Technology and the Macroeconomy

Joel Kariel

Worcester College

Department of Economics

University of Oxford

A thesis presented for the degree of

Doctor of Philosophy in Economics

Michaelmas 2021

Acknowledgements

We should always be thankful for those that are kind to, honest with, and supportive of us throughout our lives. Nonetheless, the completion of a significant project provides a moment of pause, and an opportunity to reflect on the important people that surround us. It should go without saying that writing this thesis has been a significant undertaking, which would not be possible without the support of my family, close friends, and colleagues. As is tradition, I will list all the people here anyway, but I hope that they know my gratitude already.

Firstly, I would like to thank Petr Sedláček, who has been a consistent source of support and feedback throughout my graduate studies. He has pushed me to re-evaluate research ideas, refine my writing, restructure my thoughts, and delve into new topics, even when we've been on opposite sides of the globe. His unwavering support has been essential.

Secondly, I am grateful to the faculty and students in the Economics department at Oxford for engaging on research ideas throughout the last few years, but mostly for being friendly and encouraging. Special thanks go to Hamish Low, Gabriel Ulyssea, Simon Cowan, Alexander Copestake, Lukas Freund, and Hannah Zillessen. I am particularly indebted to Richard Povey. After tutoring me as an undergraduate, he gave me the opportunity to teach, and has written me countless references over the years. I appreciate Nick Stenner, who was by my side for most of our graduate studies, and was always available to dissect a paper, have a coffee break, laugh about a mistake, or give feedback on my thoughts. I am eternally thankful to Anthony Savagar, who kindly offered me a position as a postdoctoral researcher before I completed my DPhil, and has been a wonderful mentor and colleague this past year.

Finally, I give thanks to my family and friends for just being themselves. They have been patient recipients of my repeated attempts at the 'elevator pitch' of my thesis (it turns out "are robots taking over?" is the most compelling, if not at all accurate). Gratitude goes to my grandparents, who have been interested and inquisitive, even in a field which is entirely new to them. My parents have always supported my education, and I would not have been able to do this without them. I would like to thank my flatmates, Salil Patel and Max Ahrens, who know the importance of switching off from research, with a good meal and a less good movie. And, no less important: thank you Steph Barker, for everything.

Statement of Authorship

All chapters of this thesis are single-authored. I received guidance throughout from my supervisor, Petr Sedláček. I produced the bulk of the third chapter while working as a postdoctoral researcher with Anthony Savagar. Anthony's ESRC grant on Scale Economies, alongside access to UK firm data, was necessary to pursue this project. The views expressed herein are mine and I take sole responsibility for any errors.

COVID-19 statement: I had largely completed Chapters 1 and 2 when the university buildings were closed to non-COVID-19 research. The closure of the Economics Department did not materially impact the completion of this thesis.

Word Count: Approximately 77,470 words (taking page 3 as a typical example, with 305 words and 254 non-bibliography pages).

Contents

List of Figures

List of Tables

Introduction

Technology & Economics

The emergence of new technologies has always had implications for the production and organisation of economic activity. Such innovations are adopted because they provide gains to the ultimate beneficiary of a specific economic activity, such as higher revenue, reduced costs, or less managerial oversight. However, technological improvement may not benefit all, nor will it unambiguously improve macroeconomic outcomes. The former depends on the allocation of the productivity gains supplied by the new technology. The latter is tied to the distribution of technology adoption, and the subsequent dynamic general equilibrium outcomes.

A classic historical example of technological disruption to economic activity is John Kay's 'Flying Shuttle', which significantly improved the efficiency of weaving cotton (Paulinyi 1986). The spread of this invention also led to backlash. This is now known as the Luddite movement of the early 19th century, during which English textile workers - fearful of having their work automated - destroyed machinery in protest (Autor 2014). There is some evidence that such anxiety goes back even further. In the 16th century, English inventor William Lee was refused a patent for the first knitting machine for fear it would replace workers (Pasold 1975). Even Aristotle noted in Politics I, Part IV, that "if every instrument could accomplish its own work, obeying or anticipating the will of others... workmen would not want servants, nor masters slaves." Clearly he recognised that, if tools were advanced enough, they could replace humans in the completion of some tasks, with subsequent effects on the system of work.

There is some insight in each of these concerns, yet they ignore important mechanisms. Many economists contend that machine innovations over the past centuries have been q-complements to labour – an increase in the former raises the marginal product of the latter (Autor, Levy, and Murnane 2003). To the extent that individuals are paid according to their marginal product, this leads to increasing wages. It is possible that automation can reduce human injury risk and boredom in the workplace. For example, Faunce (1965) predicted that “automation, because it represents a major change in production technology, may in the long run produce basic changes in occupational structure”, leading to less division of labour for workers and hence more meaningful and interesting work. Likewise, from 1940 – 1980 in the US, “occupational change skewed away from physically demanding, dangerous and menial work towards skilled blue- and white-collar work” (Autor 2015). Importantly, new jobs may be created by technology, for which humans are more suited than robots, such as employment as data scientists (R. Susskind and D. Susskind 2015).

It is also important to think about the macroeconomic impact from heterogeneous firms adopting emerging technologies. It is unlikely to be the case that all businesses will be impacted equally from technological progress, in the same way that globalisation and policy changes have had distributional consequences across firms (Bustos 2011; Caliendo and Rossi-Hansberg 2012; Baldwin and Okubo 2005; Costinot, Rodríguez-Clare, and Werning 2020; Zwick and Mahon 2017). To this end, the types of firms that adopt new technologies, and the subsequent effect on firm distributions, could have significant macroeconomic repercussions. It could, for example, allow for more ‘superstar firm’ dynamics - where the most productive firms in an industry grow much faster at the expense of their competitors (Suedekum and Woessner 2019; Autor, Dorn, Katz, Patterson, and Van Reenen 2020).

The recent growth in new technologies, sometimes referred to as the ‘Fourth Industrial Revolution’, or ‘Industry 4.0’, include the rising adoption of robots, Artificial Intelligence (AI) and Big Data, Cloud Computing, and other advances in digitisation.

The stock of industrial robots is plotted in Figure 1a: growth was steady until 2010, but has sped up since then. To get a sense of the *relative* impact of robots on labour markets, Figure 1b shows the growth in the number of robots per thousand workers across a range of countries since the early 1990s. This is indicative of the kind of growth - and variation across regions - in new technologies.

The substantial growth in the adoption of new technologies is well-documented (E. Brynjolfsson and McAfee 2014; R. Susskind and D. Susskind 2015; Ford 2015), yet we have only scratched the surface on the broader economic impact. If it replaces jobs, which are being impacted, and does this vary country or type of technology? Is the aggregate impact on employment positive or negative? What is the role that firms play in the adoption and spread of automation technologies? The answers to these questions will have consequences for automation-induced inequality (in income and wealth; across occupations; and between firms), and the resulting policy implications.

Thesis Outline

In this thesis, I investigate some of these questions. There are three distinct chapters, which are connected under the broad aim of analysing the relationship between new technologies and the macroeconomy. Chapter 1 studies the adoption of industrial robots on labour markets in the UK (Kariel 2021b). Chapter 2 explores adoption of automation technologies across Italian firms, and embeds automation in a firm dynamics model to understand the aggregate implications. Chapter 3 estimates revenue elasticity in the UK, and provides a theoretical foundation and empirical evidence that computer software is associated with a rise in revenue elasticity.

Robots and Labour Markets

The growth of industrial robots across the world has picked up in recent years (see Figure 1a). The UK has relatively low adoption compared to other advanced economies (see Figure 1b). Despite this, the number of robots per worker has steadily increased

Figure 1: Growth in Industrial Robots, 1993 - 2016

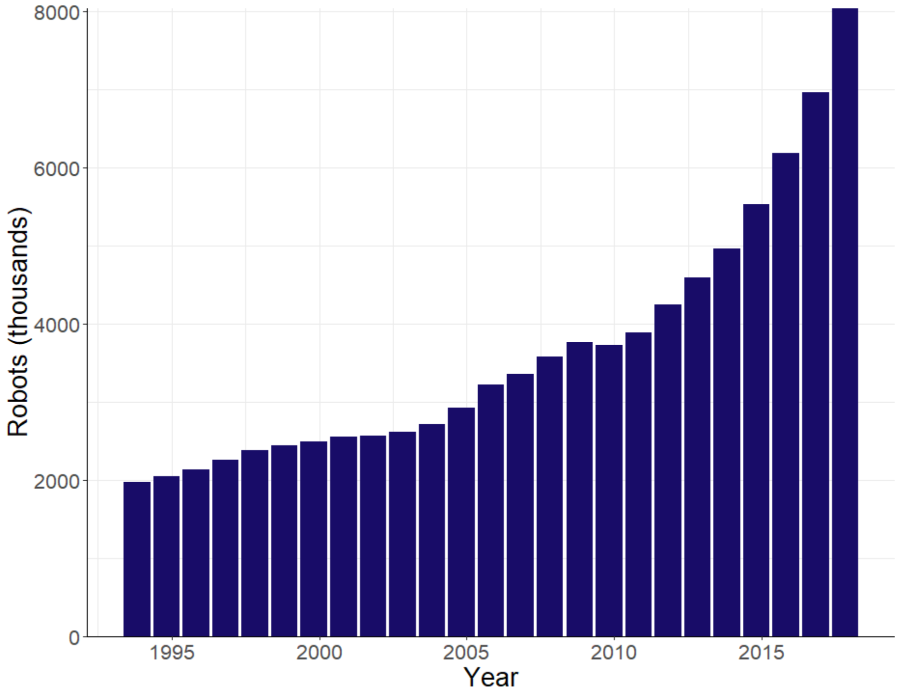

(a) Industrial Robot Stock, 1993 - 2016. Data from the International Federation of Robotics.

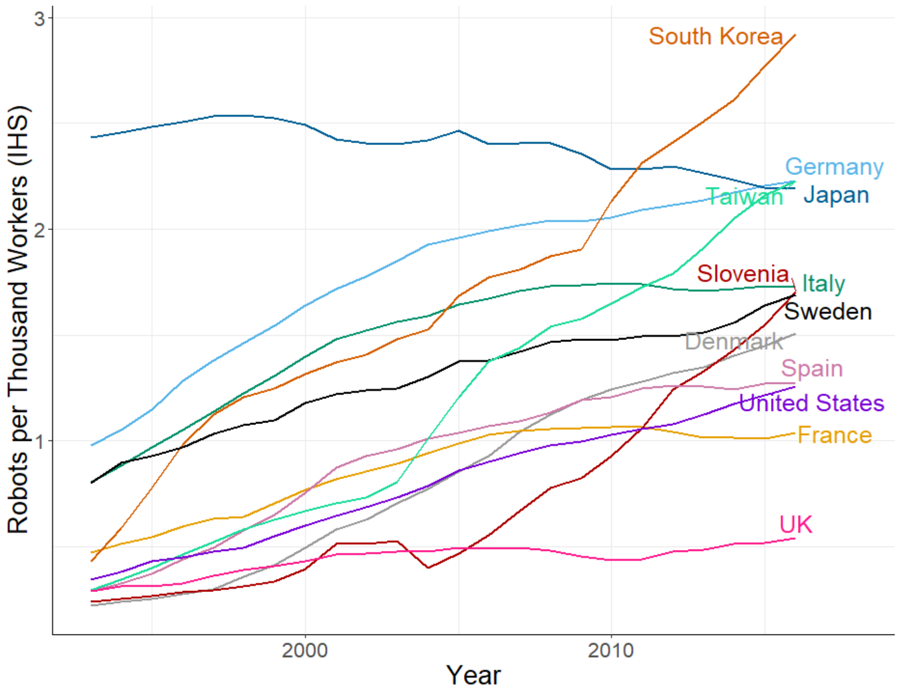

(b) Industrial Robots Adopted per Thousand Full-Time Workers (transformed by Inverse Hyperbolic Sine function). Robots data from the International Federation of Robotics. Employment data from the International Labor Organisation. Calculations by author.

over the last 30 years. Furthermore, it may be that the impact of robots differs at lower levels of adoption, or that the UK is structurally distinct from other countries.

In order to tease apart the causal impact of robots on labour markets, I construct regional measures of adoption based on historical employment shares. However, it is plausible that robots have been purchased in response to changing labour market outcomes, which obscures the relationship of interest. Therefore, I compute two instrumental variables which proxy for the technological frontier of robotics, and the supply of robot innovations. For the former, I use the rise in robots *in other EU countries*, interacted with the pre-analysis employment shares in UK regions, summed over industries. For the latter, I repeat the exercise with global robot-related patents. Such instruments should be correlated with UK robot adoption, as they represent the competition and supply channels in this market, but should not affect labour markets directly.

Contrary to studies in other countries, I find that that higher robot use is associated with increased employment. I also find some evidence of a positive effect on part-time pay, and no effect on full-time salaries, suggesting that the boost to employment is not offset by lower wages. 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 had positive effects on other sectors such as services. This result is unsurprising in the context of capital-skill complementary production processes wherein new technologies substitute for some types of workers, and complement others (E. Lewis 2011).

Firms That Automate

Chapter 2 investigates automation at the firm level, with novel analysis of Italian survey data. I study the types of firms using of a variety of cutting-edge automation technologies, such as AI, 3D Printing, and Cloud Computing. The previous evidence on which firms automate is scarce, and it is important. It will aid our understanding of the role of market structure in technology adoption, how firm distributions change with and influence automation, the impact of entrepreneurs, and the importance of regional economies (Seamans and Raj 2018a).

I find that businesses that automate are significantly larger, and grow faster. To the extent this result is causal, there are implications for the macroeconomy: if automation boosts business growth, and only larger firms can make investments in such technologies, this will raise across-firm inequality. I find evidence that automation causally increases firm size. I estimate the impact of employing automation technologies by comparing firms before and after adoption, to similar firms that do not adopt. The impact of automation adoption on firm size is in the range 1 - 3% in the first two years post-adoption, followed by 4 - 11% after six years. This holds at both the extensive and intensive margins. Importantly, though, this effect only seems to hold for 'skilled' workers. As in Chapter 1, this is indicative of capital-skill complementarity for these technologies.

These findings motivate further investigation with a model, to better understand the effects of automation on firm dynamics and macroeconomic outcomes. This is especially relevant to help reconcile my finding that firms expand with automation, despite macroeconomic evidence that new technologies have reduced employment (Acemoglu and Restrepo 2020). I extend a standard firm dynamics model to include routine and non-routine labour inputs which produce different sets of tasks, and an automation technology which can replace routine workers. The introduction of automation technology lowers the variable cost of producing, but an associated fixed cost of implementation/maintenance leads only the most-productive firms to adopt. Low-productivity firms exit, but the remaining incumbents grow. On aggregate, employment falls, due to the reduction in the number of firms. This model is consistent with my novel findings on automation boosting firm size, alongside the macroeconomic evidence of technological displacement of workers.

Scale Economies

Chapter 3 returns to studying the UK, but the focus is on firms, scale economies, and computer software. Robert Solow famously said that "you can see the computer age everywhere but in the productivity statistics" (Solow 1987). [1] One way we should expect to see the computer age in the data is the way computer software affects costs, which will influence returns to scale and revenue elasticity. Returns to scale, or the inverse cost elasticity, is the ratio of average to marginal costs. Revenue elasticity is equal to returns to scale divided by the markup. As computer software can reduce the cost of replicating various tasks, it should be associated with a fall in marginal costs, and hence a boost to both returns to scale and revenue elasticity.

In this chapter, I set out a simple theoretical framework to study the effect of computer software on revenue elasticity. Optimising firms choose computer software alongside traditional inputs, which changes variable, marginal, and fixed cost curves. If a firm invests in computer software, revenue elasticity is multiplied by a factor which depends on the level of investment in computer software, and the curvature of the software cost function. Clearly in this model revenue elasticity is increasing in computer software.

I estimate revenue elasticity from 1998 - 2014 in a services-dominated economy, using control function and cost share approaches. I find that economy-wide revenue elasticity is close to unity, and has risen over this time period. I also present results on the significant sectoral heterogeneity of revenue elasticity.

The estimation strategy allows for firm- and time-varying estimates of revenue elasticity. I use these to test the implications of the model. I find a strong positive relationship between revenue elasticity and computer software investment.

[1] This statement was widely critiqued, suggesting that perhaps computers aren't actually everywhere, measurement is crucial, and there is a lag in productivity (Triplett 1999). Furthermore, the impact of computers on TFP growth eventually seemed to be visible in the mid-1990s in the United States (Jorgenson 2005).

Chapter 1 Job Creators or Job Killers? Heterogeneous Effects of Industrial Robots on UK Employment

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

1.1 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, Roberts, and King 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 (U.S.) (Acemoglu and Restrepo 2020) and Germany (Dauth, Findeisen, Südekum, and Wößner 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.

[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).

Following Acemoglu and Restrepo (2020), 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.

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 (2020). 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 2020), 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 1.2. The empirical approach and data are introduced in Sections 1.3 and 1.4. The estimation results are presented in Section 1.5, along with further investigation of heterogeneity across industries. Supporting tables, graphs and information can be found in the Appendix.

1.2 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” (R. Susskind and D. 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 Daniel 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, Jaimovich, and Siu 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 (Maarten 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 (2020) is studying the impact of robots on labour market outcomes at a region × 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, Findeisen, Südekum, and Wößner (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.

1.3 Empirical Approach and Identification

This section outlines the underlying theoretical model of Acemoglu and Restrepo (2020) 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 2020).

1.3.1 Model

An economy is made up of commuting zones in which most adjustment to shocks takes place (Moretti 2010; Autor, Dorn, and Hanson 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}.

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

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

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 $s \le M_i$ and $s > M_i$:

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

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:

$$d\ln L_c^d = -\sum_{i\in\mathcal{I}} \ell_{ci} \frac{dM_i}{1 - M_i} - \sigma \sum_{i\in\mathcal{I}} \ell_{ci} d\ln P_{Xci} + d\ln Y_c \tag{1.1}$$

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

The first effect in equation (1.1) 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 (1.1) 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 (1.1) 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:

$$L_i = (1 - M_i)\frac{X_i}{A_i\gamma} \implies \frac{X_i}{A_i} = \frac{\gamma L_i}{1 - M_i} \tag{1.2}$$

$$R_i = M_i\frac{X_i}{A_i} \implies M_i = \frac{R_i A_i}{L_i} \tag{1.3}$$

Log differentiating the rearranged robot equation (1.3) $M_i = \frac{R_i A_i}{L_i}$ yields $dM_i \approx \frac{dR_i A_i}{L_i}$. Plugging this into the rearranged labour equation (1.2) $\frac{X_i}{A_i} = \frac{\gamma L_i}{1-M_i}$ 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 (1.1) 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] A similar relationship for the change in wages is found in Section A.2 of the Appendix, but requires the general equilibrium model, which isn't included here for parsimony.

1.3.2 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 (2020): 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:

$$X_{c,t+h} - X_{c,t} = \beta \underbrace{\sum_{i \in \mathcal{I}} \ell_{ci}^t \frac{R_{i,t+h} - R_{i,t}}{L_{i,t}}}_{\text{Exposure to Robots}_c} + \Gamma C_{c,t} + \epsilon_c \tag{1.4}$$

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. 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, Dorn, and Hanson 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 2020). 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.

1.4 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.

1.4.1 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.

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 -

[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.

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 thresholds that

[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

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 A.3 of the Appendix.

1.4.2 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 (2020) and 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]

[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

1.4.3 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 A.6 of the Appendix - data is grouped in 16 industries for the analysis. The summary data is provided in Table A.1 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 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 A.3 in the Appendix, with three of our four dependent variables showing a decline in (Local Authority) population-weighted terms.

The regional $\times$ industry census data allows computation of Exposure to Robots, introduced in Section 1.3.2: $\sum_{i \in \mathcal{I}} \ell_{ci} \frac{dR_i}{L_i}$. Each Local Authority receives a value for Exposure to Robots based on this calculation. To adapt to the data limitations, the formula is:

$$\begin{aligned} &\text{Exposure to Robots} \\ &\text{from 1993 to 2011}_c \end{aligned} = \sum_{i \in \mathcal{I}} \ell_{ci}^{1991} \left(\frac{R_{i,2011}}{L_{i,1991}} - \frac{R_{i,1993}}{L_{i,1991}} \right) \tag{1.5}$$

The distribution of Exposure to Robots over Local Authorities is summarised in Table A.2. 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 A.4. 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.

1.4.4 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 are in Sections A.6 and A.7 of the Appendix. The variable is computed in the same way as for

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

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

Exposure to Robots:

$$\begin{aligned}\text{Exposure to Automation}\\ \text{Patents from 1991 to 2011}_c\end{aligned} = \sum_{i \in \mathcal{I}} \ell_{ci}^{1991} \left(\frac{P_{i,2011}^{UK}}{L_{i,1991}} - \frac{P_{i,1991}^{UK}}{L_{i,1991}} \right) \tag{1.6}$$

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 A.2. The distribution is approximately bell-shaped with a right-skew.

1.4.5 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 A.5 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, Machin, and Meghir (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, Katz, and Kearney 2005).[12] The summary statistics are described in Table A.6. 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.

1.4.6 Instruments

For the 'world' technological frontier, I employ data from EU countries (hereafter referred to as the EU-7) (Acemoglu and Restrepo 2020). The economies which 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]:

$$\begin{aligned}&\text{Exogenous exposure to} \\ &\text{robots from 1995 to } 2011_c\end{aligned} \quad = \sum_{i \in \mathcal{I}} \ell_{ci}^{1981} \left(\frac{R_{i,2011}^{EU-7}}{L_{i,1995}^{EU-7}} - \frac{R_{i,1995}^{EU-7}}{L_{i,1995}^{EU-7}} \right) \tag{1.7}$$

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 1.3.2, 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-

[12]The 80^{th} decile is chosen over the more commonly-used 90^{th} decile due to data limitations.
[13]Unfortunately the EU-7 data only exists back to 1995.

related patent data:

$$\begin{array}{l}\text{Exposure to Global}\\ \text{Automation Patents}\\ \text{from 1981 to 1991}_c\end{array} = \sum_{i \in \mathcal{I}} \ell_{ci}^{1981} \left(\frac{P_{i,1991}^{Global}}{L_{i,1981}} - \frac{P_{i,1981}^{Global}}{L_{i,1981}} \right) \tag{1.8}$$

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 A.2 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.

1.4.7 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 2020; Dauth, Findeisen, Südekum, and Wößner 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 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, Dorn, and Hanson (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.

1.5 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 (2020) and Dauth, Findeisen, Südekum, and Wößner (2017), which allows for a cleaner comparison with results over an almost identical time frame, and a comparable methodology, in the US and Germany.

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

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.

1.5.1 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 A.4 of the Appendix.

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 $h_{c,2011}^{mean} - h_{c,1997}^{mean}$.

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 1.1, with the rest in Table A.8 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.[15] There is evidence of instrument validity across all baseline specifications.[16] There is also evidence of endogeneity of Exposure to Robots, but I report OLS estimates for completeness.[17] The evidence of endogeneity is less clear when I investigate at the industry level. All test results can be found in Table A.11 in the Appendix.

Firstly, I report the differences between the OLS and 2SLS estimates in Table 1.1. 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 A.11). 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 1.1. The estimated coefficient on robot exposure is negative and statistically significant in column (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.

[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 above 5%.

[17] Wu-Hausman tests find the 2SLS coefficients are significantly different from OLS in most of the models estimated.

Table 1.1: 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.

An important control here is 'baseline industry shares', the inclusion of which leads to the estimated coefficient becoming positive and significant, as can be seen by comparing results in columns (2) to (3), and (6) to (7). This variable is also included in similar studies in the US Acemoglu and Restrepo (2020) and Germany Dauth, Findeisen, Südekum, and Wößner (2017), although its importance is less pronounced. The results for employment at the extensive margin can be seen in Figure 1.1.

