Job Creators or Job Killers?

Heterogeneous Effects of Industrial Robots on UK Employment

Joel Kariel[†]

Department of Economics, University of Oxford, Manor Road, Oxford, OX1 3UQ, United Kingdom

Abstract

There is concern about robots taking our jobs. This analysis looks at the impact of industrial robot adoption in the UK. Using a novel instrument to deal with endogeneity of robot adoption, estimates suggest that higher robot use is associated with increased employment and some evidence of a positive effect on part-time pay, contrary to evidence from other countries. However, there is a large amount of heterogeneity across industries. The results show that industrial robots have directly replaced workers in automobile manufacturing. On the other hand, they have had positive effects on other areas of the labour market such as services.

JEL classification: J23, J24, O33
Keywords: Robots, Labour Market Effects, Automation, United Kingdom

[†]DPhil Candidate in Economics, University of Oxford, joel.kariel@worc.ox.ac.uk. I thank Petr Sedláček, Gabriel Ulyssea, participants at the Royal Economic Society Symposium of Junior Researchers, Nicholas Stenner, and anonymous referees from the University of Oxford for helpful comments. The author declares that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

> "...but I am convinced, that the substitution of machinery for human labour, is often very injurious to the interests of the class of labourers."
> David Ricardo, 1817

Economists and social commentators have long been concerned that machines will eventually replace humans in the workplace. Automation is "the substitution of labour by capital" (Lawrence et al., 2017), when technological change allows machines to perform tasks that were previously limited to the domain of human capability. However, capital can either replace humans or make them more productive, which would boost demand for labour.[1] This paper investigates the impact of the rise of automation on labour market outcomes in the United Kingdom (UK).

Recent evidence from the United States (US) (Acemoglu and Restrepo, 2017) and Germany (Dauth et al., 2017) finds that robots had a negative impact on both employment and wages, but this question has not been considered for the UK. The case of the UK is particularly interesting because while it is a large developed economy, it finds itself in a different phase of robot adoption compared to the US and Germany, with robot adoption per thousand workers at around one twelfth of that in Germany. This paper fills that gap by identifying the long-run impact of the increased use of industrial robots on employment and wages in the UK economy.

This paper makes three distinct contributions. To the best of my knowledge, it is the first to systematically analyse the impact of robots on the UK labour market. Secondly a new measure of automation is proposed to provide further evidence on the impact of such technology on the labour market. Finally, I combine industry-level data to provide evidence on the mechanism through which robot exposure affects the labour market.

Following Acemoglu and Restrepo (2017), I estimate the equilibrium impact of industrial robot adoption on employment and wages at the level of the local labour market, proxied by Local Authorities. The industry-level variation across these regions is crucial to identification, as this information is linked to national robot adoption to approximate for local robot shocks. This permits a differences-in-differences approach which compares labour markets which had high and low exposure to robots over two decades.

A new measure of automation is introduced and adds evidence to the robot adoption analysis. I leverage UK-eligible patent data related to robots and automation at a sectoral level, and construct a local labour market measure using a Bartik-type approach (Bartik, 1991). This variable describes the extent to which a region is exposed to automation innovations. For example, if a region has a high share of employment in textile manufacturing, and many patents for textile robots have been submitted, then this area is more likely to adopt automation technology over time.

[1] To the extent that more productive capital complements certain types of labour (Autor and Dorn, 2013), or as automation reduces production costs, which increases the demand for labour in non-automated tasks (Acemoglu and Restrepo, 2018).

In order to deal with potential endogeneity whereby labour market shocks affect both the use of robots and labour market outcomes, instrumental variables are used. Towards this end, two instruments are constructed. One uses the European robots data, as in Acemoglu and Restrepo (2017). A new instrument is presented which uses global patent data, which is becoming an increasingly popular tool (Mann and Püttmann, 2018). I argue that both instruments contain different information which should influence UK robot adoption, but should not directly influence labour market outcomes.

The results suggest that industrial robots have raised employment rates in the UK. My calculations suggest that an increase of one robot is associated with the employment of ten more workers. Over the period 1993 - 2011, this translates into a rise in employment of over 60,000. This is striking when compared to the results in the US, where one additional robot is estimated to *reduce* employment by six workers (Acemoglu and Restrepo, 2017), leading to job losses in the region of half a million over a similar time period. I find that the same increase in robot adoption raises part-time pay by over £500 per year, although the income data has eleven (of 348) missing observations so results should be interpreted with a little caution.

These results are in sharp contrast to those found in the US (where employment and wages fell in response to an increase in robot exposure) and Germany (where total employment was unaffected, and the impact on wages varied over skill types). In order to get a sense of the sources of these differences, I analyse employment data at a more disaggregated level, leveraging industry information. Analysis of industries suggests that while industrial robots have led to increase in overall employment, they have reduced employment shares in some areas of manufacturing, such as automobile. These results suggest that robots have directly replaced workers - especially machine operators - in automobile and metal manufacturing. In contrast, several other industries, and especially services, have experienced an increase in employment in response to higher robot exposure.

The rest of this paper will be organised as follows. The literature on the economics of automation is discussed in Section 2. The empirical approach and data are introduced in Sections 3 and 4. The estimation results are presented in Section 5, along with further investigation of heterogeneity across industries. Supporting tables, graphs and information can be found in the Appendix.

2 Related Literature

Automation of work can affect the labour market in a variety of ways. Clearly if a machine can perform the same task more efficiently, quickly, and cheaply than a human, there is a strong incentive to automate. This "displacement effect" (Acemoglu and Restrepo, 2016) is of chief concern to many, but it is not the only mechanism of importance.

It has been suggested that automation will create spillovers, such that technologically lagging sectors will see employment rise while industries that automate will experience a fall in employment (Baumol, 1967). The question as to whether the spillovers offset the direct displacement of work is crucial.

In the task-based literature, jobs are not "lumps of labour" (Susskind and Susskind, 2015) but built up from many tasks, from routine to non-routine and manual to cognitive. More recent economic models of automation consider a continuum of tasks increasing in (loosely-defined) cognitive 'difficulty'. Machines can replace humans up to some point on this continuum, and no further.

The task-based framework allows us to focus on some of the crucial effects of automation (Acemoglu and Restrepo, 2018). The displacement effect occurs when AI and robots replace workers, which reduces the labour demand, placing downwards pressure on wages and employment. The productivity effect is the expansion of output as the cost of producing automated tasks falls, which raises the demand for labour in non-automated tasks (in both the same and other sectors). The trade-off between these mechanisms is important, and is likely to change with new automation technology, the changing distribution of labour skills and the varying demand for tasks. For example, if new technologies are only marginally better than labour in the tasks they displace, the displacement effect of automation dominates the productivity effect, reducing labour demand and wages (Acemoglu and Restrepo, 2018).

While Susskind (2017) provides a bleak prediction that labour will perform a progressively shrinking set of tasks until it is squeezed out of the economy by advanced capital, Acemoglu and Restrepo (2016) are more optimistic. More complex tasks are endogenously created in their model, (akin to the continuum 'extending') which leads to a self-correcting mechanism and a stable balanced growth path.

The idea that new tasks arise is congruent with a dynamic economy that responds to automation; there are jobs people do now that didn't exist only a few decades previously, such as software engineers and data scientists. Nevertheless, automation may also shift the distribution of jobs in the economy in potentially harmful ways. There is plenty of research on job polarisation - a fall in labour demand for middle-skill, routine work - which has been attributed to automation, trade, offshoring (Autor, 2010), changes in propensity to work in such jobs (Cortes et al., 2016), and consumers favouring variety over specialisation (Autor and Dorn, 2013).

A final concern is that competition from machines will put downwards pressure on wages. One might also hypothesise that polarisation of work is the cause of wage inequality and overall depressed wages (Goos and Manning, 2007), yet polarisation itself may be driven by automation. For example, Acemoglu and Autor (2011) find that wage differentials are driven by the complementarity of computers with high-skilled workers. The evidence on the impact on wages is mixed.

Although Graetz and Michaels (2015) find robots are associated with a boost to wages and no impact on overall employment (although some crowding out of low-skilled workers), the analysis is conducted at a country level. Such macro-level results may miss significant heterogeneity, limiting the policy-usefulness of such research.

A key innovation of Acemoglu and Restrepo (2017) is studying the impact of robots on labour market outcomes at a region $\times$ industry level. They find a robust negative relationship between exposure to robots and changes in employment *and* wages at a local level in the United States, even when controlling for baseline demographics, industrial structure, trends, other labour market shocks and excluding the regions most-exposed to the robot shock. On the other hand, Dauth et al. (2017) study the same question in Germany, finding that robots have had no impact on total employment. Despite this, robots have affected the sectoral composition of employment, with the significant fall in manufacturing jobs caused by robots being offset by additional jobs in services. The lack of employment effect does come at the cost of a fall in wages, although there is considerable heterogeneity across individuals.

Other measures of automated capital have also been introduced to the literature in recent years, such as Mann and Püttmann (2018) who construct a variable describing local exposure to patents for devices that carry out a process independently. They study the US and find a positive impact on overall employment, driven by a rise in services employment compensating a fall in the manufacturing sector. Using patents may be a more all-encompassing measure of automation innovation, although it is perhaps less clear that it *directly* relates to labour markets. For this reason, my research focuses on robot adoption, but a patents measure is constructed as an instrument.

3 Empirical Approach and Identification

This section outlines the underlying theoretical model of Acemoglu and Restrepo (2017) and explains my identification approach. The mechanism of their model is that products are made by combining a set of tasks which can be fulfilled by workers or autonomous machines. Robots can operate easier tasks, whereas only labour can complete the more complex ones. The demand for labour thus depends on the set of tasks robots can perform (the 'cutoff' point), the price of products in the goods market, and aggregate demand (Acemoglu and Restrepo, 2017).

3.1 Model

An economy is made up of commuting zones in which most adjustment to shocks takes place (Moretti, 2010; Autor et al., 2013). Each commuting zone $c \in \mathcal{C}$ has preferences defined over an aggregate of consumption of the output in each of the $\mathcal{I}$ industries:

$$Y_c = \left(\sum_{i \in \mathcal{I}} \alpha_i Y_{ci}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

where $\sigma > 0$ denotes elasticity of substitution across goods produced in different industries and α_i's are share parameters of industry i's importance in the consumption aggregate, summing to one.

It is assumed that each commuting zone consumes all its own production X_{ci} so this is equal to Y_{ci} for all $c \in \mathcal{C}$ and $i \in \mathcal{I}$.[2] Let the consumption aggregate in each commuting zone be the numeraire and the price of output of industry i in c is P_{Xci}.

[2]The results here pertain to the partial equilibrium in autarky. Details of the extended general equilibrium model with trade can be found in the Appendix, but the intuition is the same.

Each industry produces output by combining tasks on the continuum $s \in [0, S]$ where $x_{ci}(s)$ is the quantity of task s used to produce X_{ci}. Tasks must be combined in fixed proportions, reflecting perfectly complementary across task inputs, so:

$$X_{ci} = A_{ci} \min_{s \in [0,S]} x_{ci}(s)$$

where A_{ci} is productivity of industry i.

The model is task-based, where there is a cutoff point below which tasks are "technologically automated" and can be operated by workers or robots. This is simply determined by the complexity of tasks which robots can currently complete, and is exogenously determined by technological breakthroughs. This cutoff point is common across all commuting zones. Let robots be substitutes for labour in a set of tasks $[0, M_i]$ across all commuting zones so the production function for task s in industry i in commuting zone c is:

$$x_{ci}(s) = \begin{cases} r_{ci}(s) + \gamma l_{ci}(s), & \text{if } s \leq M_i. \\ \gamma l_{ci}(s), & \text{if } s > M_i. \end{cases}$$

where $l_{ci}(s)$ is labour used, $r_{ci}(s)$ is the number of robots used, and γ is the productivity of labour in each task.

It is assumed that firms will employ robots in all tasks which are "technologically autmomated" (i.e. those below the 'cutoff' M_i on the task continuum). This requires an assumption that the cost savings π_c from using robots over labour are positive. If this weren't the case, firms could use either robots or labour for tasks $s \leq M_i$, and would choose whichever factor was cheaper. This assumption ensures robots complete tasks where it is technologically feasible, as it is also cost efficient.

To find the demand for labour, two steps are taken. Firstly, *output* Y_{ci} for each industry and commuting zone is chosen to minimise cost, subject to obtaining aggregate consumption Y_c. Secondly, *production inputs* $l_{ci}(s)$ and $r_{ci}(s)$ for each industry, commuting zone and task are chosen to minimise the cost of producing X_{ci}. In equilibrium, production equals consumption within each industry in every commuting zone, so $Y_{ci} = X_{ci}$ and then labour demand can be computed.

Consider cost minimisation for the final product in each industry and commuting zone Y_{ci}. Differentiate the budget constraint $\sum_{i \in \mathcal{I}} P_{Xci} Y_{ci}$ with respect to Y_{ci} subject to the aggregate of consumption Y_c. This implies that:

$$Y_{ci} = \alpha_i^\sigma P_{Xci}^{-\sigma} Y_c$$

In equilibrium, production equals consumption, so $Y_{ci} = X_{ci}$. Thus production is $X_{ci} = \alpha_i^\sigma P_{Xci}^{-\sigma} Y_c$. Cost minimisation for production is simple, by combining the production function X_{ci} with the task production function $x_{ci}(s)$ and rearranging for labour and robots for both cases $s \leq M_i$ and $s > M_i$:

$$l_{ci}(s) = \begin{cases} 0, & \text{if } s \leq M_i. \\ \frac{X_{ci}}{A_{ci}\gamma}, & \text{if } s > M_i. \end{cases} \qquad r_{ci}(s) = \begin{cases} \frac{X_{ci}}{A_{ci}}, & \text{if } s \leq M_i. \\ 0, & \text{if } s > M_i. \end{cases}$$

Therefore the demand for labour and robots in each task is production X_{ci} divided by productivity A_{ci} (and labour productivity γ for workers) for the set of tasks in which each input has a comparative advantage.

Aggregating across tasks, and then plugging in the resulting relationship from the first cost minimisation for Y_{ci} gives demand for labour and robots in each industry and commuting zone:

$$L_{ci} = (1 - M_i)\frac{X_{ci}}{A_{ci}\gamma} = (1 - M_i)\frac{\alpha_i^\sigma P_{Xci}^{-\sigma} Y_c}{A_{ci}\gamma} \qquad R_{ci} = M_i \frac{X_{ci}}{A_{ci}} = M_i \frac{\alpha_i^\sigma P_{Xci}^{-\sigma} Y_c}{A_{ci}}$$

Aggregating across industries involves simply summing over all $i \in \mathcal{I}$. Log differentiating the labour demand equation gives:

$$d\ln L_c^d = -\sum_{i\in\mathcal{I}} \ell_{ci} \frac{dM_i}{1 - M_i} - \sigma \sum_{i\in\mathcal{I}} \ell_{ci} d\ln P_{Xci} + d\ln Y_c \tag{1}$$

where ℓ_{ci} is share of labour in industry i in commuting zone c.

The first effect in equation (1) has been called the *displacement effect*: holding prices and output constant, robots displace workers and reduce labour demand as robots are more efficient in production (Acemoglu and Restrepo, 2016).

