Robots at Work*

Georg Graetz
Uppsala University†
Guy Michaels
London School of Economics‡

February 27, 2015

Abstract

Despite ubiquitous discussions of robots' potential impact, there is almost no systematic empirical evidence on their economic effects. In this paper we analyze for the first time the economic impact of industrial robots, using new data on a panel of industries in 17 countries from 1993-2007. We find that industrial robots increased both labor productivity and value added. Our panel identification is robust to numerous controls, and we find similar results instrumenting increased robot use with a measure of workers' replaceability by robots, which is based on the tasks prevalent in industries before robots were widely employed. We calculate that the increased use of robots raised countries' average growth rates by about 0.37 percentage points. We also find that robots increased both wages and total factor productivity. While robots had no significant effect on total hours worked, there is some evidence that they reduced the hours of both low-skilled and middle-skilled workers.

KEYWORDS: Robots, Productivity, Technological Change.

JEL CLASSIFICATION: E23, J23, O30.

*Arpita Patnaik provided excellent research assistance. We thank Andy Feng for pointing us towards the World Robotics database. We also thank Francesco Caselli, Alan Manning, Steve Pischke, John Van Reenen, and seminar participants at the Institute for Evaluation of Labour Market and Education Policy, Uppsala, for their helpful suggestions. We are grateful to the Centre for Economic Performance for its financial support. The usual disclaimer applies.

†Economics Department, Uppsala University, P.O. Box 513, 75120 Uppsala, Sweden. Email: georg.graetz@nek.uu.se. Tel: +46(0)18-471-1594

‡Economics Department and Centre for Economic Performance, LSE, Houghton Street, WC2A 2AE, London, UK. Email: g.michaels@lse.ac.uk. Tel: +44(0)20-7852-3518.

1 Introduction

Robots' capacity for autonomous movement and their ability to perform an expanding set of tasks have captured writers' imaginations for almost a century.[1] But more recently, robots have emerged from the pages of science fiction novels into the real world, and discussions of their possible economic effects have become ubiquitous. For example, from 1990-2000, the fraction of items on the New York Times website mentioning the word "robots" almost doubled from 0.22% to 0.39%. This figure more than doubled again from 2000-2010 (reaching 0.83%) and again from 2010-2013 (reaching 1.76%), reflecting a growth that has been faster than exponential. But despite this rapidly growing interest, there is almost no hard evidence on the global economic impact of robots.

The paucity of evidence on robots' economic impact contrasts with the accumulating evidence on other new technologies, such as information and communication technologies (ICT), including computers.[2] Substantial gains from ICT have been documented at the firm level (Basker, 2012; Bloom, Sadun, and Van Reenen, 2012; Brynjolfsson and Hitt, 2000; Doms, Jarmin, and Klimek, 2004). At the level of industries or countries it appeared initially to be difficult to detect the impact of ICT (Solow, 1987). Stiroh (2002) presents evidence that ICT production and use are associated with faster productivity growth in US industries, and O'Mahony and Timmer (2009) estimate the contribution of ICT to EU and US aggregate labor productivity growth form 1995-2005 at 0.6 and 1.0 percentage points, respectively, applying standard growth accounting. But recent work on the US finds that gains in productivity are concentrated in ICT-producing industries, and not in ICT-using industries (Acemoglu, Autor, Dorn, Hanson, and Price, 2014). At the same time, the macro literature has been concerned with the possibility that productivity gains from technology in general may have slowed down. Gordon (2012, 2014) expresses a particularly pessimistic view, and there are broader worries about secular macroeconomic stagnation (Summers, 2014; Krugman, 2014), although others remain more optimistic (Brynjolfsson and McAfee, 2014). None of these works, however, provides direct evidence on the productivity effects of robots. Against this backdrop, the main goal of our paper is to provide the first systematic evaluation of the effect of industrial robots on productivity, which we estimate using variation over time across countries and industries.

In addition to studying robots' effects on productivity, we also shed light on the concerns that they might have a negative effect on employment. Fears that technological innovations destroy jobs are not new, and the Luddites' destruction of machines during the early nineteenth century is a striking example (Hobsbawm, 1952). A growing literature has studied the effects on labor demand of ICT in general, but not of robots.[3] In recent work, Brynjolfsson and McAfee

[1] Discussions of automata and physical construction of working machines go back to the ancient world. But according to the Oxford Online Dictionary, the word "robot" comes from robota, the Czech word for 'forced labor'. The term was coined in Čapek's 1920 play R.U.R. 'Rossum's Universal Robots' (`http://www.oxforddictionaries.com/definition/english/robot`, accessed on Dec 8, 2014). Robots gained in popularity following the work of Asimov (1950).

[2] There is some overlap between ICT (software, computing and communications equipment) and robots, since the latter typically feature computing equipment for programming and control. But most of the hardware components of robots are not considered ICT.

[3] For evidence on the labor market effects of ICT, see for example Autor, Katz, and Krueger (1998), Autor, Levy,

(2014), Ford (2009), and Frey and Osborne (2013) argue that in the future robots will likely replace many existing jobs. These concerns have been exacerbated by the evidence that labor's share of national income has been falling (Karabarbounis and Neiman, 2014; Elsby, Hobijn, and Sahin, 2013). At the same time, disagreements about the potential effects of robots on the labor market are presently common even among experts in the field (Pew Research Center, 2014). The second goal of our paper is therefore to assess the impact of robots on hours worked in industries that employ them across the developed world.

Specifically, we study the impact of industrial robots, utilizing new data from the International Federation of Robotics (IFR). The IFR measures deliveries of "multipurpose manipulating industrial robots", based on the definitions of the International Organization for Standardization (ISO), which allow us to compare delivery numbers across country-industry pairs and over time. Specifically, the IFR definition refers to a "Manipulating industrial robot as defined by ISO 8373: An automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications" (IFR, 2012, see also ISO definitions here: `https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en`).[4]

Using data from the International Federation of Robotics (2006), we estimate that from 1990-2005 the price of robots in six major developed economies fell by approximately one half. When quality improvements are taken into account, the fall in prices was even steeper: by 2005 quality-adjusted robot prices were about one fifth of their 1990 level.

This rapid decline in robot prices led to increased utilization of robots (which we dub "robot densification") in a range of different industries. We use International Federation of Robotics (2012) and EUKLEMS (Timmer, van Moergastel, Stuivenwold, Ypma, O'Mahony, and Kangasniemi, 2007) data to estimate robot density (the stock of robots per million hours worked) in 14 industries in 17 countries from 1993-2007. Averaged across the 17 countries in our dataset, robot density increased over this period by more than 150 percent, from 0.58 to 1.48. Among the countries in our dataset, robot densification was particularly pronounced in Germany, Denmark, and Italy. Among the industries, transport equipment, chemicals, and metal industries led the way in increasing robot density.

Using our panel data, we find that industry-country pairs which saw more rapid increases in robot density from 1993-2007, experienced larger gains in labor productivity (value added per hour worked) and similarly sized increases in value added per worker. These results are robust to controlling for country-specific and industry-specific fixed effects, which in our long-differenced specifications control for country and industry trends. At the same time, we find that larger increases in robot density translated into increasingly small gains in productivity, suggesting that there are some congestion effects (or diminishing marginal gains) from increased

and Murnane (2003), Acemoglu and Autor (2011), Michaels, Natraj, and Van Reenen (2014), Goos, Manning, and Salomons (2014), Akerman, Gaarder, and Mogstad (2013). A growing literature analyzes theoretically the impact of increased automation on the economy (Benzell, Kotlikoff, LaGarda, and Sachs, 2015; Feng and Graetz, 2015; Hemous and Olsen, 2014), but does not provide empirical evidence on the impact of robots.

[4]Besides industrial robots, the IFR also started reporting on service robots in 2002. However, service robots were then still in their infancy, and the IFR does not provide country-industry-level data on service robots during the period we analyze.

use of robots.

We provide further evidence that robots increase productivity using a novel instrumental variable strategy. To construct our instrument we use data on "robot applications" (IFR, 2012), which classify the tasks performed by robots. We match these to data on US occupations in 1980, before robots became ubiquitous, and define occupations as "replaceable" if by 2012 their work could have been replaced, completely or in part, by robots. We then compute the fraction of each industry's hours worked in 1980 that was performed by occupations that subsequently became prone to replacement by robots. Our industry-level "replaceability" index strongly predicts robot densification when robot prices fell sharply from 1993-2007. Two-stage least squares (2SLS) estimates using replaceability as an instrument for robot densification show that increased use of robots raised both labor productivity and value added, consistent with our OLS estimates. These results are again largely unchanged when we allow for country-specific trends in outcomes.[5]

Both our OLS and 2SLS results are robust to a large set of specification checks involving alternative measures of robot use, different sub-samples, and controls for the changes in other inputs.

We calculate that on average across the 17 countries in our dataset, robot densification from 1993-2007 raised annual growth of GDP and labor productivity by 0.37 and 0.36 percentage points, respectively.[6] This figure is a conservative one as it is based on the lower range of our estimates of the impact of robots. The figure is fairly comparable to the estimated total contribution of steam technology to British annual labor productivity growth of around 0.35 percentage points, which was, however, sustained over a period that was about four times longer, from 1850-1910 (Crafts, 2004). The overall contribution of robots is less than the upper range of estimates of ICT's contribution to EU and US labor productivity growth form 1995-2005, which O'Mahony and Timmer (2009) estimate at 0.6 and 1.0 percentage points, respectively. But importantly, the total value of ICT capital services is at least five times larger than that of robot services.[7]

We next turn to our second question—how did robot densification affect employment? Our OLS and 2SLS estimates show no significant effect of robot densification on aggregate hours worked, although some of the estimates are negative and close to significant. When we look at the effects of robots on the hours worked by different skill groups, we find some evidence, especially in the 2SLS specifications, that robots reduced hours worked by low-skilled and (to a lesser extent) middle-skilled workers. We also find that robots had no effect on hours worked by high-skilled workers. These results are consistent with viewing robotics technology as skill biased, at least in relative terms. At the same time, we find that unlike ICT, robots do not polarize the labor market, since their negative effects on the least educated are no smaller than those on the middle-skilled.[8]

[5]Since our instrument only varies at the industry level, we cannot allow for industry-specific trends in our 2SLS estimates.

[6]Average annual growth in GDP and labor productivity was 3.14 percent and 2 percent, respectively.

[7]The contribution of robots to growth is also less than that of post-war road construction in the US, which Fernald (1999) estimates at 1 percent for the period 1953-1973. However, the value of the road stock is much larger than that of robots, at about one quarter of private business GDP in 1994.

[8]For related discussions of the effect of ICT on skill demand, see e.g. Michaels, Natraj, and Van Reenen (2014) and

Finally, we estimate the effect of robots on other outcomes. Our estimated effects of robot densification on total factor productivity (TFP) and average wages are positive and statistically significant in most specifications. We do not, however, find any significant effect of robot densification on the labor share.

During the period we analyze, industrial robots were used in just under a third of the economy (as averaged across the countries in our dataset) and service robots were still in their infancy. This means that there is plenty of potential for increased use of robots in new industries. Moreover, as new robot capabilities are developed, they may be used more intensively in the industries that are already using them. This suggests that the likely contribution of robots on future growth is substantial. At the same time, our finding of congestion effects in robot use suggests that ever increasing robot densification is not a panacea for growth.

2 A Model of Production Using Robots and Workers

We use a simple two-sector model to motivate our empirical analysis and guide the interpretation of the results. In a static, frictionless environment there are two sectors, one using a CES technology to combine robots and labor, and another using labor only. Let Y_R and Y_N denote the outputs of the robot-using and non-robot-using sectors, respectively. Suppose that consumer have CES utility and that all outputs are consumed, hence consumers maximize $U = \left[Y_R^{\frac{\varepsilon-1}{\varepsilon}} + Y_N^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}$ subject to a budget constraint; ε is the elasticity of substitution in consumption between the outputs of the two sectors.[9]

As we further discuss below, we find it useful to allow for two distinct interpretations of the two sectors. The two sectors could represent two different industries within one country, or we could think of them as representing the same industry located in two different countries. The two interpretations will have different implications for the value of the elasticity of substitution ε.

Denote the quantity of robots used in production by R, and suppose that robots are perfectly elastically supplied at an exogenous rental price ρ. A fixed amount L of labor is supplied perfectly inelastically and labor is perfectly mobile across sectors. The labor inputs in the two sectors are denoted by L_R and L_N, respectively. The production functions in the two sectors are $Y_R = \left[R^{\frac{\sigma-1}{\sigma}} + L_R^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$ and $Y_N = L_N$, where the elasticity of substitution between robots and labor in the robot-using industry is denoted by σ.

As we demonstrate in the Model Appendix, the model gives rise to the following predictions:

1. A fall in the price of robots leads to an increase in robot density.

2. An increase in robot density leads to an increase in average labor productivity in the robot-using sector.

Autor (2014).

[9]For simplicity, we assume homothetic preferences, which rules out the possibility that income growth affects relative demand.

3. As robot density increases, so does the output in the robot-using sector relative to the non-robot-using one.

4. The labor input in the robot-using sector increases (decreases) with robot density if $\varepsilon > \sigma$ ($\varepsilon < \sigma$). The allocation of labor across sectors does not depend on robot density when $\varepsilon = \sigma$.

The intuition for the last result may be stated as follows. A decline in robot prices induces firms to substitute robots for labor, but at the same time increases output in the robot-using sector, depressing its relative price. This in turn leads consumers to buy relatively more of the robot-using sector's output. Whether the increased output is met by the increase in the robot input or whether an inflow of workers is required, depends on whether the firm's response to the fall in the price of robots is stronger than that of consumers to the fall in the relative output price, as measured by σ and ε.

We interpret the elasticities σ and ε as reduced-form rather than structural parameters. For example, σ captures the substitution possibilities between factors at the level of a task—which may be performed by either robots or workers—as well as the substitution possibilities between the various tasks that are needed to produce a sector's output. A very small substitution elasticity between tasks may translate into a very small σ, even if robots and workers are perfect substitutes at the task level.[10]

Similarly, the value of ε may depend on which interpretation of our two-sector model is more relevant. If we think of the two sectors as two different industries in one country, say automotive and construction, then substitution possibilities may be limited and ε may be small. However, if we interpret the two sectors as the same industry located in different countries, say automotive in the US and in Germany, then substitution may be very strong and hence ε may be large.[11] We believe that both interpretations may have some empirical relevance in our context.[12]

To sum up, our simple model suggests that increases in robot density caused by a fall in the price of robots should lead to a rise in value added and labor productivity, while the effect on hours is ambiguous.