Figure 1.1: Estimated 2SLS coefficients from Exposure to Robots on Employment Rates

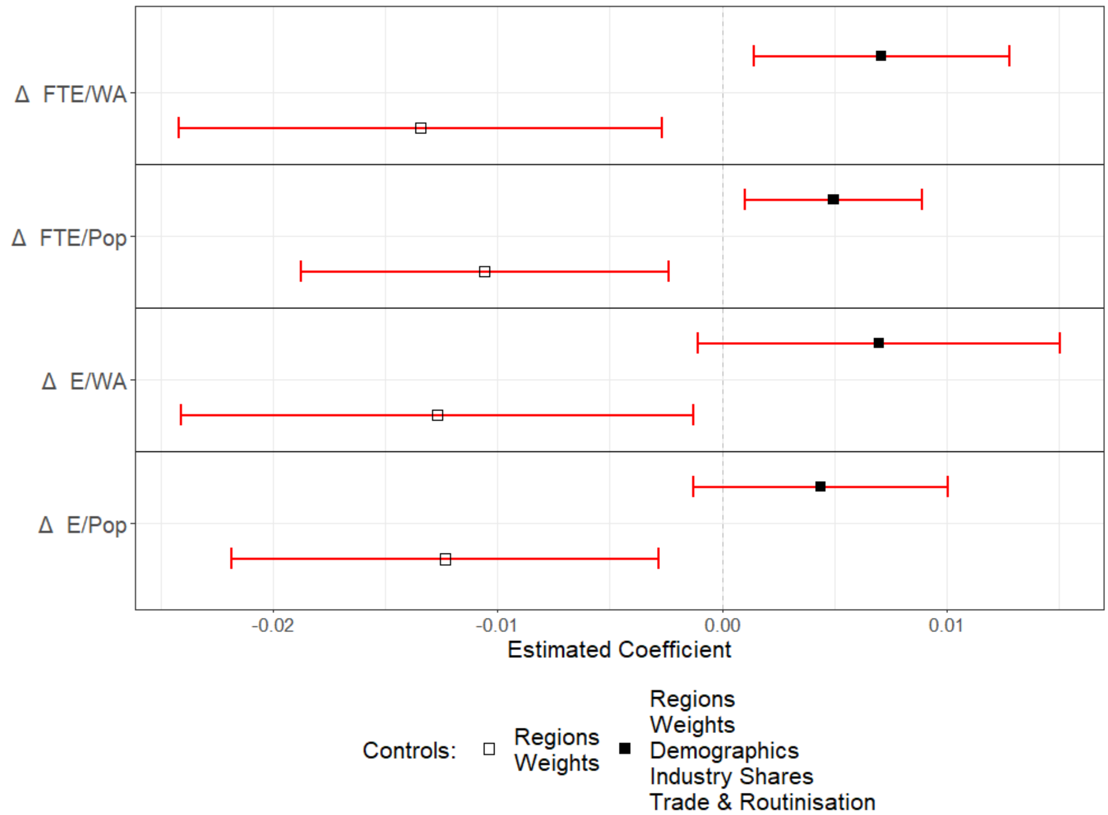

Note: The points are estimated 2SLS coefficients of Exposure to Robots on the change in employment-to-population rates (1991 - 2011). The lines indicate the 95% confidence intervals. Each panel shows the results for a different measure of the change in the employment rate, for two models with a different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

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

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, I provide a breakdown of the variance inflation factors (VIF) (Fox and Monette 1992) in the fully-specified models, with all controls. Table A.13 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 A.4 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 *reduced* employment by around 6 workers (Acemoglu and Restrepo 2020). 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, Findeisen, Südekum, and Wößner 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

[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.

wages in Section 1.5.1.

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 A.9 in the Appendix).

On the intensive margin, the results are in Panel II of Table 1.1 for part-time hours, and in A.8 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 A.11), 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.

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 A.11 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 1.1. 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 A.8. 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.

Figure 1.2: Estimated OLS coefficients from Exposure to Robots on Weekly Pay

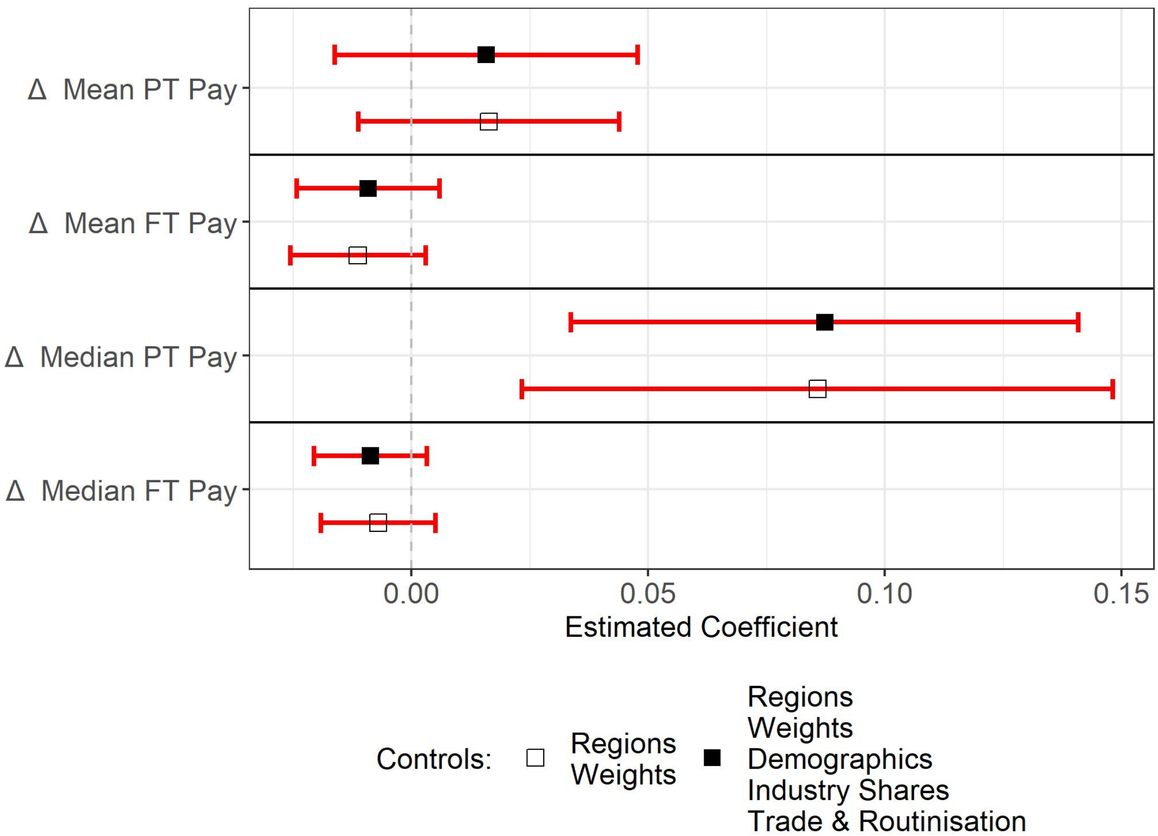

Note: The points are estimated OLS coefficients of Exposure to Robots on the change in weekly pay (1997 - 2011). The lines indicate the 95% confidence intervals. Each panel shows the results for a different measure of the change in weekly pay, for two models with a different set of covariates (the filled 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 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 2020). The estimates from Germany are much lower, with an extra robot per thousand workers reducing average pay by around €15 per year (Dauth, Findeisen, Südekum, and Wößner 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 A.9 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 1.1. 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

[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.

[21] The computation here uses the estimates in Dauth, Findeisen, Südekum, and Wößner (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 $-€15 = -0.05 \times 5,959/20$.

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 A.8).

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.

1.5.2 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.

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 A.5 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 $\Delta\text{Emp. Share}_c^i$ 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 1.3. It's important to note that employment shares sum to one, so a rise in one sector must be offset in another industry.

Figure 1.3: Impact of Exposure to Robots on Industry Employment (1991 - 2011)

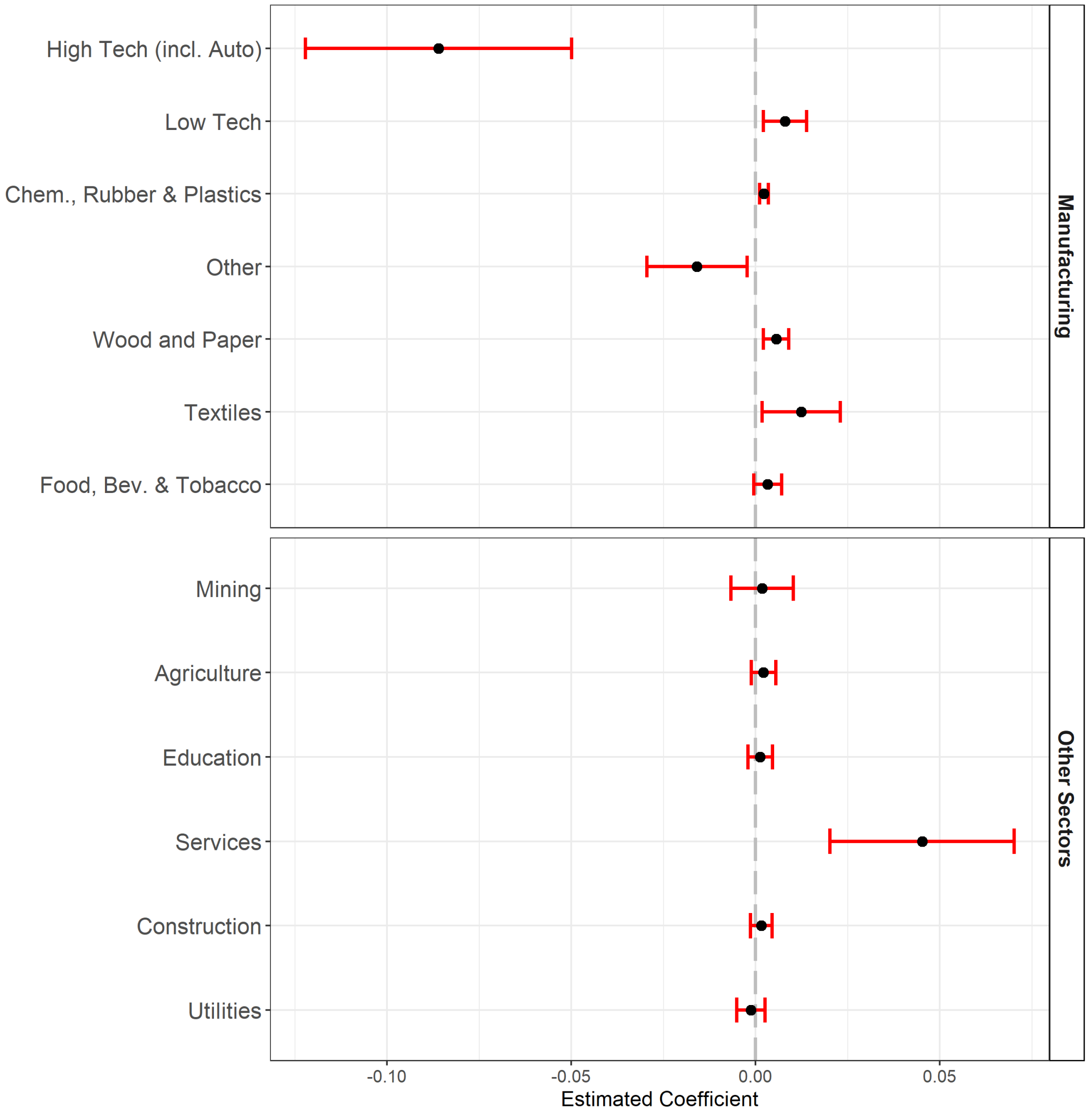

Note: The points represent estimated 2SLS coefficients of Exposure to Robots on the change in employment shares ratio (1991 - 2011) across a set of industries. The lines indicate the 95% confidence intervals. The models have the full set of controls used in baseline estimation.

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

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 A.1 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 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

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

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 across all industries. The instruments are not always valid, although the Sargan values suggest that validity does hold for the two most important sectors: High Tech Manufacturing and Services. There is some evidence of endogeneity of Exposure to Robots, but not for all industries. See Table A.12 in the Appendix for associated p-values on instrument tests. Finally, 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.

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

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

Figure 1.4: Impact of Exposure to Robots (excl. Auto) on Employment in High Tech (incl. Auto) and Services from 1991 - 2011

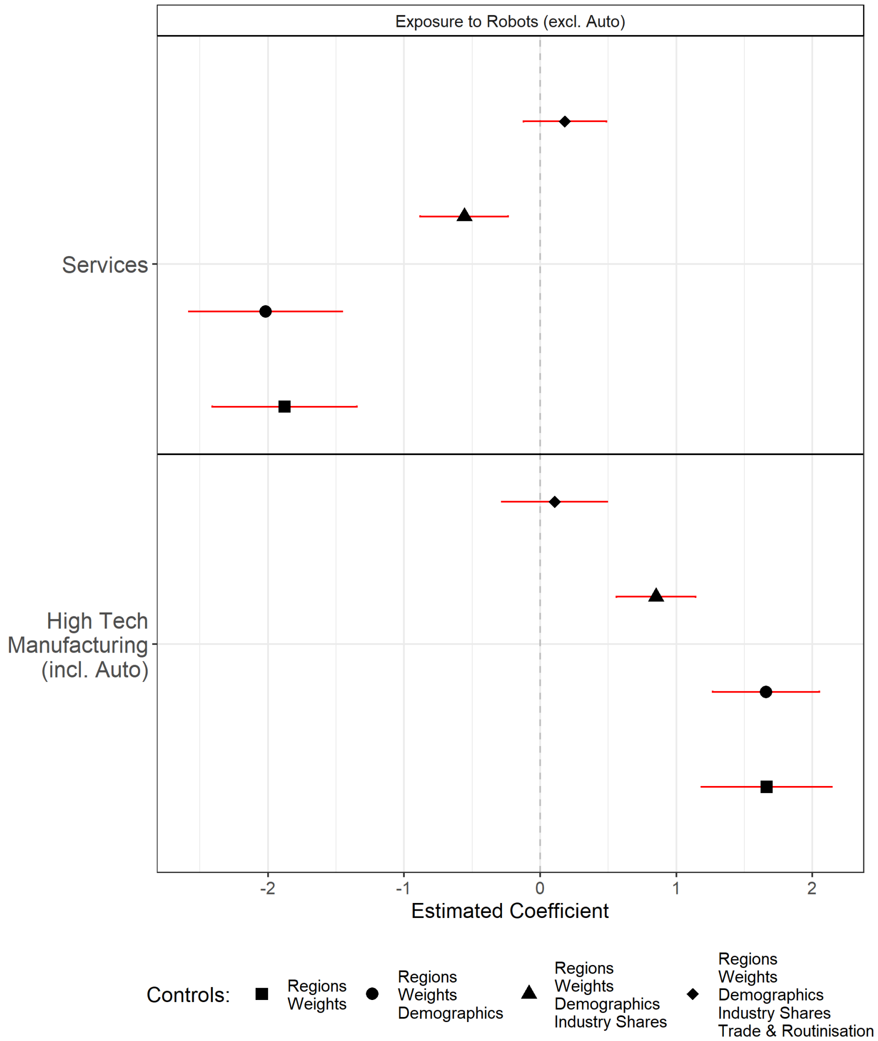

Note: Estimated coefficients and 95% confidence interval for Exposure to Robots (excl. Auto) on change in employment shares ratio (1991-2011) in High Tech Manufacturing (incl. Auto) and services, with an increasing set of controls.

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

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

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 con-

[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.

trols and instruments. The results can seen in Figure 1.4 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 1.4. 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. 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.

1.6 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 (2020) 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, Findeisen, Südekum, and Wößner (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.

Chapter 2

Firms That Automate: Evidence & Theory

Which firms are using automation technologies, and what are the effects on firms and the aggregate economy? Using Italian survey data on adoption of cutting-edge technologies, such as Artificial Intelligence and Cloud Computing, I compile a set of novel findings on firm adoption. I find that firms that automate are larger, pay higher wages, and are more productive. In addition, I find technology adopters grow faster once they start using automation technologies. I embed technology adoption in a heterogeneous firm model to investigate the equilibrium and aggregate implications of automation. The model can reconcile my firm-level evidence of technology adoption boosting firm size, with the various macro studies suggesting a negative overall effect on employment: the equilibrium effect on prices, entry and exit are crucial.

JEL classification: D22, J23, J24, L11, O14

Keywords: Automation, Firm Dynamics, Employment, Technology, Productivity.

2.1 Introduction

A recent wave of research has considered the impact of automation technologies, which can perform tasks that - until recently - were limited to the domain of human labour. The sharp increase in the capability and spread of automation technologies has been widely discussed and studied (Acemoglu and Restrepo 2016; Daniel Susskind 2017). Robots are an established automation technology, and the *growth rate* of the global stock of robots has been rising rapidly, from around 4% in 2010 to over 15% in 2017 Robotics (2014). The increased interest and use of Artificial Intelligence has been startling. For example, Fujii and Managi (2018) found the global number of AI-related patents had tripled from 2012 to 2016. However, most of the empirical analysis of such technologies has taken place at a macro level. For example, studies of how national robot adoption has impacted labour markets are quite commonplace (Graetz and Michaels 2015; Acemoglu and Restrepo 2020; Kariel 2021b).

This paper fills a noticeable gap in the literature by studying the adoption of automation technologies across firms. This research has both empirical and theoretical components. Firstly, I use a unique panel dataset of Italian firms to look at the heterogeneity in adoption, and the impact, of new automation technologies. This data has five main advantages over many existing studies on automation among firms (Kwon and Stoneman 1995; Bartel, Ichniowski, and Shaw 2007; Dinlersoz and Wolf 2018; Kromann and Sorensen 2019; Zator 2019; Acemoglu, Lelarge, and Restrepo 2020): (1) it includes a wide range of cutting-edge automation technologies, (2) there is information on *when* firms begin automating, (3) the data is from recent years, (4) the firms comprise a panel, and (5) there is a large sample of around 5,000 firms.

The evidence suggests that automation-adopting firms are already larger and more productive, but also benefit significantly once they adopt. The static results suggest automating firms are larger, pay higher wages, and are more productive. Taking a dynamic perspective, I find that these firms have higher growth rates. Finally, I find that firms grow faster once they become automaters.

There is concern that automation will disrupt labour markets, replacing the work of many across a variety of industries and occupations (R. Susskind and D. Susskind 2015). It seems that this has already happened in some contexts, with the most widely-studied automation technology to date: robots. There is growing evidence that, on aggregate, the adoption of robots has led to a fall in total employment (Acemoglu and Restrepo 2020; Chiacchio, Petropoulos, and Pichler 2018; Carbonero, Ernst, and Weber 2020). How can we reconcile this with my findings that firms using automation technologies get larger after adoption? I take a standard firm dynamics model and introduce two simple extensions: (1) firms hire routine and non-routine labour, with adjustment costs, and (2) firms can invest in automation technology which can perform the tasks of routine workers. The model replicates many of my empirical findings, and crucially it shows that while automating firms may grow, aggregate employment can shrink. It turns out that the general equilibrium effect is crucial: automating firms expand, prices fall, and low-productivity firms exit.

The model allows me to highlight the general equilibrium effects of automation with heterogeneous firms. It can be used to understand what automation among incumbents means for competition, aggregate productivity, the rate at which startups enter and at which firms shut down.

As highlighted in Seamans and Raj (2018b), economists "lack an understanding about how and when robotics and AI contribute to firm-level productivity, the conditions under which robotics and AI complement or substitute for labor, how these technologies affect new firm formation, and how they shape regional economies." This has been primarily due to a lack of available data on technological adoption at the firm level.

A number of recent studies have attempted to respond to this problem, utilising firm-level automation data from the USA (Dinlersoz and Wolf 2018), Denmark (Kromann and Sorensen 2019; Humlum 2019), Germany (Zator 2019), China (Cheng, Jia, D. Li, and H. Li 2019), Spain (Koch, Manuylov, and Smolka 2019) and the EU (Jager, Moll, Zanker, and Som 2015). However each study has its limitations, either the time period of analysis, the variety of automation technologies, or a small sample of firms. Furthermore, only Koch, Manuylov, and Smolka (2019) look at the relationship between automation and productivity.

I use the Bank of Italy's "Survey of Industrial and Service Firms" firm-level automation data to investigate the relationship between automation, productivity, firm dynamics and employment. This survey has been conducted over a long time, for a large sample of firms, asking a wide range of automation-related questions, including recent questions on "Industry 4.0". For example, it asks about the use of Artificial Intelligence (AI), the Internet of Things (IoT), and Industrial Robotics. This survey has approximately 5,000 firms in the panel, and has been run from 1984 to the present. It is an excellent dataset to examine which firms automate, how this has changed over time, and the impact of these choices. The first section analyses the distribution of automation across firms, investigates variation between industries, and provides descriptive statistics on firm-level automation. Then I leverage the panel dataset to better understand how automation affects firms.

My results are consistent with this nascent field of research on firm-level adoption of automation technologies. For example, firms that automate tend to be larger (Zator 2019; Bartelsman, Van Leeuwen, and Nieuwenhuijsen 1998; Koch, Manuylov, and Smolka 2019; Cheng, Jia, D. Li, and H. Li 2019; Acemoglu, Lelarge, and Restrepo 2020). Likewise, I find that such firms are 2 - 6% larger, even when accounting for firm age, sector, and region. I also find that firms that automate are more productive, with estimated TFP around 3% higher. Other studies have also found this positive relationship between automation and productivity, although often focusing only on robots or manufacturing (Dinlersoz and Wolf 2018; Kromann and Sorensen 2019; Zator 2019; Bartelsman, Van Leeuwen, and Nieuwenhuijsen 1998; Kwon and Stoneman 1995; Koch, Manuylov, and Smolka 2019).

I find that technology adopters grow faster than non-adopters, but there is little evidence of this difference for blue-collar employment. This suggests that some workers benefit from these technologies, but it varies across different skills and industries

(Dixon, Hong, and Wu 2020; Kariel 2021b; Dauth, Findeisen, Südekum, and Wößner 2017). By focusing on the impact *on adoption*, I continue to find significant overall employment effects, suggesting the contribution of selection effects is muted. Again, there seems to be a limited impact on unskilled workers and productivity (apart for firms using robotics). In order to fully understand what happens to firms on adoption of technologies, I conduct event studies which show a considerable boost in employment from automation technologies, in the range of 4 - 11% five years after adoption.

A model of automation is built, extending the Hopenhayn (1992) framework. This is a workhorse model of industry dynamics which lends itself to questions about the size distribution of firms, entry and exit behaviours, and the impact of policies (e.g. a robot tax). Automation is built into the production function through a task-based framework (Acemoglu and Restrepo 2016), where such technology can displace routine workers, but complements non-routine labour. The calibrated model produces predictions about the effect of automation, and can then be compared to the data.

The intuition behind the model is simple: automation technology is assumed to be cheaper than human labour, but it requires some fixed set-up and implementation costs. Thus only the most-productive and profitable firms can automate, because they can afford the fixed costs. So the effect of automation may be partly due to selection, as well as the reallocation of output towards already-productive firms, rather than due to it being a huge efficiency-saving technology.

The paper is organised as follows. Section 2.2 presents a brief literature review, covering the economics of automation, existing research on the various technologies in this study, and heterogeneous firm models. Section 2.3 introduces the data, presents results on firm dynamics, and summarises findings on automation adoption and investment. Section 2.4 focuses on the firms that automate. I investigate their growth and productivity rates, and look at the causal impact of adopting automation technologies. The model is described in Section 2.5, along with the calibration and results.

2.2 Literature Review

The fast-growing literature on the economics of automation has a variety of branches, considering different technologies, contexts, and theoretical frameworks. New technologies which automate tasks can affect the labour market in a variety of ways. Firstly, workers are most likely to benefit if they produce labour tasks which are complemented by automation technologies (Autor 2015). This is somewhat related to skill-biased technical change (SBTC), as those with higher skills - and perhaps skills more complementary to machines - have faced the greatest rise in demand (Griliches 1969). Secondly, wage gains depend on the elasticity of labour supply; if these complementary skills are abundant in the labour market and workers have elastic supply, workers may flow into the automated industry, somewhat mitigating the income growth (Autor 2015). Thirdly, output elasticity of demand and income elasticity both play a crucial role. Autor (2015) highlights the impact of the former by noting that the share of expenditure on food has fallen following productivity improvements in agriculture, while the share on healthcare has risen with gains in productivity. The importance of income elasticity is most clearly related to Baumol's 1967 distinction between "technologically lagging" sectors and those at the vanguard of automation. Rising productivity in leading sectors may boost employment in the laggards, where demand for the final goods are strongly income elastic. The spillover effects of productivity improvements essentially arise from an increase in output where larger untapped demand exists, and the effects have been found to be consistently positive (Autor and Salomons 2017).

An important strand of research has highlighted that jobs are not "lumps of labour" (R. Susskind and D. Susskind 2015) but a combination of tasks which may require distinct skills and inputs. Many recent economic models of automation embed this concept in the production function, by assuming work consists of a continuum of tasks, increasing in cognitive difficulty. For a given level of technological progress, machines can replace labour up to a certain point, but beyond that human input is required. For example, this may include tasks for which face-to-face contact is still essential, or the variety and complexity of tasks has not yet been competently replicated by AI, such as private hire driving, managing other people, or therapy. This task-based framework allows us to identify different effects of automation: the "displacement effect" whereby workers are replaced, putting downwards pressure on wages and employment; the "productivity effect" which raises demand for labour in non-automated tasks; "capital accumulation" if automation raises demand for capital, and thus for labour; "capital deepening" which is simply improved productivity in already-automated tasks (Acemoglu and Restrepo 2018). Therefore, the overall labour market impact of adopting automation technology is ambiguous.

Furthermore, it is likely to depend on the complementarity between the technologies and various labour skills, along with the varying demand for different tasks. Finally, the impact will depend on the creation of new tasks, in which humans still have a comparative advantage. In the task-based framework, this can be modelled as the continuum extending, which under certain assumptions leads to a stable balanced growth path (Acemoglu and Restrepo 2016), rather than the eventual destruction of all jobs (Daniel Susskind 2017).

The empirical research has used a variety of datasets and identification methods to consider the relationship between labour markets and automation. Many studies using national industrial robotics data have found a relationship between this technology and employment, but the effect varies across countries. There is evidence of displacement of jobs in the U.S., France, and the E.U. as a whole, while there was no overall fall in employment in Germany, but job composition was affected (Acemoglu and Restrepo 2020; Chiacchio, Petropoulos, and Pichler 2018; Aghion, Antonin, and Bunel 2019; Dauth, Findeisen, Südekum, and Wößner 2017). Similarly, there is evidence robot adoption led to changes in the skill composition of labour across Europe, alongside increases in labour productivity (Graetz and Michaels 2015). Globally, there is evidence that robots have reduced worldwide employment (Carbonero, Ernst, and Weber 2020).