The second and third terms of equation (1) describe what Acemoglu and Restrepo (2016) name the *productivity effect*: the improved efficiency of robots raises efficiency of production, raising output and thus the demand for labour. The second term is the price-productivity effect, as automation lowers the cost of production in an industry, that industry expands and thus increases its demand for labour. This rises with the elasticity of substitution σ between industries. The third term captures the scale-productivity effect as a reduction in costs expands total output, raising demand for labour in all industries (since industries are q-complements).

However it is not possible to simply estimate equation (1) as the cutoff point for automatable tasks M_i is not observable. This problem is solved below. Take the resulting equations from cost minimisation above, and integrate over commuting zones:

$$L_i = (1 - M_i)\frac{X_i}{A_i\gamma} \implies \frac{X_i}{A_i} = \frac{\gamma L_i}{1 - M_i} \tag{2}$$

$$R_i = M_i \frac{X_i}{A_i} \implies M_i = \frac{R_i A_i}{L_i} \tag{3}$$

Log differentiating the rearranged robot equation (3) $M_i = \frac{R_i A_i}{L_i}$ yields $dM_i \approx \frac{dR_i A_i}{X_i}$. Plugging this into the rearranged labour equation (2) $\frac{X_i}{A_i} = \frac{\gamma L_i}{1-M_i}$ gives the relationship $\frac{dR_i}{L_i} \approx \gamma \frac{dM_i}{1-M_i}$. This relates the industry change in robot per worker to the (normalised) change in the exogenous technological capability of robots M_i. The left side of this equation is observable, while the right side is not. Combining this equation $\frac{dR_i}{L_i} \approx \gamma \frac{dM_i}{1-M_i}$ with equation (1) gives a relationship between observable variables - the change in labour demand and the change in robots per worker:[3]

$$d\ln L_c^d = -\gamma \sum_{i\in\mathcal{I}} \ell_{ci} \frac{dR_i}{L_i} - \sigma \sum_{i\in\mathcal{I}} \ell_{ci} d\ln P_{Xci} + d\ln Y_c$$

3.2 Identification

The model yields a relationship between the change in labour market variables and the change in robots per worker, which is used for estimation. However the baseline regression equation may suffer from endogeneity, so two solutions will be offered: controls and instruments. One instrument is similar in spirit to that used by Acemoglu and Restrepo (2017): it leverages robot adoption in other countries as a proxy for the robotic 'technological frontier'. The patents instrument is part of a growing trend to use information on underlying innovations to measure technological progress (Mann and Püttmann, 2018).

The effect of robots on employment and on wages can be estimated by:

$$d\ln L_c = \beta_c^L \sum_{i\in\mathcal{I}} \ell_{ci} \frac{dR_i}{L_i} + \epsilon_c^L \quad \text{and} \quad d\ln W_c = \beta_c^W \sum_{i\in\mathcal{I}} \ell_{ci} \frac{dR_i}{L_i} + \epsilon_c^W$$

where ϵ_c^L and ϵ_c^W are unobserved shocks, and β_c^L and β_c^W are the coefficients to be estimated. The more general regression equation is:

$$X_{c,t+h} - X_{c,t} = \beta \underbrace{\sum_{i\in\mathcal{I}} \ell_{ci}^t \frac{R_{i,t+h} - R_{i,t}}{L_{i,t}}}_{\text{Exposure to Robots}_c} + \Gamma C_{c,t} + \epsilon_c \tag{4}$$

where $X_{c,t}$ is a labour market dependent variable in commuting zone c in base year t, and $C_{c,t}$ are control variables, many of which account for changes between t and $t+h$ (but some are simply baseline characteristics).

A univariate regression can be used to estimate these coefficients, by regressing the change in the employment or wages on a variable which proxies for robot adoption, hereafter referred to as 'Exposure to Robots'. This is an industry-weighted sum of the change in robots per worker in a commuting zone: $\sum_{i\in\mathcal{I}} \ell_{ci} \frac{dR_i}{L_i}$.

Despite the Bartik-instrument features of the constructed Exposure to Robots variable, there are two potential causes of endogeneity, which would lead to biased estimates. One reason for this is the omitted variable bias: industries may adopt robots in response to variables which also impact their labour demand, such as changes to the skill or demographic composition of the labour market. For example, if a local labour market experiences a long-term decline in specific skills of its workers, this may simultaneously raise unemployment and incentivise firms to purchase robots to fill this skill gap. Another possible endogeneity concern is reverse causality, where any shock to labour demand in a commuting zone affects the decision to adopt robots.

[3] A similar relationship for the change in wages is found in Section B of the Appendix, but requires the general equilibrium model, which isn't included here for parsimony.

For example, a cost-push shock to wages would affect labour demand and lead to a substitution towards robot adoption.

These problems are mitigated in several ways. Firstly, by using differences on both sides of the regression equation, I hope to deal with potentially important unobserved time-invariant regional factors which affect employment rates. For example, it may be that certain local labour markets contain far superior educational institutions, leading to better job outcomes. To the extent there are such time-invariant characteristics playing a role, the specification used here mitigates such concerns. Secondly, I include a set of local controls on potentially relevant labour market variables. The baseline share of manufacturing, share of routine employment and exposure to Chinese imports are especially likely to soak up local variation which may influence both employment growth and robot adoption - this has been highlighted in the literature (Autor, 2010; Autor et al., 2013; Autor and Dorn, 2013). My results show that controlling for such factors affects the results significantly and is crucial for identification. Finally, potential endogeneity is assuaged by using instruments which affect labour market outcomes only through robot adoption.

The 2SLS method used in the subsequent analysis instruments for Exposure to Robots by taking advantage of two instruments. The first is a proxy for the global technological frontier of robots: by using robot adoption data from *other* countries, it encodes information about how many robots exist worldwide and the industries in which they are used (Acemoglu and Restrepo, 2017). If robots become more widely used by competitor firms in other countries, that should incentivise UK adoption, but it is unlikely to directly affect the UK labour market, other than through the robot market channel. The patents instrument proxies for the supply of robot innovations by using automation-related patent data.[4] This contains information on the flow of new robot and automation ideas across industries, rather than just the current stock of robots. The mechanism here is that the quantity of such patents summarises the extent of automation and robot innovations in each industry, which should directly influence the choice to use robots, but should not affect labour market outcomes.

4 Data

The analysis for this research question requires data from a variety of sources. The various sources of data are introduced and descriptive analysis is provided.

4.1 Local Labour Markets

The local labour market is the unit of analysis used in this paper. If the boundaries set are too small, it is possible that commuting across the incorrectly-constructed markets will attenuate the results. This is because commuting is an endogenous response to labour market shocks.

[4] A similar variable was constructed by Mann and Püttmann (2018) concurrently. They use US utility patents and define patents as ‘automation-related’ with a linguistic classification algorithm.

Local Authorities are chosen as suitable proxies for stable local labour markets. The ONS produces Travel To Work Areas (TTWAs) which should provide appropriate boundaries for local labour markets. However, they are not suitable for this research for three reasons: (1) the change in the number of TTWAs over time, (2) the lack of data by TTWA, and (3) the arbitrary cutoffs for defining boundaries. Firstly, the changing number of TTWAs renders the cross-sectional differences-in-differences framework unusable, as I would be unable to compare the same labour market in 1991 and 2011. Secondly, issues of endogeneity would be more pronounced with less regional data: for example, I would not be able to control for baseline industry shares, which are likely to influence both robot adoption and subsequent changes to the labour market. Thirdly, the criteria used to determine the TTWAs, while sensible, are quite arbitrary - probably for simplicity - and may not best trade off internal integration (i.e. low commuting between TTWAs) and self-containment (i.e. maximising commuting within TTWAs).[5]

The number of Local Authorities does not change over the analysis period, there are fewer issues with data availability, and the number of observations (348 local labour markets) is not too far from the number of TTWAs in 1991.

To test whether Local Authorities are good proxies for local labour markets, two tests are undertaken, with both providing positive support. The first test looks at the stability of inter-Local Authority commuting behaviour between 1991 and 2011. This is achieved with a Mantel test which looks at the similarity between two distance matrices, which represent the 'closeness' of Local Authorities according to their commuting behaviour.[6] The null hypothesis that there is no relationship between the matrices is tested by randomly permuting one matrix and computing its correlation with the other matrix. This is compared to the *actual* correlation between the two distance matrices. If the matrices are unrelated, then the permuted matrix correlations should be more or less correlated with equal likelihood. The results suggest that commuting *between* Local Authorities has not changed in a statistically significant way.[7]

The second test checks if workers' travelling patterns can be well approximated by Local Authorities. I use a clustering algorithm on commuting flows between 8,800 wards in England and Wales (a more disaggregated geographical split) to check if the number of clusters is within the range of the 348 Local Authorities. There is no standard method to identify the 'optimal' number of clusters, so instead I compute clusters for a range of threshold values and different linkage criteria (methods for measuring 'distance' between wards). It seemed reasonable to choose a range of

[5]TTWAs are computed so that at least 75% of the resident workforce work in the area and at least 75% of the people who work in an area also live there. In addition, the area must have at least 3,500 economically active participants. For areas with a working population over 25,000, the required rates are lowered to 66.7%.

[6]The distance matrix is $D_{ij} = 1 - P_{ij}$ where P is the similarity matrix, computed by $P_{ij} = P_{ji} = \frac{f_{ij}+f_{ji}}{\min(rlf_i, rlf_j)}$ where f_{ij} = the number of people commuting from region i to region j and rlf_i = the resident labour force of region i.

[7]With 1000 permutations, the p-value is 0.000999, suggesting that the matrices have a strong relationship, with a correlation of 0.983, using the **ade4** package in R

thresholds that had a relatively small impact on the number of clusters, at the margin. The results suggest that the region of 300 - 450 clusters is a good approximation of the commuting clusters - not far from the number of Local Authorities. Full details on the commuting datasets and methodology are in Section C of the Appendix.

4.2 Robots Data

The International Federation of Robotics (IFR) provides data on industrial robots, and has recently been used by a number of researchers to answer automation-related questions (Acemoglu and Restrepo (2017); Graetz and Michaels (2015)). The IFR dataset contains the stock and flow of industrial robots by industry, country and year, based on annual surveys of robot suppliers. The IFR defines an industrial robot as "an automatically controlled, reprogrammable, and multipurpose [machine]" (IFR, 2017).[8]

The analysis that follows in this paper focuses on industrial robots, as there is no adequate data available on services robots for the UK.

The raw industrial robots data is less meaningful than that normalised by the number of workers, which will be discussed in the next section. However, some numbers will provide a bit of context. The total stock of industrial robots in the UK doubled from around 7,500 in 1993 to 15,000 in 2007, but this fell to 13,500 by 2011. In this period, between 40 - 60% of these robots were employed in the automobile sector, highlighting its relative importance for the robot industry. As a proportion of global industrial robots, the UK was home to somewhere in the range of 1-2% between 1993 and 2011. For comparison, the UK has produced 3.5-5.5% of global GDP over this time.[9]

4.3 Employment Data

Aggregate employment data is available from EUKLEMS. It is used here to analyse national trends in robots per thousand workers. For data split by industry, occupation and region, the UK census (via NOMIS) provides extensive margin data. The Annual Survey of Hours and Earnings provides detailed intensive margin data.

Overall employment figures for EUKLEMS are used for UK-wide trends. The data is annual and split by industry. A mapping between the industries in the robots and employment data is discussed in Section F of the Appendix - data is grouped in 16 industries for the analysis. The summary data is provided in Table A1 in the Appendix, for a selection of industries.

The average number of robots per thousand employees across the UK has risen from 0.38 to 0.54 from 1995 to 2011, but this hides heterogeneity at the industry and regional levels. Unsurprisingly, the biggest gains have been found in automobile manufacturing, with a leap from 11.5 to 32.0 robots per thousand employees. It

[8] The IFR notes that there are difficulties in calculating the operational stock, as not all countries take surveys of robot stocks. In cases where this data isn't available, IFR assumes an average service life of 12 years followed by immediate withdrawal.

[9] `https://data.worldbank.org/indicator/NY.GDP.MKTP.CD`

is also important to highlight that services have increasingly made up the majority of jobs, with manufacturing - especially automobile manufacturing - on the decline. Therefore the impact of industrial robots on jobs is likely to be concentrated in the manufacturing industry, which has shrunk to just below 10% of UK employment.

The UK census provides detailed extensive margin employment data at a highly disaggregated level, in terms of industry, region and occupation type. The employment data is from a 10% sample in 1991 and 2011, obtained from NOMIS. This allows various employment ratios to be computed across Local Authorities, with employment measured as Full-Time (FTE) or Total (E) employment, and normalised by either the Population (Pop) or the Working Age population (WA).

The summary statistics for these employment ratios are in Table A3 in the Appendix, with three of our four dependent variables showing a decline in (Local Authority) population-weighted terms.

The regional × industry census data allows computation of Exposure to Robots, introduced in Section 3.2: $\sum_{i \in \mathcal{I}} \ell_{ci} \frac{dR_i}{L_i}$. Each Local Authority receives a value for Exposure to Robots based on this calculation. To adapt to the data limitations, the formula is:

$$\begin{array}{l}\text{Exposure to Robots} \\ \text{from 1993 to } 2011_c\end{array} = \sum_{i \in \mathcal{I}} \ell_{ci}^{1991} \left(\frac{R_{i,2011}}{L_{i,1991}} - \frac{R_{i,1993}}{L_{i,1991}} \right) \tag{5}$$

The distribution of Exposure to Robots over Local Authorities is summarised in Table A2. There are some negative values, but the distribution is bunched close to zero with the mean and third quantile at 0.20. There are a number of large values with the maximum at over 4 additional robots per thousand workers between 1993 - 2011.

For the intensive margin, data from the Annual Survey of Hours and Earnings (ASHE) is obtained.[10] The measures used in the subsequent analysis are mean hours for full-time and part-time workers, from 1997 and from 2011 across Local Authorities. The data is available from the Office for National Statistics (ONS), and is based on a 1% sample of employee jobs from Pay As You Earn (PAYE) records, and excludes self-employed, or those for who earnings were affected by absence (e.g. sickness).

The summary statistics are in Table A4. There was a decline in the mean hours worked for both full-time and part-time workers. In fact, there was a fall in mean hours worked for part-time workers in over 66% of the 348 Local Authorities, and for full-time workers in over 83% of these regions.

4.4 Patents Data

A new measure of automation is introduced, which leverages data on UK-eligible automation-related patents. The patent data comes from the World Intellectual Property Organisation (WIPO) which provides an online platform for searching the database of global patents.[11] Full details of the search method and industry mapping

[10] https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/placeofworkbylocalauthorityashetable7

[11] https://https://ipportal.wipo.int/

are in Sections G and F of the Appendix. The variable is computed in the same way as for Exposure to Robots:

$$\begin{aligned}\text{Exposure to Automation}\\ \text{Patents from 1991 to 2011}_c\end{aligned} = \sum_{i\in\mathcal{I}} \ell_{ci}^{1991}\left(\frac{P_{i,2011}^{UK}}{L_{i,1991}} - \frac{P_{i,1991}^{UK}}{L_{i,1991}}\right) \tag{6}$$

where $P_{i,t}^{UK}$ are UK-eligible automation-related patents in industry i in year t. The distribution of Exposure to Automation Patents over Local Authorities in summarised in Table A2. The distribution is approximately bell-shaped with a right-skew.