3 Data Description

Our main source of data on robots is the International Federation of Robotics (2012), which compiles information from national robot federations on industrial robots. The IFR measures deliveries of "multipurpose manipulating industrial robots" based on the definitions of the International

[10]See for instance Zeira (1998) and its treatment in Acemoglu (2010, in particular Equation (20)).

[11]In this case the substitution elasticity in consumption would be large but not necessarily infinite, since there may still be product differentiation across countries.

[12]Increased robot use could also facilitate higher product quality and/or variety, and this could positively affect product demand and employment. A further simplification of our model is that we assume labor to be homogenous. One could imagine a model with heterogenous labor and potentially different degrees of substitutability between robots and the various types of labor. If low-skilled labor were more substitutable than high skill labor, then low-skilled workers would be more likely to see their labor input decline in the robot-using sector as a result of robots becoming cheaper.

Organization for Standardization. Their definition refers to a "Manipulating industrial robot as defined by ISO 8373: An automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications" (International Federation of Robotics, 2012, see also ISO definitions here: `https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en`). Each element of the definition is essential for a machine to be considered an industrial robot. For instance, a manipulator that is not reprogrammable or that has a single purpose is not considered an industrial robot.[13]

Typical applications of industrial robots include assembling, dispensing, handling, processing (for instance, cutting) and welding, all of which are prevalent in manufacturing industries; as well as harvesting (in agriculture) and inspecting of equipment and structures (common in power plants).

The IFR provides data on the number of robots delivered to each industry, in each country and year.[14] We construct the stock of robots based on the deliveries using the perpetual inventory method, assuming a depreciation rate of ten percent.[15] This approach is similar to the EUKLEMS procedure for computing the stock of ICT capital. We set the initial (1993) value of our stock measure equal to the corresponding estimate of the robot stock provided by the IFR.[16]

One attraction of our measure of the industrial robot stocks, is that the IFR criterion for counting these robots is fixed across industries and over time. Nevertheless, industrial robots are admittedly heterogeneous (as are workers' hours), and perhaps more importantly their technology has changed over time. To mitigate this concern, we construct an alternative measure of "robot services", which we use in our robustness checks, and in order to compare robot inputs to other capital inputs. Our robot services variable is calculated using turnover-based prices, which are available for our entire sample period but only for the US (see the Data Appendix for details). For selected countries the IFR also report price indices based on list prices, but these stop in 2005. List prices, together with data on changes in characteristics of robots, enabled the IFR to construct quality adjusted price indices, as well. We report these indices in Figure 1. Comparing the raw indices with the quality-adjusted ones reveals substantial quality improvements over time. While we are unable to directly observe how the quality of the robot stock changes, we check that our results are robust to assuming different depreciation rates when constructing the robot stock. A higher depreciation rate means that more weight is put on recent robot vintages, which may capture some of the quality growth.

[13]ISO defines a manipulator as a "machine in which the mechanism usually consists of a series of segments, jointed or sliding relative to one another, for the purpose of grasping and/or moving objects (pieces or tools) usually in several degrees of freedom... A manipulator can be controlled by an operator, a programmable electronic controller...".

[14]The IFR aims to capture the universe of robot suppliers: "The statistical data collected in the present World Robotics are based on consolidated data provided by nearly all industrial robot suppliers world-wide" (IFR, 2012, p.19).

[15]We explore different depreciation rates ranging from 5-15 percent, and find that our results are robust (Section 4.2).

[16]The IFR's estimates of robot stocks are based on the assumption that the service life of a robot is exactly twelve years—hence the estimated number of robots in use at a given point in time is equal to the sum of robots delivered over the past twelve years. While we prefer to use a measure of the robot stock that is based on more conventional assumptions about depreciation, we must rely on the IFR estimates to initialize our series of robot stocks.

A complicating feature of the IFR data is that for about half of the countries in our final sample, a breakdown of deliveries by industries is not available for the early years in the sample, when all delivered units are reported under the "unspecified" category. For these countries, we impute industry-level deliveries based on an industry's share in total deliveries during the years when the breakdown was reported (see the Data Appendix for details). As we discuss below, we find that our results are robust to excluding the countries for which we imputed robot deliveries in the base year.

Our second major source of data for this paper is EUKLEMS (Timmer, van Moergastel, Stuivenwold, Ypma, O'Mahony, and Kangasniemi, 2007). These data include information on inputs (including breakdowns of capital and labor aggregates), outputs, and prices at the industry-country-year level. We use data from the EUKLEMS March 2011 update for value added, hours worked, capital and labor compensation, breakdown of the capital input; and the EUKLEMS March 2008 release for the breakdown of the labor input.[17] IFR and EUKLEMS data use different industry classifications at varying levels of aggregation. The most detailed breakdown of EUKLEMS industries that allows us to consistently match the IFR data is shown in Appendix Table A1. Of the 28 EUKLEMS industries we are able to match 14. These include all manufacturing industries (except "machinery, not elsewhere classified") as well as agriculture, mining, utilities, construction, and "education and R&D". The IFR industries we do not use are "all other manufacturing", "all other non-manufacturing", and "unspecified". This means we lose about 16 percent of deliveries on average, mainly accounted for by the "unspecified" category.

The first year for most of our analysis is 1993, the first year covered in the IFR data, and the last year we use is 2007.[18] The IFR data end only in 2011, but coverage in the EUKLEMS data becomes extremely uneven after 2007. Furthermore, a virtue of omitting post-2007 data from the analysis is that this way our results are not influenced by the large cyclical fluctuations of the Great Recession and the subsequent recovery. The countries included in our sample are shown in Table 1.[19]

Nominal variables such as value added at current prices or compensation of labor and capital are reported in units of local currency in the EUKLEMS data. When comparing these variables across countries, we convert them to US$ using annual nominal exchange rates from the Penn World Table, Version 8.0 (Feenstra, Inklaar, and Timmer, 2013). We measure real variables in 2005 US$.

All the descriptive statistics and estimates that we report, unless noted otherwise, are weighted by an industry's initial (1993) share of hours in the country-wide amount of hours worked. We do this to ensure that our estimates reflect the relative importance of industries

[17]Our choice of EUKLEMS releases ensures that for each set of variables we use the most recent version.

[18]We use EUKLEMS data going back to 1979 for a falsification exercise.

[19]The major robot-using countries covered in the IFR data but dropped from our sample include China and Japan. While China is absent from the EUKLEMS data, basic variables such as output and hours worked can be found in World KLEMS (Wu, 2012). However, these data start only in 1999, therefore we were forced to omit China from our sample. We drop Japan because the reported deliveries and stocks feature implausible trend breaks which are due to substantial reclassification of what kinds of robots were to be included in the data, as we learned from personal communication with the IFR.

within countries. However, we give equal weight to each country. Our weighting procedure is the same as that of Michaels, Natraj, and Van Reenen (2014).

Our main dependent variables are the growth in real value added, hours worked, and productivity. More precisely, for each country-industry cell we compute the log difference of both real value added and hours between 1993 and 2007. We define labor productivity as the ratio of real value added to hours worked, and hence its growth is equal to the difference between the growth in value added and the growth in hours. Besides the three main outcomes we also consider growth of TFP and average wages, changes in the labor share, as well as growth in hours worked by three different skill groups: high, middle, and low. High-skilled workers include those with a college degree and above. While there is some variation across countries owing to differences in educational systems, the group of middle-skilled workers usually consists of high school graduates, people with some college education, and those with nonacademic professional degrees (Timmer, van Moergastel, Stuivenwold, Ypma, O'Mahony, and Kangasniemi, 2007, see pp.28-31).

The main regressor in our empirical analysis is based on our measure of robot density, which we define as the number of robots per million hours worked.[20] We refer to changes in robot density over time as "robot densification". As we discuss further below, changes in robot density are heavily concentrated at small positive values (or zero), while the distribution has a long right tail. This makes fitting a linear model using raw changes in density challenging. In our main analysis we therefore use as regressor the percentile of changes (based on the weighted distribution of changes). We verify in robustness checks that our results are mostly robust to a range of alternative functional forms.

Our instrument for robot densification is an industry-level measure that we call replaceability. We construct this instrument based on data from IFR on robot applications, the US Census occupational classifications, and the distribution of hours across occupations and industries from the 1980 US Census (Ruggles, Alexander, Genadek, Goeken, Schroeder, and Sobek, 2010). The IFR distinguishes among different applications of robots, including (among others) welding, processing, and assembling (IFR, 2012, p.33). We take the 2000 Census three-digit occupations and assign a replaceability value of one to an occupation if its name contains at least one of the IFR application categories.[21] Otherwise an occupation gets assigned a value of zero. We then map our replaceability measure into the 1990 Census occupational classification, which is available across multiple Census years, including the 1980 and 2000 censuses. If several 2000 occupations map into one 1990 occupation, then we assign the 1990 occupation a replaceability value of one if and only if at least one of the corresponding 2000 occupations has a value of one.

To measure replaceability at the industry level, we assign each individual in the 1980 IPUMS Census file a replaceability value based on their reported 1990 occupation. Furthermore, we assign each individual one of our 28 EUKLEMS industries based on a cross walk to the 1990

[20] In this definition we largely follow the IFR, except that the IFR define robot density as the ratio of the number of robots to workers. We prefer to use hours to normalize the number of robots, since workers can differ in the number of hours that they work.

[21] We also considered using older Census occupational classifications to construct our instrument. Given the changes in occupational terminology over time, we found that the matching of occupation names and robot applications works much better when using a classification that is more contemporaneous with the IFR report.

Census industry classification. We then compute the fraction of replaceable hours for each of the 14 robot-using industries by dividing the sum product of replaceability and annual hours worked by the total sum of hours worked (applying person weights both when computing the numerator and the denominator).[22]

4 Empirical Analysis

In the previous section, we have described the construction of our data containing information on value added, labor and capital inputs, and critically, robot use, for 238 country-industries between 1993 and 2007. We now turn to the empirical analysis. We first document the increased use of industrial robots across countries and industries; present our data graphically; and describe the results on the impact of robots on productivity, value added, and hours growth from our preferred OLS and 2SLS specifications in Section 4.1. We then examine the robustness of these results to using alternative functional forms, different sub-samples, and to the inclusion of additional controls in Section 4.2. We evaluate the economic importance of robots by quantifying their contribution to aggregate growth in Section 4.3. Finally, we present results for additional outcome variables, such as TFP and wage growth, and hours growth by skill group, in Section 4.4.

4.1 Main results

We begin our empirical analysis by describing the use of robots in the seventeen countries we analyze. These include the US, fourteen European countries, South Korea, and Australia. As panel A of Table 1 shows, in 1993, the average robot density (robots per million hours worked) in our sample was 0.58.[23] Robot densities, in robots per million hours worked, were highest in Germany (about 1.7), followed by Sweden (about 1.4), Belgium (1.2) and Italy (about 1.1). The figure for the US was just above two-thirds of the 17-country average. Four of the seven countries (Australia, Greece, Hungary, and Ireland) had either no industrial robots or almost none.[24] Table 1 also reports means for our main dependent variables, the logarithms of value added per million hours, value added, and millions of hours worked, and for the logarithms of the capital stock and the wage bill.

Panel B of Table 1 reports mean changes by country in robot densities from 1993-2007. The leading country was again Germany (about 2.7), followed by Denmark (about 1.6) and Italy

[22]We also computed a replaceability measure based on the wage bill. We calculated the fraction of the wage bill that is replaceable in the same way as replaceable hours, except that we substitute annual wage and salary income for hours worked. Since wage income is top-coded in the data, we assigned each top-coded observation the product of 1.5 and the maximum value of income among the non-top-coded observations. We do not report results using the wage bill instrument, however they are qualitatively and quantitatively very similar to those we do report. We prefer using replaceable hours as an instrument as this gives us a stronger first stage.

[23]Recall that the means we report are averaged within country using base period employment shares as weights, but the average across countries is unweighted.

[24]For most countries, a fraction of robot deliveries (typically less than 20 percent) is always classified as "unspecified", and is thus not part of our analysis. This means that we underestimate true robot densities. The fraction of "unspecified" deliveries is particularly large even in 2007 for Australia (82 percent) and Ireland (56 percent). These countries belong to the group for which we need to impute base-line robot stocks, and we show below that our results are robust to excluding this group.

(about 1.4). By 2007 industrial robots were employed in all the seventeen countries in our sample. The most striking fact from Table 1 is that from 1993-2007, mean robot density across the seventeen countries that we analyzed increased by more than 150 percent.

The most natural explanation for this rapid increase in robot intensity is the dramatic fall in robot prices. The International Federation of Robotics (2006) collects list prices of robots reported by surveyed companies in selected countries. Sub-figure (a) of Figure 1 shows that from 1990-2005, the price of robots in the six countries for which we have aggregate annual price data (US, France, Germany, Italy, Sweden, and UK) roughly halved.[25] Even this dramatic fall, however, does not reflect the full change in robot prices. Sub-figure (b) of Figure 1 shows that on average across the six countries in our sample, quality-adjusted robot prices fell by almost 80 percent. Even if we restrict our attention to 1993-2005, the average decline in quality-adjusted robot prices was still around 50 percent.

In light of this rapid fall in robot prices, it is unsurprising that they were widely adopted not only across countries, but also across industries. Table 2 reports the means for the same variables as Table 1, except this time for each of the fourteen industries in our dataset. In 1993 the transport equipment and metal industries led in the use of robots, with about 5.4 and 2.4 robots per million hours worked, while construction, education, mining, and utilities had negligible robot densities. From 1993-2007 the fastest increase in the number of robots per million hours worked took place in the transportation equipment (about 8.1), chemical (about 3.3) and metal (about 1.7) industries.

Tables 1 and 2 reveal an uneven distribution of robot density, with more variation across industries than across countries. Appendix Table A2 further shows the skewness of the distribution of robots across country-industry pairs, which correspond to our observations. In 1993, the median country-industry in our country had a negligible robot density of 0.004 robots per million hours worked. Panel B of the table shows that the gains in robot density from 1993-2007 were again skewed, with a median of just over 0.02 and maximum of over 28. Only in ten observations did the robot density decline over our sample period.[26] Appendix Table A2 also reports similar summary statistics for our alternative measures of robot density, and again both the levels in 1993 and the changes until 2007 were skewed to the right.

