoring, *American Economic Review*, 104, n.8, pp.2509-2526
- Goos M., Manning A., Salomons A. (2009), Job polarization in Europe, *American Economic Review*, 99, n.2, pp.58-63
- Graetz G., Michaels G. (2018), Robots at work, *The Review of Economics and Statistics*, 100, n.5, pp.753-768
- Hantrais L., Allin P., Kritikos M., Sogomonjan M., Anand P.B., Livingstone S., Williams M., Innes M. (2021), Covid-19 and the digital revolution, *Contemporary Social Science*, 16, n.2, pp.256-270
- Josten C., Lordan G. (2019), *Robots at Work. Automatable and Non Automatable Jobs*, IZA Discussion Paper n.12520, Bonn, IZA
- Katz L.F., Autor D. (1999), Changes in the wage structure and earnings inequality, in Ashenfelter O.C., Card D. (eds.), *Handbook of labor economics Vol. 3A*, New York, Elsevier, pp.1463-1555

- Katz L.F., Murphy K.M. (1992), Changes in relative wages, 1963-1987. Supply and demand factors, *The Quarterly Journal of Economics*, 107, n.1, pp.35-78
- Klenert D., Fernandez-Macias E., Anton J.I. (2020), *Do robots really destroy jobs? Evidence from Europe*, JRC Working Papers Series on Labour, Education and Technology n.1, Seville, European Commission - Joint Research Centre <<https://bit.ly/3xyKfF0>>
- Koch M., Manuylov I., Smolka M. (2021), Robots and Firms, *The Economic Journal*, article in press
- McKinsey Global Institute (2017), *A future that works. Automation, employment, and productivity*, New York, McKinsey & Company <<https://mck.co/3AMyyML>>
- Marcolin L., Miroudot S., Squicciarini M. (2018), To be (routine) or not to be (routine), that is the question. A cross-country task-based answer, *Industrial and Corporate Change*, 28, n.3, pp.477-501
- Mokyr J., Vickers C., Ziebarth N.L. (2015), The History of Technological Anxiety and the Future of Economic Growth. Is This Time Different?, *Journal of Economic Perspectives*, 29, n.3, pp.31-50
- Nedelkoska L., Quintini G. (2018), *Automation, skills use and training*, OECD Social, Employment and Migration Working Paper n.202, Paris, OECD
- Paba S., Solinas G., Bonacini L., Fareri S. (2020), Robot, ICT e globalizzazione. Gli effetti sui sistemi locali del lavoro in Italia, *L'industria. Rivista di economia e politica industriale*, XLI, n.1, pp.71-100
- Pfeiffer S. (2016), Robots, industry 4.0 and humans, or why assembly work is more than routine work, *Societies*, 6, n.2, pp.1-26
- Pratt G.A. (2015), Is a Cambrian Explosion Coming for Robotics?, *Journal of Economic Perspectives*, 29, n.3, pp.51-60
- Spitz-Oener A. (2006), Technical change, job tasks, and rising educational demands. Looking outside the wage structure, *Journal of Labor Economics*, 24, n.2, pp.235-270
- Webb M. (2020), *The impact of artificial intelligence on the labor market*, Stanford CA, Stanford University <<https://bit.ly/3AO65X2>>