4.5 Income Data

The ASHE provides wage data at the Local Authority level. The data contains the mean and median weekly earnings across Local Authorities, along with earnings deciles along the pay distribution. The data coverage for earnings deciles is somewhat lacking, but the missing observations don't seem to be heavily skewed across observables.

Each variable for weekly pay has seen an increase from 1997 to 2011 in its mean value across Local Authorities (see Table A5 in the Appendix). The increase in the mean pay variables has been greater than that of the median pay measures, implying an increased right skew in the earnings distribution, as noted in Gosling et al. (2000) for the UK. Only a few observations are missing in the first three columns (full-time pay and *mean* part-time pay), but over 100 observations are missing for median Part-Time pay, so any results for this dependent variable should be treated with caution. On further investigation, the missing observations have a very similar distribution across observables as the available data, which slightly mitigates this concern. The regions without data for median part-time pay are not outliers in robot adoption. They are also not outliers for all employment and control variables in the dataset, but these regions do tend to be smaller (i.e. have lower populations). Therefore the results should be fairly robust to the missing data. Fortunately, my overall conclusions do not hinge on this particular dependent variable.

The ASHE earnings decile data is also collected to investigate the changes in various wage percentile ratios over time. Three log wage ratios are computed to represent different aspects of the earnings distribution (Autor et al., 2005).[12] The summary statistics are described in Table A6. The arithmetic mean of the earnings ratios across Local Authorities has increased for each of these three measures of earnings inequality. However, the changes in Log 50-10 and Log 80-10 are negative and close to zero, respectively, when weighted by the 1997 employment data.

4.6 Instruments

For the 'world' technological frontier, I employ data from EU countries (hereafter referred to as the EU-7) (Acemoglu and Restrepo, 2017). The economies which

[12] The 80^{th} decile is chosen over the more commonly-used 90^{th} decile due to data limitations.

have detailed industry-level employment data back to the 1990s in EUKLEMS are Denmark, Finland, France, Germany, Italy, Spain and Sweden. In 2009, these seven countries accounted for 45% of the global stock of industrial robots, although that has fallen in recent years (International Federation of Robotics, 2017).

Aggregate EU-7 robot data is used to construct the instrument[13]:

$$\text{Exogenous exposure to robots from 1995 to } 2011_c = \sum_{i \in \mathcal{I}} \ell_{ci}^{1981} \left(\frac{R_{i,2011}^{EU-7}}{L_{i,1995}^{EU-7}} - \frac{R_{i,1995}^{EU-7}}{L_{i,1995}^{EU-7}} \right) \tag{7}$$

using the local UK industry employment shares in 1981 and the aggregated EU-7 data for the change in robots by industry and the worker normalisation. Data from 1981 is used as this was the oldest regional-industrial UK employment data available, and it holds important information on historical differences across Local Authorities.

As described in Section 3.2, a new instrument for Exposure to Robots is introduced in this paper, using data on global automation-related patents in 1981 and 1991. The patent data comes from the WIPO.

For the supply of global automation innovations, the instrument leverages automation-related patent data:

$$\text{Exposure to Global Automation Patents from 1981 to } 1991_c = \sum_{i \in \mathcal{I}} \ell_{ci}^{1981} \left(\frac{P_{i,1991}^{Global}}{L_{i,1981}} - \frac{P_{i,1981}^{Global}}{L_{i,1981}} \right) \tag{8}$$

where $P_{i,t}^{Global}$ are global automation-related patents in industry i in year t.

The summary statistics for these two instruments are in Table A2 in the Appendix. The Exogenous exposure to Robots variable has much higher values than Exposure to Robots, which is unsurprising given it is computed using data from other European countries where robot adoption is significantly higher than the UK. However, the distribution follows a similar shape. The Exposure to Global Automation-related Patents variable has an approximately bell-shaped distribution, with a right-skew.

4.7 Control Variables

There are many structural differences between Local Authorities which may confound the relationship between robots and labour market outcomes. Therefore, the following set of control variables are included in the regressions. They are chosen to be similar to existing studies in the US and Germany for better comparison of results (Acemoglu and Restrepo, 2017; Dauth et al., 2017).

Regional Dummies: controls for 9 regions across England and Wales, for structural differences across these labour markets. These regions are: East Midlands, East of England, London, North East, North West, South East, South West, Wales, West Midlands, Yorkshire and The Humber.

Demographics: changes between 1991 and 2011 in the share of the working-age population (to control for changes to the labour force), the share of the population

[13]Unfortunately the EU-7 data only exists back to 1995.

that is of white ethnicity (to control for potential changes in labour market discrimination), and the percentage change in the population size (to control for broader aggregate demand- and supply-side changes).

Broad industry shares: the 1991 baseline shares of employment in manufacturing and construction and female employment in manufacturing. As robot adoption is concentrated in these industries, it needs to be controlled for.

Trade and routinisation: exposure to Chinese imports (likely to have significant impacts on the labour market, as shown for the US in Autor et al. (2013)). The data is from Eurostat database.[14] I also control for 1991 baseline share of employment in routine jobs (to control for jobs more likely to be automated, as in Autor and Dorn (2013)).

The computed variables (robots, patents, routineness, trade) are checked for multicollinearity by calculating the partial correlations between them, controlling for demographic and broad industry shares. There is no evidence of multicollinearity apart from a high partial correlation between the Exposure to China and Exposure to Automation-related Patents. Therefore, there might be issues when introducing trade controls, which will be checked.

5 Results

This section contains the results and discussion from a range of regression models which analyse the impact of industrial robots on labour market outcomes. The estimation methods are OLS and 2SLS alongside a selection of controls, with observations at a Local Authority level across England and Wales. The results describe the impact of robot adoption on employment at the extensive and intensive margins, earnings, and changes in the income distribution. The estimated coefficients on Exposure to Robots are shown, for an increasing set of these covariates. The controls chosen are very similar to those in Acemoglu and Restrepo (2017) and Dauth et al. (2017), which allows for a cleaner comparison with results over an almost identical time frame, and a comparable methodology, in the US and Germany.

The regression specification takes the form:

$$\Delta y_c = \beta \underbrace{\sum_{i \in \mathcal{I}} \ell_{ci}^{1991} \left(\frac{R_{i,2011}}{L_{i,1991}} - \frac{R_{i,1993}}{L_{i,1991}} \right)}_{\text{Exposure to Robots}_c} + \Gamma C_{c,1991} + \epsilon_c$$

where Δy_c is the change in the dependent labour market variable. For the 2SLS estimation, Exposure to Robots will be instrumented by the two aforementioned instruments: one proxying for the global technological frontier of robots, the other for the supply of robot innovations.

[14] http://ec.europa.eu/eurostat/web/international-trade-in-goods/data/database

5.1 Economy-wide Results

The results are presented for the impact of robot adoption on both employment and wages at the aggregate level, for a range of models and outcomes. Robustness checks are mentioned where relevant, and an extended discussion on this topic can be found in Section D of the Appendix.

5.1.1 Employment and Hours

To test for a relationship at the extensive margin, the change in the employment ratio from 1991 to 2011 is regressed on the Exposure to Robots and a set of baseline controls. The dependent variable is computed as $\left(\frac{L_{c,2011}}{N_{c,2011}} - \frac{L_{c,1991}}{N_{c,1991}}\right)$, where L is an employment measure and N is a population normalisation.

On the intensive margin, mean hours worked is regressed on robot adoption, for both full-time and part-time workers. This requires defining the dependent variable as the change in mean hours worked $h^{mean}_{c,2011} - h^{mean}_{c,1997}$.

All regressions are computed with OLS and 2SLS estimation using the two instruments already introduced. Four sets of covariates are considered for each dependent variable, and each of these models can provide useful information. The baseline results for selected dependent variables can be seen in Table 1, with the rest in Table A8 in the Appendix.

The 2SLS estimation method allows testing of weak instruments, endogeneity of Exposure to Robots, and the validity of the instruments. Across all baseline regressions, the instruments are found to be strong and valid.[15,16] There is mixed support for endogeneity of Exposure to Robots, hence I also report OLS estimates.[17] All test results can be found in Table A11 in the Appendix.

Firstly, I report the differences between the OLS and 2SLS estimates in Table 1. Columns (1) - (4) are estimated by OLS, while columns (5) - (8) use 2SLS. For employment at the extensive margin in Panel I, the point estimates are fairly similar, especially in the fully-specified models. For employment at the intensive margin in Panel II, the OLS estimates are much larger than the 2SLS estimates, and statistically significant while the latter are not. This result is unsurprising due to the evidence of endogeneity of Exposure to Robots in the part-time hours regression (see Table A11). For the wage and inequality regression results contained in Panels III - V, the overall results are broadly similar between OLS and 2SLS - where the point estimates differ in magnitude, the standard errors are large enough that I cannot interpret significance. Again this is consistent with the Hausman test statistics, which show that for these regressions there is little evidence of endogeneity.

For employment at the extensive margin, consider Panel I in Table 1. The estimated coefficient on robot exposure is negative and statistically significant in column

[15] F-test on first stage rejects the null that instruments are weak with p-values always below 0.1%.

[16] Sargan tests provide support that instruments are valid such that they are not correlated with the error term, as p-values are mostly under 2%.

[17] Wu-Hausman tests find the 2SLS coefficients are significantly different from OLS in some of the models estimated, but not all.

Table 1: The estimated coefficient on Exposure to Robots on UK labour market outcomes using OLS & 2SLS estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	I: *Employment - Extensive Margin* (Δ FTE/WA, $n = 348$)							
Robot Exposure	-0.004	0.003	0.005^{**}	0.004^{*}	-0.013^{**}	-0.003	0.006^{**}	0.007^{**}
(1991 - 2011)	(0.005)	(0.003)	(0.003)	(0.003)	(0.005)	(0.005)	(0.003)	(0.003)
	II: *Employment - Intensive Margin* (Δ Part-Time Hours, $n = 348$)							
Robot Exposure	0.50^{**}	0.48^{**}	0.50^{**}	0.50^{**}	0.21	0.11	0.21	0.22
(1993 - 2011)	(0.24)	(0.23)	(0.21)	(0.22)	(0.36)	(0.36)	(0.33)	(0.33)
	III: *Wages* (Δ Ln Mean Full-Time Pay, $n = 345$)							
Robot Exposure	-0.007	-0.007	-0.009	-0.009	0.004	0.006	0.001	0.001
(1993 - 2011)	(0.006)	(0.007)	(0.006)	(0.006)	(0.014)	(0.013)	(0.008)	(0.006)
	IV: *Wages* (Δ Ln Mean Part-Time Pay, $n = 337$)							
Robot Exposure	0.086^{***}	0.084^{***}	0.084^{***}	0.087^{***}	$0.074^{\dagger}$	0.065	$0.070^{\dagger}$	0.080^{**}
(1993 - 2011)	(0.032)	(0.028)	(0.027)	(0.027)	(0.047)	(0.046)	(0.045)	(0.039)
	V: *Inequality* (Δ 80/10 Ratio, $n = 242$)							
Robot Exposure	0.024^{**}	0.012^{*}	$0.010^{\dagger}$	0.009	0.032^{***}	0.012	-0.002	0.003
(1993 - 2011)	(0.01)	(0.007)	(0.006)	(0.007)	(0.01)	(0.009)	(0.014)	(0.012)
Controls:								
Weight by population	✓	✓	✓	✓	✓	✓	✓	✓
Regional dummies	✓	✓	✓	✓	✓	✓	✓	✓
Demographics		✓	✓	✓		✓	✓	✓
Broad industry shares			✓	✓			✓	✓
Trade & routinisation				✓				✓
2SLS					✓	✓	✓	✓

Note: Long-run estimates of the impact of the exposure to robots on labour market outcomes. All regressions are weighted by baseline population, have regional dummies, and Liang and Zeger (1986) cluster-robust standard errors are reported in brackets (clustered at the regional level). Demographic controls are the changes between 1991 and 2011 in the share of working-age population, the share of the population that is of white ethnicity and the percentage change in the population size. Broad industry shares control for 1991 baseline shares of employment in manufacturing and construction and the share of female employment in manufacturing. Trade and routinisation controls for the exposure to Chinese imports and the 1991 baseline share of employment in routine jobs as defined in Autor and Dorn (2013). Instruments for 2SLS are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents. Note that results with *** are significant at a 1% confidence level, ** at a 5% level, * at a 10% level, and † at a 15% level.

(5), where controls for population and regions are included. This is consistent across the four employment ratios analysed. However, once controls are added for structural differences between Local Authorities, in terms of demographics, baseline industrial composition, trade and routinisation, the estimated coefficients flip to positive, with statistical significance across most extensive margin dependent variables. This can be seen in Figure 1.

Estimated 2SLS coefficients on Exposure to Robots on Employment Rates

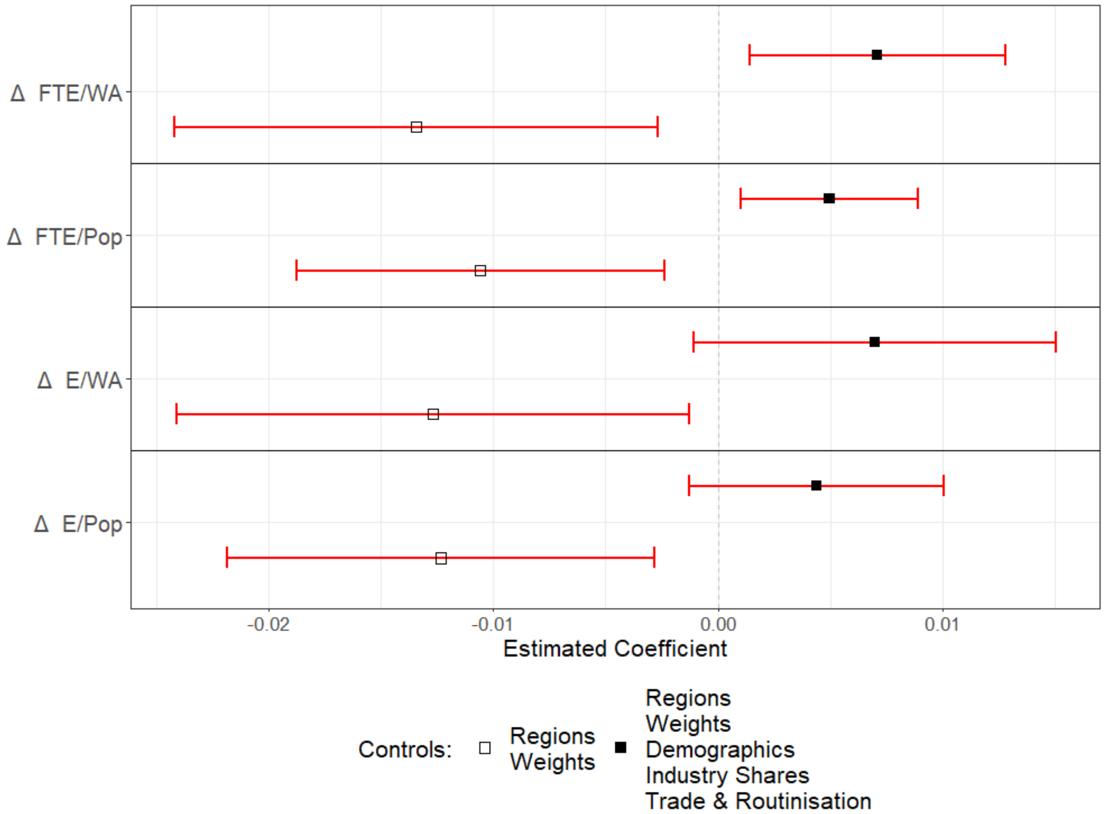

Figure 1: The points are estimated 2SLS coefficients of Exposure to Robots on the change in employment-to-population rates (1991 - 2011). The lines indicate the 95% confidence intervals. Each panel shows the results for a different measure of the change in the employment rate, for two models with a different set of covariates (the filled point contains the full set of controls).