Before presenting our regression analysis of the impact of industrial robots, we examine the relevant functional form. Figure 2 plots the change in the log of labor productivity from 1993-2007 against measures of increased robot use. In sub-figure (a), we plot the percentile of the change in robot density on the horizontal axis, as well as the fitted regression line.[27] The slope is positive and precisely estimated, and the distribution of data points around the fitted line suggests that the relationship between productivity growth and the percentile of robot densification is well approximated by a linear functional form. In sub-figure (b), we instead

[25]As we discuss in the data section, for reasons of data availability we use a turnover-based measure of prices, rather than the list-based one, in parts of our analysis.

[26]Robot stocks declined in twelve observations. In eight of these robot densities also declined. In four observations the robot stock declined but the density (weakly) increased, and in two observations the density declined but the stock (weakly) increased.

[27]Percentiles are based on the weighted distribution of changes in robot density, where within-country 1993 employment shares of a country-industry are used as weights.

plot changes in raw robot density on the horizontal axis, together with the fitted line. Here a linear functional form (though still positive and significant at conventional levels) seems much less adequate, and the estimated slope appears sensitive to a few observations near the top of the distribution of robot densification. While an approximation of the functional form in our theoretical model might suggest using changes in raw robot density, Figure 2 shows that using percentiles gives a much better fit. In most of the analysis below we use the percentile of robot densification as our main regressor, although in the robustness checks we report results using other functional forms.

In Figure 3 we examine graphically the relationship between robot intensity and changes in productivity and value added at the industry level. The figure plots simple means of the variables of interest by industry, averaged across all the countries in our data. Sub-figures (a) and (b) of Figure 3 suggest that industries with higher deciles of change in robot density experienced faster growth in productivity and value added.[28]

To support our choice of functional form, we begin our analysis by allowing the change in outcomes to vary by quartile of change in robot density. Specifically, we estimate regressions of the form:

$$\Delta Y_{ci} = \beta_1 + \sum_{j=2}^{4} \beta_{2j} \mathbb{1}\left\{\Delta(\#robots/hours)_{ic} \in (Q_{j-1}, Q_j]\right\} + \beta_3 controls_{ci} + \varepsilon_{ci}, \tag{1}$$

where ΔY_{ci} is the change in the outcome of interest, Y_{ci} in industry i in country c from 1993-2007, $\mathbb{1}\left\{\Delta(\#robots/hours)_{ic} \in (Q_{j-1}, Q_j]\right\}$ is an indicator for the change in robot density falling between quartiles $j-1$ and j, and ε_{ci} is the error term. Some of the specifications include $controls_{ci}$, which are country fixed effects and in some cases also industry fixed effects. Since the specification is estimated in changes, these fixed effects effectively absorb country (industry) specific trends.

We estimate our regressions on 1993-2007 changes, because we are interested in long-run trends. Including intermediate years would not necessarily increase the precision of our estimates since it would lower the signal-to-noise ratio. As discussed in the data section, we weight all the regressions using industries' base year shares of hours worked within each country.

Table 3 reports the results of our estimation of equation (1). Column (1) shows that without controlling for any trends, only the country-industry observations in the top quartile of robot densification experienced faster growth in labor productivity. Once we control for country fixed effects in column (2), the observations in the top quartile stand out even more, while the third quartile coefficient is positive and significant but smaller than the top quartile, and the second quartile coefficient is positive, smaller than the third quartile coefficient, and imprecisely estimated. The country-industries with fastest growth in robot density saw productivity grow about 3.2 percentage points faster per year.[29]

While the results suggest that increases in robot density are systematically associated with

[28] Figure 3 also suggests that the electronics industry is a bit of an outlier with its high growth of productivity and value added, but our results are robust to excluding this industry.

[29] 0.446 divided by 14 equals 0.032. The precise expression for calculating differences in annual growth rates is $e^{\beta/14} - 1$, which is well approximated by our simpler formula.

increases in labor productivity, the magnitudes are also interesting. As Appendix Table A2 shows, the mean change in robot density was over ten times higher in the top quartile than in the third quartile, which in turn was more than ten times higher than the average in the second quartile. Despite these large differences, our estimates which account for differences in country trends suggest that the marginal impact of increasing robot densities is diminishing.

Columns (3)-(4) of Table 3 show estimates of regressions as in columns (1)-(2), but this time using as an outcome the growth in value added. The picture is very similar to what we find for labor productivity, with larger increases in robot density corresponding to larger increases in value added. Conditional on country trends, the country-industries with fastest growth in robot density saw value added grow about 3.5 percentage points faster per year.

Finally, columns (5)-(6) show that robot densification was associated with almost no change in total hours worked: the estimated coefficients are all both small and imprecise.

Using our previous evidence on the functional form relationship between changes in robot density and the outcomes of interest, we proceed to estimate regressions of the form:

$$\Delta Y_{ci} = \gamma_1 + \gamma_2 \Delta robots_{ci} + \gamma_3 controls_{ci} + \varepsilon_{ci}, \tag{2}$$

where $\Delta robots_{ci}$ is the percentile change in robot density.

Panel A of Table 4 shows that moving from the bottom to the top of the ranking of changes in robot density distribution corresponds to an increase of about 0.36-0.57 in the logarithm of labor productivity (which translates into increases in annual growth of 2.6-4.1 percentage points), depending on the set of controls. The corresponding estimates for changes in the logarithm of value added are very similar (columns (3)-(4)), as country-industries with higher growth in robot densities experienced no significant changes in hours worked (columns (5)-(6)).

To check whether these results could be driven by unobserved differences between industries, we include industry fixed effects, which allows us to control for industry-specific trends, given that we estimate first differences. This is a very demanding specification since it may be precisely the differences in production processes *between* industries that are important in explaining increased robot use. Panel A of Appendix Table A3 shows that the relationships between productivity and value added growth, and robot densification, are still positive but less steep and less precisely estimated in the non-parametric specification. Panel B shows OLS with and without industry trends. Controlling for industry trends leads to smaller estimates of the effects of increased robot use on productivity and value added of 0.35 and 0.36, respectively. Unlike in the non-parametric specification, these estimates are precise and statistically significant. Controlling for industry trends does not change the finding that robot densification appears to have no effect on hours worked.

The robust pattern that we document in Table 3 and panel A of Table 4, where robot densification is associated with increased labor productivity and value added even after controlling for country trends (and for industry trends in Appendix Table A3), is strongly suggestive. Nonetheless, we may be concerned about the interpretation of the estimates for a few reasons. First, we might worry about attenuation bias due to measurement error in the changes in robot densities. This is not a trivial concern, given our discussion above of the data construction, and the fact that we are estimating specifications in changes, which could worsen the signal-to-noise ratio compared to regressions on levels. Second, we might worry that the estimates for labor productivity are biased because we use the change in hours to construct both our dependent variable and the regressor of interest. We note that this does not apply to the value added regressions, but it is nonetheless a concern. Finally, we might worry about reverse causality, where faster-growing industries invest more in increasing their robot densities (robot suppliers may target their products to the industries they expect to grow fastest).

To address these concerns, we use our measure of replaceable hours ($replaceable_i$) as an instrument for the changes in robot density over time. As we explain in our data section, this measure reflects the percent of hours worked in each industry in the United States in 1980, by occupations whose tasks are (at least partly) replaceable by robots. Sub-figure (c) of Figure 3 shows that this measure strongly predicts the increase in robot intensity: as robot prices fell—both in absolute terms and relative to wages—industries with higher initial replaceability increased their use of robots more than others. Sub-figures (d) and (e) show that those same industries also experienced faster increases in productivity and value added.

The bottom panel of Table 4 uses our replaceability index as an instrument for the change in robot density. As column (1) shows, we have a precisely estimated first stage, and our two stage least squares (2SLS) estimate of the effect of robot densification is roughly 50 percent larger than the OLS estimate when we allow for differential country trends in column (2). Given our concerns about measurement error, this seems quite plausible.

Moving on to columns (3) and (4) we see that robot densification increases value added in the 2SLS specifications, and the results are quite similar to those we report using OLS. Finally, columns (5) and (6) show that using our 2SLS specifications, robot densification is associated with a reduction in working hours, although the estimates are not significant at the five percent level. This result, however, is noteworthy, and we revisit it in our discussion below.

Although our 2SLS estimates address potential concerns about measurement error and reverse causality, we might still worry that replaceable industries followed different trajectories even before they began to adopt robots. To mitigate this concern, Table 5 presents falsification tests for our replaceability instrument. The table shows reduced-form regressions of the type:

$$\Delta Y_{ci} = \delta_1 + \delta_2 replaceable_i + country_c + \varepsilon_{ci} \tag{3}$$

where $country_c$ is a set of country fixed effects. Panel A presents the reduced form estimates for our full sample (our benchmark), and as before we see that during the fourteen year period from 1993-2007, replaceable industries saw large and precisely estimated increases in labor productivity and value added, and a smaller and imprecise decline in hours worked. In panel B we restrict our sample to country-industries that did not use any robots (robot "non-adopters") in 1993. The coefficients for this sample are similar but imprecisely estimated. At the bottom of the table, we report the p-value from tests for equality of coefficients across the various panels. We cannot reject that the relationship between the outcome variables and the share of replaceable hours is the same in our benchmark and in the sample of 1993 non-adopters.

Next, in panel C of Table 5 we check whether growth in the outcome variables during the fourteen-year period from 1979-1993 was systematically related to replaceability in the 1993 non-adopting country-industries. We find that replaceable industries did not experience differential growth in productivity, value added, or hours before robot adoption began. The estimated coefficients on productivity and value added are much smaller than during the period 1993-2007, and imprecisely estimated. We can reject the equality of coefficients between our benchmark and those we consider in panel C at the five percent significance level (except for hours, where we find no significant effect in any of the specifications).

Lastly, panel D of Table 5 presents a different falsification check, where we restrict the sample to industries that did not yet adopt robots in 2007. This sample is small, comprising only 27 observations, so the estimates are even less precise. Nonetheless, they again suggest that replaceable industries do not follow systematically different trends before they start to employ robots.[30]

While the results in panel C of Table 5 suggest that replaceable industries did not follow different trends before the adoption of robots, we still check the robustness of our 2SLS estimates, and the OLS estimates, to controlling for changes in outcomes (value added and hours worked) during the preceding fourteen year period (1979-1993).[31] The estimated regressions are essentially the same as those reported in Table 4 with country fixed effects, and in some OLS specifications we also add industry fixed effects. The results, reported in Appendix Table A4, show that our estimates are robust to controlling for past trends. As before, we find no significant effect on hours worked, and the 2SLS estimate, which is negative and marginally significant without controls, becomes smaller and imprecise once the controls are added.

Our replaceability instrument is constructed at the industry level and thus it does not predict within-industry variation in robot densification. We therefore consider two alternative instruments: a country-industry's initial (1993) robot density, and a shift-share instrument which interacts initial robot density at the global industry level with aggregate, country-level robot densification (see the Data Appendix for details on the construction of these variables). Appendix Table A5 shows that these alternative instruments yield strong first stages, and the 2SLS estimates of the effects of robot densification on productivity growth are positive, precisely estimated, and similar in magnitude to the estimates using the replaceability instrument. The effects on hours are estimated to be positive but small in magnitude, and they are not statistically significant. When we enter all three instruments at the same time, results are again very similar, although initial robot density has no predictive power anymore in the first stage.[32]

Before moving on to further robustness checks involving different functional forms, different subsamples, and additional controls, we perform one more check on our baseline results. Because of the relatively modest number of countries and industries, we report robust standard

[30] The fact that some country-industries did not adopt industrial robots by 1993 and some not even by 2007 suggests that these country-industries might be special, and the relationships between outcomes and replaceability might not be the same for these observations as it would be for the rest of the sample in the absence of robots. However, the similarities of results in panels A and B mitigate this concern.

[31] We use logarithms for the outcomes, so labor productivity is the sum of value added and hours worked, and adding it would result in perfect multicollinearity.

[32] Since our replaceability instrument is constructed at the industry level, we cannot estimate 2SLS specifications with industry fixed effects due to perfect multicollinearity. While the alternative instruments do feature within-industry variation, we find that they have little power once we absorb industry trends.

errors that are not clustered. Appendix Table A6, however, reports estimates using the same specifications as Table 4, except that standard errors are clustered by country. The estimated standard errors do not change by much, and all our main estimates remain statistically significant.

4.2 Further robustness checks

Having shown our main results and robustness checks, we now proceed to report estimates from further checks.

The first set of further robustness checks concerns the choice of functional form and the choice of depreciation rate used when constructing the robot stock. While Figures 2 and 3 support our use of the percentile of robot densification as the preferred functional form, in Appendix Table A7 we nevertheless explore different functional forms for the change in robot density. In panel A we report estimates using the same specifications as in Table 4, except this time using the change in the number of robots per million hours worked as an alternative measure of robot densification. The OLS estimates without controls and those that allow for country trends show positive and significant effects of robot densification on value added and productivity, with little impact on hours worked. The 2SLS estimates for these two variables are similarly positive and precisely estimated. We also note that for this functional form the first stage F-statistic is not as precise as before, suggesting that this functional form is perhaps not as suitable, although we still do not suffer from a weak instrument problem.

Appendix Table A7 also reports estimates for two other functional forms for robot densification: the change in the logarithm of $1 + robots/[million\ hours\ worked]$ and the ratio of robot services to the wage bill (panels B and C). Once again the picture is similar to before: OLS and 2SLS estimates for labor productivity and value added are all positive and precisely estimated.[33] As in most of the previous analysis, we find small and imprecise OLS estimates of the effects of robot densification on hours worked, and somewhat negative but imprecise estimates in the 2SLS specifications.[34]

As the IFR (2012) points out, there is uncertainty regarding the average service life of robots, and therefore it is important to check whether our results are robust to a range of depreciation rates around our default choice of ten percent. Using a higher depreciation rate may also capture some of the growth in robot quality, since more weight is put on recent robot vintages. In panel A, columns (1) and (2) of Appendix Table A8 we present estimates from our preferred OLS and 2SLS specification of the effects of robot densification on productivity and hours growth, where the robot stock was computed assuming a depreciation rate of 5 percent. Columns (3) and (4) repeat the baseline results and columns (5) and (6) show results for a depreciation rate of 15 percent. The estimated effects on productivity growth are all positive, statistically significant, and of similar magnitudes, while for hours we again find that a zero effect cannot be rejected in any of the regressions. Panel B shows results using changes in density rather than the percentile, and again results are robust to different depreciation rates, although the magnitudes are more

[33] In the last functional form, however, our instrument is unfortunately rather weak.

[34] When we control for industry trends in the specifications using alternative functional forms of robot densification, coefficients in the productivity and value added regressions become small and are imprecisely estimated.

sensitive.