Clearly, the validity of these results rests on the instruments successfully eliminating endogeneity issues, such as cost-push shocks which simultaneously increase the incentive to adopt robots and reduce demand for labour. One solution is to set up simultaneous equations and estimate with 3SLS to account for the two-way causal relationship between robots and labour. Jung and Lim (2020) perform such an analysis across 42 countries and find robots reduced employment, and skewed labour demand from unskilled to skilled workers. However, the use of national robots data, rather than a more disaggregated measure, may yield misleading inference, especially if there is significant heterogeneity. This may mask interesting behaviour occurring where the robots are adopted: within firms.

Thus the recent firm-level research may provide more insight. Bessen, Martin Goos, Salomons, and Berge (2019) take advantage of Dutch micro-data to show that increased expenditure on automation increases the probability of worker separation and reduces employment at the intensive margin, with no impact on wage rates. In contrast, Barth, Roed, Schone, and Umblijs (2020) find that robot adoption by Norwegian firms has led to an increase in wages, especially for those with higher education, and managers. On the flip side, Dechezleprêtre, Hémous, Olsen, and Zanella (2020) consider the opposite direction of causality, and they show their novel measure of firm-level automation innovations rises with exogenous shocks to low-skill wages. This evidence suggests firms respond to changes in factor prices by making more investments in innovations to replace low-skill workers.

Recent studies on the impact and heterogeneity of automation by firms cover a range of countries and technologies. For example, Acemoglu, Lelarge, and Restrepo (2020) investigate robot adoption by French manufacturing firms, with adopters boosting their employment at the expense of rivals, but also experiencing declines in labour share. Likewise, Koch, Manuylov, and Smolka (2019) examine robot adoption by Spanish firms, and find similar gains in output, employment, and a falling labour share, as well as reallocation towards adopters. Another large-sample study of robot adoption in Canada also finds a significant increase in firm size for technology adopters, but a fall in the number of managers, suggesting this innovation allows for reduced variation in the production process and thus changes in organisational design (Dixon, Hong, and Wu 2020). Zator (2019) shows that adoption of digitization and automation technologies is associated with local labour scarcity, although variation across technologies and industries is important. Notably, they find similar substitution effects with robots as in Acemoglu, Lelarge, and Restrepo (2020), but complementarities between digitization and labour in some service industries. This suggests that looking beyond robots will give important insights about the macroeconomic effects of new technologies. However, such data is hard to come by, especially at the firm level (Seamans and Raj 2018b), leaving economists with a limited understanding of adoption of automation, and its effects.

One such study on U.S. firms has found that adopters of advanced technologies tend to be larger and older (Zolas, Kroff, Erik Brynjolfsson, McElheran, Beede, Buffington, Goldschlag, Foster, and Dinlersoz 2020). Furthermore, they show that while digitisation (such as Cloud Computing) is quite widespread, the adoption of more advanced technologies (such as Artificial Intelligence) is rare. Finally, they find a hierarchy of technological adoption, such that firms using the more advanced technologies also tend to use the more commonly adopted innovations. The results in this paper align with the majority of their findings, suggesting that my other conclusions are possibly generalisable to contexts such as the United States.

I will briefly touch upon the theoretical and empirical work on non-robotic technologies, such as Cloud Computing and Artificial Intelligence (AI).

Cloud Computing involves using remote servers to manage data, and can encompass a range of technologies such as software (e.g. Google Docs), platforms (e.g. Microsoft Azure), and data storage (e.g. Amazon S3) (Bayrak, Conley, and Wilkie 2011). There is very little economic research on the adoption and impact of Cloud Computing. Bayrak, Conley, and Wilkie (2011) highlight this gap in the literature, and they underscore that Cloud Computing doesn't necessarily permit the completion of tasks that were unavailable with previous technologies. However, it does reduce fixed costs, and offers faster scalability than before. This insight is theoretically grounded in Etro (2009) in a DSGE framework: they conclude that Cloud Computing can raise the growth rate and employment through the entry of new small- and medium-sized enterprises (SMEs). In fact, Etro (2009) estimates that the largest medium-run impact in the EU would occur in Italy, with approximately 80,000 new jobs under fast adoption.

AI is "the capability of a machine to imitate intelligent human behavior" (Merriam Webster, 2020) and has been subject to a lot of discussion within the economics profession. It is sometimes used interchangeably with 'automation'. For example, Aghion, B. Jones, and C. Jones (2017) study the impact of AI on long-run growth, extending the Zeira (1998) model (which itself considers 'machines' that are more akin to robots) and they regularly employ the language of tasks being 'automated' or not. However, AI may not *always* simply be an automation technology, and may need to be combined with robots, Big Data, Cloud Computing, or other technologies, in order to automate a task. It seems more reasonable to think of the common tasks of AI: decision-making, prediction and anomaly detection, among others (Babina, Fedyk, He, and Hodson 2020). There are both theoretical (Aghion, B. Jones, and C. Jones 2017; Acemoglu and Restrepo 2018) and empirical (Babina, Fedyk, He, and Hodson 2020; Zator 2019) contributions on the economic impact of AI: however, as yet, there is limited *direct* evidence on firm adoption of this technology, due to a lack of data.

The Internet of Things (IoT) is a set of technologies which connect machines and the internet. They are part of an infrastructure that connects sensors and actuators to computer systems via networks, and these systems can track the performance and actions of machines through the internet (Manyika, Chiu, Bisson, Woetzel, Bughin, and Aharon 2015). There is limited economic research on the impact of these technologies, although Edquist, Goodridge, and Haskel (2019) find significant relationships between TFP growth and the number of IoT connections per person. This country-level analysis likely suffers from endogeneity concerns, as the authors note, so can really only be interpreted in terms of high-level correlations.

3D printing technology uses a digital depiction of an item alongside "additive manufacturing" techniques to create physical objects (Bechtold 2015). There is little existing research on the economic impact of this technology, and the findings are mostly descriptive, context-specific, or lacking a theoretical and empirical basis (Weller, Kleer, and Piller 2015; Bechtold 2015; Reeves and Mendis 2015). One exception is work by Abeliansky, Martínez-Zarzoso, and Prettner (2015) which introduces 3D printing in a Melitz (2003) framework such that countries facing high trade costs will be faster adopters of this technology. Their findings support this theoretical prediction, but the analysis does take place at a country level.

Closer in spirit to my research are two important papers which take a similar approach of combining microdata on firms and technology with a structural model of the macroeconomy. Research by De Ridder (2019) and Lashkari, Bauer, and Boussard (2019) look at the impact of software and Information Technology (IT), bringing together novel empirical insights on firms (from French microdata). They each build structural models to delve into the implications for market power and productivity. The research by De Ridder (2019) presents a mechanism to explain the slowdown in productivity growth, decline in business dynamism and rise in market power, all as a function of the rise in *intangibles* such as software, which reduce marginal costs and raise fixed costs. My model also embeds the idea that automation technology affects firms' cost structures. Crucially, this paper speaks to the scalability of intangible inputs, and the relative efficiency of such inputs differing across firms. I differ from this assumption by introducing a Markov process for firm productivity, and the endogenous automation choice depends on this. Somewhat differently, Lashkari, Bauer, and Boussard (2019) frame IT inputs as having marginal product rising with scale: it helps deal with organisational limits. Their model implies endogeneity of firm-level returns to scale, which leads to larger firms using more IT, and a lower labour share.

Finally, a recent paper from Hubmer and Restrepo (2021) similarly embeds automation in a standard firm dynamics model, with a fixed cost to automation, to investigate long-run shifts in the U.S. labour share and the role of endogeneous markups to study reallocation. They find that substitution of labour for machines has played an important role in the falling manufacturing labour share.

2.3 Data

The data comes from a survey of Italian firms, which I describe in this section. The reason for using this survey is threefold: it samples a large number of representative firms; most firms stay in the survey allowing for panel analysis; questions on automation technology adoption are asked regularly.

This section introduces the data, presents results on the firm dynamics, and summarises the findings on automation adoption. The firm dynamics analysis is focused around two key features of the firm dynamics literature: size and age (Haltiwanger, Jarmin, and Miranda 2013). I find that firm growth rises in size and falls in age.

I present new evidence on technology adoption and investment across firm distributions. A consistent finding is that adopters of automation technologies tend to be larger, but do not systematically vary by age. Technology adopters also have higher productivity and turnover, and pay higher wages.

2.3.1 Introduction

The data is from the Bank of Italy "Survey of Industrial and Service Firms". This survey asks a wide range of questions from a representative sample of around 5,000 firms with at least 20 employees. From 2001 and 2002 respectively, the survey has included smaller businesses and services firms, to better represent the Italian firm population. Firms in the survey are always contacted to be included in the next year, but may not be included in the sample if they switch sector or fall below the threshold size.

The survey uses a one-stage stratified sample design. The strata are selected according to sector, size (average annual number of employees), and region of head office. Firm-year weights are computed with Horvitz-Thompson estimators, and post-stratification adjusted weights are computed using outside information of certain geographic characteristics. Further details on the survey design can be found in the Appendix B.1.

The survey collects annual data on firm employment, investment, turnover, and other structural information. It also has further detail on specific issues, which may not be asked annually, such as strategies, governance, and technological factors. Firms were asked questions on automation technologies from 2015 - 2019. In odd years, they are asked about the use of distinct technologies which have the potential to replace labour inputs across some tasks. Firms give a binary response on the current use of each technology, alongside the number of years since adoption (in 2015 and 2017). These technologies are: Cloud Computing, Artificial Intelligence (AI), Big Data, the Internet of Things (IoT), Industrial Robotics, and 3D Printing. From 2016 - 2019, firms are asked the share of total investment which was spent on advanced digital technologies.

Information on the adoption of various technologies was collected in 2015, 2017, and 2019, with approximately 3,400, 3,800, and 2,100 firms responding in each year, respectively.[1]

2.3.2 Firm Dynamics

To begin, I describe some general facts about the firms in this survey. Firm age and size distributions seem to follow an approximately Pareto distribution, as has been found elsewhere in the literature (Axtell 2001; Garicano, Lelarge, and Van Reenen 2016; Segarra and Teruel 2012). For these plots, see Figure B.19 in the Appendix. The average firm is 36 years old, has 50 employees, and earns over €10.7 million in turnover. The average firm wage per worker is around €29,600.

Following Davis, Haltiwanger, and Schuh (1996), firm-level growth is computed

[1]The decline in the 2019 sample is attributed to the COVID-19 pandemic. Data is collected for 2019 at the start of the following year, from January to May 2020. The Bank of Italy reports that the response rate did not vary systematically across sectors, firm size classes, and regions. See https://www.bancaditalia.it/pubblicazioni/indagine-imprese/2019-indagine-imprese/en_statistiche_IIS_01072020.pdf?language_id=1.

as:

$$\gamma_{it} = \frac{E_{it} - E_{it-1}}{0.5 \times (E_{it} + E_{it-1})},$$

where E_{it} is the employment of firm i in period t. Over the time period 2008 - 2019, firm growth rises in firm size and falls in age, replicating the results found for Italian firms in Manaresi (2015), which analysed a panel of all private Italian firms.[2] I run a set of regressions of employment growth on firm size and/or firm age, and with year fixed effects. These results are contained in Table 2.1 below (and Table B.15 in the Appendix without binning the variables).

We can see that firms grow at a faster rate if they are in larger size bins. For example, employment growth is about twice as fast for firms with more than 500 workers compared to businesses with under 100 employees. The subdued up-or-out dynamics of Italian firms (relative to U.S. firms) is found in this data, as has been shown elsewhere (Manaresi 2015). However, it is still the case that younger firms grow faster. When we control for firm age, the estimated coefficient on size is always slightly larger (see Figure B.20 in Appendix B.4 for visual depiction). Furthermore, over the period 2008 - 2019, the unconditional average employment growth rates of younger firms are higher than that of older firms (see Figure B.21 in Appendix B.4).

These findings concur with studies in Italy (Manaresi 2015) and in the U.S. (Haltiwanger, Jarmin, and Miranda 2013), which find a positive relationship between firm size and firm growth, conditional on age.

2.3.3 Automation Technologies

Firms using automation technologies are larger, but adoption is less strongly associated with age. Firm technology adoption varies across the technologies, both in terms of sectors (e.g. Robotics is more common in manufacturing, while Cloud Computing is adopted in services) and firm characteristics (such as age and size). I estimate firm-

[2] Manaresi (2015) found that firm growth was weakly increasing in firm size, and young firms grew faster. My findings suggest stronger growth for larger firms (see Figure B.21), which is likely a feature of the survey; firms are dropped if they fall below the 20 worker size threshold. This is discussed further in Appendix B.1.

Table 2.1: *Binned Variables Firm Growth Regressions 2008 - 2019*

Dependent variable: *Employment Growth*					
Size	*50-99*	*100-199*	*200-499*	*500-999*	*1000+*
	0.014***	0.021***	0.022***	0.029***	0.028***
	(0.002)	(0.003)	(0.002)	(0.004)	(0.004)
Size (control for Age)	*50-99*	*100-199*	*200-499*	*500-999*	*1000+*
	0.014***	0.022***	0.023***	0.030***	0.029***
	(0.002)	(0.003)	(0.002)	(0.004)	(0.004)
Age	*10-20*	*20-30*	*30-40*	*40-60*	*60+*
	−0.002	−0.006	−0.008	−0.011*	−0.018**
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Age (control for Size)	*10-20*	*20-30*	*30-40*	*40-60*	*60+*
	−0.002	−0.006	−0.007	−0.011*	−0.021***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)

*Note: Coefficients represent estimated growth rates relative to omitted bin (for firm size: 20 - 49 employees; for firm age: 0 - 10 years). Regression includes year fixed effects. N = 38,702. Estimates are significant at levels of 1%: ***, 5%: **, 10% *.*

level productivity differences between automaters and non-automaters, finding that firms using these technologies are more productive.

Automation technologies are adopted by around one third of firms, but these businesses employ a significantly larger proportion of workers. The most-used advanced technology is Cloud Computing (21% of firms in 2017, and 28% in 2019). The other technologies are adopted by a smaller share of firms: the Internet of Things (13 - 16%), 3D Printing (5 - 11%), Industrial Robots (7 - 15%), AI and Big Data (9 - 13%).[3]

Overall, firms adopting *any* of these automation technologies take up a much larger share of employment: in 2015, 30% are adopters, employing 43% of workers. This rises to 33% of firms and 51% of employees in 2017, followed by 37% of businesses and 56% of labour in 2019. The graphs in Figure 2.1 split this relationship by each technology in 2019. The light blue bars show the share of firms adopting each technology, while the dark blue bars plot the share of employment in adopting businesses. Clearly the employment shares exceed the firm shares, generally by a factor of between two and three.

[3]The raw and weighted proportions of firms using these technologies is in Appendix B.2.

Further results on the adoption of automation technologies across the age and size distribution are contained in Appendix B.5, and show that the proportion of firms adopting these technologies is increasing in firm size, but not in firm age.

Figure 2.1: Technology Adopters: Proportion of Firms and Employment

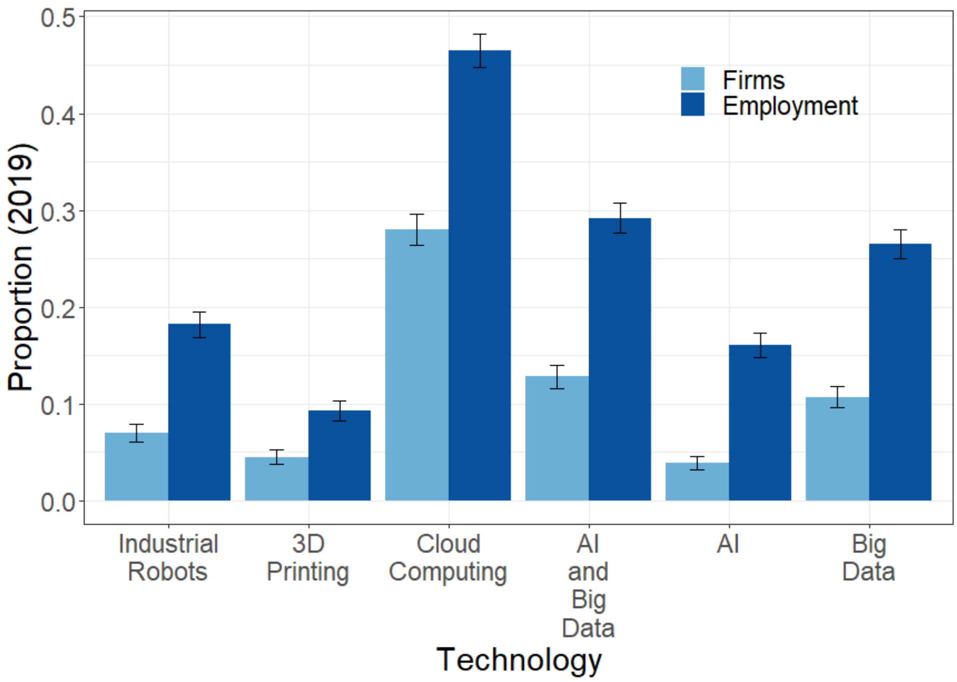

As an example, firms using AI & Big Data employ almost 63 workers on average in 2015, compared to just under 49 for non-adopters. Furthermore, despite being larger, firms using this automation technology also paid more *per* worker: such firms renumerated workers with over €31,000, compared to under €29,000 for businesses not using this technology.

The adoption of automation technologies varies widely across industries. The 'Digital' technologies (Cloud Computing, AI, Big Data, Internet of Things) are most common in Transport & Communication, Real Estate, and Wholesale & Retail. For example, Cloud Computing is most commonly found in Real Estate businesses, having been adopted by almost a third of firms in 2017. Within 'Digital' tools, there are significant sectoral differences: the Internet of Things is heavily adopted in Metal Manufacturing but not Hospitality, while this adoption pattern is somewhat reversed for the other three technologies. It is not surprising that 'Physical' technologies (Industrial Robotics and 3D Printing) are concentrated in Manufacturing. For example, Robotics are adopted in over 20% of firms in Metal Manufacturing, and over 10% of businesses in Chemicals, Rubber & Plastics. Robots are least commonly found in Services (Hotels, Restaurants, Wholesale, Retail). Similar numbers are obtained for adoption of 3D Printing. All sectoral results can be found in Appendix B.3.

Regressions of technology adoption on log employment and age show that employment is always strongly positively associated, while age is not. The estimated coefficients on log employment range from 0.03 - 0.08, and are always significant at the 1% level or below (see Tables B.18 to B.20 in the Appendix). Thus a percentage rise in employment raises the likelihood of adoption in the 3 - 8% range. These estimated coefficients barely change with the inclusion of firm age, which is, generally, small and insignificant. It is only for AI & Big Data (negatively) and Industrial Robotics (positively) that there is evidence age is associated with adoption of technology.

Firms that automate tend to adopt more than one of these advanced technologies (see Tables B.11 and B.12 in the Appendix). Firms adopting AI overwhelmingly use Big Data in 2017 (90%). These two technologies go hand-in-hand; the benefits of AI and Big Data depend significantly on use of the other. Likewise, there is significant adoption of both of the more industrial technologies: Robots and 3D Printing (84% and 86% in 2015 and 2017). This suggests a distinction between firms that automate, by adopting *any* of these technologies, and those that do not.

I find that adoption of automation technologies is reversible. A non-negligible proportion of firms that use advanced technologies in time t no longer use them in time $t+h$. Figure B.33 in the Appendix highlights this result. For example, around 15% of firms that use AI & Big Data *at some point* between 2015 and 2019 go from 'adopter' to 'non-adopter' in that time frame. A similar share go in the other direction, while less than 3% use this technology throughout. [4]

Further details on the heterogeneity in adoption across the various automation technologies can be found in Appendix B.3.

[4]It is possible that this finding is partially affected by measurement error. If different respondents fill in the survey, and are not fully aware of technology adoption within the company, I could spuriously count this as 'dropping' a technology.

2.3.4 Automation Investment

The share of total investment allocated to advanced technologies increases systematically with size, but not unconditionally with age. Figure 2.2 highlights the size relationship. Firms with under 50 employees allocate less than 0.8% of total investment towards automation technologies, while the share is over twice that for large businesses with over 1000 workers. Further graphs on automation investment shares in Appendix B.5.

Figure 2.2: Investment Share on Advanced Technology Rises with Firm Size.

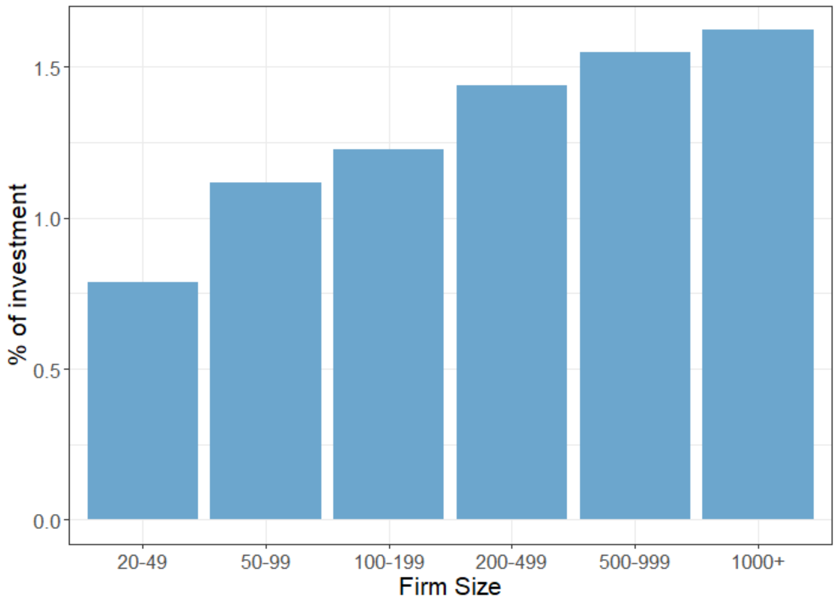

Note: Each bar shows the average share of investment spent on advanced technologies in each size bucket, averaged over 2016 - 2019.

I regress the share of investment in advanced technology on firm size and age, with sector and region fixed-effects, in each year 2016 - 2019 separately. The results are contained in Table B.17 in the Appendix.

The results highlight that larger firms invest a greater proportion in these technologies. Across each year, the estimated coefficient on firm size is large in magnitude and statistically significant, while age is not. Larger firms invest a greater share in automation technology, even controlling for age, sector, and region, with an increase of 0.25 - 0.30% for each percentage rise in employment.[5]

[5]The regression specification of $Y = \beta \ln X + \epsilon$ gives $\beta = \frac{dY}{\%\Delta X}$ which, in this case, is the change in the

2.4 Empirical Findings

This section focuses on automating firms, investigating their productivity and growth rates, relative to firms that do not automate. I show that firms that automate have productivity around 3% higher than non-adopters, although there is variation across the technologies. It is also the case that automating firms grow at faster rates, in the range of 0.6%.

These findings help build a picture of firm behaviour with regards to automation, and inform the model. However, they also present challenges in interpreting the causal relationship: does automation boost firm size, or are larger firms more able and willing to invest in such technologies? I use three approaches to answer this question: propensity score matching, two-way fixed effects models, and event studies. Each method provides strong evidence that automating firms grow faster than non-adopters. Put simply, firms that automate grow faster *when* they do so.

2.4.1 Productivity of Automating Firms

The relationship between productivity and technology use is an important question to examine. Both the correlation and the direction of causality can be informative about firm behaviour. To estimate firm-level productivity, I follow the literature on control function estimation (Olley and Pakes 1996). Full details of the methodology (including constructing firm-level capital stock), alongside the results, are included in Appendix B.10. I find that firms that adopt automation technologies are more productive than those that do not.

The results for estimated TFP across firms that do and don't adopt automation technologies are shown in Figure 2.3. For firms adopting *any* technology, they are around 3% more productive than firms which do not, on average. For the 'Physical' automation technologies - Robotics and 3D Printing - adopters are only slightly more productive, at under 1% and 0.1% respectively. The productivity of firms using

share of investment in advanced technologies with a percentage change in employment.

Figure 2.3: Adopters have higher TFP than Non-Adopters, 2010 - 2018

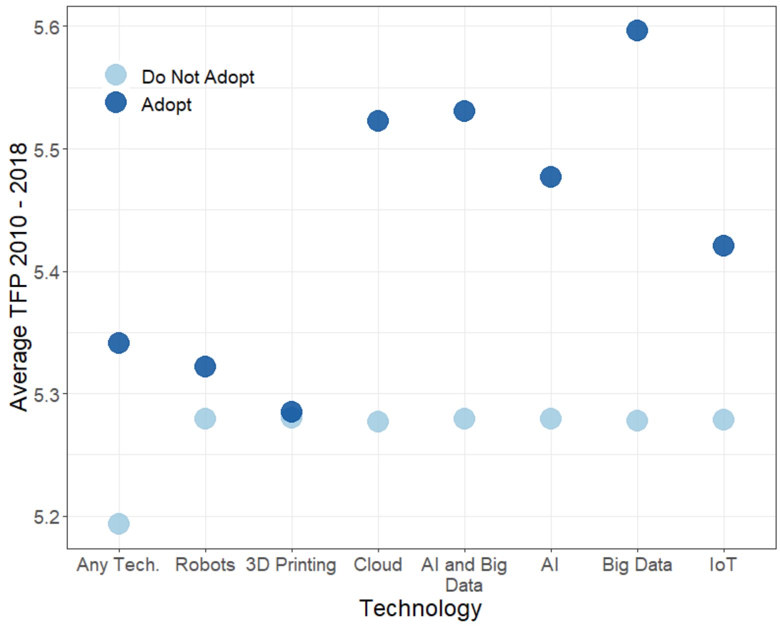

Note: Each point shows the employment-weighted average estimated TFP for firms that do or do not automate, over the time period 2010 - 2018. The TFP estimation procedure follows standard control function approaches.