The statistically significant negative coefficients prior to introducing most covariates shows that simple correlations between robots and employment indicate robots are replacing jobs. However, the careful use of relevant covariates can reverse this result.[18]

It is reasonable to consider *why* the coefficient flips when further controls are added, as this might raise concerns about multicollinearity. To examine this further,

[18]Furthermore, these included control variables are almost always statistically significant at very low thresholds.

I provide a breakdown of the variance inflation factors (VIF) (Fox and Monette, 1992) in the fully-specified models, with all controls. Table A13 in the Appendix contains Generalised VIFs and all values are far below the rules-of-thumb (5, or 10) considered in the literature (O'brien, 2007), suggesting that issues of multicollinearity are not substantial.

I found non-negligible partial correlations between the instruments and Exposure to Chinese Trade (when controlling for demographic and broad industry shares). Therefore, a further check for multicollinearity is provided by excluding the trade control in Section D of the Appendix. However, the magnitude and significance of the estimated coefficients do not change in an economically meaningful way.

So what best explains the change in the sign of the estimated coefficient for employment at the extensive margin? For the 2SLS regressions, the flip occurs when I control for broad industry shares. These variables have wide-ranging, and sometimes substantial, correlations with the dependent variable (employment ratios) and the variable of interest (Exposure to Robots). Crucially, the correlations go in different directions for the dependent and independent variable, for each of these 1991 baseline employment shares. The most stark example is the manufacturing employment share in 1991, which has a 22.2% correlation with Exposure to Robots, and a -19.7% correlation with the change in full-time employment per person. Essentially, the results suffer from significant omitted variable bias without the controls for industry shares. Notice also that for full-time employment rates, the inclusion of these control variables changes the coefficient from insignificant to significant at the 5% level. This evidence suggests that (1) baseline industry shares are important for identification, and (2) their opposing relationships with the dependent variable (employment rate) and Exposure to Robots lead to drastic changes in the sign - and magnitude - of the estimated coefficients.

Importantly, these results suggest that robot adoption raises employment rates. This can also be expressed as one additional robot increasing employment by around 10 workers.[19] In the UK between 1993 - 2011, the total number of robots increased by 6,165. Thus the estimates suggest a rise in employment of over 60,000 as a result of robot adoption over this 20 year period. Simply using the average rise in Exposure to Robots of 0.20 multiplied by the point estimates suggests a rise in employment rates of 0.09 - 0.14%, which translates to 40,000 - 50,000 more jobs. For context, total employment increased by around 1.15 million between 1991 and 2011. Therefore my results can account for around 3 - 5% of the rise in employment. It is important to note that these estimates are subject to significant estimation uncertainty, so they should be considered instructive for gauging the broad impact of robots, but are by no means precise.

This result is in stark contrast to the findings in the US, where an extra robot

[19] One extra industrial robot per thousand workers over 1991 to 2011 is equivalent to raising the full-time employment-to-population ratio by approximately 0.005, per the regression coefficients. Dividing this by the number of robots per person yields the change in employment per robot. One extra robot per thousand workers is equivalent to $24,873 (= 24,872,843/1000)$ more robots, or roughly $0.47 (= 24,873/52,983,335 \times 1000)$ robots per thousand people. Thus one extra robot raises employment by $10.65 (= 0.005 \times 1000/0.47)$ workers.

reduced employment by around 6 workers (Acemoglu and Restrepo, 2017). Over a similar time period, this translates into hundreds of thousands of lost jobs in the US (estimates range from 360,000 - 670,000). On the other hand, results from Germany found no aggregate employment losses, but an extra robot did reduce *manufacturing* employment by 2 workers, and this led to a fall of around 275,000 manufacturing jobs (Dauth et al., 2017). Notice that both countries had greater adoption than the UK, with an increase of 1.0 (US) and 5.6 (Germany) robot(s) per thousand workers compared to 0.2. Therefore the effect on overall employment rates was a reduction of 0.18 - 0.34% in the US, and no effect in Germany, in contrast to the 0.09 - 0.14% rise in the UK.

Notice that the statistical significance is stronger for full-time employment rates (FTE) than total employment rates (E). The latter includes part-time employment, which might be affected differently by the technological shock. Further evidence to support this hypothesis is provided when considering the effect of robot adoption on wages in Section 5.1.2.

Further investigation into the Local Authorities most-affected by robot adoption suggests that it is these labour markets that drive the positive employment effect. By removing the 5% of Local Authorities with the highest Exposure to Robots, the estimated coefficients become small, negative and statistically insignificant (see Table A9 in the Appendix).

On the intensive margin, the results are in Panel II of Table 1 for part-time hours, and in A8 for full-time hours. The OLS estimates are positive, and statistically significant for part-time and full-time hours worked. However, there is evidence of endogeneity (see Table A11), so the 2SLS estimates are preferred. The 2SLS estimated coefficients on Exposure to Robots are not found to differ significantly from zero for either part-time or full-time employment at the intensive margin. The point estimates are positive for both full-time and part-time hours worked, and quite large for the latter. However, the robust clustered standard errors are also large, so the null hypothesis of no relationship between robot adoption and hours worked cannot be rejected. These results are robust to removing the most-affected Local Authorities.

5.1.2 Wages and Inequality

The regression is also computed for changes in mean and median log weekly pay, and for changes in the log earnings decile ratios. For the wages and inequalities data, there is little evidence of endogeneity (see Table A11 in the Appendix) so the OLS results are more efficient and preferred.

The results for the impact of robot exposure on *mean* wages can be seen in Panels III and IV of Table 1. The OLS estimated coefficients are negative for the change in mean full-time pay, but they are very small and statistically insignificant. However for mean part-time pay, the results are larger and there is statistical significance across all the models. In fact for mean part-time pay, the estimated coefficient is largest and has the lowest p-value when all controls are included. The evidence suggests that robot adoption led to an increase in part-time wages but not full-time pay. The results therefore suggest that the technological shock affects full-time workers on the quantity side (i.e. employment) but part-time workers through prices (i.e. wages).

The regressions with the change in *median* pay as the dependent variable show no statistically significant relationship with robot adoption, as seen in the Table A8. Firstly, I note that these regressions - especially for part-time pay, had more missing observations due to data availability. Perhaps more importantly, the lack of impact on median earnings might tell us something about the effect of robot adoption on the *distribution* of pay. The mean is more likely to be affected by changes at the extremes of the distribution than than the median. This motivates an investigation into the impact of robots on inequality.

Estimated OLS coefficients on Exposure to Robots on Weekly Pay

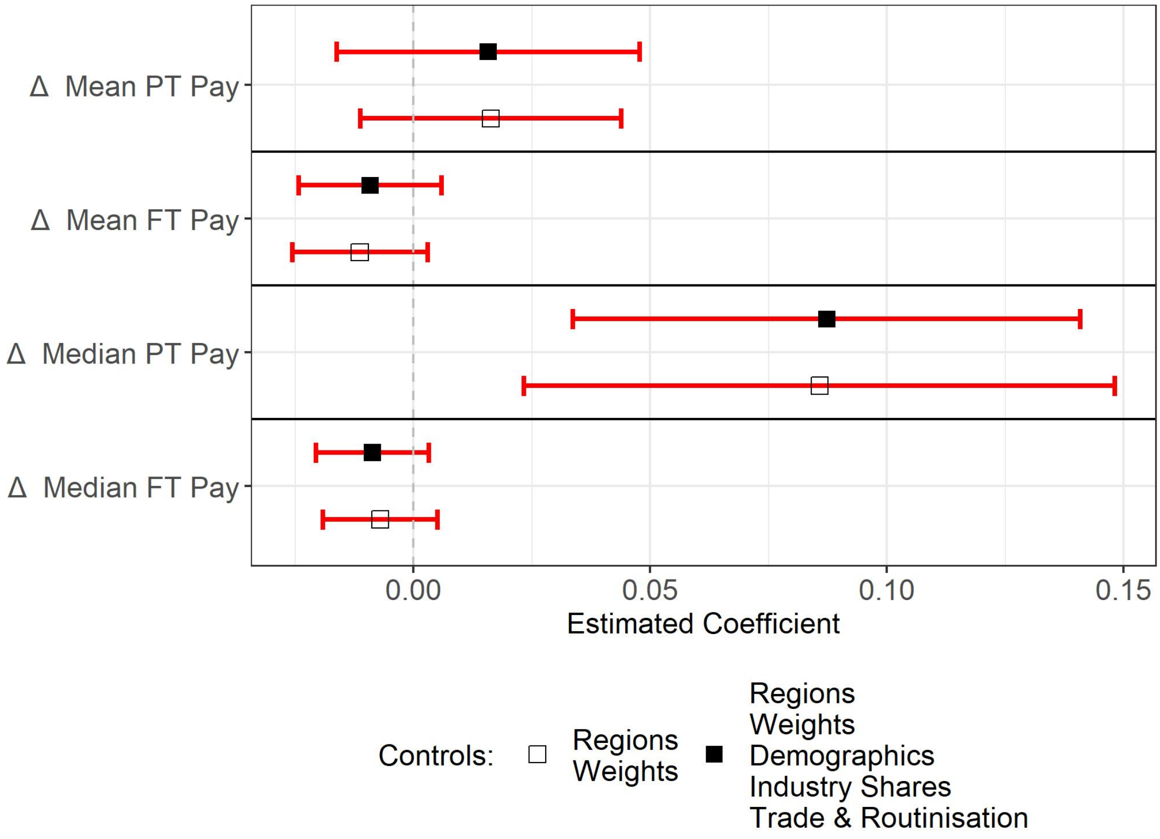

Figure 2: The points are estimated OLS coefficients of Exposure to Robots on the change in weekly pay (1997 - 2011). The lines indicate the 95% confidence intervals. Each panel shows the results for a different measure of the change in weekly pay, for two models with a different set of covariates (the filled point contains the full set of controls).

This evidence suggests that adopting one more robot per thousand workers increases part-time pay in the range of £537 per year for the average part-time worker.[20] In the UK between 1991 - 2011, robot adoption increased by 0.20 on average across

[20] From the baseline of £113.7 in 1997, the estimated coefficient of 0.087 suggests that one more robot per thousand workers raises part-time weekly pay by $9.1\% = e^{0.087} - 1$. This leads to a weekly wage of $124.0 (= 113.7 \times 1.091)$ which is an increase of £10.33 per week, or around £537 over a year.

the 348 Local Authorities, leading to an increase in annual pay of over £107 for part-time workers. An increase in robot exposure of 0.20 is associated with an increase in mean part-time pay of 1.82%. For context, the average increase in part-time pay from 1997 - 2011 was 72.7%. Thus the use of industrial robots can explain over 2.5% of the change in mean part-time pay.

In contrast, an extra robot per thousand workers in the US is estimated to *reduce* wages by 0.25 - 0.5% (note this is for all employed persons, so the levels are much higher) which translates to a fall in yearly pay of around $200 (Acemoglu and Restrepo, 2017). The estimates from Germany are much lower, with an extra robot per thousand workers reducing average pay by around €15 per year (Dauth et al., 2017).[21]

It should be noted that this relationship between robot adoption and wages seems to be driven by the most-affected Local Authorities. When these regions are excluded, the estimated coefficients are negative and statistically significant, for both mean and median full-time *and* part-time wages (see Table A9 in the Appendix). This suggests that the effect of robots on labour income are heterogeneous across labour markets.

The effect of robot exposure on inequality is estimated by regressing the change in log earnings decile ratios on Exposure to Robots. The results are in Panel V of Table 1. The estimated coefficient on the 80/10 inequality ratio is found to be positive and significant where the only controls are regional dummies and population weights. However, as controls are added, the magnitude of the coefficient falls close to zero and there is no statistical significance. It seems that the introduction of demographic covariates, especially the change in the proportion of the population that is white, plays a crucial role. It seems likely that this covariate is a proxy for structural factors that are influencing changes in inequality, and that the robot shock is unimportant. The results are similar for the 80/50 ratio, while there is no statistical significance for the change in the 50/10 ratio (see Table A8).

Overall, there is substantial evidence that robot adoption has led to an increase in full-time employment rates and part-time wages in the UK over this time period, suggesting that distinct types of workers are differentially affected by the technology. This contrasts with findings in the US and Germany, which should not be surprising given the substantially different structure of their labour markets and the contrasting uptake of industrial robots over this period.

5.2 Industry Analysis

The automobile industry is of crucial importance when analysing the impact of industrial robots in the UK. On average from 1993 - 2011, this industry accounted for over 55% of robots in the UK. But even more important - at least when looking at the *change* in exposure to automation - is that the automobile industry is responsible for almost 80% of the change in robots over this time period.

[21] The computation here uses the estimates in Dauth et al. (2017). The daily fall in wages is $-€0.05 = (e^{\frac{-0.0417}{100}} \times 120.7) - 120.7$ and the average worker is employed for 5,959 over 20 years, which is almost 300 days per year. Thus the total annual fall in wages is $-€15 = -0.05 \times 5,959/20$.

This section aims to look more carefully about the relationship between industrial robots and employment across industries. The findings suggest that robot adoption has reduced employment in some manufacturing sectors, while boosting employment in services. The displaced workers seem to mostly be machine operators (see Section E of the Appendix for occupation results). More detailed analysis suggests that robots purchased for the automobile industry have played a crucial role in the overall employment effects.

The regression specification used in this section takes the form:

$$\Delta \text{Emp. Share}_c^i = \beta \underbrace{\sum_{i \in \mathcal{I}} \ell_{ci}^{1991} \left(\frac{R_{i,2011}}{L_{i,1991}} - \frac{R_{i,1993}}{L_{i,1991}} \right)}_{\text{Exposure to Robots}_c} + \Gamma C_{c,1991} + \epsilon_c$$

where ΔEmp. Share$_c^i$ is the change in the employment share for industry i in Local Authority c between 1991 and 2011. The change in employment shares are computed for each industry between 1991 and 2011. 2SLS estimation is performed on the change of employment shares across 14 industries with the full set of controls.[22] The point estimate and 95% confidence intervals are shown for the industry breakdown in Figure 3.

There is evidence that robot adoption impacted employment shares for around half the industries considered. In particular, the estimated coefficient is large and statistically significant for High Tech Manufacturing and services.[23] There is also evidence that industrial robots have affected employment shares in some other manufacturing sectors.