Another concern about the data, as mentioned in Section 3, is that constructing the robot stock for half of the countries in our sample required imputing deliveries in 1993, since for these countries deliveries were only reported at the aggregate level. Appendix Table A9, however, shows that excluding these countries from our analysis does not substantially change the results reported in Table 4.

There is potential heterogeneity in our sample not only because of differences in data quality, but also because of differences between industries in the nature of their products. In particular, our sample contains both tradable and non-tradable goods producing industries. In Appendix Table A10 we check whether restricting our sample to tradable industries affects our results. Columns (5) and (6) suggest that the impact of robot densification on value added growth is somewhat larger in tradable industries than in the full sample, both in the OLS and 2SLS specifications. While the effects on productivity appear reduced in magnitude (column (2)) and the effects on hours are estimated to be more positive (column (10)) in the tradable industries for our baseline specifications, controlling for industry trends in the OLS regressions yields no meaningful differences between the full sample and tradable industries, with positive effects on productivity and value added, and small and statistically insignificant effects on hours (columns (3)-(4), (7)-(8), (11)-(12)).

As one further check on whether the impact of robot densification is heterogeneous across countries and industries, we have run our baseline OLS and IV specifications excluding each industry at a time, and each country at a time. The results are available on request. They are highly robust except that in the regressions of changes in hours, the coefficients on robot densification are negative and statistically significant when we exclude agriculture. However, the coefficient in the OLS regression becomes attenuated and statistically insignificant when controlling for industry trends.

A potential concern that we have not addressed so far is that robot densification might be associated with changes in the use of other inputs at the industry level, which may affect our estimates. For instance, if robot densification involves skill upgrading, then there could be a spurious correlation between robot densification and productivity growth. Similarly, robot densification may be correlated with increased use of other types of capital, which affects productivity and value added growth. However, if other inputs change as a result of the fall in robot prices, then these changes are not pre-determined with respect to robot densification, and including them on the right-hand side is not without problems. Nevertheless, it is an interesting question whether our results are robust to including these additional controls.

In Appendix Table A11 we add to our main specifications from Table 4 controls for the changes in the fractions of high-skilled and low-skilled workers in total hours. Furthermore, we add changes in the log of the average wage as a control. This extreme robustness check is a way of making sure that the productivity gains associated with increased robot densification are not driven by the increased hiring of workers that are more productive along skill dimensions that we cannot otherwise measure directly. However, the wage is clearly an endogenous variable, and in fact we use it as an outcome in Section 4.4.

As columns (1)-(3) show, adding these controls leads to somewhat lower estimates of the impact of robot densification on productivity, but the estimates are still precise and statistically significant both for OLS and 2SLS specifications (in the OLS specifications, this holds even when including industry fixed effects, see columns (4)-(6)). The estimated effects on value added are not affected by controling for changes in the skill mix or the wage (columns (7)-(12)).

We further control for changes in the ratio of the value of capital services to the wage bill, and the share of ICT capital services in total capital services. Given that the EUKLEMS data aim to measure the entire capital stock, there is an overlap between our esimtates of the robot and capital inputs, and a partial overlap between robots and ICT in particular. Nonetheless, robots make up a small part of capital as a whole, and the overlap with ICT is far from perfect (since robots are largely made up of hardware that is not considered ICT), so we check that our results are not driven by changes to these larger input categories. As Appendix Table A12, columns (1)-(2) and (4)-(5) show, the estimated effects of robot densification on labor productivity and value added remain positive and statistically significant, and their magnitude is also robust. This remains true when we control for changes in the compositions of the capital and labor inputs at the same time (columns (3) and (6)).

Appendix Table A13 suggests that our OLS estimates are smaller in magnitude and imprecisely estimated when we control for both industry trends and changes in other inputs. However, as a comparison of columns (1) and (5) with columns (2) and (6) shows, this decline in magnitude and loss of precision is due to the reduction in sample size that comes from including measures of other capital, rather than due to the inclusion of industry trends.

4.3 Magnitudes

Having presented our main results and examined their robustness, we next consider the implications for aggregate labor productivity and value added of our estimates of the effects of robot densification. We consider a counterfactual scenario in which robot densities (robots per million hours worked) in 2007 would have remained the same as in 1993. We calculate how much lower labor productivity and value added would have been in this case.

To calculate counterfactual productivity and value added for this scenario, we proceed as follows. We first compute the 'zero-percentile', the percentile of changes in robot density that corresponds to no change, q_0. Let q_{ci} denote the actual percentile of the change in robot density in country c and industry i. For each outcome $Y \in \{\text{VA/H,VA}\}$ we then calculate its counterfactual log change as

$$(\Delta \ln Y_{ci})^{cf} = \Delta \ln Y_{ci} - \widehat{\beta}_Y \left(q_{ci} - q_0\right),$$

where $\widehat{\beta}_Y$ is the preferred estimate of the effect of robot densification on the outcome of interest. Using $(\Delta \ln Y_{ci})^{cf}$, we compute the counterfactual log values and levels of productivity and value added in 2007 for each country-industry. We then aggregate levels of productivity and value added to the country level, simply summing value added, but weighting productivity by an industry's 2007 share in total hours in the given country, obtaining $Y^{cf}_{c,2007}$. By comparing these numbers to the actual 2007 levels, we obtain an estimate of how much lower productivity and value added would have been in the absence of robot densification. In particular, we calculate the percentage loss as $100 \times (1 - Y^{cf}_{c,2007}/Y_{c,2007})$.

We base our analysis on the OLS estimates from the specifications that allow for both country and industry trends, setting $\widehat{\beta}_{\text{VA/H}} = 0.352$ and $\widehat{\beta}_{\text{VA}} = 0.358$.[35] Since these estimates are lower than our IV estimates, the results reported here may be viewed as conservative. The bottom row in Appendix Table A14 shows that the counterfactual loss in labor productivity for the robot-using industries implied by the OLS estimate is on average about 16 percent across countries, and similarly for value added. We calculate that countries with more rapid robot densification would experience a larger loss in productivity and value added in the absence of robot densification. The loss in both productivity and value added would have been highest for Germany and lowest for Hungary.

What effect have robots had on the aggregate economies? Assuming that no robots are used in the industries excluded from our sample, we obtain the loss in economy-wide productivity and value added by multiplying our figures for the robot-using industries by the share in value added of the robot-using industries in 2007. This share is typically around one third or less, and hence our estimates of losses in productivity and value added drop substantially. Still, we find that productivity and value added would have been about 5 percent lower in the absence of robot densification. This implies that robot densification increased annual growth of GDP and labor productivity by about 0.37 percentage points. This figure is fairly comparable to the estimated total contribution of steam technology to British annual labor productivity growth, which was around 0.35 percentage points, but was sustained over a period that was about four times longer, from 1850-1910 (Crafts, 2004). The overall contribution of robots is less than the upper range of estimates of ICT's contribution to EU and US labor productivity growth form 1995-2005, which O'Mahony and Timmer (2009) estimate at 0.6 and 1.0 percentage points, respectively. However, the total value of ICT capital services likely exceeds that of robot services by a factor of at least five.[36]

In sum, the contribution of robot densification to growth has been substantial, especially given the small share of robots in total capital.

4.4 Additional results

We have so far discussed our main set of results, showing that robot densification increased labor productivity and value added. We now turn to estimating the effect of robot densification

[35] These are taken from panel B, columns (2) and (4) of Appendix Table A3.

[36] Averaged across countries and the years 1993 and 2007, the share of robot services in total capital services is 0.64 percent (2.25 percent in robot-using industries), compared to 11 percent for ICT services (13 percent in robot-using industries). However, the IFR (2012, p.11) point out that their data on the value of the robot stock "do not include the cost of software, peripherals and systems engineering", and that the true value of the robot stock may be three times as large. A further difficulty in this context is that EUKLEMS data break down the capital stock into ICT and non-ICT, but robots are made of both ICT and non-ICT components.

The contribution of robots to growth is also less than that of post-war road construction in the US, which Fernald (1999) estimates at 1 percent for the period 1953-1973. However, the value of the road stock is much larger than that of robots, at about one quarter of private business GDP in 1994.

on other outcomes: total factor productivity (TFP), wages, the labor share, and on different skill groups' hours worked and wage bills.

We begin in Table 6 by reporting the results from estimating the same specifications as in Table 4 using the change in the logarithm of TFP as an outcome variable. The OLS estimates without controls and with country trends suggest that robot densification increased TFP. The 2SLS estimates—both with and without country trends—are also statistically significant and somewhat larger than the OLS estimates. These results are consistent with a large and positive effect of robots on productivity at a relatively low cost.

The estimated effects of robot densification on the change in the logarithm of mean wages is positive and significant in the OLS specifications (except the one that does not allow for country trends), and in both 2SLS specifications. Some of the productivity gains from robot densification appear to be shared with workers.

Table 6 also shows results using the change in the labor share as a dependent variable. The point estimates are negative and of non-trivial magnitudes, but given a lack of precision, we cannot reject a zero effect of robot densification on the labor share.

The results we discussed so far consider labor as one input. But did robot densification have different effects for workers with different levels of skill? To address this question, we report in Table 7 estimates using the hours and wage bills of high-skilled, middle-skilled, and low-skilled workers as dependent variables. The OLS estimates that allow for country trends suggest that the hours worked (and possibly also the wage bill) of skilled workers may have increased faster, and the 2SLS estimates suggest that growth in hours worked and the wage bill of low-skilled workers may have suffered from robot densification. Middle-skilled workers may have been adversely affected as well, but to a lesser extent than low-skilled workers. This result is noteworthy given the recent findings in the literature that some technological change is biased against middle-skilled workers (see for example the discussion of the effect of ICT in Michaels, Natraj, and Van Reenen (2014) and more general discussions in Goos, Manning, and Salomons (2014), Goos and Manning (2007) and Autor (2014)).

Appendix Table A15 provides further suggestive evidence that in relative terms the main losers from robot densification have been low-skilled workers. Panels A and B show that all the OLS and 2SLS estimates of robots' effects on low-skilled workers' share of hours are negative. The estimates for the two higher skill groups are mostly noisy (with the exception of one positive and significant estimate for high-skilled workers), making it difficult to tell precisely which workers gained in relative terms from robot densification.

Panels C and D of Appendix Table A15 paint a starker picture when it comes to wage bill shares. Here the majority of the estimates for high- and low-skilled workers are significant, providing evidence that robot densification shifts demand from the low-skilled towards the high-skilled. As before, there is no significant effect of robot densification on the wage bill share of middle-skilled workers.

5 Conclusions

We estimate, for the first time, the effects of industrial robots on economic outcomes. Using a panel of industries from 17 countries from 1993-2007, we find that industrial robots increased labor productivity and value added. We find that the contribution of increased use of robots to economic growth is substantial, and calculate using conservative estimates that it comes to 0.37 percent, accounting for just over one tenth of aggregate growth. This finding is robust to including various controls, for country and industry trends, for past growth, and for changes in the composition of labor and in other capital inputs. We also examine the effects of industrial robots on other economic outcomes. Specifically, we find that robot densification increased both total factor productivity and wages. While we find no significant effect of industrial robots on overall employment, there is some evidence that they crowd out employment of low-skilled and, to a lesser extent, middle-skilled workers.

As of 2007 industrial robots accounted for only around 2.25 percent of the capital stock in affected industries, and they penetrated only a limited part of the developed economies that we examine. If the quality-adjusted prices of robots keep falling at a rate similar to that observed over the past decades, and as new applications are developed, there is every reason to believe that they will continue to increase both labor productivity and value added. Recently, the development of robots has been increasingly directed towards services. Areas that are experiencing a particularly rapid expansion include medical robots, factory logistic systems, and unmanned aerial vehicles, popularly known as drones (IFR 2012, p.19).

Our analysis focused (due to data limitations) on developed economies. But recent evidence (Financial Times, 2014; International Federation of Robotics, 2014) shows that robots are increasingly used also in developing countries, and China may already be the worlds leading buyer of robots. So the contribution of robots to worldwide growth in the upcoming decades can be even larger.

At the same time, the evidence suggests that there are congestion effects, marginal returns on increased robot densification seem to diminish fairly rapidly. We also caution that the rise of robots is not a blessing for all: we find that low-skilled and middle-skilled workers in particular may lose out. And in the longer run, the findings of Frey and Osborne (2013) also demonstrate potential risks for a growing set of occupations.

References

ACEMOGLU, D. (2010): "When Does Labor Scarcity Encourage Innovation?," *Journal of Political Economy*, 118(6), 1037 – 1078.

ACEMOGLU, D., AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," vol. 4 of *Handbook of Labor Economics*, chap. 12, pp. 1043–1171. Elsevier.

ACEMOGLU, D., D. AUTOR, D. DORN, G. H. HANSON, AND B. PRICE (2014): "Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing," *American Economic Review*, 104(5), 394–99.

AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2013): "The Skill Complementarity of Broadband Internet," IZA Discussion Papers 7762, Institute for the Study of Labor (IZA).

ASIMOV, I. (1950): *I, Robot*. Gnome Press.

AUTOR, D. H. (2014): "Polanyis Paradox and the Shape of Employment Growth," NBER Working Papers 20485, National Bureau of Economic Research, Inc.

AUTOR, D. H., L. F. KATZ, AND A. B. KRUEGER (1998): "Computing Inequality: Have Computers Changed The Labor Market?," *The Quarterly Journal of Economics*, 113(4), 1169–1213.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content Of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118(4), 1279–1333.

BASKER, E. (2012): "Raising the Barcode Scanner: Technology and Productivity in the Retail Sector," *American Economic Journal: Applied Economics*, 4(3), 1–27.

BENZELL, S. G., L. J. KOTLIKOFF, G. LAGARDA, AND J. D. SACHS (2015): "Robots Are Us: Some Economics of Human Replacement," Working Paper 20941, National Bureau of Economic Research.

BLOOM, N., R. SADUN, AND J. VAN REENEN (2012): "Americans Do IT Better: US Multinationals and the Productivity Miracle," *American Economic Review*, 102(1), 167–201.

BRYNJOLFSSON, E., AND L. M. HITT (2000): "Beyond Computation: Information Technology, Organizational Transformation and Business Performance," *The Journal of Economic Perspectives*, 14(4), pp. 23–48.

BRYNJOLFSSON, E., AND A. MCAFEE (2014): *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.

CRAFTS, N. (2004): "Steam as a General Purpose Technology: A Growth Accounting Perspective," *The Economic Journal*, 114(495), pp. 338–351.