'Digital' technologies such as Artificial Intelligence is much higher, relative to non-adopters: it sits between 2.7 - 6.1%.

2.4.2 Automation & Firm Growth

I leverage the panel to label firms as technology adopters, and non-adopters, from 2010 - 2018. Firms that use advanced technologies grow consistently faster than those that do not, as seen in Table B.16 in the Appendix. The difference in growth rates for adopters and non-adopters is compared across firm size and age distributions. I find that the average growth rates of automating firms are higher than that of non-adopters, across the age distribution (see Figure B.22 in the Appendix). This finding also holds across the size distribution for small- and medium-sized firms (Figure B.23 in the Appendix) and when I separate results across technologies (Figure B.24 in the Appendix).

To get a sense of the magnitude of these growth differences, I regress firm growth on size, age, year, and an indicator for whether the firm ever adopts *any* automation technology, from 2010 - 2018. The resulting estimated coefficient represents the extra growth associated with the adoption of automation technology. Table 2.2 contains the results that indicate firm growth is around 0.6 - 0.7% higher for automating firms. Given that the difference in growth rates for *all* firms with 1000+ employees compared to those with 20 - 49 is around 3%, it is clear that automation technology is associated with significantly higher firm growth, akin to being a substantially larger firm.

Table 2.2: Estimates of technology adoption coefficient from firm growth regressions.

Coefficient	0.0068**	0.0065**	0.0065**	0.0062**
SE	(0.0031)	(0.0031)	(0.0031)	(0.0031)
Age		✓		✓
Year			✓	✓
N	25,093	25,076	25,093	25,073

Note: Robust standard errors clustered at firm level. Coefficients labelled by statistical significance at: *** 1%, ** 5%, * 10%.

2.4.3 What Happens When Firms Automate?

It is clear that automation is more common in larger firms, and that they grow faster. Therefore, identifying the *impact* of automation on firm size is tricky, as we face an endogeneity problem: the sample of firms using advanced technologies seem fundamentally different pre-adoption to those that never adopt.

Matching Automaters and Non-Automaters

The first method to account for this problem is propensity-score matching (PSM). Propensity scores for the probability of adopting automation technology are computed using a logit with a wide range of potentially useful explanatory variables (see Section B.7 of the Appendix). Then various methods can be used to 'match' automation adopters to non-adopters, based on the propensity scores. I use the 'matched' set of firms to compare those using technologies to similar non-adopters. I regress log of firm employment on technology adoption for these matched sets. This provides evidence that technology adopters are significantly larger than similar non-adopters, as seen in Table 2.3.

Table 2.3: Propensity Score Matching Regression Results, 2015

Dependent variable: *Log Employment*						
	Any Tech.	*Cloud*	*AI & Big Data*	*IoT*	*Robotics*	*3D Printing*
Tech. Adoption (nearest)	0.461*** (0.06)	0.822*** (0.07)	0.623*** (0.11)	0.475*** (0.08)	0.370*** (0.10)	0.330** (0.11)
N	1914	1376	674	1042	720	524
Tech. Adoption (full)	0.586*** (0.05)	0.400*** (0.06)	0.818*** (0.07)	0.583*** (0.06)	0.535*** (0.07)	0.537*** (0.08)
N	2554	2580	2547	2541	2544	2538
Tech. Adoption (optimal)	0.493*** (0.06)	0.046 (0.06)	0.823*** (0.11)	0.680*** (0.08)	0.648*** (0.09)	0.606*** (0.11)
N	1914	1376	674	1042	720	524

*Note: Estimates are significant at levels of 0.1%: ***, 1%: **, 5% *. Coefficients are for log employment regressed on (binary) technology adoption, with matching described in brackets: "nearest" performs greedy nearest neighbour matching; "full" performs optimal full matching, so treated and control observations are assigned to a class and each receives at least one a match; "optimal" is similar to "nearest", but aims to minimise the mean of the absolute pair distances in the matched sample.*

Fixed Effects

I leverage the large set of firms over the panel to analyse the relationship between technology adoption and firm size. There are around 3,500 firms that can be labelled as adopters or non-adopters of advanced technologies. Compiling a panel of these firms over the period 2010 - 2018 gives around 25,000 firm-year observations. I investigate how technology adopters differ over time, and how firms change before and after adoption.[6]

I employ two-way fixed-effects (TWFE) models, to estimate the difference in firm outcomes between adopters and non-adopters. This model has unobserved individual- and time-specific effects, which is important in this context. Firstly, some firms may be inherently more productive and have greater growth potential, for a host of unobserv-

[6] Given that there is potentially non-random non-response to the questions on automation adoption, I use the class non-response adjustment factor to re-weight my panel as in Dargatz and Hill (1996). For further information on the extent of non-randomness of non-responders, and the method used to re-weight the panel, refer to Section B.6 of the Appendix.

able reasons such as aggregate conditions at entry (Sedláček and Sterk 2017), ex-ante heterogeneity (Sterk, Sedláček, and Pugsley 2021), or managerial practices (Bloom, Sadun, and Reenen 2016). Secondly, firm outcomes are likely to be tightly linked to the prevailing aggregate conditions of the macroeconomy and labour market.

However, the TWFE model relies on the underlying assumption of linear additivity of the two unobserved confounders. Crucially, this leads to 'treated' units being compared to observations described as 'mismatches' in the literature (Imai and I. S. Kim 2019). The intuition here is that we *should* be evaluating the causal treatment effect by comparing each 'treated' firm to an average of control units from the same firm (within-unit), plus the average of control firms from the same year (within-time), adjusting for the average outcome across these two control groups. However, in constructing these control sets, we may match a firm to one with the same treatment status (e.g. has also adopted automation technology). Imai and I. S. Kim (2019) show it is impossible to avoid this issue, even with a weighted TWFE estimator which attempts to eliminate mismatches in the within-unit and within-time sets.

Therefore, I will also employ an event study design to show how firm outcomes change before and after adoption. This further allows me to look at the differential effects of adopting automation technology over time, and permits testing of selection by observing pre-trends. The results from the TWFE models and event studies both show broadly the same results, both in direction, and magnitude.

The TWFE model is deployed to investigate the difference in employment, blue-collar employment, and labour productivity (proxied by turnover per worker) between firms that do and do not automate. These variables are regressed on firm and time fixed-effects, along with a control for firm age, and the variable of interest: a binary technology adoption variable (i.e. does firm i ever adopt the technology?).

$$\ln Y_{it} = \mu_i + \gamma_t + \delta X_{it} + \beta \mathbb{1}\text{Tech}_{it} + \epsilon_{it} \tag{2.1}$$

The set of β are shown in Table 2.4.

The results in Table 2.4 provide evidence that firm size rises upon adoption of

Table 2.4: Estimated β from homogeneous effect TWFE model: the % change in variables when adopting technology, relative to non-adopters.

		Cloud	*AI & Big Data*	*IoT*	*Robotics*	*3D Printing*
Employment	Coeff	0.020^{***}	0.052^{***}	0.051^{***}	0.042^{***}	0.056^{***}
	SE	(0.0043)	(0.0061)	(0.0048)	(0.0062)	(0.0066)
Blue-collar Emp.	Coeff	-0.036^{*}	-0.030	0.0008	0.048	-0.025
	SE	(0.015)	(0.027)	(0.021)	(0.027)	(0.028)
Turnover per worker	Coeff	0.0057	-0.017	0.017^{*}	0.065^{***}	0.019
	SE	(0.0066)	(0.0096)	(0.0075)	(0.0097)	(0.010)

Note: Robust standard errors clustered at firm level. Coefficients labelled by statistical significance at: *** 0.1%, ** 1%, * 5%.

each of these advanced technologies, in the range of 2 - 6 %. Although the point estimates for the effect of adoption on blue-collar employment are negative for Cloud Computing, AI & Big Data, and 3D Printing, the standard errors are large enough for the latter two that I cannot interpret these as statistically significant. Likewise, there is little evidence that using these advanced technologies raises turnover per worker, apart for Industrial Robotics and IoT. For robots, the estimate is large (at 6.5%) and statistically significant. The result for IoT is smaller (at 1.7%) and only holds at a 5% level of significance.

Further TWFE models show similar patterns: a range of results can be found in Appendix B.8, where I estimate coefficients by year, and provide instrumental variable results.

Event Studies

Knowledge of *when* firms adopt technologies also permits an event study specification in which the firm outcomes are regressed on a set of dummy variables which indicate the time relative to adoption year. The specification takes the form:

$$\ln Y_{it} = \mu_i + \gamma_t + \delta X_{it} + \sum_{j=\underline{j}, j\neq -1}^{\bar{j}} \beta_j \mathbb{1}(D_{it} = j) + \epsilon_{it}$$

where $D_{it} = t - A_i$ is the 'relative time', or the number of periods relative to when firm i adopted technology in year A_i (Borusyak and Jaravel 2017). The outcomes of interest Y_{it} are employment, hours, wages, and turnover of firm i in year t. The regression includes both firm and year fixed effects. The time-invariant controls are sector and region, and the time-varying control is firm age.

I find some evidence of pre-trends from this event study specification. Results from this exercise are presented in Figures B.35 and B.36 of the Appendix. Where these pre-adoption estimates are negative and statistically significant, this is suggestive that there is a general pre-adoption trend in firm outcomes. The linear pre-trends in log employment, log hours, log wages and log turnover[7] are noticeable, especially for Industrial Robotics and 3D Printing.

Therefore, I follow Borusyak and Jaravel (2017) and deal with this by identifying the deviation from a linear pre-trend. This is done by dropping two 'reference' periods, rather than just one. I drop the earliest period (6 years before adoption) and continue to exclude the period prior to adoption. The results for Artificial Intelligence and Big Data are presented in Figure 2.4. I focus on one technology for ease of exposition. For all technologies, the results can be found in Figures B.37 and B.38 in the Appendix. Although the linear pre-trends are still somewhat apparent, they are shrunk towards zero and mostly not statistically significant.

The results from the event studies highlight a clear sustained rise in employment in the years following adoption of these automation technologies, both at the extensive and intensive margins. There is evidence that turnover also rises on adoption, but not for Cloud Computing, nor AI & Big Data. The results for wages have much larger standard errors, showing no clear pattern, and no statistically significant effects across all technologies.

The boost to employment seems most pronounced for AI & Big Data, 3D Printing, and the Internet of Things. This is consistent with the previous results which

[7] I actually use the Inverse Hyperbolic Sine (IHS) transformation (Bellemare and Wichman 2020) for turnover, rather than log, due to a small number of zeros in the data. But unlike log, the IHS approximation is defined at zero, and the coefficients can be interpreted analogously.

Figure 2.4: *Estimates from Event Study Regression on AI and Big Data*

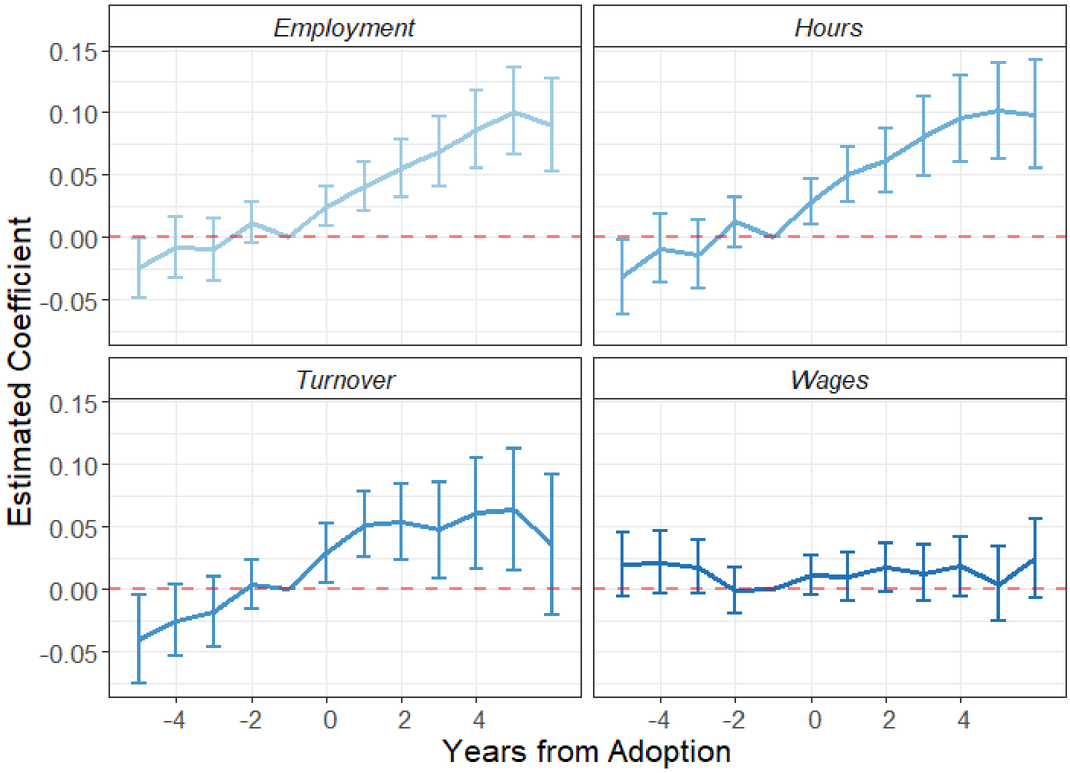

Note: Regression with 3,270 firms and 24,544 firm-year observations. Standard errors clustered at the firm level.

considered the impact of adopting automation technologies using TWFE models.

More specifically, most of the point estimates prior to the reference year are insignificant across all technologies, apart from 5 years prior to adoption in some cases. This arises from the issue of linear pre-trends, as previously discussed. There is some mild evidence of a pre-trend in employment for adopters of 3D Printing. Overall I am reasonably confident that the event study identifies how firm outcomes change with technology adoption. However, the existence of a pre-trend implies that my post-adoption point estimates for 3D Printing are biased *downwards*. The interpretation of such pre-trends is that adopters of 3D Printing grow more slowly than non-adopters *prior to adoption*. Thus, the impact of investing in 3D Printing would actually be larger than my results suggest.

The results show a statistically significant relationship between technology adoption and firm size, in the range of 1 - 3 % in the first two years post-adoption, and rising to 4 - 11% in the final two years of the sample. These effects are statistically significant across all technologies, and the magnitudes are not trivial. Furthermore, these results hold at the intensive margin, with a rise in hours of 2 - 4 % in the first two years, which increases to 4 - 12 % in the fifth and sixth years post-adoption. Overall, the employment effects are most muted - but still positive - for Mobile & Cloud Computing.

It seems that firms using automation technologies also experience a large and significant rise in turnover: in the first two years, there are increases of 2 - 8 %, with smaller gains for firms adopting the Digital technologies (Cloud, AI etc.) and bigger returns for those using Physical technologies (Robots and 3D Printers). After five years, the effect on turnover is in the range of 8 - 12 % for IoT and Physical technologies. The medium-run effect for AI & Big Data is around 6%, and although there are large standard errors, the estimated effect is still generally positive. However, for Mobile & Cloud Computing, the estimated effect on turnover is not distinguishable from zero.

Finally, firms' average wage doesn't seem to change significantly after implementing automation technologies. For Physical technologies, it seems that adopters' wages have essentially no pre-trend - coefficients are close to zero and insignificant - but the point estimates *are* positive after adoption. Sometimes these are statistically significant, especially for firms using 3D Printing technology. Overall, though, the evidence is not especially strong that automating firms' wages rise after adoption. This may not be entirely surprising: firms may grow through using new technologies, but the average wage will only grow to the extent the composition of workers adjusts. If firms that start using Machine Learning algorithms already had a higher-than-average share of skilled technical employees, and expand by continuing to hire in that ratio, the average wage may not rise.

2.4.4 Empirical Takeaways

I have highlighted important novel findings on firms that automate. Firstly, I find that firms adopting automation technologies are already larger and grow faster than non-adopters. Secondly, firms grow even faster *after* they adopt these technologies. Taken together, it seems that automation increases 'inequality' across firms. This could have important consequences for competition among firms, such as increased up-or-out dynamics, rising market power, or 'superstar firm' effects.

These facts motivate further investigation with a model, to help understand the effects of automation on firm dynamics and macroeconomic outcomes.

2.5 Model

I introduce a model of heterogeneous firms that can endogenously choose to automate. The aim of this model is twofold. Firstly, it allows me to analyse the impact of automation when firms differ and have the option to automate tasks, taking into account the trade-off between marginal and fixed costs from new technologies (De Ridder 2019; Lashkari, Bauer, and Boussard 2019). Automation lowers marginal costs, making firms more productive, but there is a selection effect, as only *more productive firms* can afford the associated fixed costs. The impact on overall productivity is ex-ante unclear, and depends on both selection and reallocation. The effect on employment is also not obvious, as it depends on which firms expand, contract, enter, and exit. Secondly, I am interested in the partial equilibrium effects. Automation influences the equilibrium price by affecting input costs. This will impact firm entry and exit. Furthermore, changes in demand for labour caused by automation technologies will lead to a shift down the labour supply curve, with wage and employment effects.

I take a standard Hopenhayn (1992) heterogeneous firm dynamics model with adjustment costs on labour. I extend it to include a task-based production function (Zeira 1998; Acemoglu and Restrepo 2016), routine and non-routine labour which produce different sets of tasks, and automation technology which can replace routine workers.

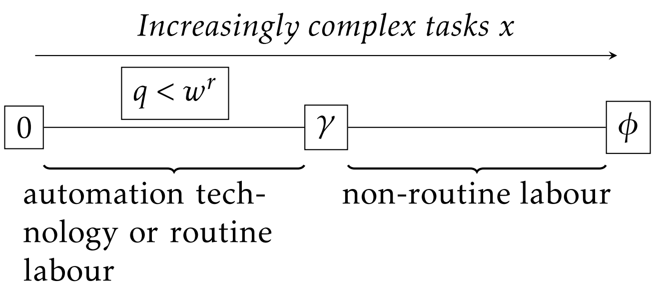

Figure 2.5: Structure of Task-Based Production Function

2.5.1 Task-based Production Function with Automation

In the standard Hopenhayn (1992) model, firms are heterogeneous in productivity z and output is determined by one labour input n with decreasing returns to scale. The production function is $y = zn^{\alpha}$.

Taking inspiration from Acemoglu and Restrepo (2016), the new production function depends on productivity and production over a set of tasks x of increasing 'difficulty'. This allows for jobs that are made up of a variety of tasks, which might be performed by different inputs depending on their complexity. Thus firm output is:

$$\ln y = \ln z + \int_0^{\phi} \ln y(x) dx \quad \text{where } \phi < 1 \text{ for DRTS} \tag{2.2}$$

The production of each task $y(x)$ is determined by the set of tasks which can be performed by different inputs. I assume that routine labour n^r can be perfectly substituted by automation technology R, such as industrial robots, for some set of tasks. However, non-routine labour n^n cannot be substituted for, among the more complex tasks. This is depicted in Figure 2.5, where $q < w^r$ highlights that, should a firm have access to automation technology, the marginal cost of technology is less than the routine labour wage.

Of course these sets of tasks will change over time, as advanced technology improves. Production of a task is determined:[8]

$$y(x) = \begin{cases} r(x) = R(x) + n^r(x) & \text{for } x \in [0, \gamma) \\ n^n(x) & \text{for } x \in [\gamma, \phi) \end{cases} \tag{2.3}$$

The prices of inputs are w^n, w^r, q for non-routine labour, routine labour, and automation technology. This price q is the per-unit cost of using the technology, such as the electricity or computing power associated. I assume that $\frac{q}{w_r} < 1$ such that automating is beneficial. However firms must pay a fixed cost c_a in each period to use automation technology R. This can be thought of as the cost of paying a company to maintain the robots, or an annual payment to access cloud computing services. This does not scale with the amount of technology used, as does the technology input cost q. The production function is therefore:

$$\ln y = \ln z + (\phi - \gamma) \ln n^n + \gamma \ln r \tag{2.4}$$

$$y = z(n^n)^{\alpha} r^{\gamma}, \tag{2.5}$$

where $r = n^r + R$, and $\alpha = \phi - \gamma$. This yields a production function that has capital-skill complementarity, which is an idea going back to at least Griliches (1969).

Heterogeneous Firm Model

Firms produce with decreasing returns to scale and choose inputs n^n, n^r, R to maximise profits (dropping the subscript t for ease of notation):

$$\pi(z, n_{-1}) = pz(n^n)^{\alpha}(n^r + R)^{\gamma} - w^n n^n - w^r n^r - qR - g(n, n_{-1}) - c_f \tag{2.6}$$

where $n = n^n + n^r$ is the total labour input, n_{-1} is the previous period choice of labour inputs and $g(.,.)$ is an adjustment cost function. Firms differ in productivity z,

[8]Choosing to split labour inputs into routine and non-routine is inspired by a rich literature on the substitutability of automation-related technologies with labour inputs. For example, Jaimovich and Siu (2012) look at the changes in employment shares by occupation group in the US, and find this distinction to be important.

which follows AR(1) process:

$$\ln z_{t+1} = \ln \bar{z} + \rho \ln z_t + \sigma_z \tag{2.7}$$

The value function of a firm is the largest of the 'automating' and 'non-automating' value functions, taking into account the fixed automation cost c_a. These value functions, and the overall value function, are expressed below:

$$v_t^a(z_t, n_{t-1}) = \max_{R_t, n_t^n, n_t^r \geq 0} \left\{ \pi(z, n_{-1}) + \beta \max \left\{ \int v_{t+1}(z_{t+1}, n_t) dF(z_{t+1}|z_t), -g(0, n_t) \right\} \right\} \tag{2.8}$$

$$v_t(z_t, n_{t-1}) = \max_{n_t^r, n_t^n \geq 0, R=0} \left\{ \pi(z, n_{-1}) + \beta \max \left\{ \int v_{t+1}^a(z_{t+1}, n_t) dF(z_{t+1}|z_t), -g(0, n_t) \right\} \right\} \tag{2.9}$$

$$\tilde{v}(z_t, n_{t-1}) = \max \{ v_t^a(z_t, n_{t-1}) - c_a, v_t(z_t, n_{t-1}) \} \tag{2.10}$$

As in Hopenhayn (1992), there is a cut-off level of productivity z^* such that all firms with $z \geq z^*$ do not exit. Furthermore, there will be a cut-off level of productivity z^a such that all firms with $z \geq z^a$ will automate. The equilibrium price is pinned down by the free-entry condition:

$$v^e(z) = \int_z v(z) dG(z) = c_e \tag{2.11}$$

Finally the law of motion for the distribution of firms can be computed:

$$\mu_{t+1}(z') = \underbrace{\int_{z \geq z^*} \mu_t(z) dF(z'|z)}_{\text{Surviving incumbents}} + \underbrace{MG(z')}_{\text{Entrants}} \tag{2.12}$$

where $\mu_t(z)$ is the mass of firms with productivity z at time t, $M > 0$ is a constant mass of potential entrants, and $G(z)$ is the productivity distribution for entrants, which is simply the stationary distribution of the AR(1) process. I solve for M in stationary equilibrium by exploiting the linearity of $\mu_t(z)$, as in Hopenhayn (1992). Effectively, the mass of entrants is chosen so that equation (2.12) holds, and there is no (net) entry nor exit.

To close the model, I assume downwards-sloping demand for final output $D(p) = \frac{\bar{D}}{p}$, where $\bar{D} > 0$ is exogenous. Price is endogenous, and satisfies the free-entry condition in equation (2.11). In equilibrium, output is equal to demand, so $Y = \int y(z)\mu(z)dF(z) = D(p)$. This equilibrium condition shows that adjusting $\bar{D}$ simply scales the values of all equilibrium objects up or down. $\bar{D}$ is chosen to clear the labour market. The labour supply curve is upwards-sloping with Frisch elasticity λ, following Clementi and Palazzo (2016): $L^s(w^n) = (w^n)^\lambda$, where w^n is the non-routine labour wage (and the routine labour wage is set as the numeraire).[9]

This model is partial equilibrium, so the labour market is not endogenised from optimisation on the household side. However, an upwards-sloping labour supply curve allows variation in labour demand to lead to changes in the equilibrium wage and employment. If instead a Melitz (2003) framework were chosen, where labour supply is fixed, any changes to labour demand would not affect aggregate employment. This would limit the model from speaking to aggregate employment effects, as all the adjustment would occur through the wage channel.

To summarise, the equilibrium conditions are as follows,

$$D(p) = \frac{\bar{D}}{p}, \quad Y = \int y(z)\mu(z)dF(z), \quad L^s(w^n) = (w^n)^\lambda, \quad L^d = \int n(z)\mu(z)dF(z) \qquad (2.13)$$

[9]As an example, raising $\bar{D}$ shifts the mass of firms across the productivity distribution up in equal proportion. Labour demand is firm-level demand multiplied by the mass of firms, integrated over the productivity distribution: $\int n(z)\mu(z)dF(z)$. This is clearly impacted by $\bar{D}$, which scales the distribution. In equilibrium, labour demand equates to labour supply, which is the non-routine wage to the power of the Frisch elasticity. Given the calibration sets the non-routine wage at 1.2 times the numeraire routine wage, I choose $\bar{D}$ such that the labour market equilibrium holds in the non-automation calibration. When there are changes to labour demand, for example induced by automation, this will lead to changes in the real wage and employment.