So how can these results be interpreted? Firstly, it is important to note that *most* industrial robots are employed in manufacturing industries, especially the automobile industry. Thus, it is not surprising this is the industry that is most affected. Secondly, the results are not sensitive to excluding the Local Authorities most-exposed to the robot shock (see Figure A1 in the Appendix). It wouldn't necessarily be a problem if that were the case, but this suggests that the aggregate industry effects aren't driven by a small proportion of manufacturing-heavy regions. Furthermore, these results fit with existing evidence (Acemoglu and Autor, 2011; Autor and Dorn, 2013) and theory (Acemoglu and Restrepo, 2016) that robots will replace routine, manual manufacturing tasks, while boosting labour demand through the productivity effect for workers that are complementary - or, at least, not substitutable - with such machinery. Finally, one could interpret these results as highlighting structural shifts in employment towards services. However, I believe this conclusion is false. These results are from models with the full set of controls, including baseline industry shares. Regions with high manufacturing employment in 1991 might be *expected* to see a greater reduction in this sector over time, as the UK moved towards a service sector economy. By controlling for this variation, my estimates isolate how exposure to robot

[22] Limitations of 2011 data necessitated merging some industries.

[23] High Tech Manufacturing includes computers, electrical equipment, machinery, automobiles and other vehicles.

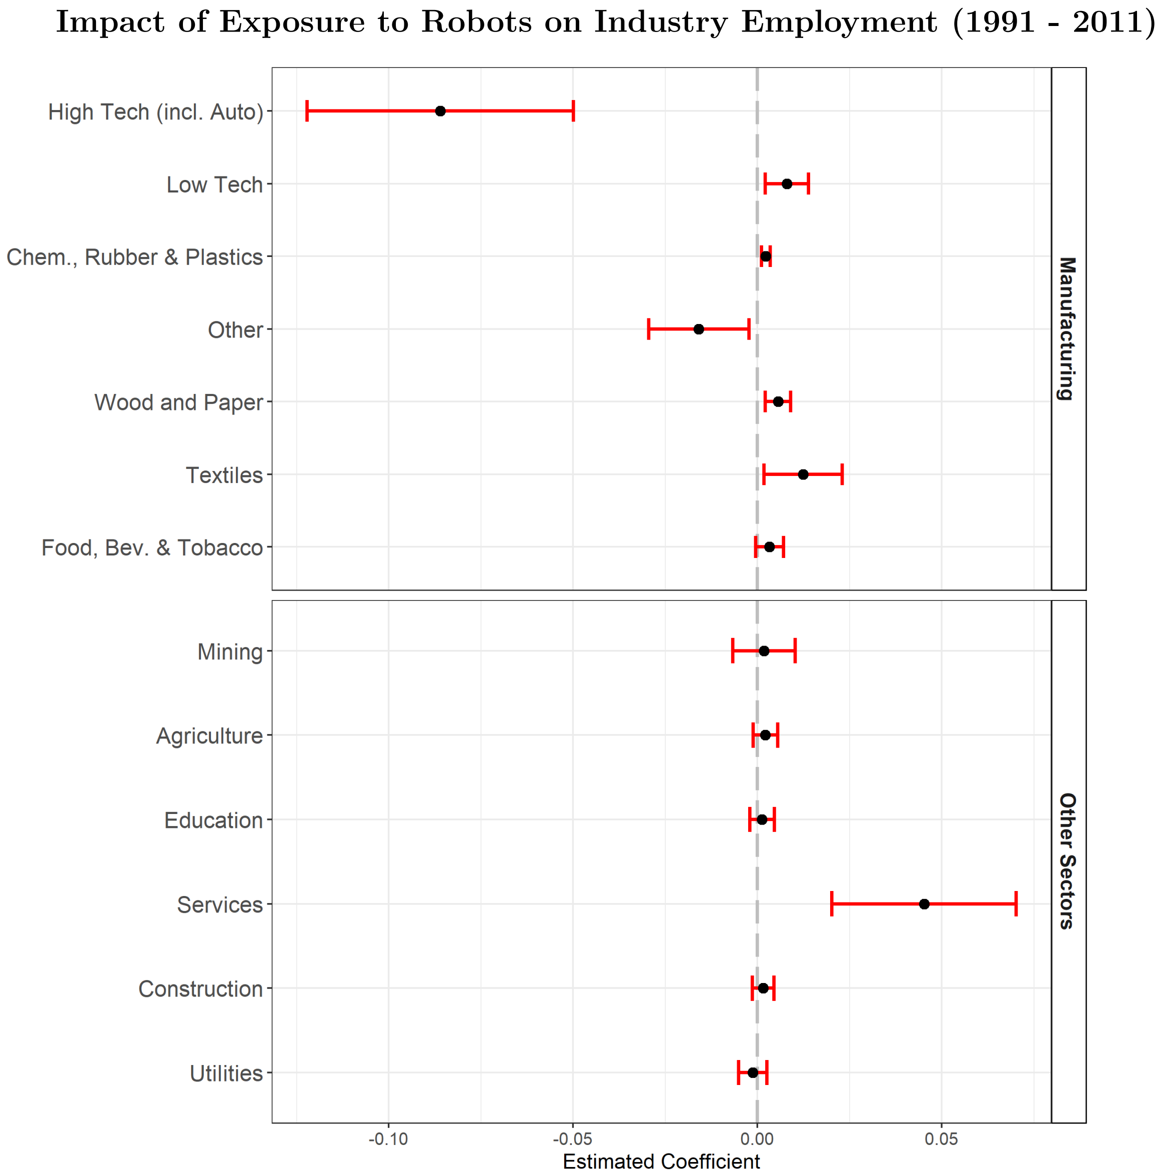

Figure 3: The points represent estimated 2SLS coefficients of Exposure to Robots on the change in employment shares ratio (1991 - 2011) across a set of industries. The lines indicate the 95% confidence intervals. The models have the full set of controls used in baseline estimation.

adoption affects industry employment shares, independent of the initial observed structure of local labour markets.

However, despite my confidence in these controls, there is still the possibility the results pick up regression to the mean. For example, in Local Authorities with a (relatively) large share of workers in the automobile industry, you would expect the employment share in High Tech (incl. Auto) to fall if there is a factory closure. On the flip side, if a region contains no automobile plant but one opens up over this period, we get the opposite result. Given the significant correlation between Exposure to Robots and the automobile employment share, the estimates may absorb this regression to the mean effect.

Evidence that this problem might be significant can be seen by running the following regression: the change in the share of automobile manufacturing employment between 1981 and 1991 on the baseline share in 1981 (including regional fixed effects and demographic controls).[24] The estimated coefficient is negative, large and significant.[25] This means that Local Authorities with a high automobile employment share in 1981 saw a greater fall in this share, even accounting for demographic changes and regional effects.

The presence of regression to the mean doesn't necessarily invalidate my conclusions. If the 2SLS approach is well-specified, then the instrumental variables should alleviate this problem. More concretely, I require evidence that the instruments are informative and valid, and that the instruments are not highly correlated with the automobile employment share. The instruments are strong and valid across all industries with significant estimates (there is mixed evidence on some of the null results), with plenty of support for endogeneity of Exposure to Robots (see Table A12 in the Appendix). Furthermore, the instruments have much lower correlation with the baseline automobile employment share than does Exposure to Robots (0.31 and 0.60, compared to 0.9).

Given the evidence that robot exposure reduces employment in High Tech Manufacturing (which includes the automobile sector), while raising it in services, the variable Exposure to Robots is recomputed to answer the question: is it only robots in the automobile industry which are driving the employment effects? The answer is yes.

To do this, I compute a slightly different Exposure to Robots variable:

$$\sum_{i \in \tilde{\mathcal{I}}} \ell_{ci}^{1991} \left(\frac{R_{i,2011}}{L_{i,1991}} - \frac{R_{i,1993}}{L_{i,1991}} \right)$$

where $\tilde{I}$ is the full set of industries *excluding* the automobile sector.[26] Employment changes in both High Tech Manufacturing (incl. Auto) and in services are regressed on Exposure to Robots *excluding the automobile industry*, along with the standard controls and instruments. The results can seen in Figure 4 and are robust to removing outliers. These coefficients can highlight the importance of robots in the automobile industry: if they are no longer significant, or differ in sign/magnitude, then clearly the automobile robots play a critical role.

First consider the lowest square points on each panel on Figure 4. These represent the estimated coefficients on Exposure to Robots (excl. Auto) for the change in services and High Tech Manufacturing (incl. Auto) employment, when only controlling for regions and population weights. These suggest robot adoption *reduces* employment in services and *boosts* employment in High Tech Manufacturing (incl.

[24] The regression is ΔAuto Emp Share$_{1981-1991,c} = \alpha + \beta$Auto Emp Share$_{1981,c} + \Gamma C_c + \epsilon_c$

[25] Estimated coefficient $\beta = -0.454$ with cluster-robust standard error 0.0267.

[26] Note that the correlation between this and the standard Exposure to Robots variable is just 0.13, so this measure represents something that varies differently across Local Authorities.

Impact of Exposure to Robots (excl. Auto) on Employment in High Tech (incl. Auto) and Services from 1991 - 2011

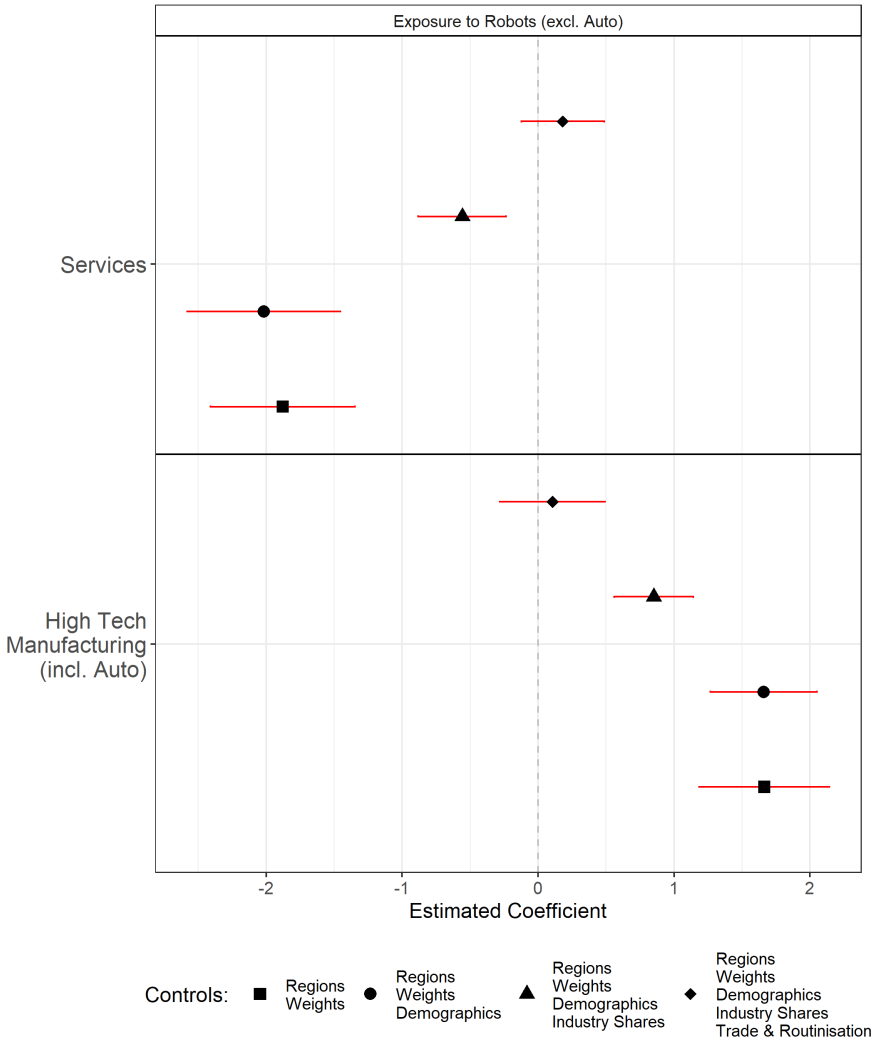

Figure 4: Estimated coefficients and 95% confidence interval for Exposure to Robots (excl. Auto) on change in employment shares ratio (1991-2011) in High Tech Manufacturing (incl. Auto) and services, with an increasing set of controls.

Auto). It seems that robots outside of the automobile sector affect employment in the opposite way to the previous findings. The inclusion of demographic controls does little to affect this conclusion, but industry share controls shrink the estimated coefficients sharply towards zero. Finally, the highest diamond points on each panel represent the model with the full set of controls. Neither provides evidence that robots outside the automobile sector affected employment shares in services or High Tech Manufacturing (incl. Auto). This exercise provides evidence that robot adoption in the automobile sector - where a large proportion of the UK robot stock is used - drives the results on changes in sectoral employment.

6 Conclusion

There is little doubt that automation, through robots and Artificial Intelligence, has the potential to alter the structure of labour markets in the UK. There has been an increase in the adoption of industrial robots, although this varies across industries and regions, and is still significantly lower than other advanced economies.

However, there has been no study of the impact of automation in the UK thus far. This paper analyses the effect of industrial robot adoption in England and Wales on employment and wages in local labour markets over the period 1991 - 2011. Local labour markets are approximated with Local Authorities, of which there are 348 in England and Wales - yielding significant cross-sectional variation and many covariates for identification by OLS and 2SLS in the differences-in-differences framework.

Analysis of industries suggests that industrial robots have reduced employment shares in some areas of manufacturing, such as automobile. However, the overall effect on employment rates is positive. The conclusion that reconciles these results is that robots have directly replaced workers - especially machine operators - in automobile and metal manufacturing. However, robots had a positive impact on other areas of the labour market such as services, so the overall effect on employment is positive. There is also evidence of a positive effect on part-time pay, but this result should be treated with a little caution due to a few missing observations.

On aggregate, there is a positive impact of industrial robots on employment at the extensive margin, but no effect at the intensive margin. Part-time wages are somewhat higher due to exposure to robots, but there is no impact on full-time pay. This contrasts with research by Acemoglu and Restrepo (2017) in the US, who estimate a robust negative relationship between exposure to robots and employment and wage changes at a local level.

There seem to be future avenues of research which can build on the analysis conducted in this paper. Firstly, industrial robots and automation-related patents are just two measures of automation. Data on other types of automation - even at an aggregate level - would be useful to extend this analysis. Secondly, it would be instructive to control for the potential heterogeneous regional impact of the Global Financial Crisis on the relationship between robot adoption and labour market outcomes.[27] Finally, the availability of linked employee-employer data (as in Dauth et al. (2017)) would be extremely useful, as it is clear that the robot impact in the UK is somewhat geographically and sectorally unequal, so following the impact on relevant individuals and firms would be insightful.

[27] This was not pursued here due to data limitations.

A Appendix - Tables & Figures

	Robots per thousand employees					
	Total	Auto. & Other Vehicles	Food & Bev.	Rubber & Plastics	Construction	Wood, Furniture & Paper
1995	0.376	11.49	0.29	0	0	2.16
2011	0.541	32.01	2.19	4.41	0.12	0.69
Change	0.165	20.51	1.90	4.41	0.12	-1.47

Table A1: Robot stock per thousand employees in 1995, 2011, and the change over that time period, for selected industries. Data from IFR, EUKLEMS.