DOMS, M., R. JARMIN, AND S. KLIMEK (2004): "Information technology investment and firm performance in US retail trade," *Economics of Innovation and New Technology*, 13(7), 595–613.

ELSBY, M., B. HOBIJN, AND A. SAHIN (2013): "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, 47(2 (Fall)), 1–63.

FEENSTRA, R. C., R. INKLAAR, AND M. P. TIMMER (2013): "The Next Generation of the Penn World Table," available for download at www.ggdc.net/pwt.

FENG, A., AND G. GRAETZ (2015): "Rise of the Machines: The Effects of Labor-Saving Innovations on Jobs and Wages," CEP Discussion Papers dp1330, Centre for Economic Performance, LSE.

FERNALD, J. G. (1999): "Roads to Prosperity? Assessing the Link between Public Capital and Productivity," *American Economic Review*, 89(3), 619–638.

FINANCIAL TIMES (2014): "China becomes largest buyer of industrial robots," June 1.

FORD, M. (2009): *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future*. CreateSpace Independent Publishing Platform.

FREY, C. B., AND M. A. OSBORNE (2013): "The future of employment: how susceptible are jobs to computerisation?," *Sept*, 17, 2013.

GOOS, M., AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 89(1), 118–133.

GOOS, M., A. MANNING, AND A. SALOMONS (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104(8), 2509–26.

GORDON, R. J. (2012): "Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds," NBER Working Papers 18315, National Bureau of Economic Research, Inc.

——— (2014): "The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections," NBER Working Papers 19895, National Bureau of Economic Research, Inc.

HEMOUS, D., AND M. OLSEN (2014): "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," CEPR Discussion Papers 10244, C.E.P.R. Discussion Papers.

HOBSBAWM, E. J. (1952): "The Machine Breakers," *Past & Present*, (1), pp. 57–70.

INTERNATIONAL FEDERATION OF ROBOTICS (2006): "World Robotics Inudstrial Robots 2006," report.

——— (2012): "World Robotics Inudstrial Robots 2012," report.

——— (2014): "World Robotics Inudstrial Robots 2012," report.

KARABARBOUNIS, L., AND B. NEIMAN (2014): "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 129(1), 61–103.

KRUGMAN, P. (2014): *Four observations on secular stagnation* chap. 4, pp. 61–68, Secular Stagnation: Facts, Causes and Cures. CEPR Press.

MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years," *The Review of Economics and Statistics*, 96(1), 60–77.

O'MAHONY, M., AND M. P. TIMMER (2009): "Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database," *Economic Journal*, 119(538), F374–F403.

PEW RESEARCH CENTER (2014): "AI, Robotics, and the Future of Jobs," Available at: http://www.pewinternet.org/2014/08/06/future-of-jobs/.

RUGGLES, S., J. T. ALEXANDER, K. GENADEK, R. GOEKEN, M. B. SCHROEDER, AND M. SOBEK (2010): "Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]," Discussion paper, Minneapolis: University of Minnesota, 2010.

SOLOW, R. M. (1987): "We'd Better Watch Out," *New York Times*, review of Manufacturing Matters: The Myth of the Post-Industrial Economy, by Stephen S. Cohen and John Zysman.

STIROH, K. J. (2002): "Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say?," *American Economic Review*, 92(5), 1559–1576.

SUMMERS, L. H. (2014): *Reflections on the 'New Secular Stagnation Hypothesis'* chap. 1, pp. 27–38, Secular Stagnation: Facts, Causes and Cures. CEPR Press.

TIMMER, M., T. VAN MOERGASTEL, E. STUIVENWOLD, G. YPMA, M. O'MAHONY, AND M. KANGASNIEMI (2007): "EU KLEMS Growth and Productivity Accounts Version 1.0," mimeo, University of Groningen.

WU, H. X. (2012): "CIP Database Round 1.0: Technical Notes," mimeo, Institute of Economic Research, Hitotsubashi University.

ZEIRA, J. (1998): "Workers, Machines, And Economic Growth," *The Quarterly Journal of Economics*, 113(4), 1091–1117.

Figure 1: The Price of Robots in Six Countries 1990-2005

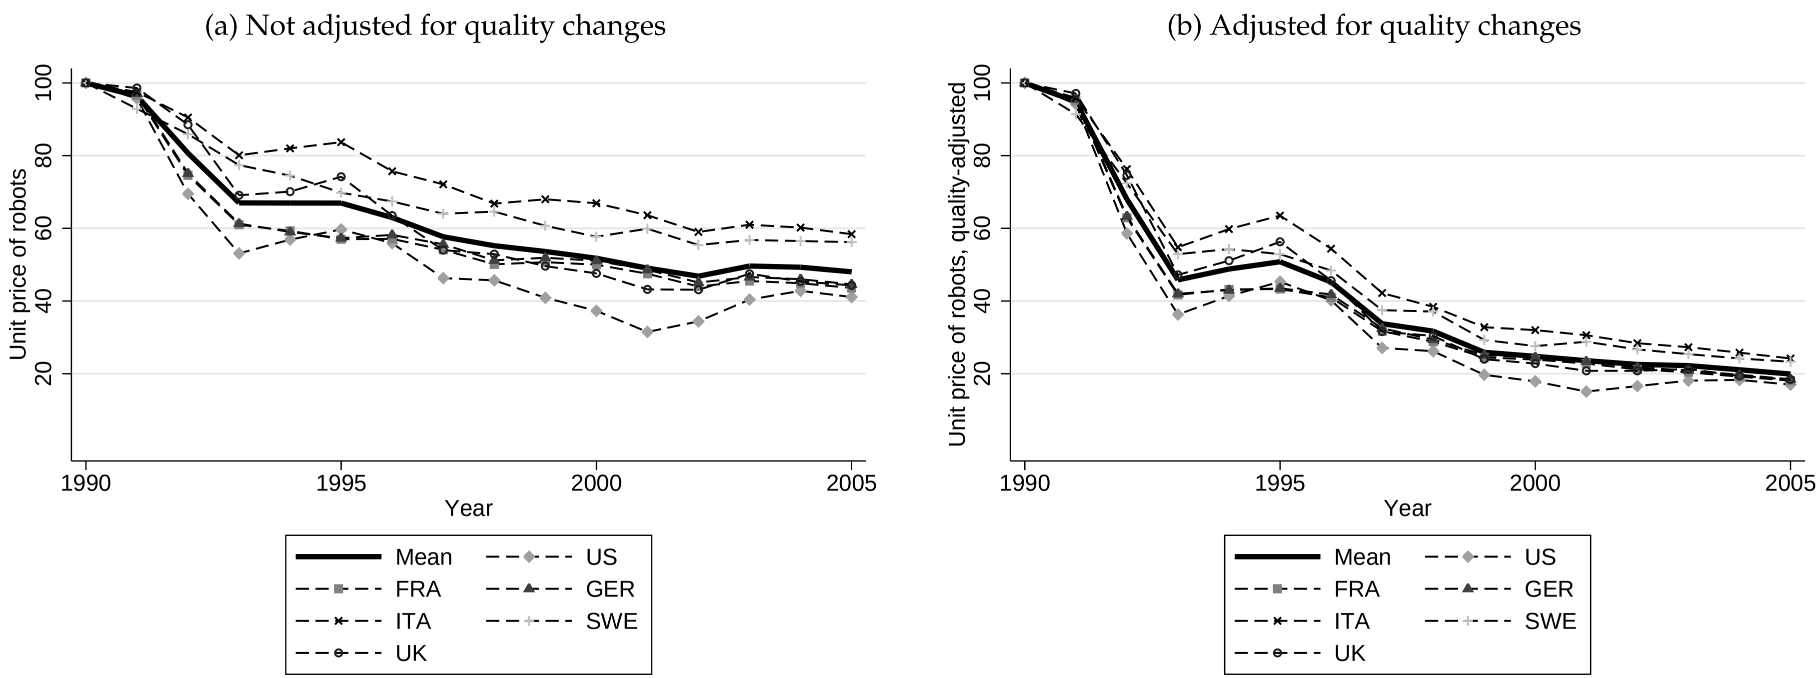

Source: International Federation of Robotics (2006, Chapter III). Price indices are nominal. Indices are based on listed prices reported by surveyed firms. Quality adjusted indices are based on assumptions about how improved characteristics of robots would have affected production cost in the base year (see International Federation of Robotics (2006, Annex C) for details). Annual surveys on robot characteristics were not carried out for the years 1991-1998, hence values of the quality adjusted indices for these years have been imputed by the International Federation of Robotics. "Mean" refers to the unweighted arithmetic average across the six countries shown. For comparison, nominal wages grew on average 58 percent in these six countries (96 percent on average in all countries included in our sample).

Figure 2: Growth of Productivity and Robots 1993-2007

(a) Percentile of Change in Robot Density

(b) Change in Robot Density

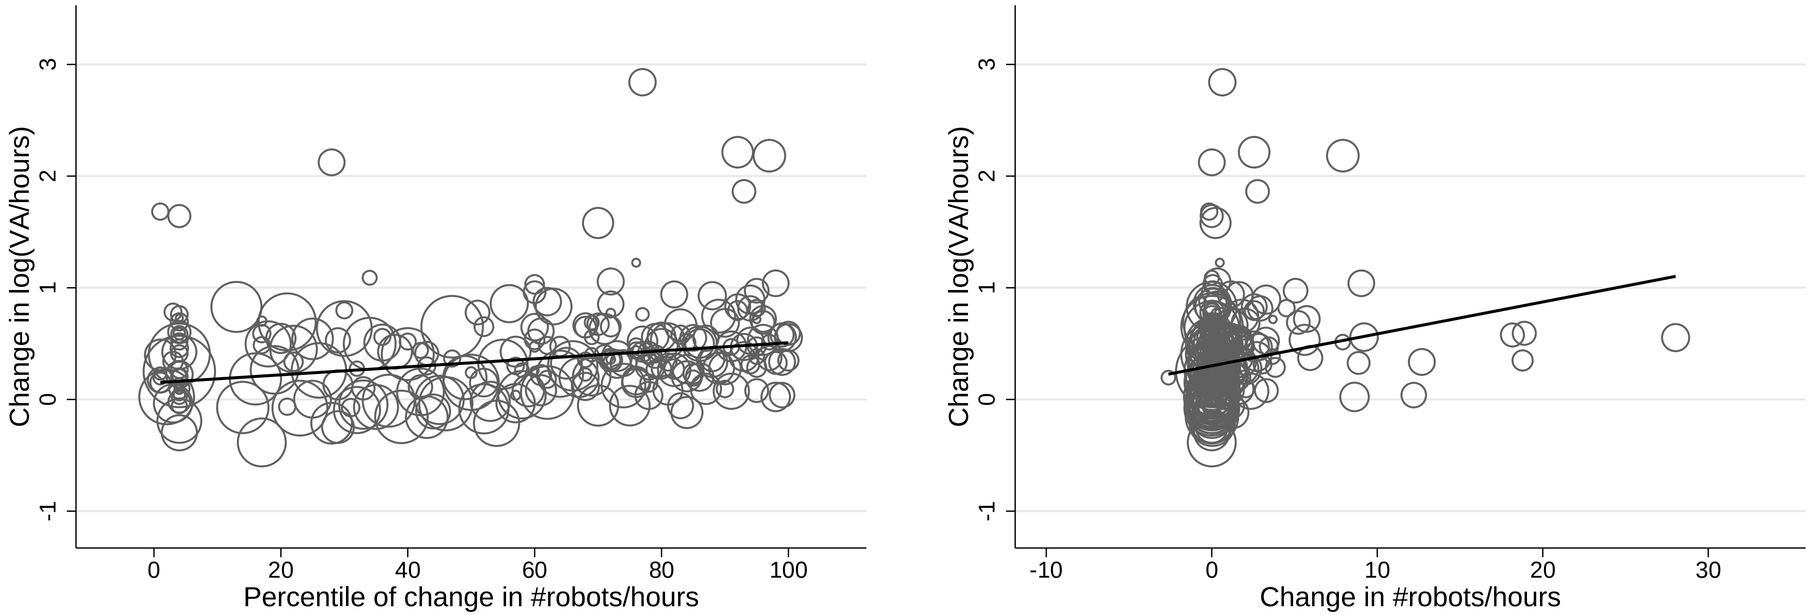

Observations are country-industry cells. The size of each circle corresponds to an industry's 1993 within-country employment share. Fitted regression lines are shown. In panel (a), the estimated slope is 0.36 with a robust standard error of 0.11. In panel (b), the estimated slope is 0.029 and the robust standard error is 0.012.