2.5.2 Solution Method

To solve this problem, I take first-order conditions (FOCs) of the static profit maximisation problem without labour adjustment costs, for firms that automate, and those that do not.[10] The presence of the labour adjustment cost prevents a closed-form solution for optimal input demands, yielding two potential solution approaches. One is to solve these numerically, but this was computationally intensive, which slowed down the calibration process. The second is to approximate lagged employment across a grid,[11] creating a labour cost adjustment matrix for all possible combinations between time periods. A brief outline of the solution algorithm is presented here:

1. Solve optimal input demands across grid of productivity and lagged employment.

2. Use the free-entry condition to find equilibrium price, both with and without the presence of automation technology.

   (a) Guess price p_0

   (b) Take FOCs of static firm problem to find maximised profit $\pi(p_0; z, n_{-1})$.

   (c) Apply value function iteration on this starting point until convergence.

   (d) Compute expected value of entry, and check free-entry condition holds.

   (e) Adjust guess of price until free-entry condition holds.

3. Find cutoff productivity level for entry, and for automation.

4. Compute firm choices and distributions.

2.5.3 Calibration

The chosen labour adjustment cost function is quadratic, such that $g(n_t, n_{t-1}) = \frac{\tau}{2}(n_t - n_{t-1})^2$ where n_t is the sum of routine and non-routine labour hired in period t. The

[10] Dropping adjustment costs as a starting point allows for closed-form solutions for the static FOCs, as we ignore the dynamic element of the labour input choice.

[11] I ran the algorithm without automation technology, to give upper and lower bounds for total employment demand.

purpose of this adjustment cost is to better match the data, as justified by long-standing literature (Kydland and Prescott 1991; Hamermesh 1989; Cogley and Nason 1995). The labour adjustment cost is for *total* labour hired: firms can move workers from routine to non-routine tasks, and it is the hiring and firing of workers which incurs a cost.

Table 2.5: Model Parameters

	Parameter	Value	Target
pre-set			
β	Discount rate	0.95	Match annual IR over period
c_e	Entry cost	$0.82 \times c_f$	Barseghyan and DiCecio (2011)
$\frac{w_n}{w_r}$	Wage ratio	1.23[12]	Vannutelli, Scicchitano, and Biagetti (2021)
q	Automation price	0.975	By assumption
γ	Routine share	0.66 - α	Residual labour share
ρ	AR1 coefficient	0.8	
σ	AR1 stdev	0.2	
λ	Frisch elasticity	2.0	Clementi and Palazzo (2016)
calibrated			
α	Non-routine share	0.40	Match avg firm size
$\bar{z}$	AR1 intercept	2.06	Match avg labour productivity
c_f	Fixed cost	35.6	Match exit rate in Manaresi (2015)
c_a	Automation cost	6.23	Match % firms that automate
τ	Labour adjustment cost	20.0	Match avg growth rates
$\bar{D}$	Demand scalar	4.16	Clear labour market without automation

The pre-set calibration parameters are fairly standard, or taken from existing literature. The automation price is set below the routine wage by assumption. This somewhat arbitrary assumption is relatively innocuous; it means at least *some* firms will have an incentive to automate, and I calibrate the fixed automation cost to match the share of firms using the technology. Thus the decision to adopt automation technology is governed by the fixed automation cost, and the automation price is simply set such that some firms automate. It doesn't seem the results are particularly sensitive to this 'automation price'.[13] The calibrated parameters are chosen to target observables which are sensitive to that parameter. The non-routine share is chosen to match

[12] Difference in log wage for non-routine to routine workers is 0.21. This implies $\frac{w^n}{w^r} = e^{0.21} = 1.23$.

[13] Calibration with other values of q simply leads to different values of c_a to the target share of automated firms, and the other calibrated parameters do not change in any significant way.

the average firm size, as this will determine the labour share in highly productive, automating firms. The AR(1) intercept is chosen to match the average labour productivity, as it determines the level of the productivity distribution from which firms draw. The fixed and automation costs are naturally chosen to match the exit rate and the share of automating firms. The adjustment cost on labour is chosen to match the average growth rate of firms, as this affects the speed of readjustment towards the 'optimal' labour choice.

The calibration is achieved with a pattern search algorithm (Koziel and Yang 2011; R. Lewis and Torczon 2002; Madić and Radovanović 2014); an optimisation method that doesn't require computation of a gradient. It is exploratory in that it considers the objective function at a pattern of points, known as a mesh, to improve the objective function value (Madić and Radovanović 2014). The algorithm chooses the calibrated parameters to find the global minimum of the weighted average of the absolute distance between the model and targeted moments, where the weights are the inverse of the targets themselves.[14]

2.5.4 Results

I report a set of targeted and non-targeted moments in the model and data in Table 2.6. The model moments are either steady-state targets, or computed from a simulation of 200,000 firms over 20 periods, where automation is introduced and firms adjust their behaviour over time to this technological change.[15] When necessary, I choose the simulated moments instead of steady-state moments, where I need to track firms that do or do not automate.

Whenever possible, the moments are compared to the data from the Italian firm

[14]Such weights force the algorithm to be 'unit-agnostic', else it would fit larger targets better than smaller ones.

[15]After calibrating the model, I take the resulting parameters and simulate 200,000 firms over 20 time periods. This allows me to introduce automation to the model, and analyse the long-run effects as firms respond: by optimally choosing inputs each period, deciding whether to automate, and exiting the market. I can follow firms over time, identifying if and when they choose to automate, and the resulting impact. I shocks firms' productivity each period following the AR(1) process, and firms choose the optimal inputs on the productivity grid. They automate only if their productivity is equal or greater than z^a, and exit if it is below z^*. If they exit, the firm stays inactive. I do not allow firms to enter for computational ease; this does not change the results which are ratios and time-averages.

survey introduced in Section 2.4. However, firms may leave the panel if they become too small, which does not *necessarily* imply firm exit. Hence I use the exit rate in Manaresi (2015) which has the full sample of private Italian firms.

For the non-targeted moments, the routine share of employment of 43% is taken from an Italian study by Vannutelli, Scicchitano, and Biagetti (2021), while all other moments from the data are computed from the survey in this paper.

The model fit is good for some moments, but not all. The firms are a bit larger in the model, on average, and grow at a slightly higher rate. The share of firms that automate is a few percent below the data. However, the exit rate is about twice as high as the data, and average labour productivity is much lower than the data (although the units make this a less useful target).

The model does well at matching the non-targeted moments. The routine share of employment matches the data very well, while automating firms are slightly larger than the data in terms of employment (48% in model; 42% in data), but about right in output (53% in model; 55% in data). The relative growth rates, relative exit rates, and relative productivity for automating firms have the right sign in the model, but the magnitudes don't match the data.

Relative growth rates of automating firms are high in the model compared to the data. In the simulation, the difference is almost 5 p.p. greater growth for automaters, compared to firms that don't automate. In contrast, this difference is under 1 p.p. in the data. Automating firms exit less frequently than non-automaters in the model, as we see in the data. However, the difference in exit rates is over 17 p.p. in the data, compared to just 9 p.p. in the model. Finally, productivity of automating firms is 19% higher in the model, compared to just 3% in the data.

Another exercise to check the model describes the data well is to calibrate the non-automated model as a starting point. To this end, I set the nonroutine to routine wage ratio at rate consistent with the data 'pre-automation' in the 1980s, at 1.1 (Basso 2019). Then I back out the labour supplies implied by the model. Automation is introduced the model, and wages are computed to clear the labour market. The nonroutine to

Table 2.6: Model Results

	Model	Data
Non-targeted		
Routine share of employment	0.44	0.43[1]
Share of employment in automating firms	0.48	0.42*
Share of output in automating firms	0.53	0.55*
Difference in growth rates for automating firms (p.p.)	0.047†	0.007*
Difference in exit rates for automating firms (p.p.)	-0.089†	-0.176*
Relative productivity for automating firms (p.p)	0.19†	0.03*
Targeted		
Average firm size	62.3	50.0*
Exit rate	0.198	0.098[2]
Average labour productivity	61.79	135.5*[3]
% firms that automate	0.27	0.30*
Average growth rate	0.021†	0.015*

†: moments estimated off long-run averages from firm simulation exercise.
*: data from Italian firm panel described in Section 2.4.
[1]Vannutelli, Scicchitano, and Biagetti (2021)
[2]Manaresi (2015)
[3]Simply computed as average turnover per worker hour. In 2017, the mean turnover per worker was €221,894. The average weekly hours were 31.5, so over a year we get 1,638. Thus average turnover per worker hour is €135.47.

routine wage ratio the model gives is 1.285, which is slightly higher than the data, but in the right direction. Overall it seems that this baseline model does a reasonable job of matching the data.

2.5.5 What is the Impact of Automation?

I initially solve the model when automation technology unavailable to firms.[16] Then I solve the model with automation, and compare the outcomes. I document the aggregate and firm-level changes when automation technology becomes available in this partial equilibrium framework. Table 2.7 contains the results.

It is important to underscore that the magnitude of the effects presented here do depend on the elasticities of product demand and labour supply. Automation technology affects firm decisions. Aggregation across firms will shift the aggregate output supply curve and labour demand curve. The extent to which this affects partial equi-

[16]Imagine that initially the per-unit price of automation q_t is greater than that of routine labour, or that automation technology is too undeveloped to produce any tasks that humans perform.

librium outcomes depends on the aforementioned elasticities.

Automation technology permits highly-productive firms to grow, and pushes down the equilibrium output price via free-entry, reducing the number of operating firms. A fall in total labour demand leads to a shift down the upwards-sloping labour supply curve, leading to a fall in aggregate employment. The extent of the fall in employment and the real wage depend on the elasticity of labour supply.

Table 2.7: Percentage point change relative to 'No Automation' model

Aggregates:	**Employment**	**-2.49**
	Price	-0.02
	# firms	-8.51
	Output-weighted productivity	+1.34
	Exit rate	+0.10
	Real wage	-1.24
Firm Level:	**Employment per firm**	**+6.58**
	Output per worker	+0.16
	% firms that automate	+27.4

The key result in Table 2.7 is that automation leads to an aggregate fall in employment, alongside a firm-level rise in employment. These two results have been found in the data, and can be reconciled in this heterogeneous-firm framework.

When cheaper automation technology is available to firms, they purchase this input if they are productive enough. These firms expand due to the low-cost input, and earn significantly greater profits. The equilibrium effect is a lower price, as determined by the free-entry condition. This is because automating firms can produce output more cheaply, raising the value of producing output. The lower output price reduces returns to low-productivity firms, so they exit. Overall, the reallocation towards more-productive firms leads to a rise in output-weighted productivity.

At a firm level, the average firm hires more workers and produces more, due to the greater skew of output towards highly productive firms. Firms that automate are larger, as they are able to expand with the new technology. The rise in average firm size over time can be seen in the model simulation in Figure 2.6, which plots the average

Figure 2.6: Average firm-level employment from simulation when automation is introduced.

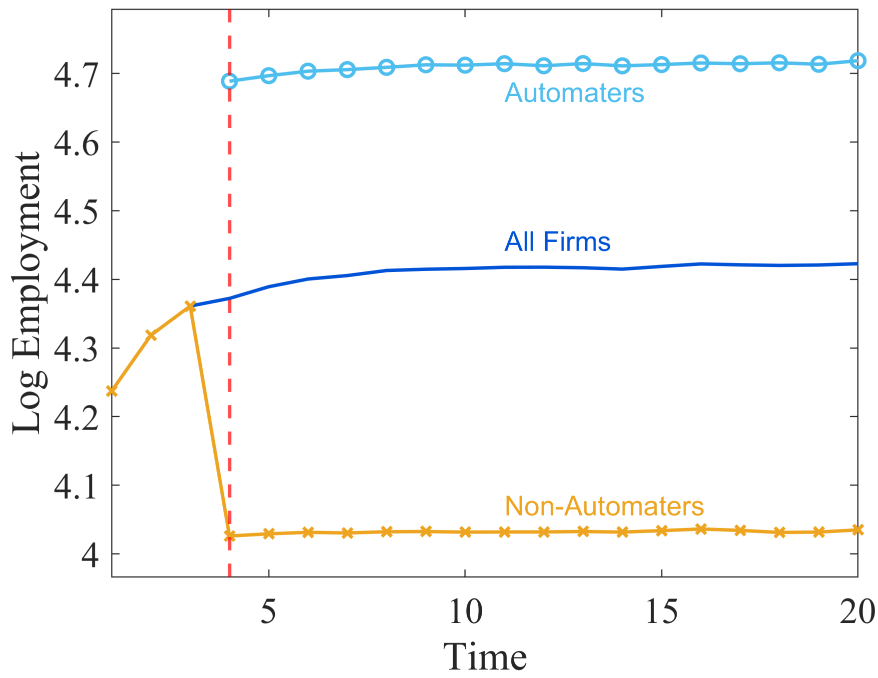

firm-level employment once automation technology is introduced.[17]

However, there is a fall in aggregate labour demand (firm-level employment multiplied by the mass of firms). This is due to the fall of the total mass of firms. As labour demand falls, the equilibrium shifts down the upwards-sloping labour supply curve. Thus, the introduction of automation technology in production leads to a fall in aggregate employment and the real wage. The fall in employment is skewed towards routine labour, which takes up a smaller share of total employment.

Figure 2.6 highlights an important dynamic uncovered in the empirical section. Firms that are already large and productive will invest in automation technology, and grow larger - this is why automaters start much larger than their technologically-lagging counterparts. This also explains why the average size of non-automaters falls: the largest of this group leave, upgrading their technology, which reduces the average size of non-automaters. In sum, the average size of all firms increases, due to the combination of the higher exit rate, and the reallocation of firms between the two groups.

I track firms that automate, before and after automation is introduced, and produce

[17] Automation technology is made available in the fourth period of the simulation, but the dynamics are very similar whenever it is introduced.

a within-model event study to compare to the data. This is presented in Figure 2.7, which shows:

- *Data:* event study regression on employment for **adopters of AI & Big Data**, with linearized pre-trends (Borusyak and Jaravel 2017). The filled in points are statistically significant at the 1% level.
- *Model:* log difference in employment for automating firms relative to non-automating firms, normalized to initial difference.

Figure 2.7: Model Event Study for Automating Firms

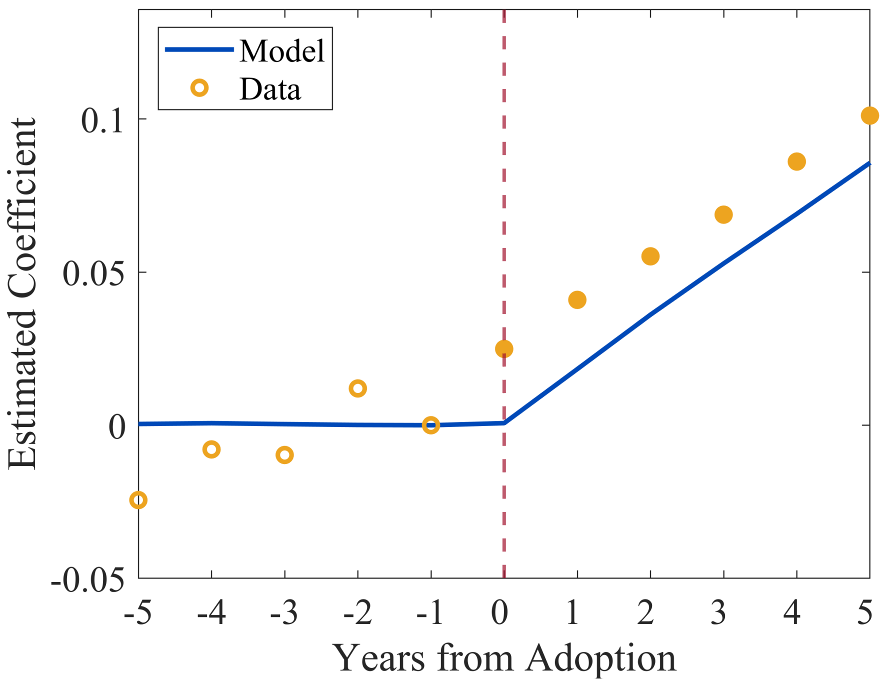

The model reproduces the results from the regression very well. As in the empirical exercise, automating firms are the same size as non-adopters pre-adoption (when controlling for a set of characteristics), but grow faster after implementing automation technologies. That the model accurately produces the shape and magnitude of the event study is promising, as this is an outcome of the mechanisms of the model, and not directly targeted. It is further evidence that the model can account for the findings from the data.

This model can also reconcile the aggregate findings that automation technology reduces aggregate employment (Acemoglu and Restrepo 2018; Dauth, Findeisen, Südekum, and Wößner 2017), with my firm-level evidence that automaters grow on adoption. This model highlights the important role for heterogeneity, reallocation, and equilibrium effects.

There are several important extensions to the model which are being worked into the next version of this paper. Firstly, I will introduce capital to the model, and use data on investment rates (correlations and inaction rates) to better pin down firm behaviour in the calibration. This should improve the model fit because firms adjust capital more slowly than labour, so this information reflects the dynamic response of firms to productivity shocks. Secondly, I intend to embed this model in a general equilibrium framework. I will introduce an endogenous labour supply, to allow the labour market channel to affect firm behaviour. Put simply, this will allow firms to flexibly adjust their optimal input choices to changing labour market conditions. This is especially important in the context of labour market polarisation, potentially induced by changes in technology and automation (Jaimovich and Siu 2012). Thirdly, I would like to introduce policy parameters to the model, to assess the impact of government interventions such as investment in Artificial Intelligence research, or a tax on automation technology. Finally, I want to quantitatively assess the balance of automation on directly boosting productivity, versus reallocating towards more-productive firms.

2.6 Conclusion

This paper seeks to understand the adoption of automation technologies among firms, and the equilibrium effects thereof. On the empirical side, I take a novel dataset on firm use of advanced technologies, and I show that automating firms tend to be larger, but also grow faster, both before and after they adopt. The employment growth doesn't show up in blue-collar workers, suggesting that complementarities between automation technologies and skilled workers are strong and important. Furthermore, adopters of new technologies are more productive, but I do not find that productivity is boosted by adoption.

Taking stock of these findings, I extend Hopenhayn (1992)'s seminal heterogeneous firm model to include two types of labour - routine and nonroutine - alongside automation capital which can be substituted for the former. In the model, automation is achieved only by highly productive firms, due to the fixed cost of using this technology. These firms benefit most from automation, allowing them to expand. The equilibrium effect lowers prices, forcing low-productivity firms out of the market. Unsurprisingly, workers lose out significantly, especially those in routine occupations. Importantly, such a model can reconcile the growing evidence that firms grow when they adopt advanced technologies, but the aggregate effect on employment is often negative.

Chapter 3 Scale Economies, Productivity, and Computer Software

Computers have changed the way firms produce. For example, information technology and software have affected firm cost structures, by introducing different fixed and marginal costs. This can affect output and revenue elasticities, which are closely linked to returns to scale. I estimate revenue elasticity with a large dataset covering firms of all sizes, and across all sectors in the UK economy. I document the relationship of software to the changing patterns of revenue elasticities. A theory is developed which highlights how technology endogenously affects scale economies, and thus allows firms to grow faster.

JEL classification: E32, E23, D21, D43, L13

Keywords: Returns to Scale, Production Functions, Software, Productivity, Market Structures.

3.1 Motivation

Scale economies describe the relationship between firm size and its costs of production. Increasing returns to scale reflect faster growth in outputs than inputs (Basu 2008), holding all else constant. Returns to scale is computed as the ratio of average to marginal costs. Thus it describes where a firm operates on its cost curves, and is indicative of the underlying market structures.

Returns to scale can be shown as the function of variables which describe industry competitiveness - markups and profit shares - or as the function of parameters underlying the production technology - the slope of marginal cost and size of fixed cost. Therefore, changes to market competition or production technologies will be reflected in returns to scale. In the long run, returns to scale will determine the capacity for firm growth, and which firms can survive. The former is important for how firms affect aggregate income, and the latter for the productivity distribution within industries. Thus, returns to scale is important for understanding market competition, production technology, long run growth and productivity distributions.

Factors which affect fixed costs or marginal costs will change scale economies. One such factor is technological change. There has been a proliferation in the adoption of intangible technologies such as Information Technology (IT) and computer software, which should boost the ability to turn inputs into output (Haskel and Westlake 2017). A crucial feature of such technologies is high fixed costs and low marginal costs (De Ridder 2019; Kariel 2021a). There is still debate on the macroeconomic impact of computers, especially on productivity growth (Acemoglu, Autor, Dorn, Hanson, and Price 2014). Indeed it has been suggested that, while computers are visible in almost every workplace, the impact on various macroeconomic indicators may be negligible (Solow 1987).

Returns to scale is the sum of output elasticities, which is equal to the markup multiplied by the sum of revenue elasticities. In this paper, I will compute revenue elasticities, and the markup is assumed to be constant due to a monopolistic competitive framework and constant elasticity of substition (CES) demand.

This paper focuses on the relationship between output and revenue elasticities, productivity and computer software. I describe the theory of returns to scale, and how it relates to measurable objects in the data. Specifically, I use methods to estimate the revenue elasticity, but cannot separately identify returns to scale and markups. I provide estimates on revenue elasticities across a services-dominated economy, and describe the industry heterogeneity. In addition, I introduce a simple theory to connect revenue elasticities and software, which affects both fixed and marginal costs. Finally, I take this theory to the data.

The data available only contains revenue for firms, not quantities. This presents a challenge to consistent estimation of the output elasticity of the flexible input which is required to compute markups by the De Loecker and Warzynski (2012) approach. Output prices may vary significantly across firms, but is unobserved. Hence it is necessary to assume that firms compete monopolistically with constant CES demand curves, such that markups are fixed across firms and over time (Bond, Hashemi, Kaplan, and Zoch 2021).

I follow Gao and Kehrig (2020) and estimate revenue elasticities using control function methods on production function estimation. I assume markups are fixed, due to CES demand and monopolistic competition. The production function parameters I estimate are revenue elasticities.

Estimating revenue elasticities is not straightforward, due to substantial issues of endogeneity and identification of production function parameters. I present estimates of revenue elasticities using Ackerberg, Caves, and Frazer (2015) and Gandhi, Navarro, and Rivers (2020). The production function is Cobb-Douglas, and the dependent variable is either value-added or gross-output, respectively for each estimation approach. I also present results with a translog production function, in order to obtain firm-level estimates.

I use firm-level data from the UK to (1) estimate revenue elasticities on a full set of representative firms, (2) investigate the drivers of changing revenue elasticities, specifically the relationship with adoption of software, (3) present a model which highlights the mechanisms relating scale, technology, and productivity.

Using ONS data, I estimate long-run trends of revenue elasticities in the UK over the last few decades. I obtain estimates at the industry-level, and the significant heterogeneity therein; the rise and plateau of firm-level productivity; the increase in computer software investments. Economy-wide revenue elasticity is close to unity, and has risen from 1998 - 2014. I present evidence that where more computer software has been adopted, revenue elasticities have risen more rapidly. I highlight two mechanisms which link revenue elasticity and productivity. Firstly, industries with higher scale require firms to have higher productivity in order to survive. Secondly, more productive firms are larger, which reduces their fixed cost share and this lowers revenue elasticity at the firm level. Combined, these two effects make the relationship between revenue elasticity and productivity more complicated than in standard models. I show that productivity and revenue elasticity have a strong negative relationship *within firms, over time*. On the other hand, *across industries*, I find a positive relationship. Reconciling these results with the aggregate trends provides a paradox for future research.

3.2 Literature Review

Returns to scale is typically an underlying parameter γ (or set of parameters) in production functions, and is assigned in line with the literature (Edmond, Midrigan, and Xu 2021), calibrated, directly estimated (Ruzic and Ho 2019), or assumed to be constant (Jaimovich and Floetotto 2008; De Ridder 2019; Burstein, Carvalho, and Grassi 2020).

Even if it is not explicitly discussed when presenting a model involving firms, returns to scale always plays a role in the dynamics. It is tightly linked to productivity, firm growth, and firm survival (Gao and Kehrig 2020). It is important when evaluating how firms respond to technology shocks or changes in policy (Basu and Fernald 1996;

Lashkari, Bauer, and Boussard 2019; Baqaee and Farhi 2020). Unsurprisingly it also a driver of trends in aggregate productivity, as the spread of returns to scale across heterogeneous firms will affect firm-size distributions (Ruzic and Ho 2019; Baqaee and Farhi 2020).

The standard theory for understanding and measuring returns to scale owes much to Basu and Fernald (1996), although estimation of scale economies goes back much further (e.g. Tintner 1944). I will present an outline of the theoretical framework in the next section. On the measurement side, more modern approaches to estimating production functions (e.g. Olley and Pakes 1996; Levinsohn and Petrin 2003; Wooldridge 2009; Ackerberg, Caves, and Frazer 2015; Gandhi, Navarro, and Rivers 2020) help alleviate concerns of endogeneity, whereby estimates of factor elasticities are biased due to unobserved firm-level productivity. However production function estimation is still plagued with issues, and many assumptions are required to make progress (Flynn, Traina, and Gandhi 2019; Gandhi, Navarro, and Rivers 2020; Bond, Hashemi, Kaplan, and Zoch 2021).