	Robot Variables						
	Mean	Standard Deviation	Min	Q1	Q2	Q3	Max
Exposure to Robots 1993 - 2011	0.20	0.45	-0.21	0.01	0.06	0.20	4.02
Exogenous Exposure to Robots 1993 - 2011	2.53	3.16	0.15	0.88	1.39	2.75	22.84
Exposure to Automation Patents 1991 - 2011	0.16	0.08	0.01	0.10	0.15	0.21	0.54
Exposure to Global Automation Patents 1981 - 1991	0.27	0.09	0.11	0.20	0.25	0.31	0.61

Table A2: These summary statistics describe the robot and automation variables data across the 348 Local Authorities in England and Wales.

	Employment Statistics - Extensive Margin			
	E/Pop	E/WA	FTE/Pop	FTE/WA
Employment Ratios 1991	0.44	0.66	0.30	0.45
	(0.04)	(0.06)	(0.04)	(0.05)
Employment Ratios 2011	0.45	0.62	0.28	0.38
	(0.04)	(0.05)	(0.03)	(0.04)
Total change in employment ratios (1991 - 2011)	0.012	−0.037	−0.019	−0.062
	(0.02)	(0.03)	(0.02)	(0.03)
Observations	348	348	348	348

Table A3: These summary statistics describe the employment ratios for the whole population and the 1991 population-weighted change across the 348 Local Authorities in England and Wales. Standard errors are reported in brackets. Data from 1991/2011 UK Census.

	Employment Statistics - Intensive Margin	
	Full-Time Mean	Part-Time Mean
Hours Worked 1997	40.01	18.81
	(1.02)	(1.31)
Hours Worked 2011	39.08	18.02
	(0.895)	(1.10)
Change in hours worked	−0.899	−0.725
(1997 - 2011)	(1.52)	(0.776)
Observations	348	348

Table A4: These summary statistics describe the mean hours worked dependent variables for the whole population and the 1997 employment-weighted change across the 348 Local Authorities in England and Wales. Standard errors are reported in brackets. Data from ASHE.

	Income Statistics - Gross Weekly Pay			
	Full-Time		Part-Time	
	Mean Pay	Median Pay	Mean Pay	Median Pay
Gross Weekly 1997	374.1	323.3	113.7	80.3
	(79.0)	(57.2)	(24.0)	(37.0)
Gross Weekly 2011	607.5	512.2	196.4	149.2
	(154.1)	(113.7)	(42.5)	(32.9)
Change in Weekly Pay	221.7	180.5	83.6	51.3
(1997 - 2011)	(72.0)	(57.7)	(35.7)	(38.4)
Observations	345	344	337	218

Table A5: These summary statistics describe the arithmetic mean for weekly pay and the 1997 employment-weighted change across the 348 Local Authorities in England and Wales. Standard errors are reported in brackets. Data from ASHE.

	Income Statistics - Earnings Inequality		
	Log 80-50	Log 50-10	Log 80-10
1997 Ratios	0.36	0.56	0.88
	(0.14)	(0.14)	(0.33)
2011 Ratios	0.40	0.57	0.94
	(0.12)	(0.08)	(0.26)
Change in Ratio (1997 - 2011)	0.015	−0.018	0.003
	(0.056)	(0.083)	(0.087)
Observations	242	332	242

Table A6: These summary statistics describe various log earnings distribution ratios and the 1997 employment-weighted change across the Local Authorities in England and Wales. Standard errors are reported in brackets. Data from ASHE.

	Control Variables						
	Mean	Standard Deviation	Min	Q1	Q2	Q3	Max
Change in working age ratio 1991 - 2011	0.057	0.014	0.017	0.049	0.056	0.064	0.13
Change in white ratio 1991 - 2011	−0.059	0.063	−0.361	−0.08	−0.032	−0.017	−0.006
Exposure to China 1991 - 2011	1.769	0.851	0.104	1.163	1.603	2.255	4.755
Routine Share 1991	0.482	0.065	0.195	0.441	0.483	0.529	0.689
Population Growth 1991 - 2011	0.129	0.095	−0.069	0.07	0.117	0.183	0.781
Automobile Share 1991	0.010	0.022	0.000	0.001	0.003	0.009	0.195
Manufacturing Share 1991	0.179	0.065	0.051	0.132	0.169	0.214	0.412
Construction Share 1991	0.074	0.014	0.013	0.066	0.074	0.082	0.110
Female Manufacturing Share 1991	0.301	0.051	0.180	0.265	0.295	0.331	0.451

Table A7: These summary statistics describe the covariates data across the 348 Local Authorities in England and Wales.

Table A8: The estimated coefficient on Exposure to Robots on UK labour market outcomes using OLS & 2SLS estimation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Employment - Extensive Margin (Δ FTE/Pop, $n = 348$)							
Robot Exposure	-0.003	0.002	0.004^{**}	0.003^{*}	-0.011^{**}	-0.002	0.004^{**}	0.005^{**}
(1993 - 2011)	(0.004)	(0.002)	(0.002)	(0.002)	(0.004)	(0.003)	(0.002)	(0.002)
	Employment - Extensive Margin (Δ E/Pop, $n = 348$)							
Robot Exposure	-0.004	0.002	0.004	0.003	-0.012^{**}	-0.003	0.003	$0.004^{\dagger}$
(1993 - 2011)	(0.005)	(0.003)	(0.002)	(0.002)	(0.005)	(0.004)	(0.003)	(0.003)
	Employment - Extensive Margin (Δ E/WA, $n = 348$)							
Robot Exposure	-0.003	0.003	$0.006^{\dagger}$	0.005	-0.013^{**}	-0.004	0.005	0.007^{*}
(1993 - 2011)	(0.006)	(0.004)	(0.004)	(0.003)	(0.006)	(0.006)	(0.004)	(0.004)
	Employment - Intensive Margin (Δ Full-Time Hours, $n = 348$)							
Robot Exposure	0.10	0.13^{*}	0.14^{*}	0.13^{*}	0.02	0.04	0.07	0.09
(1993 - 2011)	(0.091)	(0.076)	(0.072)	(0.075)	(0.19)	(0.19)	(0.18)	(0.18)
	Wages (Δ Ln Median Full-Time Pay, $n = 344$)							
Robot Exposure	$-0.011^{\dagger}$	-0.009	-0.01	-0.009	-0.004	0.001	-0.001	0.001
(1993 - 2011)	(0.007)	(0.008)	(0.008)	(0.008)	(0.017)	(0.018)	(0.013)	(0.01)
	Wages (Δ Ln Median Part-Time Pay, $n = 218$)							
Robot Exposure	0.016	0.018	0.018	0.016	0.012	0.014	0.009	0.018
(1993 - 2011)	(0.014)	(0.016)	(0.015)	(0.016)	(0.037)	(0.039)	(0.039)	(0.036)
	Inequality (Δ 80/50 Ratio, $n = 242$)							
Robot Exposure	0.019^{**}	$0.015^{\dagger}$	$0.014^{\dagger}$	0.012	0.024^{**}	$0.019^{\dagger}$	0.012	0.012
(1993 - 2011)	(0.009)	(0.01)	(0.009)	(0.009)	(0.012)	(0.012)	(0.01)	(0.011)
	Inequality (Δ 50/10 Ratio, $n = 332$)							
Robot Exposure	0.006	-0.002	-0.003	-0.002	0.004	-0.008	-0.012	-0.008
(1993 - 2011)	(0.01)	(0.008)	(0.008)	(0.009)	(0.015)	(0.014)	(0.014)	(0.013)
Controls:								
Weight by population	✓	✓	✓	✓	✓	✓	✓	✓
Regional dummies	✓	✓	✓	✓	✓	✓	✓	✓
Demographics		✓	✓	✓		✓	✓	✓
Broad industry shares			✓	✓			✓	✓
Trade & routinisation				✓				✓
2SLS					✓	✓	✓	✓

Note: Long-run estimates of the impact of the exposure to robots on labour market outcomes. All regressions are weighted by baseline population, have regional dummies, and Liang and Zeger (1986) cluster-robust standard errors are reported in brackets (clustered at the regional level). Demographic controls are the changes between 1991 and 2011 in the share of working-age population, the share of the population that is of white ethnicity and the percentage change in the population size. Broad industry shares control for 1991 baseline shares of employment in manufacturing and construction and the share of female employment in manufacturing. Trade and routinisation controls for the exposure to Chinese imports and the 1991 baseline share of employment in routine jobs as defined in Autor and Dorn (2013). Instruments for 2SLS are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents. Note that results with *** are significant at a 1% confidence level, ** at a 5% level, * at a 10% level, and † at a 15% level.

Table A9: The estimated coefficient on Exposure to Robots on UK labour market outcomes using OLS and 2SLS estimation, removing outliers

Model	OLS Estimate	2SLS Estimate	n
ΔE/Pop	-0.01^{*}	-0.01	330
ΔFTE/Pop	-0.01^{*}	0.001	330
ΔE/WA	-0.01^{*}	-0.01	330
ΔFTE/WA	-0.01^{*}	0.002	330
Δmean_FT_hrs	-0.02	-0.57	330
Δmean_PT_hrs	0.13	-2.56^{*}	330
Δmedian_FT_wages	-0.05^{***}	0.04	326
Δmedian_PT_wages	-0.04	-0.16	207
Δmean_FT_wages	-0.05^{***}	0.07	327
Δmean_PT_wages	-0.07^{**}	-0.48^{**}	319
cp80.50	0.04	0.04	229
cp50.10	-0.06^{**}	-0.27^{**}	314
cp80.10	-0.04	-0.20^{**}	229

Note: Long-run estimates of the impact of the exposure to robots on labour market outcomes, removing the Local Authorities most-exposed to robots. All regressions are weighted by baseline population, and Liang and Zeger (1986) cluster-robust standard errors are reported in brackets (clustered at the regional level). Models estimated with full set of controls (detailed in the main text): demographic; broad industry shares; trade and routinisation. Instruments for 2SLS are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents. Note that results with *** are significant at a 1% confidence level, ** at a 5% level, * at a 10% level, and † at a 15% level.

Table A10: The estimated coefficient on Exposure to Robots on prior employment trends using 2SLS estimation

	(1)	(2)	(3)	(4)
	Δ FTE/Pop 1981 - 1991			
Robot Exposure	-0.002	0.003	0.01	0.013
(1993 - 2011)	(0.003)	(0.004)	(0.021)	(0.021)
	Δ FTE/WA 1981 - 1991			
Robot Exposure	-0.003	0.003	0.014	0.018
(1993 - 2011)	(0.005)	(0.006)	(0.031)	(0.031)
	Δ E/Pop 1981 - 1991			
Robot Exposure	-0.001	-0.005	0.01	0.003
(1993 - 2011)	(0.004)	(0.004)	(0.021)	(0.019)
	Δ E/WA 1981 - 1991			
Robot Exposure	-0.002	-0.008	0.014	0.007
(1993 - 2011)	(0.006)	(0.006)	(0.03)	(0.026)
	Δ Manufacturing Share 1981 - 1991			
Robot Exposure	0.023^{**}	0.027^{***}	-0.023	$0.051^{\dagger}$
(1993 - 2011)	(0.009)	(0.007)	(0.035)	(0.032)
Controls:				
Weighted by population	✓	✓	✓	✓
Regional dummies	✓	✓	✓	✓
Demographics	✓	✓	✓	✓
Broad industry shares	✓	✓	✓	✓
Trade		✓		✓
Remove outliers			✓	✓
Observations	348	348	330	330

Note: Long-run estimates of the impact of the exposure to robots on prior employment trends 1981 - 1991. All regressions are computed with 2SLS, weighted by baseline population, and Liang and Zeger (1986) cluster-robust standard errors are reported in brackets (clustered at the regional level). Demographic controls are the changes between 1981 and 1991 in the share of working-age population, the share of the population that is of white ethnicity and the percentage change in the population size. Broad industry shares control for 1981 baseline shares of employment in manufacturing and construction and the share of female employment in manufacturing. Trade controls for the exposure to Chinese imports. Instruments for 2SLS are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents. Note that results with *** are significant at a 1% confidence level, ** at a 5% level, * at a 10% level, and † at a 15% level.

Table A11: p-values for 2SLS tests for all dependent variables

Dependent Variable	Weak Instruments	Sargan	Wu-Hausman
Δ E/Pop	0.00	0.00	0.59
Δ FTE/Pop	0.00	0.00	0.26
Δ E/WA	0.00	0.00	0.59
Δ FTE/WA	0.00	0.00	0.28
Δ mean Full-Time hours	0.00	0.01	0.65
Δ mean Part-Time hours	0.00	0.01	0.14
Δ median Full-Time pay	0.00	0.01	0.71
Δ median Part-Time pay	0.00	0.02	0.79
Δ mean Full-Time pay	0.00	0.01	0.43
Δ mean Part-Time pay	0.00	0.00	0.88
Δ 80/50 ratio	0.00	0.14	0.40
Δ 50/10 ratio	0.00	0.00	0.51
Δ 80/10 ratio	0.00	0.04	0.98

Note: Test results from 2SLS estimation on baseline regression models, with full set of controls. Instruments are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents.

Table A12: p-values for 2SLS tests for all industry employment rates

Dependent Variable (Δ emp. share)	Weak Instruments	Sargan	Wu-Hausman
Food, Bev. & Tobacco Manuf.	0.00	0.58	0.41
Textiles Manuf.	0.00	0.00	0.00
Wood & Paper Manuf.	0.00	0.00	0.00
Other Manuf.	0.00	0.00	0.00
Chem., Rubber & Plastics Manuf.	0.00	0.09	0.00
Low Tech Manuf.	0.00	0.00	0.00
High Tech Manuf. (incl. Auto)	0.00	0.00	0.08
Utilities	0.00	0.26	0.15
Construction	0.00	0.89	0.93
Services	0.00	0.00	0.34
Education	0.00	0.82	0.51
Agriculture	0.00	0.03	0.03
Mining	0.00	0.81	0.70

Note: Test results from 2SLS estimation on industry share regression models, with full set of controls. Instruments are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents.