Figure 3: Cross-Industry Variation in Growth of Value Added, Productivity, and Robots, and the Replaceability of Labor

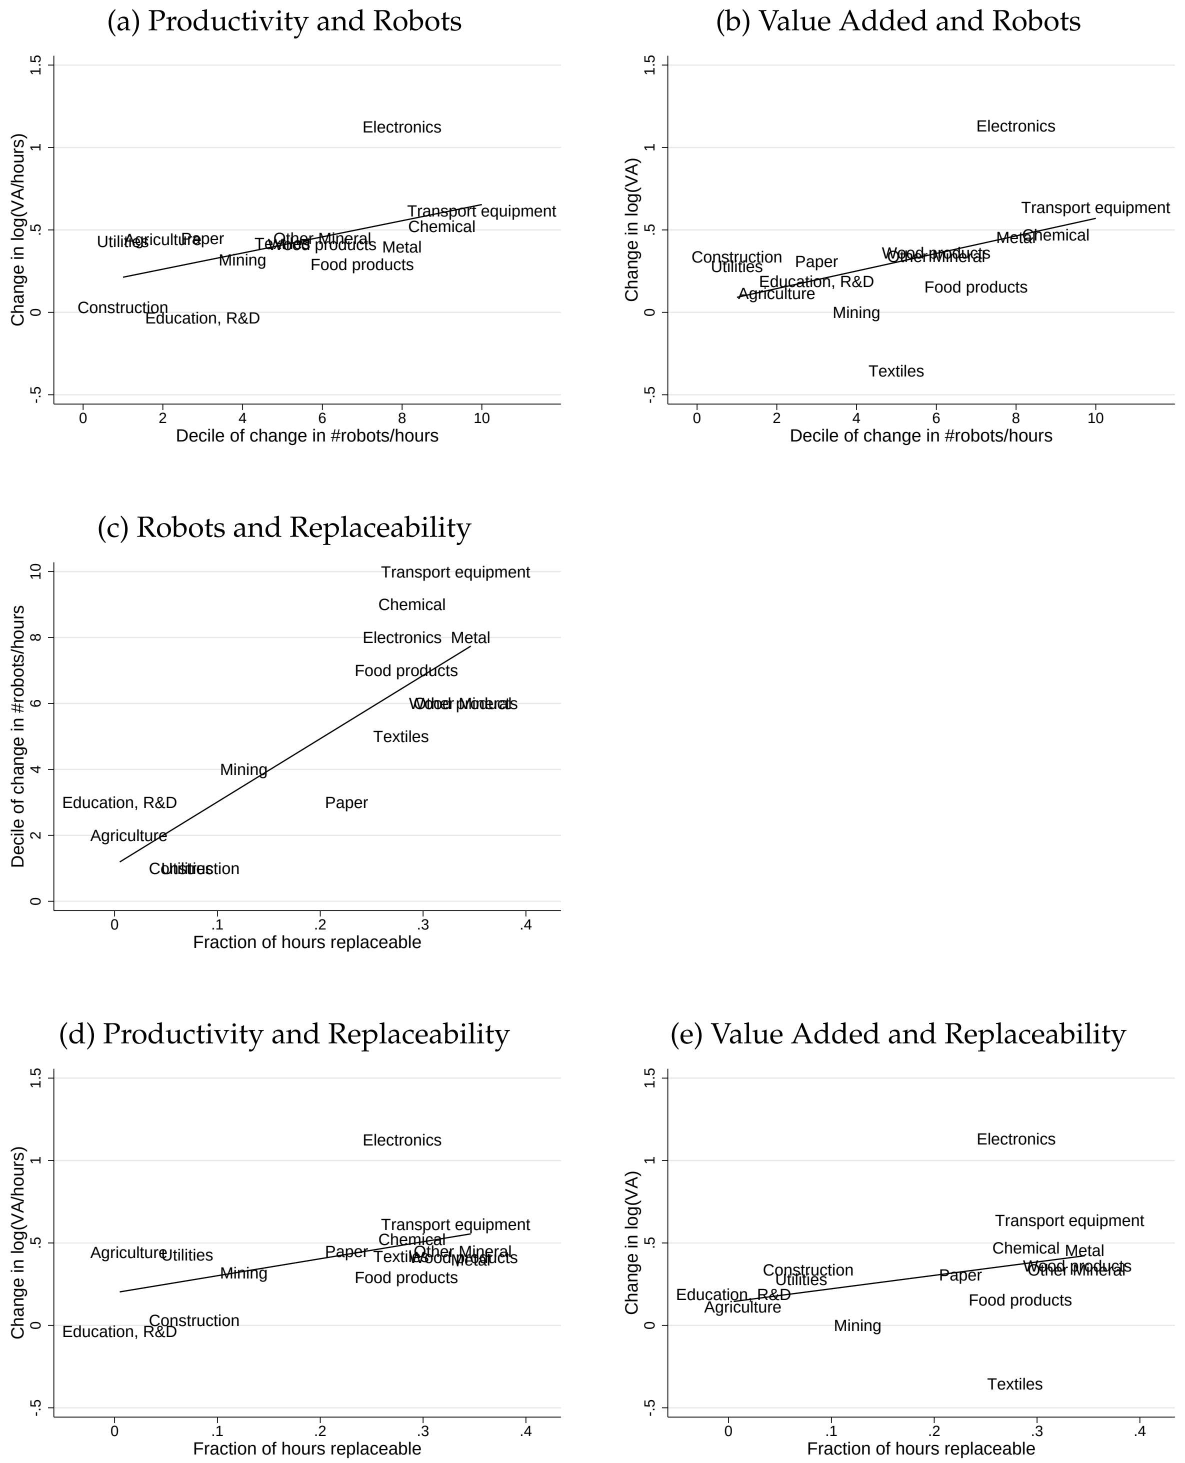

Table 1: Summary Statistics by Country

A. 1993 Levels Averaged by Country

	#robots/H	ln(VA/H)	ln(VA)	ln(H)	ln(K)	ln(wH)
Australia	0.07	3.30	9.56	6.26	8.08	9.24
Austria	0.63	3.09	8.64	5.55	7.84	8.66
Belgium	1.20	3.72	8.94	5.22	7.40	8.87
Denmark	0.42	3.52	8.41	4.89	6.15	8.34
Finland	0.68	3.15	8.29	5.14	6.45	8.29
France	0.79	3.37	10.63	7.26	8.89	10.56
Germany	1.71	3.38	11.00	7.63	9.36	11.01
Greece	0.00	2.53	8.76	6.23	.	8.50
Hungary	0.05	1.68	7.50	5.82	.	.
Ireland	0.00	3.26	8.05	4.79	6.69	7.99
Italy	1.13	3.17	10.54	7.37	8.70	10.37
Netherlands	0.25	3.60	9.35	5.75	7.80	9.16
South Korea	0.28	1.90	9.76	7.86	.	10.11
Spain	0.36	3.21	10.12	6.91	8.86	9.66
Sweden	1.39	3.21	8.69	5.47	6.85	8.78
United Kingdom	0.50	3.38	10.62	7.24	8.09	10.57
United States	0.41	3.39	12.27	8.88	10.59	12.08
Mean	0.58	3.11	9.48	6.37	8.03	9.51

B. Changes from 1993-2007 Averaged by Country

	Δ(#robots/H)	Δ ln(VA/H)	Δ ln(VA)	Δ ln(H)	Δ ln(K)	Δ ln(wH)
Australia	0.12	0.22	0.34	0.12	0.72	0.15
Austria	0.61	0.51	0.32	-0.19	0.02	-0.15
Belgium	1.23	0.29	0.20	-0.09	.	.
Denmark	1.57	0.19	0.17	-0.02	0.51	0.00
Finland	1.05	0.43	0.39	-0.04	0.40	0.01
France	1.20	0.29	0.14	-0.15	0.26	-0.06
Germany	2.73	0.28	0.02	-0.26	0.04	-0.24
Greece	0.03	0.16	0.04	-0.12	.	0.05
Hungary	0.08	0.56	0.37	-0.20	.	.
Ireland	0.10	0.44	0.65	0.20	0.94	0.28
Italy	1.39	0.17	0.10	-0.06	0.43	-0.04
Netherlands	0.54	0.24	0.19	-0.05	0.39	0.04
South Korea	1.31	0.71	0.45	-0.26	.	-0.16
Spain	1.21	0.13	0.31	0.18	0.48	0.26
Sweden	0.80	0.43	0.46	0.04	0.78	0.07
United Kingdom	0.34	0.26	0.14	-0.12	0.54	-0.03
United States	0.97	0.27	0.28	0.01	0.60	0.05
Mean	0.90	0.33	0.27	-0.06	0.47	0.02

H stands for million hours worked. Value added (VA), capital services (K), and the wage bill (wH) are measured in millions of 2005 US$, converted from local currencies using 2005 nominal exchange rates where applicable. Country-level and overall means are weighted by each industry's 1993 share of hours within a country.

Table 2: Summary Statistics by Industry

A. 1993 Levels Averaged by Industry

	#robots/H	ln(VA/H)	ln(VA)	ln(H)	ln(K)	ln(wH)
Agriculture	0.01	2.34	9.24	6.90	8.18	9.52
Chemical	1.16	3.72	9.40	5.68	8.89	9.15
Construction	0.01	3.30	10.26	6.96	8.40	10.15
Education, R&D	0.02	3.45	10.18	6.72	7.27	10.18
Electronics	0.95	2.78	8.38	5.60	7.76	8.94
Food products	0.34	3.35	9.32	5.97	8.41	9.05
Metal	2.37	3.23	9.09	5.86	8.13	9.07
Mining	0.07	4.27	8.22	3.95	7.75	7.32
Other Mineral	0.34	3.27	8.07	4.80	7.23	8.04
Paper	0.06	3.36	8.89	5.53	7.96	8.84
Textiles	0.12	2.79	8.34	5.55	6.45	8.43
Transport equipment	5.36	3.14	8.41	5.27	7.25	8.65
Utilities	0.00	4.30	9.13	4.83	9.02	8.44
Wood products	0.77	2.77	7.36	4.59	6.35	7.49

B. Changes from 1993-2007 Averaged by Industry

	Δ(#robots/H)	Δ ln(VA/H)	Δ ln(VA)	Δ ln(H)	Δ ln(K)	Δ ln(wH)
Agriculture	0.03	0.44	0.11	-0.33	0.07	-0.26
Chemical	3.33	0.52	0.47	-0.05	0.42	0.01
Construction	0.02	0.03	0.34	0.30	0.71	0.35
Education, R&D	0.06	-0.03	0.19	0.22	0.98	0.29
Electronics	1.32	1.13	1.13	0.00	0.62	0.09
Food products	1.21	0.29	0.16	-0.14	0.26	-0.04
Metal	1.67	0.40	0.45	0.06	0.30	0.14
Mining	0.29	0.32	0.00	-0.32	0.42	-0.19
Other Mineral	0.81	0.45	0.34	-0.11	0.25	-0.03
Paper	0.14	0.45	0.31	-0.14	0.45	-0.07
Textiles	0.30	0.42	-0.35	-0.77	-0.13	-0.68
Transport equipment	8.07	0.61	0.64	0.02	0.47	0.06
Utilities	0.02	0.43	0.28	-0.15	0.26	-0.07
Wood products	0.84	0.41	0.36	-0.05	0.40	0.00

H stands for million hours worked. Value added (VA), capital services (K), and the wage bill (wH) are measured in millions of 2005 US$, converted from local currencies using 2005 nominal exchange rates where applicable. Means are not weighted.

Table 3: Changes in Robots Input and Growth in Productivity, Value Added, and Hours Worked, 1993-2007: Non-Parametric Specification

	$\Delta \ln(VA/H)$		$\Delta \ln(VA)$		$\Delta \ln(H)$	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.265		0.195		-0.070	
	(0.064)		(0.066)		(0.088)	
Δ(#robots/hours), quartile 2	-0.060	0.124	0.014	0.233	0.074	0.109
	(0.091)	(0.110)	(0.088)	(0.119)	(0.120)	(0.111)
Δ(#robots/hours), quartile 3	0.040	0.200	0.002	0.201	-0.038	0.001
	(0.084)	(0.093)	(0.083)	(0.105)	(0.103)	(0.102)
Δ(#robots/hours), quartile 4	0.272	0.446	0.279	0.491	0.008	0.045
	(0.090)	(0.109)	(0.095)	(0.114)	(0.091)	(0.086)
Country trends	No	Yes	No	Yes	No	Yes

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 238.

Table 4: Changes in Robots Input and Growth in Productivity, Value Added, and Hours Worked, 1993-2007: OLS and IV Estimates

	$\Delta \ln(VA/H)$		$\Delta \ln(VA)$		$\Delta \ln(H)$	
	(1)	(2)	(3)	(4)	(5)	(6)
A. OLS						
Percentile of Δ(#robots/hours)/100	0.359	0.572	0.336	0.602	-0.023	0.030
	(0.106)	(0.118)	(0.117)	(0.121)	(0.114)	(0.099)
B. IV, replaceable hours						
Percentile of Δ(#robots/hours)/100	0.833	0.873	0.545	0.607	-0.289	-0.266
	(0.188)	(0.157)	(0.155)	(0.143)	(0.169)	(0.155)
First-stage F statistic	93.7	152.6	93.7	152.6	93.7	152.6
Country trends	No	Yes	No	Yes	No	Yes

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 238.

Table 5: Falsification Tests for the Automation Propensity Instrument

	$\Delta \ln(VA/H)$ (1)	$\Delta \ln(VA)$ (2)	$\Delta \ln(H)$ (3)
A. Growth in outcome 1993-2007 (benchmark)			
Share of hours replaceable	1.153	0.802	-0.351
	(0.206)	(0.210)	(0.204)
Observations	238	238	238
B. Growth in outcome 1993-2007, non-adopters (1993)			
Share of hours replaceable	0.912	0.872	-0.041
	(0.605)	(0.792)	(0.811)
Observations	72	72	72
C. Growth in outcome 1979-1993, non-adopters (1993)			
Share of hours replaceable	-0.132	-0.200	-0.068
	(0.558)	(0.505)	(0.419)
Observations	72	72	72
D. Growth in outcome 1993-2007, non-adopters (2007)			
Share of hours replaceable	-0.286	-0.154	0.132
	(1.030)	(1.469)	(1.229)
Observations	27	27	27
p-value of test for equality, A versus B	0.679	0.927	0.687
p-value of test for equality, A versus C	0.019	0.049	0.455
p-value of test for equality, A versus D	0.021	0.262	0.498

Results from OLS regressions are shown. All regressions include country fixed effects. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. Panel *A* shows the reduced form for the full sample. Panel *B* shows the reduced form for country-industry cells that had zero robots in 1993 (non-adopters in 1993) and for which data on prior outcomes is non-missing. In panel *C* the LHS variables are prior changes in the outcomes for the same sample as in panel *B*. Panel *D* shows the reduced form for country-industry cells that did not use any robots in 1993 or 2007 (non-adopters in 2007). Tests for equality of coefficients were performed using Stata's *suest* command.

Table 6: Further Outcomes: TFP, Average Wages, Labor Share

	$\Delta \ln(\text{TFP})$		$\Delta \ln(\text{average wage})$		$\Delta(\text{labor share})$	
	(1)	(2)	(3)	(4)	(5)	(6)
A. OLS						
Percentile of Δ(#robots/hrs)/100	0.366	0.430	0.004	0.045	-0.067	-0.061
	(0.104)	(0.117)	(0.021)	(0.012)	(0.100)	(0.058)
B. IV, replaceable hours						
Percentile of Δ(#robots/hrs)/100	0.633	0.663	0.057	0.070	-0.153	-0.134
	(0.170)	(0.152)	(0.029)	(0.016)	(0.121)	(0.089)
First-stage F statistic	87.6	121.7	93.1	142.8	88.2	141.2
Observations	182	182	210	210	224	224
Country trends	No	Yes	No	Yes	No	Yes

The labor share falls between 0 and 1. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country.