Returns to scale is equal to the sum of output elasticities on all inputs, as this describes the change in output when all inputs are scaled slightly. Consistent estimation of output elasticities presents a range of challenges. Even with quantity data on both inputs and output, Bond, Hashemi, Kaplan, and Zoch (2021) show the conditions under which identification of output elasticities is possible with various methods and under assumptions of different market structures. To use the GMM estimator of Blundell and Bond (2000), persistent variation across firms in the input price of a perfectly flexible factor allows for consistent estimation of the output elasticity on that factor. However in many cases only revenue data is available, and in such situations only revenue elasticity can be consistently estimated (Bond, Hashemi, Kaplan, and Zoch 2021). If there is no markup heterogeneity across firms, then it may be possible to obtain revenue elasticity. Revenue elasticity is equal to the output elasticity divided by the markup, so is closely related to returns to scale. Gorodnichenko (2007) describes revenue elasticity as "returns to scale in the revenue function."

Basu and Fernald (1996) find constant or slightly decreasing returns to scale in U.S. production from 1959 - 1989. Using the same data, Altug and Filiztekin (1999) highlight that estimates of returns to scale differ when using prime or dual equations for estimation. Research on returns to scale in the UK have often analysed data from the mid-20th century, used older estimation methods with endogeneity issues, and focused only on manufacturing (e.g. Oulton 1996; Harris and Lau 1998; Girma and Görg 2002). Estimates from this research suggested approximately constant or slightly decreasing returns to scale.

Existing research commonly use *implied* revenue elasticity (which is the ratio of returns to scale over the markup) $\zeta = \gamma/\mu = 0.85$ where $\mu = 1$ (e.g. Atkeson and Kehoe 2005; Restuccia and Rogerson 2008; Hopenhayn 2014; Barseghyan and DiCecio 2016).[1] This $\zeta = 0.85$ implies that the economic profit share in revenue is 15%. This figure for the revenue elasticity is not typically specified directly by authors, but is an indirect outcome of specifying the markup and output elasticity. Typically one of these two is assumed be one. Either because authors assume constant returns to scale and some market power ($\gamma = 1,\ \mu > 1$) or because authors assume perfect competition and decreasing returns ($\mu = 1,\ \gamma < 1$). Atkeson and Kehoe (2005) specify $\mu = 1.11$, $\zeta = 0.95$ based on estimates in the literature, where ζ represents diminishing returns in variable factors. On the other hand, in a model with perfect competition and no markups ($\mu = 1$), revenue elasticity and output elasticity are equivalent (e.g. Restuccia and Rogerson (2008) set $\gamma = 0.85$ thus implying $\zeta = 0.85$).

In recent work, Gao and Kehrig (2020) estimate returns to scale at $\gamma = 0.96$ on U.S. manufacturing firms, with markups assumed to be constant. The authors highlight the variation in returns to scale across industries, ranging from 0.86 to 1.3. Similarly, Ruzic and Ho (2019) estimate returns to scale (with non-constant μ) to be $\gamma = 0.96$ in U.S. manufacturing industries in 2007, having experienced a gradual decline from over 1.2 in 1982. Over the same time, the standard deviation in returns to scale across industries rose by almost 40%. Another paper which investigates changing returns to

[1] Atkeson, Khan, and Ohanian (1996) investigate this parameter specification.

scale, and specifically the role of technology, is Lashkari, Bauer, and Boussard (2019), which posits a theory allowing for endogenous firm-level returns to scale depending on Information Technology itself. They find returns to scale ranging from 0.75 to 1.06, with smaller firms obtaining larger scale economies.

Why is technology so important? De Ridder (2019) gives a compelling argument along the following lines. There is evidence of a slowdown in productivity growth over the last fifteen years (Syverson 2017; Haskel, Goodridge, and Wallis 2015). Wide dispersion in firm-level productivity is well-documented (Gorodnichenko, Revoltella, Svejnar, and Weiss 2018). How could we reconcile these facts with the productivity puzzle? It is plausible that *scale economies* - the relationship between firm size and its costs of production - have changed due to new technologies, and this creates a wedge between firm-level productivity growth and the aggregate.

To be concrete, software has low marginal cost and high fixed cost, leading to a rise in returns to scale (De Ridder 2019). The evidence of high fixed costs associated with intangible capital - at least, relative to physical capital - is both intuitive, and has been found in the U.S. economy (Chiavari and Goraya 2021). Furthermore, Lashkari, Bauer, and Boussard (2019) show a correlation between firm size and the relative marginal product of IT, enabled by this technology allowing firms to deal with organisational complexity. Thus, adopters of such technologies can raise returns to scale, leading to a reallocation among firms, and a resulting productivity slowdown. In addition, there is evidence that markups have risen in the UK (Hwang and Savagar 2020) and the welfare and efficiency implications of higher markups can be severe (Edmond, Midrigan, and Xu 2021; Hsieh and Klenow 2009). The tight relationship between markups and market power encourages an investigation into the evolving market structures, and the role of technological change in this trend.

This paper fits firmly at the nexus of three recent papers already mentioned: De Ridder (2019), Lashkari, Bauer, and Boussard (2019), and Gao and Kehrig (2020). Firstly, the deep investigation into industry-level revenue elasticity across a range of estimation methods has much in common with Gao and Kehrig (2020). Following De

Ridder (2019), I look to understand how computer software has affected firms, but focus on output and revenue elasticities. To this end, I use more general production technology (without constant returns to scale, and including capital to match the estimation procedure), and simplified the model dynamics, excluding R&D and ignoring variation in the efficiency of adopting software. This allows us to focus more clearly on the relationship between optimising firms adopting software, and returns to scale in output and revenue. Finally, I do not build a fully structural model, but highlight the relationships between software adoption, output and revenue elasticities, and productivity. Returns to scale in output and revenue are endogenous, as in Lashkari, Bauer, and Boussard (2019), but without the non-homotheticity embedded in their model.

3.3 Theory

This section presents the theory behind returns to scale in a clear and general framework. I highlight that returns to scale can be expressed in two ways: as parameters of a homogeneous production function, or as a function of firm outcomes determined from optimising behaviour, in a competitive factor market.[2]

3.3.1 Returns To Scale

Definition 1. *Returns to scale* (RTS) *are the inverse cost elasticity. The inverse cost elasticity is the ratio of average cost to marginal cost. Thus,*

$$\mathrm{RTS} \equiv \left(\frac{\partial \mathcal{C}}{\partial y}\frac{y}{\mathcal{C}}\right)^{-1} = \frac{\mathcal{AC}}{\mathcal{MC}}, \tag{3.1}$$

where AC $\equiv \mathcal{C}/y$ *and* MC $\equiv \partial\mathcal{C}/\partial y$.

We typically do not observe total and marginal costs, so directly computing returns

[2]Basu (2008) and Savagar (2021) explain these two approaches to acquiring returns to scale. Our expression of RTS in terms of technological factors, and that of Basu (2008), assumes the fixed cost is output denominated. This would change if the fixed cost were factor denominated. But the broad interpretation would be the same. Returns to scale is explained by a component that relates to the elasticity of variable production (slope of marginal cost) and a component that relates to the fixed cost shares.

to scale using equation 3.1 is generally not possible. Thus we use theory to obtain different expressions for returns to scale. Consider the following general framework.

Firms minimise costs and maximise profits. The factor market is competitive, so factor prices are given. The competitive environment for output is monopolistic. Net output is given:

$$y = zF(k, \ell) - \phi, \tag{3.2}$$

where y is net output, z is Hicks-neutral productivity, k is capital, ℓ is labour, and ϕ is an output-denominated fixed cost.[3] Let $F(\cdot)$ be twice continuously differentiable, strictly concave, and homogeneous of degree ν.

In linearized form net output is given by:

$$\hat{y} = \hat{z} + \varepsilon_{yk}\hat{k} + \varepsilon_{y\ell}\hat{\ell} \tag{3.3}$$

where $\hat{x}$ represents the percentage change in variable x, and ε_{yx} is the output elasticity with respect to factor x. The sum of the two coefficients represent the change in net output from a change in both inputs. Intuitively this represents *returns to scale*. If they sum to one, such that a change in inputs causes a proportional change in output, then there are constant returns to scale.

The output elasticities are:

$$\varepsilon_{y\ell} \equiv \frac{\partial y}{\partial \ell}\frac{\ell}{y} = zF_{\ell}(k, \ell)\frac{\ell}{y}$$

$$\varepsilon_{yk} \equiv \frac{\partial y}{\partial k}\frac{k}{y} = zF_{k}(k, \ell)\frac{k}{y}.$$

Using the output elasticties and Euler's homogeneous function theorem I can show:

$$\nu(1 + s_{\phi}) = \varepsilon_{yk} + \varepsilon_{y\ell}, \tag{3.4}$$

[3]This type of fixed cost is crucial to endogenise returns to scale. It represents lost output that is not taken to the market. It is a fixed cost paid by active incumbents, but the cost does not vary with firm size. For example, a firm may pay a flat fee for access to Cloud services, but not fully utilise this production technology. There may also be extra costs which do scale with usage. In this case, the fixed cost for Cloud computing, per unit of output produced, would be very high. Thus the firm would have high returns to scale. Other examples include licensing fees, set-up costs, or HR costs.

where s_ϕ is the fixed cost share in net output, and ν is the degree of homogeneity of the production function, and is the slope of the marginal cost curve. This represents returns to scale with parameters of the production function.

Equation 3.4 highlights the importance of the fixed cost in net output, and the resulting endogeneity of returns to scale. Larger firms will have lower returns to scale, as the fixed cost share in output falls. These firms 'use up' the fixed cost over many units of output, moving down their cost curves, resulting in lower returns to scale.

Finally, returns to scale can be represented as markups multiplied by revenue elasticities, by rearranging the profit function:

$$\frac{\mathcal{AC}}{\mathcal{MC}} = \mu(1 - s_\pi) \tag{3.5}$$

where $\mu = \frac{P}{MC}$ is the markup, and $s_\pi = \frac{\pi}{py}$ is the share of profit in revenue.

It can be shown that *revenue elasticity* is equal to returns to scale divided by the markup, and thus is equal to $(1 - s_\pi)$ (D. Kim and Savagar 2021). Revenue elasticity will be equal to:

$$\text{Revenue Elasticity} = \frac{\nu}{\mu}\left(1 + s_\phi\right) \tag{3.6}$$

3.3.2 Optimisation

Firms solve the standard two-stage problem: they choose inputs to minimise variable costs subject to a production function, and select output to maximise profits subject to a demand schedule.

Cost Minimisation

Cost minimising firms solve:

$$\mathcal{C} := \min_{k,\ell} w\ell + rk \quad \text{s.t.} \quad y \geq zF(k,\ell) - \phi.$$

This yields the solution:

$$w = z\lambda \frac{\partial F}{\partial \ell}, \qquad r = z\lambda \frac{\partial F}{\partial k}.$$

And the variable costs are:

$$\mathcal{C} = w\ell + rk = z\lambda y \left(\frac{\partial F}{\partial \ell} \frac{\ell}{y} + \frac{\partial F}{\partial k} \frac{k}{y} \right) = \lambda y \left(\varepsilon_{y\ell} + \varepsilon_{yk} \right)$$

Given that marginal costs $\lambda = \frac{\partial \mathcal{C}}{\partial y}$, rearranging this expression gives the sum of net output elasticities equal to the inverse cost elasticity, which is the definition of returns to scale in equation 3.1:

$$\left(\frac{\partial \mathcal{C}}{\partial y} \frac{y}{\mathcal{C}} \right)^{-1} = \frac{AC}{MC} = \varepsilon_{y\ell} + \varepsilon_{yk}.$$

Thus, returns to scale is the sum of net output factor elasticities.

Profit Maximisation

Firms maximise profits subject to demand and compete monopolistically, such that each firms' price does not depend on the prices charged by competitors. The demand and inverse demand functions are:

$$y = \mathcal{D}(p) \quad \text{and} \quad p = \mathcal{P}(y).$$

p is the price, $\mathcal{D}$ is the inverse function of $\mathcal{P}$ which is monotone and twice continuously differentiable. I assume firms have market power, and hence their price depends on output. Profit maximising firms solve:

$$\pi := \max_{p,y} \; py - \mathcal{C}(y; w, r) \quad \text{s.t.} \quad p = \mathcal{P}(y).$$

This yields the familiar solution relating inverse markups to price elasticity:

$$\left(1 + \frac{\partial \mathcal{P}}{\partial y} \frac{y}{p} \right) = \frac{MC}{p} = \mu^{-1}.$$

Revenue Elasticity

Price elasticity of demand is given by:

$$-\frac{\partial \mathcal{D}}{\partial p}\frac{p}{y} = -\left(\frac{\partial \mathcal{P}}{\partial y}\frac{y}{p}\right)^{-1}.$$

Firms' revenue elasticity will depend both on the direct effect of changing an input X, as well as the indirect effect of changing input X on demand. Revenue elasticity with respect to input X is defined:

$$\frac{\partial R}{\partial X}\frac{X}{py} = \left[\frac{\partial \mathcal{P}}{\partial y}\frac{\partial y}{\partial X}y + p\frac{\partial y}{\partial X}\right]\frac{X}{py} = \left[\left(\frac{\partial \mathcal{D}}{\partial p}\frac{p}{y}\right)^{-1} + 1\right]\frac{\partial y}{\partial X}\frac{X}{y}.$$

where R is revenue, X is an input in the production function.

Given that the markup is the inverse of the final term in square brackets, it follows that revenue elasticity of input X is equal to output elasticity of input X divided by the markup. Returns to scale is the sum of output elasticities, so revenue elasticity is equal to returns to scale divided by the markup.

3.3.3 Productivity & Scale Economies

The relationship between productivity and scale is important, but the theoretical link is not simple. The classic interpretation (e.g. Gao and Kehrig 2020) posits that when returns to scale are high, average productivity is higher, as only more productive firms can exist in an environment with greater average-to-marginal cost ratios.

However, there is another important mechanism at play when output-denominated fixed costs are present, which seems under-emphasised in the literature.[4] More productive firms are typically larger, which lowers average costs, which lowers both output and revenue elasticities. This can be clearly seen in Equation 3.4, as higher y reduces the fixed-cost share s_ϕ. The intuition here is subtle, yet simple: more pro-

[4]Such fixed costs include licensing fees, human resources, or defective output. The idea is that some output is lost or destroyed in each period, and this represents a fixed cost to the firm that is unrelated to scale. It seems reasonable that such costs are fairly fixed over time and do not depend on firm productivity or factor input choices.

ductive firms spread the fixed-cost ϕ over a larger number of units, reducing their average cost and hence lowering their returns to scale and revenue elasticity. The cost curves shift with firm-level productivity, and the cost-minimising position that heterogeneous firms choose is associated with different ratios of the average and marginal cost curves.

Take the expression for returns to scale in Equation 3.4, and differentiate with respect to productivity z:

$$\frac{\partial \gamma}{\partial z} = -\frac{\nu\phi}{y^2}\frac{\partial y}{\partial z} \tag{3.7}$$

As long as more productive firms are larger ($\frac{\partial y}{\partial z} > 0$), returns to scale is falling in productivity. Why? Simply, average costs fall faster than do marginal costs, as productivity increases. In this framework with fixed markups, revenue elasticity is also falling in productivity.

Notice also that the elasticity of returns to scale with respect to productivity is:

$$\frac{\partial \gamma}{\partial z}\frac{z}{\gamma} = -\frac{s_\phi}{1+s_\phi}\varepsilon_{yz},$$

where ε_{yz} is the elasticity of output to productivity. Given that elasticity can be approximated with logs, a regression of log-RTS on log-productivity across firms within industries should yield: (1) negative coefficients and (2) correlation between coefficients and the fixed cost shares.

How can we reconcile the positive theoretical relationship between returns to scale and productivity (as in Gao and Kehrig (2020)) with this negative relationship? The key difference between these two results is *between* and *within* comparisons. For a given fixed cost ϕ, we should expect that returns to scale will fall as the average firm grows in size: this is the *within* effect. On the other hand, the productivity cut-off required for a firm to survive will differ *between* industries for different fixed costs ϕ.

3.3.4 Model Extension: Software

The above section sets out a standard environment to describe scale economies. To allow returns to scale in output and revenue to endogenously change with software adoption, I extend the model slightly. The intuition is that computer software scales down costs, but requires a fixed cost for adoption.[5] This differs in a key way from an important paper in this field by Lashkari, Bauer, and Boussard (2019), which instead introduced a non-homothetic production function to allow software to endogenously affect scale economies. My exposition is more simple, at the cost of some additional moments which can be matched in the data. I draw on some ideas from De Ridder (2019), but the production technology is more general: I do not impose constant returns to scale, and include capital as an input to better map onto the empirical work.

This model has two 'fixed costs'. For a firm to operate, they must pay ϕ, which is the same for all firms. This makes output and revenue elasticity endogenous: larger firms will have lower elasticity. This fixed cost is paid out of gross output in the production function, as in Equation (3.2). It is value that cannot be recouped in the output market, because it represents "lost output" from total firm production. This fixed cost is taken from gross output, as opposed to the cost function, following Savagar (2021). The main reason for this decision is to obtain an expression for returns to scale as in Equation (3.4) that depends only on the slope of the marginal cost curve ν and the fixed cost share s_ϕ. I show in Appendix C.6 that an output-denominated fixed cost in the cost function introduces the markup μ into this equation for returns to scale. This chapter uses a framework of monopolistic competition with CES demand which yields a constant markup. However, in other competitive environments with heterogeneous markups, the relationship between returns to scale and productivity becomes more complicated if the markup enters this expression for scale.

[5] Software could be considered as a direct input in the production function. I do not pursue this approach for two reasons. Firstly, this would require computing a software capital time series, and the combination of many zeros for software investment alongside high estimated depreciation rates would lead to many dropped observations in the empirical section. Secondly, the purpose of this model is simply to highlight the relationship between returns to scale, productivity, and computer software. Allowing software to directly affect costs allows for this in a clear way, and links cleanly to the intuition that software allows for scaling up at low cost.

The second 'fixed' cost is related to software adoption. This appears in the cost function, as it scales with software investment. This will be discussed in more detail, but it permits a separate mechanism for endogenous revenue elasticity, through computer software. In sum, heterogeneous firms will have different revenue elasticity, due to the fixed cost ϕ. Furthermore, those that adopt software will have an additional channel through which firm decisions affect output and revenue elasticity.

Software

Firms can use one of two production technologies: (1) produce with just labour and capital, or (2) invest in software, with the associated software fixed cost and reduced marginal costs, as in De Ridder (2019) and Kariel (2021a). As in Altomonte, Favoino, Morlacco, and Sonno (2021), let $s \in [0,1)$ be the reduction in marginal cost from using intangible software technology, with an associated 'fixed' cost $f(s)$.[6] Firms choose to use these technologies at both the extensive and intensive margin.

Let the 'fixed' cost $f(s)$ depend on the intensive margin software investment, with $f', f'' > 0$, and it satisfies both $f(0) = 0$ and $\lim_{s \to 1} f(s) = \infty$. Thus the 'fixed' cost is not paid when firms don't invest in software, and the fixed cost to eliminate marginal costs is infinite. Firms produce net output $y(z) = zF(k, \ell) - \phi$. They hire capital and labour in a cost minimisation procedure, and factor prices equal their marginal products. Firms turn to the market and maximise profits by choosing a price and a level of intangible investment s.

Software expenditure reduces the cost of employing given amounts of inputs, instead of raising output for a given quantity of inputs, as pursued by De Ridder (2019). If firms are cost minimising, either assumption will affect total costs in the same way. I show this equivalence in Appendix C.9. The cost function of cost-minimising firms will always be $\mathcal{C} = \lambda \nu y + f(s)$, where λ are marginal costs, ν is the slope of the marginal cost curve, y is output, and $f(s)$ is the cost function for software s.

[6]Note that this isn't *really* a fixed cost in the typical sense, as it depends on the amount of software s. However, as s is not an input, $f(s)$ isn't a typical marginal cost which enters the cost minimisation problem.

The decision to focus on the cost side in this paper is twofold. Firstly, returns to scale is the inverse cost elasticity, so it felt natural to focus on how software influences costs. Secondly, it seems intuitive that investing in software would scale down the cost of replicating tasks. However, for an alternative framework where software boosts output for a given quantity of inputs employed, see the model in Appendix C.9.

Demand

The demand side follows Melitz (2003), with Cobb-Douglas preferences over sectors $j \in \mathcal{J}$, and CES production over varieties ω within each sector.[7] This yields variety demand in each sector $y_j(\omega) = A_j p_j(\omega)^{\sigma_j}$ and the CES price index P, which each firm on the continuum takes as given. This can be solved to show that demand for each firm $y(z)$ can be written as a function of the relative price:

$$y(z) = \left(\frac{p(z)}{P}\right)^{\sigma} Y.$$

Cost Minimisation

When firms invest s in software, variable costs are scaled down by the factor $(1-s)$, with an associated software 'fixed' cost. The cost minimisation problem is:

$$\mathcal{C} := \min_{k,\ell} (1-s)(w\ell + rk) + f(s) \quad \text{s.t.} \quad y \geq zF(k,\ell) - \phi.$$

The Lagrange function is defined as follows:

$$\mathcal{L}(k,\ell,\lambda) = (1-s)(w\ell + rk) + f(s) - \lambda\left[zF(k,\ell) - \phi - y\right],$$

[7]Full details in Appendix C.8.

where λ is the Lagrange multiplier. The optimization conditions are

$$
\begin{aligned}
y &= zF(k,\ell) - \phi, \\
w &= \frac{z\lambda}{1-s}\frac{\partial F}{\partial \ell}, \\
r &= \frac{z\lambda}{1-s}\frac{\partial F}{\partial k}.
\end{aligned}
$$

We obtain minimised variable costs as in a standard model, with the additional fixed software cost:

$$\mathcal{C} = \nu\lambda(y+\phi) + f(s).$$

To obtain the marginal cost $\lambda = \frac{\partial \mathcal{C}}{\partial y}$, note that changes in output ∂y affect variable costs $\mathcal{C}$ through λ, y, and $f(s)$ as the optimal choice of software will depend on firm size. This will be shown and solved at the profit maximisation decision.

Following De Ridder (2019), consider the following function for the 'fixed' software cost:

$$f(s) = \left(\frac{1}{1-s}\right)^{\psi} - 1.$$

The marginal cost is:

$$\frac{\partial \mathcal{C}}{\partial y} = \lambda + \frac{\partial f(s)}{\partial s}\frac{\partial s}{\partial y}.$$

This expression requires obtaining $\frac{\partial s}{\partial y}$, which we compute in the following section.

Profit Maximisation

The profit function is:

$$\pi(z,s) = y(z)p(z) - (1-s)(wl + rk) - f(s).$$

Rewriting firm demand as a function of its price, the firm maximises:

$$\max_{p(z),s} \left(\frac{p(z)}{P}\right)^{-\sigma} Yp(z) - \nu\lambda\left(\left(\frac{p(z)}{P}\right)^{-\sigma} + \phi\right) - f(s)$$

This yields the following first-order conditions with respect to $p(z)$ and s:

$$\frac{\sigma-1}{\sigma}p(z) = \lambda\nu$$
$$\frac{\nu\lambda(y+\phi)}{1-s} = f'(s)$$

The first first-order condition rearranges to yield the result that price over marginal cost equals a constant markup:

$$\frac{p(z)}{\lambda} = \frac{\sigma}{\sigma-1}\nu$$

Given the derivative of the 'fixed' software cost is:

$$f'(s) = \psi\left(\frac{1}{1-s}\right)^{\psi-1}.$$

then the second first-order condition yields optimal software investment as:

$$s = 1 - \left[\frac{1}{\psi}\left(\frac{y+\phi}{z}\right)^{\frac{1}{\nu}} G(w,r,1)\right]^{\frac{1}{2-\psi}}. \tag{3.8}$$

where $G(w,r,1)$ is an arbitrary function independent of output y, as in Savagar (2021). It is obtained by integrating over the partial differential relationship between the cost function C and the marginal cost λ that is the result of cost minimisation. Further details are in Appendix C.5.

Software investment therefore depends on firm size, and the derivative is:

$$\frac{\partial s}{\partial y} = \frac{\lambda}{\psi(\psi-2)}(1-s)^{\psi-1}.$$

Full derivation in Appendix C.8.

And the expression for firm's marginal costs:[8]

$$\frac{\partial \mathcal{C}}{\partial y} = \lambda\left(\frac{\psi - 1}{\psi - 2}\right)$$

Scale and Software

Returns to scale is the ratio of average to marginal costs. In this environment with software, the expression is as follows:

$$\frac{\mathcal{AC}}{\mathcal{MC}} = \frac{\nu\lambda(1 + s_{\phi}) + \frac{f(s)}{y}}{\lambda\left(\frac{\psi-1}{\psi-2}\right)}.$$

Substitute the equation for the 'fixed' software cost $f(s)$, and the FOC from profit maximisation with respect to software s:

$$\frac{\mathcal{AC}}{\mathcal{MC}} = \nu(1 + s_{\phi})\left[\left(\frac{\psi - 2}{\psi(\psi - 1)}\right)\left(\psi + 1 - (1 - s)^{\psi}\right)\right] \tag{3.9}$$

Equation 3.9 highlights that returns to scale will be equal to the expression when software is unavailable, multiplied by the term in square brackets. This term depends on the amount of software adopted s (which itself is determined by productivity z), and the curvature of the software 'fixed' cost ψ.