Table A13: Generalised Variance Inflation Factors for all control variables across all models

	Employment				Hours		Wages				Inequality		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Construc.% 1991	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.39	1.33	1.34	1.38	1.32	1.38
Δ % White	1.62	1.62	1.62	1.62	1.62	1.62	1.63	1.60	1.63	1.62	1.65	1.62	1.65
Δ WA	1.22	1.22	1.22	1.22	1.22	1.22	1.22	1.26	1.22	1.23	1.31	1.24	1.31
Exposure to China	1.96	1.96	1.96	1.96	1.96	1.96	1.96	2.11	1.96	1.96	1.94	1.94	1.94
Exposure to Robots	1.30	1.30	1.30	1.30	1.30	1.30	1.30	1.34	1.30	1.30	1.30	1.31	1.30
F. Manuf.% 1991	1.21	1.21	1.21	1.21	1.21	1.21	1.21	1.21	1.21	1.21	1.25	1.21	1.25
Manuf.% 1991	2.71	2.71	2.71	2.71	2.71	2.71	2.70	2.89	2.71	2.71	2.70	2.69	2.70
Pop. Growth	1.54	1.54	1.54	1.54	1.54	1.54	1.54	1.66	1.54	1.54	1.71	1.55	1.71
Region	1.16	1.16	1.16	1.16	1.16	1.16	1.16	1.18	1.16	1.16	1.18	1.17	1.18
Routine % 1991	1.79	1.79	1.79	1.79	1.79	1.79	1.80	1.83	1.80	1.81	1.88	1.79	1.88

Note: Models (1) to (13) cover the different dependent variables in this paper. Models (1) - (4) are Δ E/Pop, Δ FTE/Pop, Δ E/WA and Δ FTE/WA. Models (5) and (6) are Δ Full-Time hours and Δ Part-Time hours. Models (7) and (8) are changes in mean Full-Time and Part-time pay, while (9) and (10) are the same for median wages. Models (11) - (13) are the change in the 80/50, 50/10 and 80/10 income ratios.

Impact of Exposure to Robots on Industry Employment (1991 - 2011), removing regions most-exposed to Robots

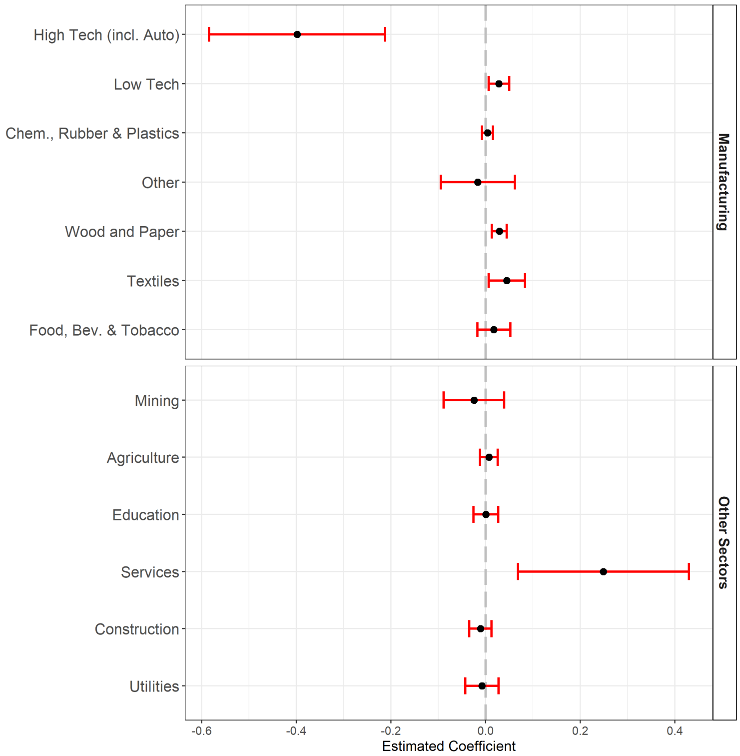

Figure A1: Estimated 2SLS coefficients and 95% confidence interval for Exposure to Robots on change in employment shares ratio (1991-2011) across a set of industries, conditional on the full set of controls used in baseline estimation. Instruments are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents. Regions most-exposed to robots are removed.

Impact of including Chinese Trade control

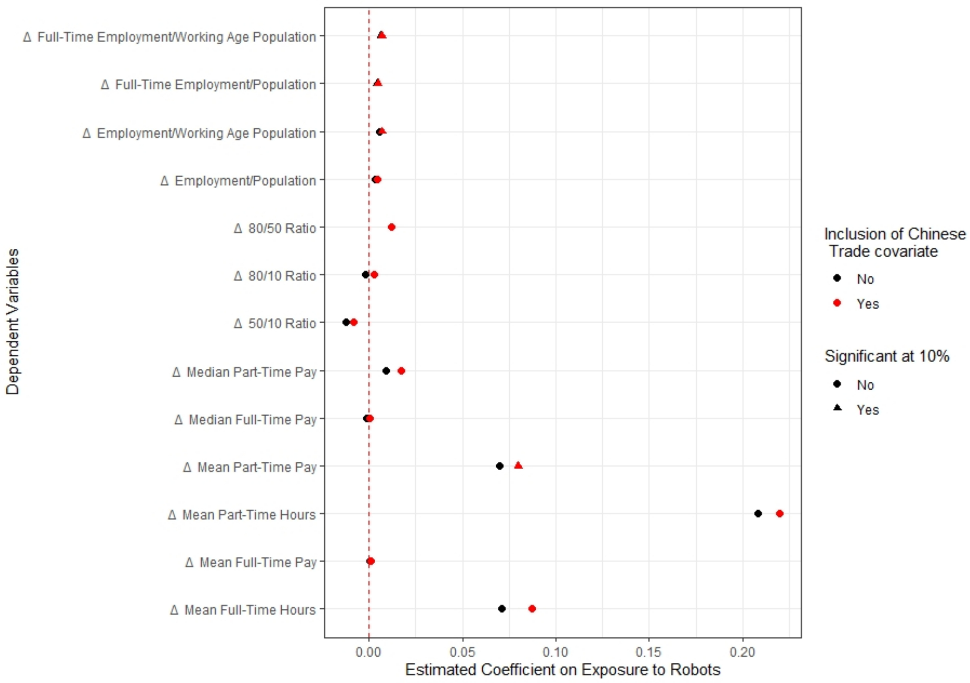

Figure A2: Estimated 2SLS coefficients on Exposure to Robots on models with all controls, when including and excluding covariate which controls for Chinese Trade. Instruments are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents.

B Appendix - General Equilibrium and Trade Model Extensions

This section presents a brief summary of the intuition and results from the general equilibrium and trade versions of the model in the Acemoglu and Restrepo (2017) paper.

Changes in prices and output depend on the changes in prices and quantities of robots and labour as well as the tasks which can be automated M_i. The proof can be found in the appendix of Acemoglu and Restrepo (2017).

The general equilibrium impact is given by:

$$d\ln L_c = -\frac{1+\eta}{1+\epsilon}\sum_{i\in\mathcal{I}}\ell_{ci}\frac{dM_i}{1-M_i} + \frac{1+\eta}{1+\epsilon}\pi_c\sum_{i\in\mathcal{I}}\ell_{ci}\frac{s_{icL}}{s_{cL}}\frac{dM_i}{1-M_i} \tag{9}$$

$$d\ln W_c = -\eta\sum_{i\in\mathcal{I}}\ell_{ci}\frac{dM_i}{1-M_i} + (1+\eta)\pi_c\sum_{i\in\mathcal{I}}\ell_{ci}\frac{s_{icL}}{s_{cL}}\frac{dM_i}{1-M_i} \tag{10}$$

where s_{cL} is the labour share of total output in commuting zone c, s_{icL} is the labour share of industry i output in commuting zone c, η and ϵ are elasticities of supply for robots and labour respectively, and π_c is the cost saving from using robots over labour.

The first term in equation (9) is the general equilibrium *displacement effect*, while the second is the *productivity effect*, expressed as a function of the changes in robotics technology, i.e. which tasks can be automated. The impact on employment could be negative. The magnitude of the productivity effect depends on π_c, which is the cost saving from substituting robots for human labour. If this is small, the productivity effect will be small.

The key extension to the partial equilibrium result is the appearance of the elasticities of supply for robots and labour, which is intuitive as changes in the demand for robot and labour influence the changes in output and prices, which will be determined by the intersection of demand and supply in the robot and labour markets.

Equation (10) represents the relationship between the change in wages and the change in the technological capability of robots. This extension from the partial equilibrium model yields a relationship which can then be tested with the data.

Links between commuting zones are crucial as lower costs of production in one zone (for example, due to the adoption of robots) will expand trade with other zones. Trading between zones can be modelled by assuming X_{ci} is exported to all commuting zones. The preferences in each commuting zone are defined by the same aggregate over consumption goods, but now these goods are aggregates of varieties sourced from all commuting zones.

The demand for labour in the trading equilibrium satisfies the following equation:

$$d\ln L_c^d = -\sum_{i\in\mathcal{I}}\ell_{ci}\frac{dM_i}{1-M_i} - \lambda\sum_{i\in\mathcal{I}}\ell_{ci}d\ln P_{Xci} + (\lambda-\sigma)\sum_{i\in\mathcal{I}}\ell_{ci}d\ln P_{Yi} + d\ln Y \tag{11}$$

We assume that varieties of the same good from different commuting zones are more substitutable than different products in the consumption aggregator (i.e. $\lambda > \sigma$).

Labour demand in commuting zone under trading equilibrium (equation (11)) is different from autarky (equation (1)). We still have the same displacement effect in equation (11) as in equation (1). Now we have three terms making up the productivity effect. The price-productivity effect, which is more powerful than under autarky due to higher substitutability of the same good from different zones. When an industry lowers its costs and hence price, it can raise its market share. The productivity effect is somewhat dampened because the greater use of robots in industry i reduces the cost of production in all commuting zones - this is the spillover. The scale-productivity effect is still present but works through expansion of total output in the economy rather than just in the commuting zone.

C Appendix - Local Authority Analysis

This section provides more detail on the two tests on the suitability of Local Authorities as proxies for stable local labour markets. The Mantel test analyses the similarity of the commuting flows across Local Authorities between 1991 and 2011. This gives information on the stability of these regions with regards to commuting behaviour. Secondly, a clustering algorithm is applied to more disaggregated geographies (as in Tolbert and Sizer (1996)) - between 8,800 wards in England and Wales - to check if the number of clusters is within the range of the 348 Local Authorities. This checks if workers' travelling patterns can be well approximated by Local Authorities.

C.1 Data

Commuting data is obtained from the Web-based Interface to Census Interaction Data (WICID) which permits access to UK Census data on various migration and commuting flows, between different levels of regional aggregation. The Special Workplace Statistics (SWS) are used, which provide counts of flows of employed and self-employed between their usual residence and their workplace.[28] All data is for England and Wales, provided at a 10% sample. The first two datasets are at the Local Authorities level, used for the first test. The final dataset is at the ward level, used for the second test.

- 1991 SWS Set C - commuting across UK Interaction Data Districts (2001).[29]
- 2011 WF01EW - commuting flows across (merged) UK Local Authorities.[30]
- 1991 SWS Set C (2001 geog.) - commuting between 2001-defined Standard Table wards.

C.2 Stability of Local Authorities

I take the commuting data from 1991 and 2011 and drop all entries not in both flows matrices (i.e. removing authorities not common to both datasets). This leaves 333 authorities in both matrices. Distance matrices D_{ij} are constructed as in Tolbert and Sizer (1996).

The relevant test is to check that commuting patterns between the Local Authorities in 1991 and in 2011 are not too dissimilar. To statistically analyse the difference in commuting patterns over time, a Mantel test is used, which tests for the similarity between two distance matrices. It deals with the lack of independence between observations by calculating the correlation between the matrices many times after

[28] For all usual residents aged 16+ and employed in the week before the Census.

[29] This is a 376 x 376 matrix as data isn't available by Local Authorities. The difference is that some Local Authorities are broken into smaller parts. For example, Bedford and Central Bedford are the Local Authorities, whereas Bedford, Mid Bedfordshire and South Bedfordshire are the Interaction Data Districts.

[30] This is a 346 x 346 matrix as Westminster and Scilly are merged - they are small authorities with control over certain services.

randomly permuting the rows and columns of one matrix. The null hypothesis that there is no relationship between the matrices is therefore tested by comparing the randomly permuted matrix correlation to the actual correlation. If the matrices are unrelated, then the permuted matrix correlations should be more or less correlated with equal likelihood. With 1000 permutations, the p-value is 0.000999, suggesting that the matrices have a strong relationship, with a correlation of 0.983.[31]

C.3 Clustering of Wards

I follow Tolbert and Sizer (1996) in performing hierarchical clustering on wards.[32] This assigns each ward to a cluster, and subsequently groups clusters together based on the distance between them. The 'distance' between clusters is defined by the linkage criterion, and I consider three popular criteria. Tolbert and Sizer (1996) choose a cutoff point for between-cluster distances, and justified this decision as producing "reasonable and consistent results across the wide variety of U.S. counties." However, there is no standard method to work out the 'optimal' number of clusters. Instead of this approach, I compute the number of clusters for a range of 'threshold' values and consider how quickly the number of clusters falls as the 'threshold' value is changed. I find it reasonable to choose a threshold that had a relatively small impact on the number of clusters, at the margin. Therefore the number of clusters is computed when the slope of this curve (clusters against threshold) reaches some small proportion p of the maximum slope. While this may seem arbitrary, this 'optimal' number of clusters is found for a range of p values, with the aim of finding "reasonable and consistent results" as in Tolbert and Sizer (1996). Local Authorities seem to be a reasonable approximation for the commuting behaviour of workers in 1991. The following table summarises the results. It shows the number of clusters chosen by the algorithm, for a variety of popular linkage criteria, specifying that the slope of clusters against threshold reaches proportion p of the maximum slope.

	Linkage Criteria		
p	Single	Complete	Average
.01	134	207	199
.02	326	342	294
.03	410	464	367
.04	410	764	412
.05	542	764	663

Table A14: Resulting clusters from applying hierarchical clustering to 1991 commuting data between 8,800 UK wards, for three linkage criteria, with threshold being varied to ensure relative 'stability' of results.

[31] Using the **ade4** package in R

[32] Using the **scipy.cluster.hierarchy** package in Python

D Appendix - Robustness Checks

The baseline results are checked for robustness to the Chinese trade control, to prior trends, and to sensitivity to outliers. The findings suggest that the results are not particularly sensitive to Chinese trade or prior trends, but are sensitive to the regions most-exposed to robots. This highlights that the heterogeneity across regions plays an important role in aggregate results.

D.1 Control for Trade from China

It is possible that the results are confounded by the inclusion of the trade regressor, as the instruments have non-negligible partial correlations with Exposure to Chinese Trade (when controlling for demographic and broad industry shares). This is likely because a significant proportion of the increase in robot patents, robots exposure, and Chinese trade exposure has been in manufacturing.

Figure A2 shows the impact of the inclusion of the trade control: generally, the point estimate on Exposure to Robots is slightly greater for each dependent variable. Furthermore, the inclusion of the trade control reduces the p-values across almost all models, leading to stronger evidence of statistical significance for Full-Time Employment and Mean Part-Time Pay. This suggests that multicollinearity is not a problem for this control variable, as such an issue would typically raise the standard error of the coefficient on Exposure to Robots.

Crucially, even if the inclusion of the trade regressor might be problematic, the overall results do not change significantly. The sign, magnitude and statistical significance of the estimated 2SLS coefficients on Exposure to Robots do not change enough to eliminate the validity of the baseline results.

D.2 Prior Trends

One concern with the analysis is perhaps regions adopting more robots did so due to pre-existing employment trends. For example, it may be that a region experiencing sharp manufacturing employment decline prior to 1991 is more likely to integrate robots. Rather than Exposure to Robots over the period 1993 - 2011 having a relationship with employment changes from 1991 - 2011, it may be that it is related to employment changes from the decade prior.