Table 7: Further Outcomes: Hours and Wage Bill Growth by Skill Group

	high skill		middle skill		low skill	
	(1)	(2)	(3)	(4)	(5)	(6)
A1. Hours, OLS						
Percentile of Δ(#robots/hrs)/100	-0.025	0.179	-0.179	-0.051	-0.165	-0.129
	(0.121)	(0.102)	(0.132)	(0.082)	(0.120)	(0.097)
A2. Hours, IV						
Percentile of Δ(#robots/hrs)/100	-0.030	0.035	-0.330	-0.259	-0.460	-0.431
	(0.213)	(0.168)	(0.171)	(0.118)	(0.192)	(0.146)
First-stage F statistic	98.8	154.3	98.8	154.3	98.8	154.3
B1. Wage bills, OLS						
Percentile of Δ(#robots/hrs)/100	-0.068	0.182	-0.221	-0.057	-0.078	-0.100
	(0.127)	(0.112)	(0.156)	(0.085)	(0.128)	(0.096)
B2. Wage bills, IV						
Percentile of Δ(#robots/hrs)/100	-0.027	0.040	-0.321	-0.231	-0.410	-0.360
	(0.229)	(0.179)	(0.190)	(0.116)	(0.243)	(0.147)
First-stage F statistic	98.8	154.3	98.8	154.3	98.8	154.3
Country trends	No	Yes	No	Yes	No	Yes

The dependent variables are the changes from 1993-2005 in the logs of hours and the wage bill for each group as indicated. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 224 (information on skills groups is missing for Hungary in 1993).

Online Appendix for "Robots at Work"

Georg Graetz, *Uppsala University*
Guy Michaels, *London School of Economics*

Data Appendix

Imputation of the robot stock for a subset of countries. As mentioned in Section 3, a complicating feature of the IFR data is that for half of the countries in our final sample, a breakdown of deliveries by industries is not available for the early years in the sample, when all delivered units are reported under the "unspecified" category. These countries (and the year that the breakdown by industries first becomes available) include Australia (2006), Austria (2003), Belgium (2004), Denmark (1996), Greece (2006), Hungary (2004), Ireland (2006), Korea (2001, but not in 2002, then again from 2003 onwards), Netherlands (2004), and the US (2004). For this group of countries, we impute industry-level deliveries by multiplying the number of robots reported as "unspecified" by the average share of an industry's deliveries in total deliveries during the years when the breakdown was reported in the data. To compute the share of deliveries we use all the years available in the IFR data, up to and including 2011. Similarly, for these countries we multiply the stock reported by IFR as "unspecified" in 1993 by the average share of deliveries. We then apply our perpetual inventory method to compute the stock for all subsequent years.

Data on robot prices. The IFR report two measures of prices: one that is based on the total turnover of the robots producing industries, and one that is based on list prices of surveyed firms. However, the IFR do not report price data for all countries and years. Turnover-based prices are calculated as the ratio of the total turnover of the robots industries and the number of robots delivered. They are available throughout our sample period for the US only, and can be found in IFR (2005) and IFR (2012). For each country-industry-year cell, we compute robot services as the product of the turnover-based US price of robots and our measure of the robot stock, multiplied by 0.15 which is the sum of a depreciation rate of ten percent and a real interest rate of five percent. (This procedure is based on the neoclassical theory of investment, see e.g. Timmer, van Moergastel, Stuivenwold, Ypma, O'Mahony, and Kangasniemi (2007, p.33) for a discussion and application to EUKLEMS capital data.)

As the IFR point out, turnover-based prices are problematic as the total turnover also includes peripherals, customer services, etc., and is affected by volume discounts. For selected countries the IFR also report price indices based on list prices, but these stop in 2005. List prices, together with data on changes in characteristics of robots, enabled the IFR to construct quality adjusted price indices, as well. We report these indices in Figure 1.

Construction of alternative instruments shown in Appendix Table A5. "Percentile of #robots/hrs in 1993/100" is the percentile of a country-industry in the 1993 distribution (weighted by 1993 within-country employment shares) of robot density, divided by 100 to let this variable range from zero to one.

The "shift-share" instrument is the product of two terms. The first term is a given industry's percentile in the 1993 distribution of global robot density. For industry i, global 1993 robot density is calculated as $[\sum_c \#robots_{ic,1993}]$ / $[\sum_c hours_{ic,1993}]$, where the sum is over countries indexed by c, but excluding countries for which deliveries are not reported at the industry level in 1993 (see the note to Appendix Table A9 for a list of countries included). The second term is the percentile of the country-level change in robot density between 1993 and 2007. Country-level robot density in year t is calculated as $[\sum_i \#robots_{ic,t}]$ / $[\sum_i hours_{ic,t}]$ where the sum is taken over all industries (only here we include the "all other manufacturing", "all other non-manufacturing", and "unspecified" industries). Percentiles are un-weighted and were divided by 100 to let the shift-share instrument range from zero to one.

Model Appendix

Here we show how to solve the simple model introduced in Section 2 and derive the predictions stated there. Start by normalizing the price of Y_R to one and let p be the price of Y_N relative to that of Y_R. Profit maximization in the robot-using sector implies

$$\left(\frac{Y_R}{R}\right)^{\frac{1}{\sigma}} = \rho, \qquad \left(\frac{Y_R}{L_R}\right)^{\frac{1}{\sigma}} = w \tag{A1}$$

where w is the wage. Combining the last two equations results in

$$\left(\frac{R}{L_R}\right)^{-\frac{1}{\sigma}} = \frac{\rho}{w}. \tag{A2}$$

Profit maximization in the non-robot-using sector implies $p = w$.

Consumer's utility maximization subject to the budget constraint $Y_R + pY_N = I$ (where I denotes total income) results in

$$\left(\frac{Y_R}{Y_N}\right)^{-\frac{1}{\varepsilon}} = \frac{1}{p} \tag{A3}$$

In general equilibrium the income accounting identity $I = wL + \rho R$ applies.

To analyze the effect of a fall in the rental price of robots, ρ, on a variety of outcomes, we make use of the recursive structure of our equilibrium conditions. In particular, notice that robot density determines labor productivity in the robot-using sector, as follows from dividing both sides of the production function by L_R. In turn, labor productivity in the robot-using sector determines the ratio of outputs of the two sectors,

$$\left(\frac{Y_R}{Y_N}\right)^{-\frac{1}{\varepsilon}} = \left(\frac{Y_R}{L_R}\right)^{-\frac{1}{\sigma}}, \tag{A4}$$

as well as the ratio of labor inputs,

$$\frac{L_R}{L_N} = \left(\frac{Y_R}{L_R}\right)^{\frac{\varepsilon}{\sigma}-1}. \tag{A5}$$

By determining how a fall in ρ affects robot density, we immediately reveal its effects on labor productivity in the robot-using sector, the ratio of outputs, and the ratio of labor inputs.[37] In particular, we see that if a fall in ρ raises robot density, then labor productivity in the robot-using sector increases (production function), which in turn raises value added relative to the non-robot-using sector (see (A4)), but the effect on the labor input depends on the relative values of the elasticities of substitution (see (A5)).

It remains to verify that a fall in ρ indeed raises robot density. Combine (A2) with (A3), using $p = w$, and (A4) to obtain $(R/L_R)^{-1/\sigma} = \rho(Y_R/L_R)^{-1/\sigma}$. Using the production function yields $(R/L_R)^{-1/\sigma} = \rho[(R/L_R)^{(\sigma-1)/\sigma} + 1]^{-1/(\sigma-1)}$. By implicit differentiation,

$$\frac{\partial(R/L_R)}{\partial\rho} = -\frac{\sigma}{\rho}\frac{R}{L_R}\left[\left(\frac{R}{L_R}\right)^{\frac{\sigma-1}{\sigma}} + 1\right] < 0.$$

[37] To derive (A4), use $p = w$ to substitute out the wage in the second equation of (A1), and plug the result into (A3). To derive (A5), plug the production function $Y_N = L_N$ into (A4) to obtain $(Y_R/L_R/(L_N/L_R))^{-1/\varepsilon} = (Y_R/L_R)^{-1/\sigma}$ and rearrange.

Appendix Tables

Table A1: List of All EUKLEMS Industries

Code	Included in Robotics data	Label	Code description
AtB	✓	Agriculture	Agriculture, hunting, forestry, and fishing
C	✓	Mining	Mining and quarrying
15t16	✓	Food products	Food products, beverages and tobacco
17t19	✓	Textiles	Textiles, textile products, leather and footwear
20	✓	Wood products	Wood and products of wood and cork
21t22	✓	Paper	Pulp, paper, paper products, printing and publishing
23t25	✓	Chemical	Chemical, rubber, plastics and fuel
26	✓	Other mineral	Other non-metallic mineral products
27t28	✓	Metal	Basic metals and fabricated metal products
29			Machinery, not elsewhere classified
30t33	✓	Electronics	Electrical and optical equipment
34t35	✓	Transport equipment	Transport equipment
36t37			Manufacturing not elsewhere classified; recycling
50			Sale, maintenance and repair of motor vehicles and motorcycles; retail sale of fuel
51			Wholesale trade and commission trade, except of motor vehicles and motorcycles
52			Retail trade, except of motor vehicles and motorcycles; repair of household goods
60t63			Transport and storage
64			Post and telecommunications
70			Real estate activities
71t74			Renting of machinery and equipment and other business activities
E	✓	Utilities	Electricity, gas, water supply
F	✓	Construction	Construction
H			Hotels and restaurants
J			Financial intermediation
L			Public administration, defence, and compulsory social security
M	✓	Education, R&D	Education
N			Health and social work
O			Other community, social and personal services

Note: Industry M in the World Robotics data includes research and development in addition to education, whereas research and development are included in industry 71t74 in the EUKLEMS data.

Table A2: Summary Statistics for Robots Variables

A. 1993 Levels

	Mean	Stdev	Min	Median	Max
#robots/hours	0.582	1.773	0.000	0.004	15.697
$\ln(1 + \#\text{robots/hours})$	0.245	0.514	0.000	0.004	2.815
$1,000 \times$ robot services/wage bill	0.361	1.032	0.000	0.004	10.287

B. Changes from 1993-2007

	Mean	Stdev	Min	Median	Max	Mean 1st qrtl	Mean 2nd qrtl	Mean 3rd qrtl	Mean 4th qrtl
Δ(#robots/hours)	0.898	2.795	-2.617	0.024	28.028	-0.018	0.012	0.130	3.479
$\Delta \ln(1 + \#\text{robots/hours})$	0.199	0.341	-0.697	0.024	1.806	-0.007	0.012	0.103	0.700
$\Delta(1,000 \times$ robot services/wage bill$)$	0.123	0.678	-3.298	0.006	5.069	-0.206	0.003	0.022	0.691

The variable 'hours' refers to million hours worked. The number of robots was computed from annual investment data using the perpetual inventory method and assuming a depreciation rate of ten percent. The initial value was taken from the World Robotics database. Robots services equal 0.15 times the price of robots, times the number of robots. The adjustment factor of 0.15 reflects depreciation at ten percent and an interest rate of five percent. The price of robots is the average unit price of robots in the US in the relevant year, expressed in 2005 US$. Reported statistics are weighted by each industry's 1993 share of hours within a country.

Table A3: Robustness to Controlling for Industry Trends

	$\Delta\ln(VA/H)$		$\Delta\ln(VA)$		$\Delta\ln(H)$	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Non-parametric specification						
Δ(#robots/hours), quartile 2	0.124	0.157	0.233	0.168	0.109	0.011
	(0.110)	(0.070)	(0.119)	(0.084)	(0.111)	(0.053)
Δ(#robots/hours), quartile 3	0.200	0.192	0.201	0.174	0.001	-0.018
	(0.093)	(0.083)	(0.105)	(0.084)	(0.102)	(0.060)
Δ(#robots/hours), quartile 4	0.446	0.184	0.491	0.225	0.045	0.041
	(0.109)	(0.132)	(0.114)	(0.141)	(0.086)	(0.071)
B. OLS						
Percentile of Δ(#robots/hours)/100	0.572	0.352	0.602	0.358	0.030	0.006
	(0.118)	(0.144)	(0.121)	(0.144)	(0.099)	(0.094)
Country trends	Yes	Yes	Yes	Yes	Yes	Yes
Industry trends	No	Yes	No	Yes	No	Yes

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 238.

Table A4: Robustness to Controlling for Prior Changes in Outcome Variables

	Δ ln(value added/hours)				Δ ln(value added)				Δ ln(hours)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. OLS												
Pctile of Δ(#robots/hrs)/100	0.59	0.45	0.38	0.33	0.61	0.46	0.44	0.35	0.01	0.01	0.06	0.02
	(0.12)	(0.08)	(0.14)	(0.13)	(0.12)	(0.10)	(0.13)	(0.12)	(0.10)	(0.08)	(0.09)	(0.09)
Δ ln(value added), lagged		0.57		0.25		0.45		0.30		-0.13		0.05
		(0.10)		(0.10)		(0.13)		(0.11)		(0.09)		(0.07)
Δ ln(hours), lagged		-0.71		-0.28		0.02		-0.16		0.73		0.11
		(0.08)		(0.11)		(0.10)		(0.12)		(0.09)		(0.10)
B. IV, replaceable hours												
Pctile of Δ(#robots/hrs)/100	0.88	0.67			0.58	0.55			-0.30	-0.12		
	(0.15)	(0.11)			(0.13)	(0.13)			(0.15)	(0.11)		
Δ ln(value added), lagged		0.53				0.43				-0.10		
		(0.10)				(0.13)				(0.09)		
Δ ln(hours), lagged		-0.70				0.03				0.73		
		(0.07)				(0.09)				(0.09)		
First-stage F statistic	154.3	187.3			154.3	187.3			154.3	187.3		
Country trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry trends	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes

Prior changes in outcome variables refer to the years 1979-1993. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 224 (prior changes in outcome variables are missing for Hungary).

Table A5: Robustness to Using Alternative Sets of Instruments

	$\Delta \ln(\text{value added/hours})$				$\Delta \ln(\text{hours})$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. IV								
Percentile of Δ(#robots/hours)/100	0.873	0.808	0.676	0.780	-0.266	0.063	0.053	-0.108
	(0.157)	(0.188)	(0.136)	(0.138)	(0.155)	(0.150)	(0.125)	(0.129)
B. First stage								
Replacable hours	1.321			0.757	1.321			0.757
	(0.107)			(0.186)	(0.107)			(0.186)
Percentile of #robots/hrs in 1993/100		0.515		0.026		0.515		0.026
		(0.057)		(0.101)		(0.057)		(0.101)
Shift-share			0.898	0.475			0.898	0.475
			(0.063)	(0.124)			(0.063)	(0.124)
First-stage F statistic	152.6	82.9	203.0	84.7	152.6	82.9	203.0	84.7

"Shift-share" is the interaction of an industry's initial global robot density with the country-level change in robot density, see the Data Appendix for details. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. All regressions include country fixed effects. The number of observations is 238.