An alternative way to express returns to scale is:

$$\frac{\mathcal{AC}}{\mathcal{MC}} = \left(\frac{\psi - 2}{\psi - 1}\right)\left[\nu(1 + s_{\phi})\left(1 + \frac{f(s)}{\mathcal{C} - f(s)}\right)\right] \tag{3.10}$$

which highlights the importance of the 'fixed' software cost to variable cost ratio.

The revenue elasticity is the output elasticity divided by the markup. In this environment the revenue elasticity will be:

$$\text{Revenue Elasticity} = \left(\frac{\psi - 2}{\psi - 1}\right)\left[\frac{\nu}{\mu}(1 + s_{\phi})\left(1 + \frac{f(s)}{\mathcal{C} - f(s)}\right)\right] \tag{3.11}$$

[8] Because marginal costs must be positive, I will hereafter assume that $\psi > 2$. In De Ridder (2019), ψ is calibrated to 2. This assumption is equivalent to the the software cost increasing slightly more rapidly as firms invest in more software.

Bot returns to scale and revenue elasticity are increasing in software investment s. This can be seen clearly in Equations (3.9) and (3.11), with software s in the bracketed term.

Proposition 1. *Revenue elasticity is increasing in software investment.*

$$\frac{\partial Revenue\ Elasticity}{\partial s} = \frac{\nu}{\mu}(1+s_{\phi})\left(\frac{\psi-2}{\psi-1}\right)(1-s)^{\psi-1} > 0.$$

Secondly, there is an additional channel through which productivity affects returns to scale and revenue elasticity when software is available. Higher productivity firms adopt more software, as in Equation 3.8. The effect on variable costs is zero, as firms optimally increase their inputs to offset the reduction in costs afforded by software. However, fixed software costs $f(s)$ will rise. Thus, software allows higher productivity firms to obtain greater returns to scale and revenue elasticity than if software were not available.

Proposition 2. *Software investment mediates the relationship between revenue elasticity and productivity.*

3.4 Data & Estimation

This section contains information on the Annual Respondents Database X (ARDx) which is used to estimate revenue elasticities and revenue total factor productivity (TFPR). I describe the data coverage, the data cleaning process, and construction of the capital stock. I summarise the two approaches to identifying production functions, following Ackerberg, Caves, and Frazer (2015) and Gandhi, Navarro, and Rivers (2020). There are many production functions estimated, at various levels of aggregation and following different sets of assumptions. Finally I address the distinction between the model and estimation approaches, in that the latter does not assume fully-flexible inputs.

3.4.1 UK Data

The data is from the ARDx, which is a UK firm-level research dataset constructed from two ONS surveys, the Annual Business Inquiry (ABI; 1998 - 2008) and the Annual Business Survey (ABS; 2009 onwards), combined with the Business Register and Employment Survey (BRES; 2009 onwards). It contains financial and employment data for around 50,000 firms each year, with stratification by industry, size, and region to ensure a representative sample.

The ARDx is essentially a census and a survey: the former for large firms, which are repeatedly included, and the latter for small firms, which are sampled with specific rules on inclusion to reduce the administrative burden. There are approximately 11 million workers covered by the businesses in the ARDx.

Firm gross output is converted into real values using the ONS experimental industry deflators.[9] Material inputs are deflated with the ONS producer price inflation data.[10]. Finally, the constructed capital stock is deflated with the ONS gross fixed capital formation deflator.[11] For the purpose of production function estimation, I exclude certain non-market sectors: Agriculture, Public Sector, Finance & Insurance, Education, and Health.[12] I remove firms if there are fewer than 100 firms in a sector, as sector-specific production functions have too few observations to be estimated precisely. To ensure compatibility of SIC pre- and post-2007, when the classification is changed, a simple re-coding is undertaken. For SIC codes post-2007, the number is divided by 1000 to match with pre-2007 codes.

To reduce the influence of outliers, which may represent measurement or recording errors in the surveys, I winsorise firms with the top and bottom 0.1% of factor shares (M/Y, K/Y, L/Y) in each year. In addition, the top and bottom 0.1% growth rates of gross revenue, capital, labour and materials are removed. Table 3.1 contains number

[9] https://www.ons.gov.uk/economy/inflationandpriceindices/datasets/experimentalindustrydeflatorsuknonseasonallyadjusted

[10] https://www.ons.gov.uk/economy/inflationandpriceindices/datasets/producerpriceindex

[11] https://www.ons.gov.uk/economy/grossdomesticproductgdp/timeseries/ybfu/ukea

[12] 2-digit Standard Industrial Classification (SIC) 2007 codes: A, K, O, P, Q.

of firms at each stage of the data cleaning process, along with the final number of observations for estimation. Summary statistics for the variables used in production function estimation are in Table C.15 of the Appendix.

Table 3.1: Data Cleaning: Firms Dropped

	# Firms
All ARD firms	854,732
Drop if no 2-digit sector	852,424
Drop if < 100 firms in sector	852,331
Drop non-market sectors	761,348
Take logs of regression variables	539,368
Drop outliers (factor shares and growth rates)	527,813

3.4.2 Production Function Estimation

Firm-level productivity estimation is computed using control function methods. These require firm-level capital stock. This section explains the Perpetual Inventory Method (PIM) used to construct a measure of capital stock, followed by the productivity estimation approaches.

The Perpetual Inventory Method (PIM) allows construction of firm-level capital stocks when such data is unavailable, but investment data is present. The method here follows Martin (2002). The PIM is constructed using the following equation:

$$K_t = (1 - \delta)K_{t-1} + i_t.$$

where K_t is the capital stock in period t, and i_t is investment in period t. However to use this method, K_0 is required. This is the initial capital stock of a firm, and is not available in this survey. To construct this series, each firm's K_0 is a revenue-weighted share of the industry-level capital stock in the first year that firm appears in the panel, as in (Hwang and Savagar 2020). Capital stock is then constructed for all future years with the above equation, with missing investment data interpolated. The depreciation rate is taken to be 18.195%, which is a weighted average of ONS depreciation rates for the three different capital categories: Building, Vehicles, Other.

A number of papers propose techniques with a set of assumptions for recovering production function parameters from firm panel data (Olley and Pakes 1996; Levinsohn and Petrin 2003; Ackerberg, Caves, and Frazer 2015; Doraszelski and Jaumandreu 2018; Gandhi, Navarro, and Rivers 2020; Bond, Hashemi, Kaplan, and Zoch 2021). When using sales data, we can only identify the revenue elasticity. Under the assumption of fixed markups across firms and over time, as in the model, the revenue elasticity is equal to returns to scale divided by the markup, but we cannot separately identify these components.

I estimate production functions using Ackerberg, Caves, and Frazer (2015) (ACF) and Gandhi, Navarro, and Rivers (2020) (GNR).[13] The former method uses 'control function' approach. A key assumption of this method is that the idiosyncratic productivity shock at time t does not affect the choice of state variables chosen by the firm in previous periods, but does affect the decision on free variables. The latter method uses the 'cost share' approach which leverages the first-order condition on the flexible input to identify the parameters of the production function. Here I discuss the assumptions, moment conditions, and over-identifying restrictions. More detail is provided in Appendix C.1.

The relationship between outputs and inputs takes the form:

$$Y_{it} = F(l_{it}, k_{it}, m_{it}) e^{\epsilon_{it}}$$

$$y_{it} = f(l_{it}, k_{it}, m_{it}) + \omega_{it} + \eta_{it}$$

where $y_{it}, l_{it}, k_{it}, m_{it}$ are logs of output, labour, capital and materials. $f()$ is a production function which is continuously differentiable. ω_{it} is the component of productivity that firms observe prior to making decisions in period t, while η_{it} is an ex-post productivity shock.

[13] I use **prodest** in STATA to estimate using Ackerberg, Caves, and Frazer (2015), and I use **gnrprod** in R to estimate using Gandhi, Navarro, and Rivers (2020).

The productivity shock ω_{it} follows a Markov process, such that it is in the firm's information set $\mathcal{I}_{it}$ when it makes input and output choices. This Markov process means that firms form expectations on future shocks using only the current period ω_{it} and the known productivity distribution P_ω, so $P_\omega(\omega_{it}|\mathcal{I}_{it-1}) = P_\omega(\omega_{it}|\omega_{it-1})$. Due to the Markov process, we can write this period's productivity shock as $\omega_{it} = g(\omega_{it-1}) + v_{it}$ where $g(\omega_{it-1})$ is the predictable component of the productivity process, and v_{it} is an unpredictable 'surprise' to firms.

Capital is chosen in period $t-1$, while labour is chosen from $t-1$ to t. I assume labour is chosen after capital, but other assumptions may permit additional moment conditions. Materials input is chosen at the same time or after labour.

Control function methods such as Ackerberg, Caves, and Frazer (2015) assume a fully-flexible non-dynamic 'proxy' variable such as materials input m_{it} that is strictly monotonic in only one unobservable ω_{it}. This allows this unobserved productivity shock to be expressed as a function of observables and parameters. Estimates of ω_{it} can be computed for given guesses of coefficients on inputs, and the Markov process on productivity is exploited to estimate the 'surprise' component v_{it}. This yields a set of moment conditions for the first and second stages respectively (more detail in Appendix C.1):

$$\mathbb{E}(\eta_{it}|\mathcal{I}_{it}) = 0 \tag{3.12}$$

$$\mathbb{E}(v_{it} + \eta_{it}|\mathcal{I}_{it-1}) = 0 \tag{3.13}$$

In practice, Equation (3.13) means that this period capital k_{it} and lagged labour l_{it-1} can be used in the second stage to identify production function coefficients.

Notice that the moment conditions allow for extra lags for additional moments. Implied v_{it} are uncorrelated with further lags of capital and labour, as these are in $\mathcal{I}_{it-1}$. This can add efficiency to the estimator, but at the cost of reducing the effective sample for estimation. In STATA with the package **prodest**, I can estimate production functions with overidentifying restrictions following the approach of Wooldridge (2009). Previous lags are valid instruments given Equation 3.13.

Gandhi, Navarro, and Rivers (2020) use the firm's first-order condition to non-parametrically obtain the coefficient on the intermediate input from a gross output production function. The first-order condition for materials from the profit maximisation problem yields:

$$P_t \frac{\partial F(l_{it}, k_{it}, m_{it})}{\partial M_{it}} e^{\omega_{it}} \mathbb{E}(\eta_{it} | \mathcal{I}_{it}) = P_t^M \tag{3.14}$$

where P_t is the output price, $\mathcal{I}_{it}$ is the information set of firm i in period t, and P_t^M is the materials input price. Demand for materials depends on capital and labour (which are chosen before materials) and the productivity shock ω_{it}. Taking logs of Equation (3.14) and substituting in the production function yields a 'cost-share' equation which can be used to identify the flexible input coefficient:

$$s_{it} = \ln\left(\frac{\partial f(l_{it}, k_{it}, m_{it})}{\partial m_{it}}\right) + \ln \mathbb{E}(\eta_{it} | \mathcal{I}_{it}) - \eta_{it} \tag{3.15}$$

where $s_{it} = \ln \frac{P_t^M M_{it}}{P_t Y_{it}}$ is the log of the materials cost share in revenue. Gandhi, Navarro, and Rivers (2020) exploit $\mathbb{E}(\eta_{it} | \mathcal{I}_{it}) = 0$ to recover the elasticity on materials and η_{it} by regressing the log materials cost share s_{it} on a vector of capital, labour, and materials inputs. The rest of the production function can be obtained non-parametrically, with full details available in Appendix C.1.

3.4.3 Estimation Details

For both ACF and GNR I estimate a gross-ouput Cobb-Douglas production function of the following form:

$$Y_{it} = K_{it}^{\beta_k} L_{it}^{\beta_l} M_{it}^{\beta_m} e^{\epsilon_{it}} \tag{3.16}$$

where $Y_{it}, K_{it}, L_{it}, M_{it}$ represent revenue, capital stock, employment, and materials respectively, while the β's are the production elasticities. Productivity shocks $\epsilon_{it} = \omega_{it} + \eta_{it}$ where ω_{it} are ex-ante shocks and η_{it} are ex-post shocks. Taking logarithms yields a regression equation with the log of gross revenue on the left-hand side, and logs of the factor inputs on the right-hand side.

I also consider the "value-added" or Leontief production function for ACF:

$$Y_{it} = \min\{e^{\omega_{it}} K_{it}^{\beta_k} L_{it}^{\beta_l}, M_{it}^{\beta_m}\} e^{\eta_{it}} \tag{3.17}$$

Taking logarithms yields the "value-added" production function:

$$y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \omega_{it} + \eta_{it} \tag{3.18}$$

Firms draw productivity ω_{it} which is unobserved by the econometrician, leading to potential omitted variable bias, as clearly the optimal firm input choices will be correlated with this variable. As Ackerberg, Caves, and Frazer (2015) discuss, imposing this structure on the production side is meaningful, and implies that the materials input is proportional to Y_{it}.

I estimate the production functions with technology that is Cobb-Douglas (with ACF and GNR) and translog (with ACF only, due to software limitations).[14] The former gives fixed elasticity estimates, while the latter provides firm- and time-varying elasticity estimates. Hence Cobb-Douglas results for revenue elasticity provide just one estimate for the time period of estimation, while translog results provides a distribution of revenue elasticities across firms in each year.

All production functions estimated with ACF include year-specific fixed effects. Estimates using GNR do not have year dummies, due to software limitations.

Cobb-Douglas production functions are estimated in several ways. Firstly, I provide aggregate and macro-sector estimates over 1998 - 2014. Secondly, I estimate production functions separately by 2-digit sectors over the whole time period. Finally, I split the data into sub-periods (1998 - 2001, 2002 - 2005, 2006 - 2009, 2010 - 2014) and re-estimate production functions at the aggregate and industry levels. Revenue elasticity is computed by summing the elasticities obtained from production function

[14] **gnrprod** does not permit a translog specification, nor extra instruments to test for over-identification, nor the inclusion of year fixed effects.

estimation.

The moment conditions set out in Section 3.4.2 allow for extra instruments with further lags of capital and labour. The STATA package **prodest** permits production function estimation with overidentifying restrictions, following the approach in Wooldridge (2009). I estimate gross-output and value-added production functions with one additional lag in capital and labour as instruments. This permits Hansen overidentification tests to jointly test the validity of all the instruments.

3.4.4 Flexibility of Inputs

There is one important distinction between the model and the estimation procedure that requires comment. In the model, all inputs are flexible. Firms choose the optimal factor demand in the cost minimisation procedure, with no dynamic processes for inputs. In production function estimation, timing assumptions on firm input choices are crucial. Firms accumulate capital according to a dynamic process, while labour and materials are flexible inputs. Therefore, it is only these fully flexible inputs which are considered in the firm's cost minimisation procedure. I show in Appendix C.7 that introducing convex adjustment costs on capital to the standard model leads to an extra term in the expressions for returns to scale and revenue elasticity.

Both returns to scale and revenue elasticity can be higher or lower in the presence of adjustment costs. Investment in capital at the margin now depends both on the cost of capital r and the (positive or negative) gap between today's choice and the previous period choice. This pushes firms away from the one-period optimal choice of capital, and short-run cost curves are equal to or above long-run cost curves. Whether average or marginal costs rise more will depend on whether firm's capital stock is growing or shrinking. If firms are growing on average, returns to scale and revenue elasticity will be lower than if adjustment costs were absent. As more productive firms will grow more quickly in the model, adjustment costs introduce an extra channel through which productivity affects returns to scale and revenue elasticity. This effect is in the same direction as the theoretical predictions arising from fixed costs to produce output and to adopt software.

The model could include adjustment costs on capital as in Appendix C.7. The exclusion of this feature from the model does not prevent the theory highlighting how software affects revenue elasticities and scale. The inclusion of capital as a state variable with adjustment costs would add some complexity without changing the main relationships between revenue elasticity, productivity, and computer software. The identification of revenue elasticity relies on capital adjusting more slowly than the flexible inputs. I sum the estimated coefficients to obtain revenue elasticities, which is not equal to costs over revenue if adjustment costs are present (see Appendix C.7). There is a discrepancy between the model and the estimation procedure, but I do not believe it invalidates the results.

3.5 Empirical Results

3.5.1 Revenue Elasticity Estimation

In this section, I present results on revenue elasticity in the UK economy between 1998 and 2014. These estimates are computed by summing the estimated coefficients from the production function estimation. The estimated coefficients are either labour, capital, and materials (gross output production functions) or labour and capital (value added production functions). I present results where elasticities are constant over time (Cobb-Douglas) and time-varying (translog). In addition, revenue elasticity is estimated and reported at different levels of aggregation. As in Hall (1988), I find estimates differ across levels of aggregation. I report results where each estimated input elasticity is between zero and one.

The main findings are that:

1. Revenue elasticity in the UK is close to constant.
2. Revenue elasticities are heterogeneous across sectors (0.53 - 1.52).
3. Revenue elasticity in the UK has risen over time.

Table 3.2 presents the estimates of revenue elasticities over the whole period at the aggregate level using a Cobb-Douglas production function. The three estimates follow the methodologies of Gandhi, Navarro, and Rivers (2020) and Ackerberg, Caves, and Frazer (2015) with a gross-output production function, and a value-added production function for the latter approach. These estimates suggest that revenue elasticity in the UK from 1998 - 2014 is close to unity.

Table 3.2: Revenue Elasticity: Cobb-Douglas production function, 1998 - 2014

	ACF (VA)	ACF (GO)	GNR (GO)
β_l	0.545	0.177	0.329
	(0.020)	(0.024)	(0.006)
β_k	0.505	0.041	0.181
	(0.028)	(0.006)	(0.003)
β_m	-	0.758	0.514
	-	(0.017)	(0.002)
Revenue Elasticity	1.051	0.976	1.024
N	527,813	527,813	527,813

Note: Revenue elasticity and estimated coefficients on labour, capital and materials from Cobb-Douglas production functions, with gross output (GO) or value-added (VA). Standard errors in brackets. Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF) or Gandhi, Navarro, and Rivers (2020) (GNR).

I find that revenue elasticity is greatest in the UK in Manufacturing, and lowest in Services. Table 3.3 presents these results. These findings hold across all three estimation methods. More specifically, estimates of revenue elasticity are higher for Manufacturing and Construction than for Wholesale, Trade & Transport and Services. The underlying coefficients are contained in Table C.1 in the Appendix.

Table 3.3: Revenue Elasticities by Macro Sector: Cobb-Douglas production function, 1998 - 2014

	ACF (VA)	ACF (GO)	GNR (GO)
	Manufacturing		
Revenue Elasticity	1.143	1.189	1.034
N	120,712	120,712	120,712
	Construction		
Revenue Elasticity	1.192	0.874	1.044
N	51,784	51,784	51,784
	Wholesale/Trade/Transport		
Revenue Elasticity	0.926	0.852	1.016
N	181,985	181,985	181,985
	Services		
Revenue Elasticity	1.067	0.798	1.015
N	173,332	173,332	173,332

Note: Revenue elasticity from Cobb-Douglas production functions, with gross output (GO) or value-added (VA). Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF) or Gandhi, Navarro, and Rivers (2020) (GNR).

I also estimate revenue elasticities at the 2-digit industry level. These results are contained in Table C.2 in the Appendix, with underlying coefficients in Tables C.3, C.4 and C.5. The estimates follow Ackerberg, Caves, and Frazer (2015) and Gandhi, Navarro, and Rivers (2020) with a value-added or gross-output Cobb-Douglas production function. It is not always possible to obtain estimates of revenue elasticity for each industry. If the estimated input coefficients are below zero or above one, I do not compute the industry revenue elasticity. For GNR with gross-output, this only occurs for two industries. Using ACF with gross-output, there are nine industries for which revenue elasticity cannot be computed. With ACF and value-added, seven industries are missing.

The results show a wide range of revenue elasticities, from 0.53 to 1.52. For example, industries with low revenue elasticities include Legal & Accounting Services (SIC 69), Scientific & Research Development (SIC 72), and Creative, Arts & Entertainment

Activities (SIC 90). Industries with high revenue elasticities are Manufacture of Wearing Apparel (SIC 14), Manufacture of Chemicals (SIC 20), Manufacture of Furniture (SIC 31), and Information Service Activities (SIC 63). These results are consistent with the idea that firms can scale up activities more easily in manufacturing than in specialised services, as well as in industries where computers and software are important.

I also estimate production functions with an extra set of instruments. The estimates of revenue elasticity at the aggregate, macro sector and 2-digit SIC level are contained in Tables C.6, C.7 and C.8. These results produce systematically lower estimates of revenue elasticity. In addition, the Hansen test statistics always provide evidence that all instruments are jointly valid.

At the aggregate level, the overidentified production function yields revenue elasticities substantially below unity, at 0.76 and 0.91 for gross-output and value-added specifications respectively (see Table C.6). This result also holds for macro sectors (see Table C.7), with revenue elasticity sitting between 0.72 and 0.89. Table C.8 presents the estimates for 2-digit industries. The unweighted average revenue elasticity across sectors is 0.82 (gross-output) and 0.92 (value-added).

However, the extra set of instruments comes at the cost of significantly smaller sample sizes. For example, in the baseline estimates in Table C.2, the average number of observations used in each 2-digit SIC gross-output production function is over 8,500. In contrast, when extra lags are used as instruments, the average number of observations in each industry is just over 2,600. For this reason, especially given I cut the sample into smaller sub-periods, I proceed with the baseline production function estimation.

Rising Revenue Elasticity

I estimate production functions on four shorter sub-periods, in order to track changes in revenue elasticity over time. Table 3.4 presents these revenue elasticity estimates following Gandhi, Navarro, and Rivers (2020). On aggregate, there is some evidence of a rise in scale over time, from slightly below one, to a little above unity. The underlying coefficients can be found in Tables C.10 and C.11.

Table 3.4: Changing Revenue Elasticity: Cobb-Douglas production function, 1998 - 2014

	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
	Ackerberg, Caves, and Frazer (2015) (VA)			
Revenue Elasticity	0.988	1.081	1.046	1.061
N	153,874	144,465	108,619	120,855
	Ackerberg, Caves, and Frazer (2015) (GO)			
Revenue Elasticity	0.713	1.071	0.976	1.311
N	153,874	144,465	108,619	120,855
	Gandhi, Navarro, and Rivers (2020) (GO)			
Revenue Elasticity	0.994	1.028	1.032	1.025
N	153,874	144,465	108,619	120,855

Note: Revenue elasticity from Cobb-Douglas production functions, with gross output (GO) or value-added (VA). Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF) or Gandhi, Navarro, and Rivers (2020) (GNR).

Revenue elasticity has also increased across all macro sectors between 1998 - 2014. Table 3.4 shows a steady increase in estimates of revenue elasticity in each sector. The largest rise occurs in Manufacturing, from 0.997 to 1.223. Services experienced the smallest rise, from 0.982 to 1.062. The underlying estimated coefficients are in Table C.12. Estimates using Ackerberg, Caves, and Frazer (2015) are available in Tables C.13 and C.14.

Table 3.5: Changing Revenue Elasticity by Macro Sector: Cobb-Douglas production function (GNR), 1998 - 2014

	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
		Manufacturing		
Revenue Elasticity	0.997	1.016	1.053	1.223
N	39,876	34,678	24,011	22,147
		Construction		
Revenue Elasticity	1.017	1.013	1.091	1.139
N	13,484	13,416	10,210	14,674
		Wholesale/Trade/Transport		
Revenue Elasticity	1.008	1.074	1.035	1.129
N	53,814	50,631	37,906	39,634
		Services		
Revenue Elasticity	0.982	1.009	1.044	1.062
N	46,700	45,740	36,492	44,400

Note: Revenue elasticity from a Cobb-Douglas gross-output production function. Estimation approach follows Gandhi, Navarro, and Rivers (2020) (GNR).

The rise in revenue elasticity over time is more apparent when I estimate at the 2-digit industry level. Figures 3.1 plots a comparison of revenue elasticity in 1998 - 2001, compared to 2010 - 2014, across sectors.[15] It is clear in both cases that most industries experienced an increase in revenue elasticity, as the majority of points sit above the 45 degree line.

For time-varying revenue elasticity estimates, we require a translog production function.[16] This yields firm-year factor elasticities which are summed to obtain firm-year revenue elasticity. I remove estimates where labour or capital elasticities are below zero or greater than one, leaving estimates for over 300,000 firms from 1998 to 2014. A summary of translog revenue elasticity estimates are contained in Table C.9 in the Appendix. The average is 0.940, with a standard deviation of 0.344. Figure 3.2 plots averages of firm-level revenue elasticity in each year.

[15]I remove industries where estimated factor elasticities are below zero or above one.

[16]These estimates are computed following the Ackerberg, Caves, and Frazer (2015) methodology with a value-added production function, as **prodest** permits a translog production function, while **gnrprod** does not yet have this functionality.