To test for this, the employment extensive margin dependent variables over the period 1981 - 1991 are regressed against the Exposure to Robots over 1993 - 2011. The baseline 2SLS estimation is repeated, including control variables recomputed over 1981 - 1991 with minor adjustments due to data availability.[33]

The results provide strong evidence that robot adoption over 1993 - 2011 is not related to employment changes in the decade prior. In case trade with China or outliers confound the results, the regressions are run with and without these two

[33] Unable to control for share of routine work or share of population of white ethnicity, as the UK 1981 Census doesn't have the relevant data.

elements. Table A10 shows that the null hypothesis of no relationship between the two variables cannot be rejected, across all measures of employment rates.

However, there is some evidence that robot adoption from 1993 - 2011 is related to the change in *manufacturing share* from 1981 - 1991, but this is driven by the outlying Local Authorities which increased the use of industrial robots the most (see Table A10). The estimated coefficients are positive, implying that regions experiencing a rise in manufacturing employment share from 1981 - 1991 adopted more industrial robots in the subsequent decade.

D.3 Sensitivity to Outliers

The industry results show an enormous amount of heterogeneity. In addition, the automobile industry is dominant in terms of employment in regions where the robot shock was largest. An important question is if the aggregate results are driven by the changes to the automobile industry.

Therefore, the baseline regressions are repeated without the regions most-exposed to robots, to check for the importance of these Local Authorities. The results are presented in Table A9, for OLS and 2SLS estimation with the full set of controls.

The impact on the extensive margin of employment is stark: all OLS estimates are negative, while the 2SLS estimates are not significantly different from zero. The estimated coefficients on the intensive margin vary greatly between OLS and 2SLS, reflecting significant endogeneity concerns. There is some evidence that - without the most-exposed Local Authorities - Part-Time hours *fall* with Exposure to Robots.

Outliers also have an impact on the wage estimates compared to the baseline results. Given limited evidence of endogeneity for these variables, the OLS estimates are preferred and suggest significant *negative* effects on wages for both full-time and part-time workers.

Finally, although the baseline regressions showed little support for robot adoption affecting earnings inequality, this seems to be sensitive to outliers. When removed, there is strong evidence of reduced pay inequality with statistically significant and negative coefficients on the 80/10 and 50/10 ratios.

The baseline results are sensitive to outliers. It seems that by removing the areas most-exposed to the technology shock, robots reduce Part-Time hours and wages, and also reduce the wage inequality ratios. This suggests that a small number of Local Authorities which are highly-exposed to robot adoption are driving the baseline results, of increased full-time employment rates and part-time pay.

E Appendix - Occupation Analysis

Robots have had a heterogeneous impact across different occupations. In the UK, the industry and occupation codes have some overlap, but codify quite distinct information. In other words, there are significant proportions of each of the nine occupation categories in each industry, and vice versa.[34]

Therefore, employment shares are computed for each occupation for 1991 and 2011. Then 2SLS estimation is repeated on the 9 standard occupation categories (SOCs) with the full set of controls. The estimated coefficients which have both a large magnitude and statistical significance are on (8) Process plant and machine operatives and (9) Elementary occupations.

This evidence suggests that industrial robots have affected jobs differently. In particular, the lower employment share (SOC 8) is in jobs that require humans to operate, test, and maintain machinery. The higher employment share (SOC 9) is in work that typically requires "knowledge and experience" on "mostly routine tasks" but which involve machinery that may require dexterity that automated robots are yet to achieve (e.g. operating winches, cleaning animal quarters and sheering sheep).

Impact of Exposure to Robots on Occupation Employment (1991 - 2011)

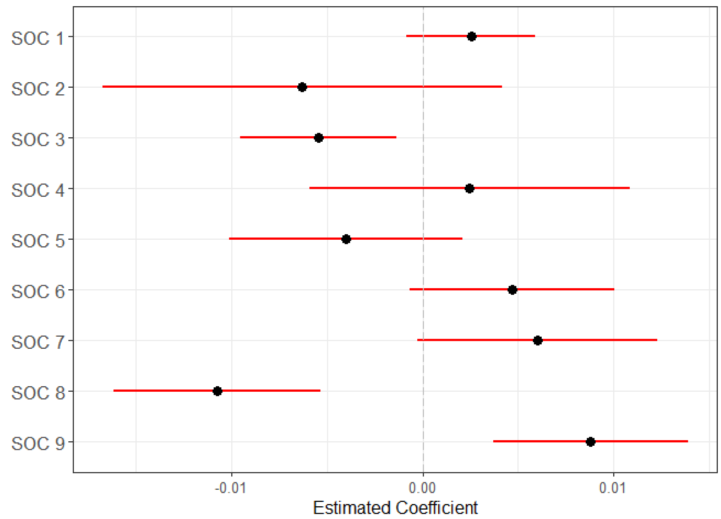

Figure A3: The points are estimated 2SLS coefficients of Exposure to Robots on the change in employment shares (1991 - 2011) across a set of occupations. The lines indicate the 95% confidence intervals. The models have the full set of controls used in baseline estimation. SOC 1: Managers and Senior Officials. SOC 2: Professional Occupations. SOC 3: Associate Professional and Technical Occupations. SOC 4: Administrative and Secretarial Occupations. SOC 5: Skilled Trade Occupations. SOC 6: Personal Service Occupations. SOC 7: Sales and Customer Service Occupations. SOC 8: Process, Plant and Machine Operatives. SOC 9: Elementary Occupations.

[34]Computation of Cramér's V, Theil's Index, and the Uncertainty Coefficient support this conclusion. Although some occupations are more common in certain industries, there is different information in the distribution of workers across occupations than in the distribution of workers over industries.

F Appendix - Industry Mapping

Across this paper, an industry-level mapping is used, which allows the linking of robot data and employment data. The EUKLEMS dataset contains national employment figures by year and industry. The industry breakdown is in accordance with the industry classification ISIC Rev. 4/NACE Rev 2. In Chapter 1 of the IFR World Robotics 2017, it is stated that all data is now consistent with the ISIC Rev. 4. Therefore the robotics data can be mapped to the national employment data. However for the *regional* employment data (UK census) uses the UK Standard Industry Classification (SIC) 1980. This necessitated the creation of a mapping between the classifications for later analysis. The industry mapping is shown in Table A15 in the Appendix, resulting in 16 industry groupings. Note that for mapping robots data to the EUKLEMS employment data, Automotive and Other Vehicles are combined, as is Wood and Furniture with Paper, yielding just 14 industries.

For construction of the Routineness measure for each Local Authority, the share of employment in 'routine' occupations is calculated (Autor and Dorn, 2013). The mapping from occupations (defined by SOCs) to routineness is shown in Table A17.

Industry	UK SIC 1980	ISIC Rev. 4
Agriculture, forestry & fishing	0	A-B
Electricity, gas & water supply	1	E
Mining & quarrying	2	C
Construction	5	F
Non-manufacturing	6, 7, 8, 91, 92, 95, 96, 97, 98, 99, 00	90
Metal manufacturing	31	24-25
Electrical manufacturing	34	26-27
Automotive manufacturing	35	29
Other vehicles manufacturing	36	30
Food & beverages manufacturing	41	10-12
Textiles manufacturing	43	13-15
Wood & furniture manufacturing	46	16
Paper manufacturing	47	17-18
Rubber & plastics manufacturing	48	22
Other manufacturing	32, 33, 37, 44, 45, 49	19, 20-21, 23, 28, 91
Education and R&D	93, 94	P

Table A15: The mapping of the 16 industries in set $\mathcal{I}$, used to construct Exposure to Robots, between UK SIC 1980 and ISIC Rev. 4

Industry	CPA 2002
Agriculture, forestry & fishing	1, 2, 5
Electricity, gas & water supply	-
Mining & quarrying	10, 13, 14
Construction	-
Non-manufacturing	72, 74, 92, 93
Metal manufacturing	27, 28
Electrical manufacturing	30, 31, 32
Automotive manufacturing	34
Other vehicles manufacturing	35
Food & beverages manufacturing	15
Textiles manufacturing	17
Wood & furniture manufacturing	20, 361
Paper manufacturing	21, 22
Rubber & plastics manufacturing	25
Other manufacturing	16, 18, 19, 23, 24, 26, 29, 33, 362 - 366
Education and R&D	-

Table A16: The mapping of the 16 industries in set $\mathcal{I}$, used to construct Exposure to Trade, between UK SIC 1980 and CPA 2002

Standard Occupational Classification	Autor and Dorn (2013)	RTI
1. Managers and administrators	Managers/prof/tech/finance/public safety	$-$
2. Professional occupations	Managers/prof/tech/finance/public safety	$-$
3. Associate professional and technical occupations	Managers/prof/tech/finance/public safety	$-$
4. Clerical and secretarial occupations	Clerical/retail sales	$+$
5. Craft and related occupations	Production/craft	$+$
6. Personal and protective service occupations	Service occupations	$-$
7. Sales occupations	Clerical/retail sales	$+$
8. Plant and machine operatives	Machine operators/assemblers	$+$
9. Other occupations	*Several categories*	$+/-$

Table A17: Mapping SOC to Autor and Dorn (2013) Routine Task Intensity

G Appendix - Patent Search

On the WIPO IP Portal, the PATENTSCOPE database search is used.[35] The Advanced Search functionality was used, which offers boolean search on the patent database. Following research on robotics-related patents and intellectual property by Wunsch-Vincent et al. (2015), the abstracts were searched for the terms 'robot', 'robotic', 'automate' and 'automation'. Given the analysis required automation-related patents from 1991 to 2011, the search was restricted to these years. The search query for automation-related patents in 1991 was:

EN_ALLTXT:(robot OR robotic OR automate OR automation) AND DP:(1991)

where 'EN_ALLTXT' searched the abstract for specified terms, and 'DP' denotes publication year.

For UK patents, the search was limited to patents in the UK and European Patent Office (EPO). The latter is included because a patent granted under the European Patent Convention (EPC) is "treated like a granted domestic patent."[36] Patents were downloaded for the UK and globally (the latter for the instrumental variable). Most patents were assigned IPC codes, sometimes up to 50. The IPC codes were manually matched to industrial NACE Rev. 2 classifications using the Eurostat IPC to Nace Rev. 2 concordance.[37] The next step was to match the NACE Rev. 2 to ISIC Rev. 4 using the Eurostat correspondence table.[38] It was now straightforward to link to the industries in set $\mathcal{I}$ using Table A15.

Clearly some patents were mapped to multiple industries. Any patent with multiple IPC codes was assigned equal weight to each code. This meant a patent could be 1/3 Textiles Manufacturing and 2/3 Automobile Manufacturing. Now the following variable could be computed by summing patents across industries in a given year:

$$\begin{array}{l}\text{Exposure to Robot Patents}\\ \text{from 1991 to } 2011_c\end{array} = \sum_{i \in \mathcal{I}} \ell_{ci}^{1991} \left(\frac{P_{i,2011}^{Robots}}{L_{i,1991}} - \frac{P_{i,1991}^{Robots}}{L_{i,1991}} \right)$$

where $P_{i,t}^{Robots}$ are robot-related patents in industry i in year t.

[35] https:/https://patentscope.wipo.int

[36] https://www.gov.uk/guidance/manual-of-patent-practice-mopp/section-77-effect-of-european-patent-uk

[37] https://circabc.europa.eu/webdav/CircaBC/ESTAT/infoonstatisticsofsti/Library/methodology/patent_statistics/IPC_NACE2_Version2\%200_20150630.pdf

[38] http://ec.europa.eu/eurostat/ramon/relations/index.cfm?TargetUrl=LST_LINK&StrNomRelCode=NACE\%20REV.\%202\%20-\%20ISIC\%20REV.\%204&StrLanguageCode=EN

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *National Bureau of Economic Research, Working Papers*, 4(16082).

Acemoglu, D. and Restrepo, P. (2016). The race between machine and man: Implications of technology for growth, factor shares and employment. *National Bureau of Economic Research, Working Paper*, (22252).

Acemoglu, D. and Restrepo, P. (2017). Robots and jobs: Evidence from U.S. labor markets. *National Bureau of Economic Research, Working Paper*, (23285).

Acemoglu, D. and Restrepo, P. (2018). Artificial intelligence, automation and work. *National Bureau of Economic Research, Working Paper*, (24196).

Autor, D. (2010). The polarization of job opportunities in the U.S. labor market: Implications for employment and earnings. *The Center for American Progress and The Hamilton Project.*

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, 103(5):1553–97.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–68.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2005). Rising wage inequality: The role of composition and prices. (11628).

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. *The American Economic Review*, 57(3):415–426.

Cortes, G. M., Jaimovich, N., and Siu, H. E. (2016). Disappearing routine jobs: Who, how, and why? *National Bureau of Economic Research, Working Paper*, (22918).

Dauth, W., Findeisen, S., Südekum, J., and Wößner, N. (2017). German Robots - The Impact of Industrial Robots on Workers. *CEPR Discussion Papers*, (12306).

Fox, J. and Monette, G. (1992). Generalized collinearity diagnostics. *Journal of the American Statistical Association*, 87(417):178–183.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics*, 89(1):118–133.

Gosling, A., Machin, S., and Meghir, C. (2000). The Changing Distribution of Male Wages in the U.K. *The Review of Economic Studies*, 67(4):635–666.

Graetz, G. and Michaels, G. (2015). Robots at Work. *CEP Discussion Papers*, (1335).

International Federation of Robotics (2017). World Robotics: Industrial Robots.

Kirsten Jäger (The Conference Board) (2017). EU KLEMS Growth and Productivity Accounts 2017 Release, Statistical Module.

Lawrence, M., Roberts, C., and King, L. (2017). Managing automation: Employment, inequality and ethics in the digital age. *IPPR Commission on Economic Justice, Discussion Paper.*

Liang, K.-Y. and Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1):13–22.

Mann, K. and Püttmann, L. (2018). Benign effects of automation: New evidence from patent texts.

Moretti, E. (2010). Local labor markets. *National Bureau of Economic Research, Working Paper*, (15947).

Office for National Statistics (2016). 2011 Census aggregate data. UK Data Service. `http://dx.doi.org/10.5257/census/aggregate-2011-1`.

Office of Population Censuses and Surveys (1997). 1991 Census aggregate data. UK Data Service. `http://dx.doi.org/10.5257/census/aggregate-1991-1`.

Office of Population Censuses and Surveys (2000). 1981 Census aggregate data. UK Data Service. `http://dx.doi.org/10.5257/census/aggregate-1981-1`.

O'brien, R. (2007). A caution regarding rules of thumb for variance inflation factors. *Quality & Quantity: International Journal of Methodology*, 41(5):673–690.

Susskind, D. (2017). A model of technological unemployment. *University of Oxford, Department of Economics Discussion Paper Series*, (819).

Susskind, R. and Susskind, D. (2015). *The Future of the Professions: How Technology Will Transform the Work of Human Experts.* Oxford University Press.

Tolbert, C. M. and Sizer, M. (1996). U.S. commuting zones and labor market areas: a 1990 update. *[Washington, D.C.]: Economic Research Service, Rural Economy Division.*

Wunsch-Vincent, S., Raffo, J., and Andrew Keisner, C. (2015). Breakthrough technologies – robotics, innovation and intellectual property. *WIPO Economic Research Working Paper*, 30.