Table A6: Main OLS and IV Results, Standard Errors Clustered by Country

	$\Delta \ln(VA/H)$		$\Delta \ln(VA)$		$\Delta \ln(H)$	
	(1)	(2)	(3)	(4)	(5)	(6)
A. OLS						
Percentile of Δ(#robots/hours)/100	0.359	0.572	0.336	0.602	-0.023	0.030
	(0.129)	(0.117)	(0.138)	(0.132)	(0.102)	(0.106)
B. IV, replaceable hours						
Percentile of Δ(#robots/hours)/100	0.833	0.873	0.545	0.607	-0.289	-0.266
	(0.116)	(0.126)	(0.120)	(0.132)	(0.118)	(0.123)
First-stage F statistic	126.9	112.0	126.9	112.0	126.9	112.0
Country trends	No	Yes	No	Yes	No	Yes

Robust standard errors, clustered by country, in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 238.

Table A7: Alternative Measures of Robots Input

	$\Delta\ln(VA/H)$		$\Delta\ln(VA)$		$\Delta\ln(H)$	
	(1)	(2)	(3)	(4)	(5)	(6)
A1. OLS						
Δ(#robots/hrs)	0.029	0.032	0.029	0.037	0.001	0.005
	(0.012)	(0.010)	(0.012)	(0.011)	(0.004)	(0.005)
A2. IV, replaceable hours						
Δ(#robots/hrs)	0.138	0.144	0.090	0.100	-0.048	-0.044
	(0.035)	(0.032)	(0.029)	(0.028)	(0.028)	(0.026)
First-stage F statistic	33.3	34.5	33.3	34.5	33.3	34.5
B1. OLS						
$\Delta\ln(1+\#robots/hours)$	0.348	0.406	0.317	0.385	-0.031	-0.021
	(0.106)	(0.097)	(0.113)	(0.110)	(0.061)	(0.065)
B2. IV, replaceable hours						
$\Delta\ln(1+\#robots/hours)$	0.753	0.769	0.493	0.535	-0.261	-0.234
	(0.157)	(0.139)	(0.142)	(0.134)	(0.147)	(0.133)
First-stage F statistic	123.1	127.5	123.1	127.5	123.1	127.5
C1. OLS						
$\Delta(1{,}000 \times$ robot services/wage bill$)$	0.133	0.126	0.120	0.127	-0.013	0.002
	(0.069)	(0.053)	(0.079)	(0.070)	(0.021)	(0.032)
C2. IV, replaceable hours						
$\Delta(1{,}000 \times$ robot services/wage bill$)$	1.199	1.241	0.764	0.842	-0.435	-0.399
	(0.516)	(0.518)	(0.364)	(0.375)	(0.294)	(0.275)
First-stage F statistic	5.7	5.7	5.7	5.7	5.7	5.7
Country trends	No	Yes	No	Yes	No	Yes

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 238 (panels *A* and *B*) and 210 (panel *C*, data on wage bills is missing for Belgium and Hungary).

Table A8: Robustness to Alternative Depreciation Rates

	$\delta = 0.05$		Benchmark: $\delta = 0.10$		$\delta = 0.15$	
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta \ln(\text{VA/H})$	$\Delta \ln(\text{H})$	$\Delta \ln(\text{VA/H})$	$\Delta \ln(\text{H})$	$\Delta \ln(\text{VA/H})$	$\Delta \ln(\text{H})$
A1. OLS						
Percentile of Δ(#robots/hours)/100	0.593	0.071	0.572	0.030	0.395	-0.053
	(0.126)	(0.107)	(0.118)	(0.099)	(0.122)	(0.094)
A2. IV, replaceable hours						
Percentile of Δ(#robots/hours)/100	0.824	-0.251	0.873	-0.266	1.107	-0.337
	(0.150)	(0.146)	(0.157)	(0.155)	(0.218)	(0.195)
First-stage F statistic	230.3	230.3	152.6	152.6	68.9	68.9
B1. OLS						
Δ(#robots/hours)	0.021	0.004	0.032	0.005	0.047	0.005
	(0.006)	(0.003)	(0.010)	(0.005)	(0.016)	(0.008)
B2. IV, replaceable hours						
Δ(#robots/hours)	0.090	-0.027	0.144	-0.044	0.236	-0.072
	(0.020)	(0.016)	(0.032)	(0.026)	(0.056)	(0.042)
First-stage F statistic	38.0	38.0	34.5	34.5	27.9	27.9

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. All regressions include country fixed effects. The number of observations is 238.

Table A9: Main OLS and IV Results, Excluding Countries for which Baseline Robot Stocks Are Imputed

	$\Delta \ln(VA/H)$		$\Delta \ln(VA)$		$\Delta \ln(H)$	
	(1)	(2)	(3)	(4)	(5)	(6)
A. OLS						
Percentile of Δ(#robots/hours)/100	0.376	0.402	0.398	0.484	0.022	0.082
	(0.129)	(0.153)	(0.135)	(0.141)	(0.136)	(0.134)
B. IV, replaceable hours						
Percentile of Δ(#robots/hours)/100	0.706	0.747	0.532	0.620	-0.174	-0.127
	(0.264)	(0.253)	(0.213)	(0.194)	(0.234)	(0.226)
First-stage F statistic	42.9	48.4	42.9	48.4	42.9	48.4
Country trends	No	Yes	No	Yes	No	Yes

The estimation sample includes Finland, France, Germany, Italy, Spain, Sweden, and the UK. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 98.

Table A10: Robustness to Excluding Non-Tradable Industries

	$\Delta \ln$(value added/hours)				$\Delta \ln$(value added)				$\Delta \ln$(hours)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. OLS												
Pctile of Δ(#robots/hrs)/100	0.57	0.28	0.35	0.33	0.60	0.78	0.36	0.38	0.03	0.50	0.01	0.05
	(0.12)	(0.12)	(0.14)	(0.19)	(0.12)	(0.14)	(0.14)	(0.20)	(0.10)	(0.07)	(0.09)	(0.10)
B. IV, replaceable hours												
Pctile of Δ(#robots/hrs)/100	0.87	0.15			0.61	0.77			-0.27	0.62		
	(0.16)	(0.12)			(0.14)	(0.15)			(0.15)	(0.09)		
First-stage F statistic	152.6	153.6			152.6	153.6			152.6	153.6		
Non-tradable excluded	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Country trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry trends	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
Observations	238	187	238	187	238	187	238	187	238	187	238	187

Non-tradable industries include construction, education/research/development, and utilities. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country.

Table A11: Robustness to Controlling for Changes in Skill Mix

	Δ ln(value added/hours)						Δ ln(value added)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. OLS												
Pcntle of Δ(#robots/hrs)/100	0.59	0.48	0.40	0.38	0.39	0.34	0.61	0.61	0.61	0.44	0.44	0.41
	(0.12)	(0.11)	(0.10)	(0.14)	(0.14)	(0.15)	(0.12)	(0.11)	(0.11)	(0.13)	(0.12)	(0.13)
Δ(high skill hrs/total hrs)		1.91	0.59		1.86	1.12		0.98	0.80		1.13	0.53
		(1.17)	(1.01)		(0.73)	(0.73)		(0.94)	(0.83)		(0.73)	(0.71)
Δ(low skill hrs/total hrs)		-1.49	-0.49		0.28	0.45		0.70	0.95		-0.11	-0.06
		(0.62)	(0.75)		(0.39)	(0.43)		(0.54)	(0.71)		(0.58)	(0.66)
Δ log(wage)			3.71			1.51			0.38			1.21
			(0.98)			(0.69)			(1.45)			(0.89)
B. IV, replaceable hours												
Pcntle of Δ(#robots/hrs)/100	0.88	0.73	0.58				0.58	0.61	0.60			
	(0.15)	(0.15)	(0.15)				(0.13)	(0.13)	(0.13)			
Δ(high skill hrs/total hrs)		1.54	0.45					0.97	0.81			
		(1.08)	(0.98)					(0.90)	(0.79)			
Δ(low skill hrs/total hrs)		-1.21	-0.39					0.70	0.94			
		(0.61)	(0.74)					(0.53)	(0.67)			
Δ log(wage)			3.36						0.40			
			(1.00)						(1.43)			
First-stage F statistic	154.3	135.4	127.6				154.3	135.4	127.6			
Country trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry trends	No	No	No	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 224.

Table A12: Robustness to Controlling for Changes in Other Capital and Skill Mix

	$\Delta\ln(VA/H)$			$\Delta\ln(VA)$		
	(1)	(2)	(3)	(4)	(5)	(6)
A. OLS						
Percentile of Δ(#robots/hours)/100	0.566	0.545	0.371	0.492	0.477	0.454
	(0.128)	(0.127)	(0.104)	(0.125)	(0.118)	(0.107)
Δ(capital services/wage bill)		0.096	0.084		0.101	0.088
		(0.054)	(0.049)		(0.060)	(0.052)
Δ(ICT capital services/total capital services)		0.112	0.004		-0.445	-0.605
		(0.228)	(0.222)		(0.151)	(0.179)
Δ(high skill hours/total hours)			0.311			1.806
			(1.116)			(1.068)
Δ(low skill hours/total hours)			-0.251			0.838
			(1.038)			(0.768)
$\Delta\log$(wage)			3.863			0.057
			(1.126)			(1.173)
B. IV, replaceable hours						
Percentile of Δ(#robots/hours)/100	0.877	0.848	0.628	0.568	0.480	0.487
	(0.171)	(0.165)	(0.159)	(0.143)	(0.132)	(0.135)
Δ(capital services/wage bill)		0.087	0.080		0.101	0.087
		(0.049)	(0.045)		(0.057)	(0.049)
Δ(ICT capital services/total capital services)		0.099	0.016		-0.445	-0.603
		(0.204)	(0.204)		(0.144)	(0.170)
Δ(high skill hours/total hours)			0.047			1.773
			(1.050)			(1.022)
Δ(low skill hours/total hours)			-0.062			0.862
			(1.004)			(0.742)
$\Delta\log$(wage)			3.363			-0.007
			(1.150)			(1.093)
First-stage F statistic	121.7	125.1	113.4	121.7	125.1	113.4
Country trends	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. The number of observations is 182.

Table A13: Robustness to Controlling for Changes in Other Capital, Changes in Skill Mix, and Industry Trends

	Δ ln(value added/hours)				Δ ln(value added)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Percentile of Δ(#robots/hours)/100	0.384	0.261	0.236	0.221	0.441	0.314	0.288	0.263
	(0.144)	(0.157)	(0.156)	(0.159)	(0.127)	(0.140)	(0.135)	(0.133)
Δ(capital services/wage bill)			0.106	0.093			0.112	0.102
			(0.057)	(0.051)			(0.061)	(0.057)
Δ(ICT capital services/total capital services)			0.003	-0.070			-0.369	-0.423
			(0.171)	(0.166)			(0.168)	(0.162)
Δ(high skill hours/total hours)				1.036				0.939
				(0.827)				(0.868)
Δ(low skill hours/total hours)				0.629				0.207
				(0.596)				(0.676)
Δ log(wage)				0.998				1.050
				(0.844)				(0.923)
Country trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	224	182	182	182	224	182	182	182

Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. Columns (1) and (5) show results using the sample that contains information on skill groups, while the remaining columns show results from restricting the sample further to observations with information on non-robot capital.

Table A14: Percentage Losses in 2007 Value Added per Hour and Value Added for the Counterfactual Scenario of No Increase in Robots

	Robot-Using Ind.		All Ind.	
	VA/H	VA	VA/H	VA
Australia	8.1	8.2	2.8	2.8
Austria	19.0	19.3	6.5	6.6
Belgium	19.3	19.6	5.8	5.8
Denmark	20.4	20.7	5.9	6.0
Finland	20.2	20.5	7.6	7.8
France	17.2	17.5	4.5	4.5
Germany	22.9	23.3	6.9	7.0
Greece	11.2	11.4	3.3	3.4
Hungary	7.1	7.2	2.7	2.7
Ireland	9.9	10.1	4.2	4.2
Italy	16.1	16.3	4.9	4.9
Netherlands	13.8	14.0	3.8	3.8
South Korea	17.7	17.9	8.3	8.4
Spain	17.7	18.0	6.1	6.2
Sweden	17.0	17.2	5.3	5.3
United Kingdom	17.0	17.3	4.7	4.8
United States	13.6	13.8	3.6	3.6
Mean	15.8	16.0	5.1	5.2

The percentage loss in variable $Y \in \{\text{VA/H}, \text{VA}\}$ is given by $100 \times (1 - Y^{cf}_{c,2007}/Y_{c,2007})$. See the text for details of how the counterfactual outcome $Y^{cf}_{c,2007}$ was calculated. The figures for the entire economy were obtained by multiplying the numbers reported in the first four columns by the share in value added of the robots-using industries in a given country in 2007. This amounts to assuming that no robots were used in the industries not included in our sample. In fact, the average share of the excluded industries ("all other manufacturing" and "all other non-manufacturing") in total robots deliveries across countries in 2007 was 0.5 percent.

Table A15: Further Outcomes: Skill Shares

	high skill		middle skill		low skill	
	(1)	(2)	(3)	(4)	(5)	(6)
A1. Hours, OLS						
Percentile of Δ(#robots/hrs)/100	0.97	2.57	1.33	1.28	-2.30	-3.85
	(1.58)	(1.16)	(2.30)	(2.02)	(1.99)	(1.52)
A2. Hours, IV						
Percentile of Δ(#robots/hrs)/100	2.51	2.72	4.65	5.73	-7.16	-8.45
	(2.27)	(1.68)	(3.54)	(3.10)	(3.44)	(2.52)
First-stage F statistic	98.8	154.3	98.8	154.3	98.8	154.3
B1. Wage bills, OLS						
Percentile of Δ(#robots/hrs)/100	1.99	4.62	-2.99	-1.93	1.00	-2.70
	(1.52)	(1.29)	(3.50)	(2.19)	(3.00)	(1.48)
B2. Wage bills, IV						
Percentile of Δ(#robots/hrs)/100	6.34	6.72	-0.45	1.17	-5.89	-7.88
	(2.43)	(1.94)	(4.44)	(3.35)	(4.32)	(2.72)
First-stage F statistic	98.8	154.3	98.8	154.3	98.8	154.3
Country trends	No	Yes	No	Yes	No	Yes

The dependent variables are the changes from 1993-2005 in the shares in hours and the wage bill for each group as indicated. Shares were scaled to fall between 0 and 100. Robust standard errors in parentheses. Regressions are weighted by each industry's 1993 share of hours within a country. Instrumental variables regression use the replaceable fraction of the wage bill as an instrument. The number of observations is 224 (information on skill groups is missing for Hungary in 1993).