Figure 3.1: Revenue Elasticity in 1998 - 2001 vs. 2010 - 2014

(a) Revenue Elasticity (Ackerberg, Caves, and Frazer (2015))

(b) Revenue Elasticity (Gandhi, Navarro, and Rivers (2020))

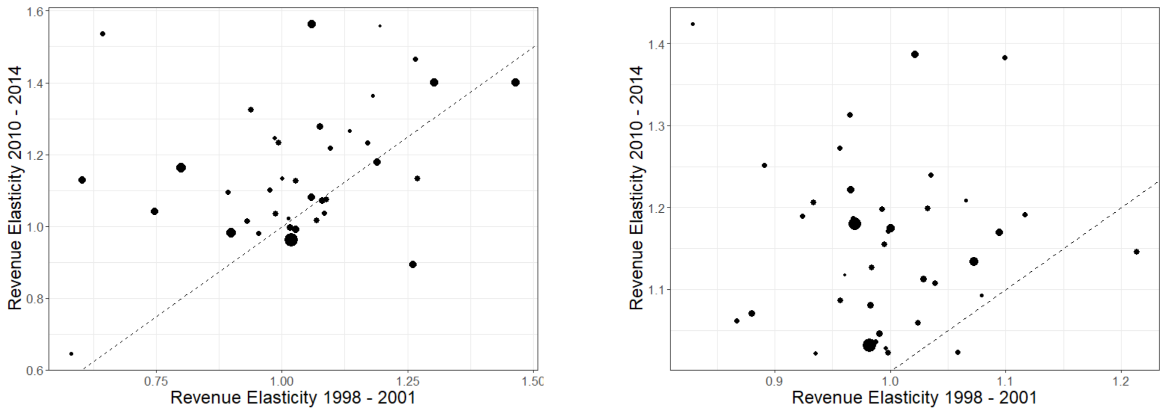

Note: Comparison of revenue elasticity at 2-digit SIC level, from 1998 - 2001 to 2010 - 2014. Size of points represents the average number of firms in that sector in each period. Dotted line is 45 degree line: points above that line are consistent with a rise in revenue elasticity.

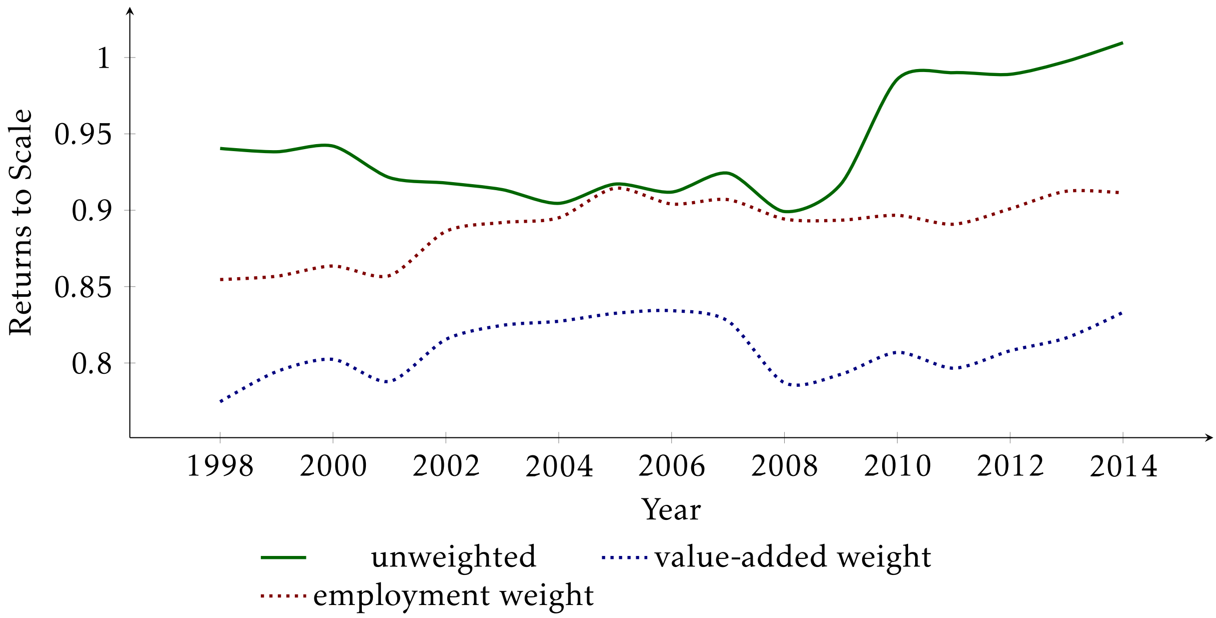

Figure 3.2: Aggregate Revenue Elasticity from translog production function.

Figure 3.2 presents estimates of revenue elasticity across the whole economy, where firm-level estimates are averaged in each year. The estimates show revenue elasticity just below unity, rising above one towards the end of the sample. When firm-level revenue elasticities are weighted by the value-added or employment, the average is lower and the increase in revenue elasticity is less pronounced. This is due to the negative relationship between firm-level revenue elasticity and firm size. Hence the unweighted average shows the change in firm-level revenue elasticity most clearly.

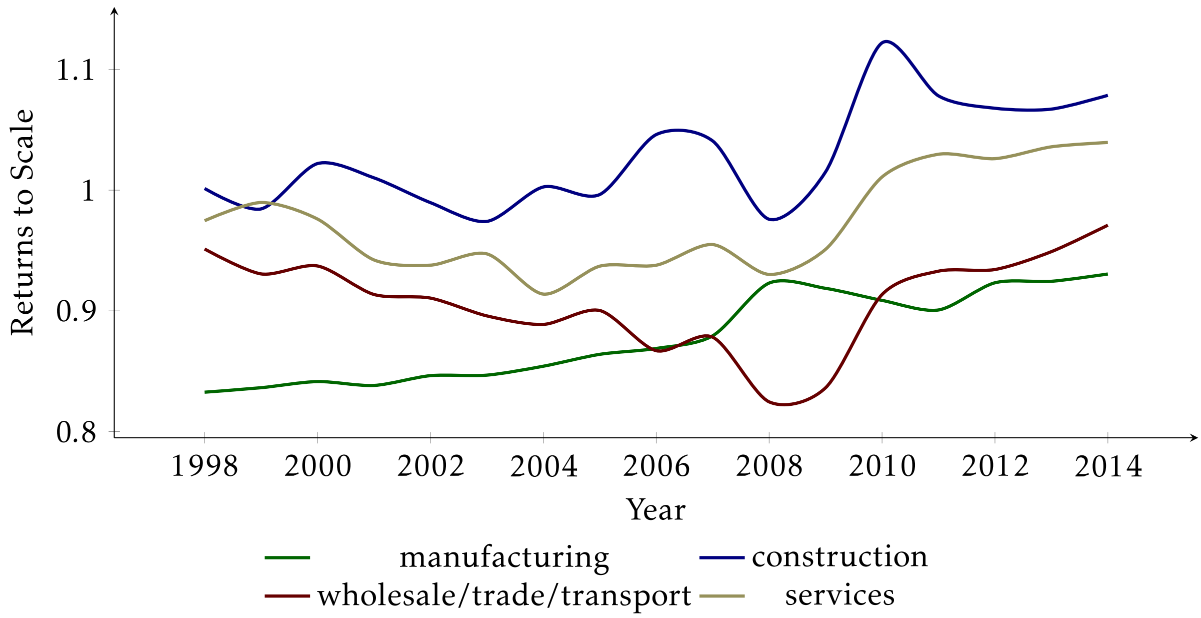

Figure 3.3: Revenue Elasticity by sector from translog production function.

For different sectors, the results are provided in Figure 3.3. These results are unweighted, and production functions were estimated at the industry level. These results show a similar increase in revenue elasticity across industries, but different levels between sectors. The exception is Wholesale, Trade & Transport, which saw revenue elasticity fall until the Financial Crisis, before rebounding.

The translog estimates also allow us to study the distribution of revenue elasticity across firms, over time. Figure C.1 shows these distributions in 1998 and 2014 for all firms in the dataset. This plot shows the increased dispersion in revenue elasticity over time. There is a significant increase in the mass of firms with revenue elasticity greater than unity. I also plot percentiles of revenue elasticity over time in Figures C.2 and C.3, which highlight the same trends. There is an increase in the level of all percentiles along the revenue elasticity firm distribution. In other words, the distribution of revenue elasticity shifts right. This is most pronounced in Manufacturing and Services.

3.5.2 Productivity Estimation

Revenue Total Factor Productivity (TFPR) is estimated using the control function methods described in Section 3.4.2. Figure 3.4 presents the results for estimates of TFPR in the cases of a gross-output Cobb-Douglas and value-added translog production functions (both following Ackerberg, Caves, and Frazer (2015)). It shows the steady rise in TFPR up to the Financial Crisis, followed by the plateauing from the late 2000s, and an uptick from around 2012. This result is robust across alternative estimation methods, as seen in Figure C.4 in the Appendix.

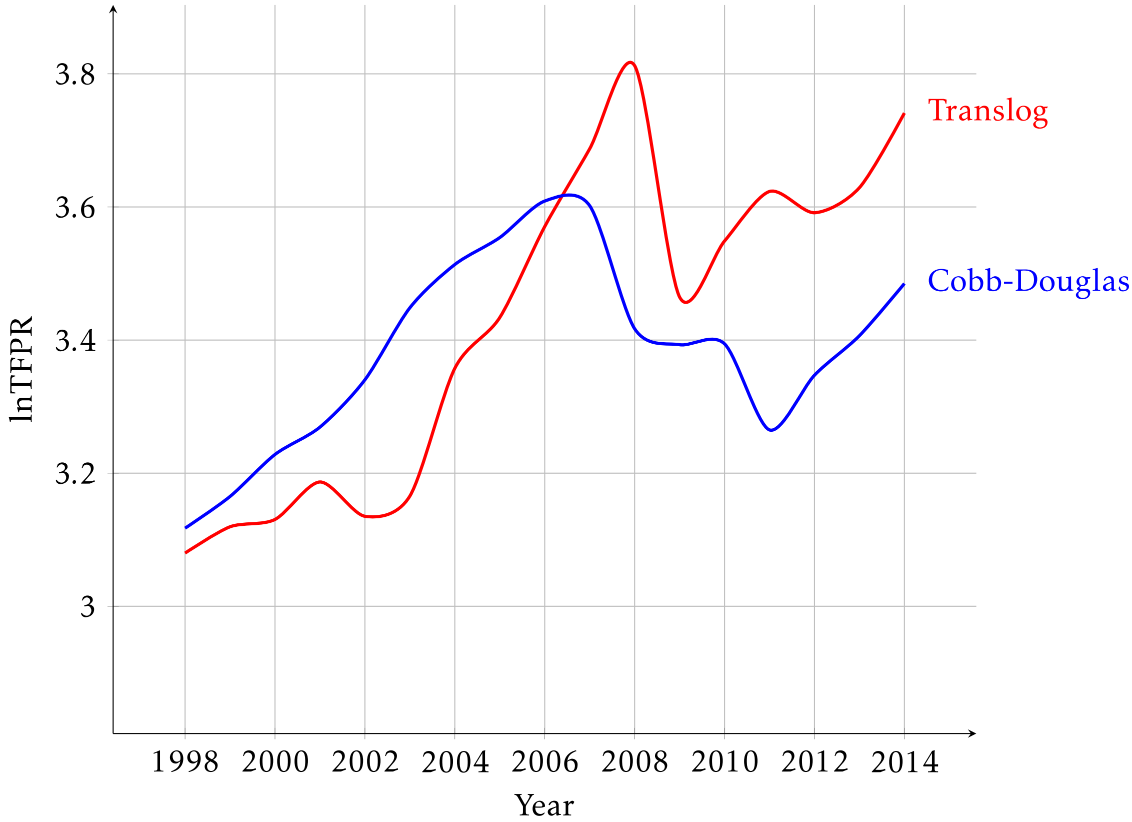

Figure 3.4: Aggregate ln TFPR.

Alongside a rise in the level of log TFPR between 1998 and 2014, the UK has also experienced an increase in the spread of log TFPR across firms. Figure C.5 plots the standard deviation of log TFPR estimated from Cobb-Douglas production functions with two different estimation methods. There is a clear rise in the standard deviation of log TFPR by around 10 percentage points from 1998 to 2007. It falls before rebounding towards the end of the sample period.

Figures C.6 and C.7 provide estimates of log TFPR by sector from Cobb-Douglas and translog production functions, respectively. Productivity is highest for Services and Wholesale/Trade/Transport, while it is lowest for Manufacturing. Each sector experienced a steady rise in log TFPR from 1998 until the Financial Crisis. By 2014 productivity exceeded its pre-2007 level in some sectors, but noticeably not in Services. Given the importance of the Services sector in the UK, it is therefore not surprising that aggregate log TFPR is still below its pre-2007 level in 2014 in Figure 3.4.

On aggregate, both revenue elasticity and productivity has increased in the U.K. between 1998 - 2014. However in order to disentangle the relationship, I look at the relationship at the 2-digit sector and firm-level. As discussed in Section 3.3.3, the theoretical link between productivity and revenue elasticity is not straightforward. For a given fixed cost ϕ, the relationship should be negative. If fixed costs are unchanged within firms, we should expect a negative relationship between revenue elasticity and TFPR *within firms*. However, fixed costs may not be unchanged. Furthermore, higher revenue elasticity may lead to a higher productivity cut-off.

I estimate time-varying revenue elasticity and productivity for each firm. This allows the computation of regressions on this relationship at the firm level, with year and firm fixed effects, as well as removing outlier observations. These results present a negative relationship: industries with greater scale also have higher TFPR. Table 3.6 reports the results from a firm-level panel regression of revenue elasticity on log productivity. I include year fixed-effects to control for the trends in log TFPR and revenue elasticity over time. I include firm fixed-effects to control for time-invariant unobserved firm-specific characteristics that may be correlated with productivity.

Table 3.6 presents evidence that firm-level revenue elasticity is negatively related to productivity. The first column simply describes a negative correlation between these two variables. Moving along the columns, I add year fixed effects, firm fixed effects, and then both fixed effects. The estimated coefficient falls in magnitude, but stays negative and statistically significant. The final column removes outlier firms (in terms of revenue elasticity), and this does not materially affect the result. Table 3.6

Table 3.6: *Regression: Revenue Elasticity on Productivity at Firm Level*

	Dependent variable: Revenue Elasticity				
Log TFPR	-0.055^{***}	-0.055^{***}	-0.022^{***}	-0.019^{***}	-0.020^{***}
	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)
N	449,484	449,484	278,391	278,391	271,357
Year FE:		✓		✓	✓
Firm FE:			✓	✓	✓
Remove outliers:					✓

*Note: Estimates statistically significant at levels of 0.1%: ***, 1%: **, 5%: *. Robust standard errors in brackets, clustered at the level of the fixed effects included. Revenue elasticity and log TFPR estimated following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function. Regressions weighted by value-added at the firm level. Outliers are the top and bottom 1% of firms by revenue elasticity.*

provides evidence that increases in productivity are associated with a fall in revenue elasticity at the firm level.

This confirms the prediction from the endogenous revenue elasticity theory: revenue elasticity is falling in productivity, if fixed costs are held constant. I find that *within firms*, those that become more productive experience falling revenue elasticity over time. The interpretation of the coefficient in the final column is that a 10% increase in TFP is associated with a fall in revenue elasticity of 0.002. I find that *within firms*, those that become more productive experience falling revenue elasticity over time.

However, there is evidence that revenue elasticity and productivity are positively correlated *across industries*. This is in line with findings by Gao and Kehrig (2020), and standard static models of heterogeneous firms with a productivity cut-off and exogenous returns to scale. I also find this result, shown in Table 3.7 for translog revenue elasticity estimates averaged at the 2-digit SIC level. Industries with higher average productivity also exhibit higher average revenue elasticity. This result holds with 2-digit SIC and year fixed effects, as well as removing outlier observations of revenue elasticity. Thus, the final two columns control for the fact that some industries generally exhibit higher revenue elasticity, and the trend towards increased revenue elasticity in the UK over time. I also confirm that there is positive relationship between

Cobb-Douglas revenue elasticity and log TFPR in Table C.16, although the sample is much smaller as there is just one estimate per sector.

Table 3.7: Regression: Revenue Elasticity on Productivity at Industry Level

	Dependent variable: Revenue Elasticity				
Mean log TFPR	0.014	0.015^{*}	0.145^{***}	0.134^{***}	0.128^{***}
	(0.009)	(0.009)	(0.017)	(0.019)	(0.019)
N	1,036	1,036	1,036	1,036	1,002
2-digit SIC FE:			✓	✓	✓
Year FE:		✓		✓	✓
Remove outliers:					✓

*Note: Estimates statistically significant at levels of 1%: ***, 5%: **, 10%: *. Robust standard errors in brackets, clustered at the level of the 2-digit SIC. Revenue elasticity and log TFPR estimated following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function, and aggregated to the 2-digit SIC level. Regressions weighted by number of firms in each 2-digit SIC and year. Outliers are sector-year estimates of revenue elasticity in the top and bottom 1%.*

Therefore, we face a paradox! Within firms, becoming more productive is associated with a fall in revenue elasticity. Across industries, this relationship is flipped. And at the aggregate level, a slowdown in the productivity *growth* has coincided with a gradual rise in revenue elasticity. This suggests that reallocation across firms has had significant aggregate implications. I leave deeper exploration of this mechanism to forthcoming research.

3.5.3 Relationships with Computer Software

Computer software investment data is obtained from the ARDx. I use software share in total investment as the preferred measure of 'software intensity', as it represents the extent to which new capital purchases are in low marginal cost technology.

Software investment shares are computed as computer software investment divided by total investment. All values outside of the range zero to one are removed.

This gives 292,706 firm-year observations of investment in computer software over this period. The average value of computer software investment share is 0.095 with a standard deviation of 0.17.

Computer software investment share varies substantially by sector. At the macro sector level, Services firms invest the most in software at 14.79% of total investment. Manufacturing and Construction firms have the lowest share of investment in software, at 6.02% and 6.16% respectively. For firms in Wholesale/Trade/Transport, the investment share in computer software sits at 8.47%. Computer software investment shares at the 2-digit SIC level is contained in Table C.17. Figure 3.5 presents the variation in software investment share across firms, with the vast majority below 20%, but there are a few firms investing very large shares in computer software.

Figure 3.5: Histogram of computer software investment share over all firms from 1998 to 2014.

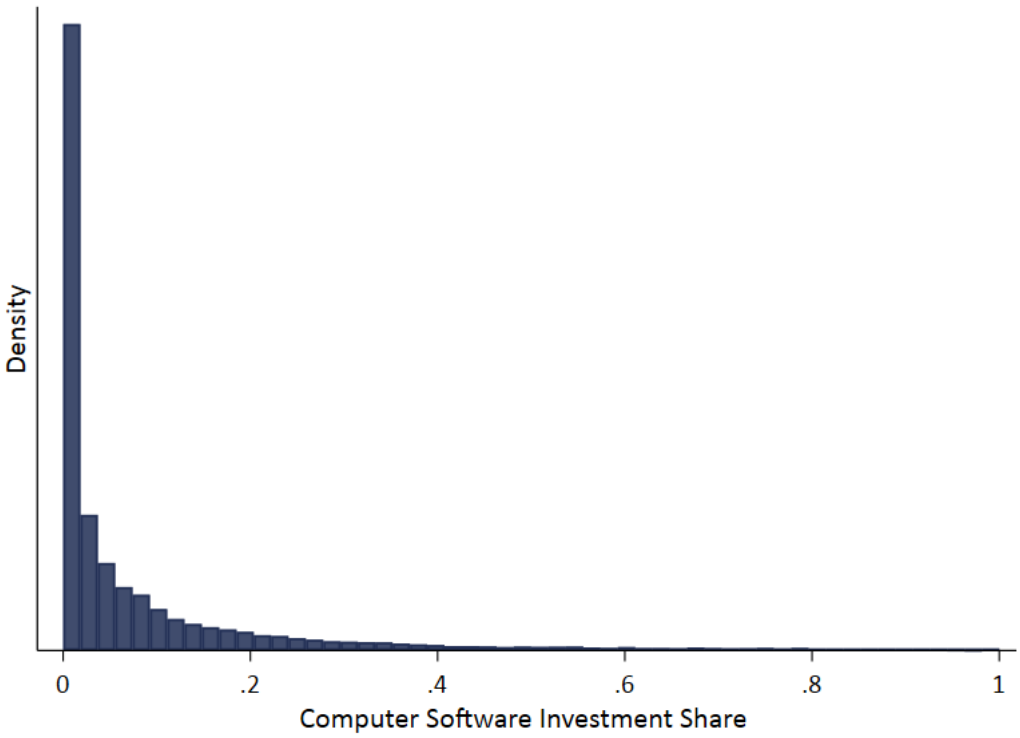

As a starting point, I plot revenue elasticity against computer software investment shares across 2-digit sectors. Figure 3.6 shows a weak positive relationship, but there is little evidence here that revenue elasticity and computer software are strongly related.

Figure 3.6: Revenue Elasticity and Computer Software Investment Share

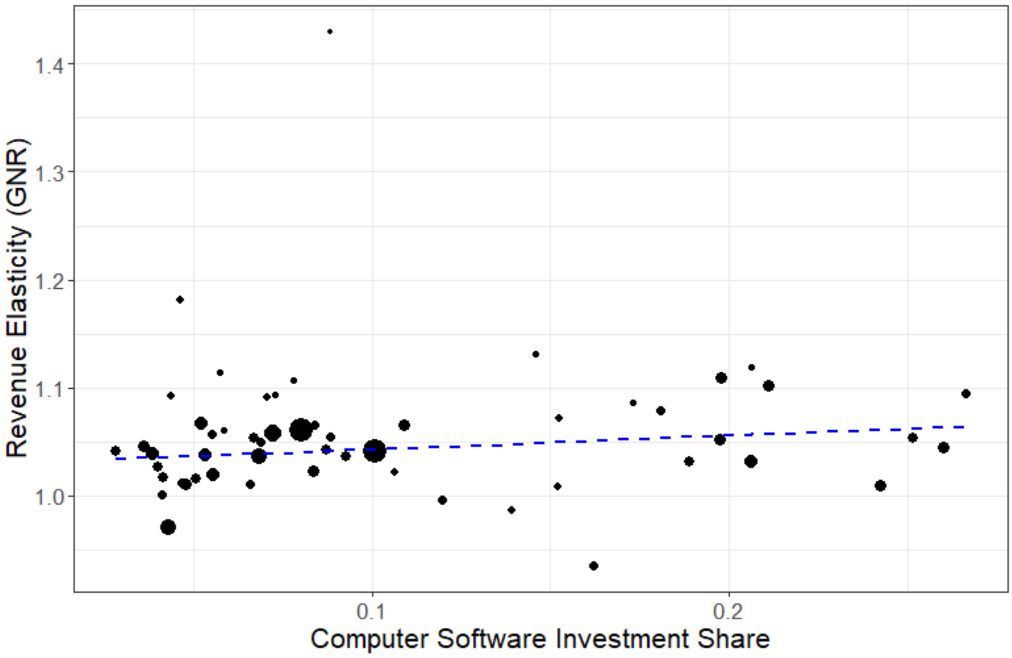

Scatter of revenue elasticity at the 2-digit SIC, estimated with a gross-output Cobb-Douglas production function following Gandhi, Navarro, and Rivers (2020), against average computer software investment share over 1998 - 2014. Points represent number of firms in each sector.

In order to test the two propositions of the model, I investigate the relationships between revenue elasticity, computer software investment, and productivity at the firm level. The subsequent fixed effect regressions use estimated revenue elasticity following Ackerberg, Caves, and Frazer (2015) with a value-added translog production function, to obtain time-varying revenue elasticity, at the 2-digit SIC and firm levels. All regressions are weighted, include robust clustered standard errors, and I show the impact of removing outliers for firm-level regressions.

The theory presented in this paper indicates that revenue elasticity increases when firms invest in computer software. I run panel regressions of revenue elasticity on computer software investment share, with fixed effects. Table 3.8 presents the results at the firm level. The evidence indicates that firms with greater computer software investment tend to have higher revenue elasticity.

The first two columns of Table 3.8 show a strong positive relationship between firm-level revenue elasticity and computer software investment share, with and with-

Table 3.8: Regression: Revenue Elasticity on Software Investment

	Dependent variable: Revenue Elasticity				
Software Investment Share	0.0342***	0.0315***	0.0039**	0.0027**	0.0024*
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
N	215,862	215,862	166,920	166,920	163,705
Year FE:		✓		✓	✓
Firm FE:			✓	✓	✓
Remove outliers:					✓

*Note: Estimates statistically significant at levels of 1%: ***, 5%: **, 10%: *. Robust standard errors in brackets, clustered at the level of the fixed effects included. Revenue elasticity estimated following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function. Regressions weighted by value-added at the firm level. Outliers are the top and bottom 1% of firms by revenue elasticity.*

out year fixed effects. Once I control for firm fixed effects in the third column, the magnitude of the estimated coefficient falls sharply. However even controlling for year *and* firm fixed effects, there is still some evidence of a strong relationship between software and revenue elasticity. The final column of Table 3.8 shows that removing 'outlier' firms (by revenue elasticity) reduces the strength of this relationship to some extent. In other words, firms with very high or very low revenue elasticity are important in driving this positive association.

I also split the data halfway in 2007, to investigate the relationship between revenue elasticity and computer software in each half of the sample. Table C.18 presents evidence that the association is much stronger in the second half of the sample, both in magnitude and statistical significance. The role of software in raising returns to scale in the revenue function has increased over time. One interpretation is that computer software has got better at scaling up production, by keeping marginal costs low and allowing for replication of revenue-generating activities.

The model also suggests that computer software investment provides an additional channel to link revenue elasticity and productivity. I regress revenue elasticity on the log of TFPR and the computer software investment share. The results are presented in Table 3.9. When productivity is included, the positive relationship between revenue elasticity and computer software is weakened. This suggests that, conditional on a firm's level of TFPR, computer software is not associated to revenue elasticity. This is evidence against the second proposition of the model.

Table 3.9: Regression: Revenue Elasticity, Productivity, and Software Investment

	Dependent variable: Revenue Elasticity				
Software Investment Share	0.0229***	0.0206***	0.0032**	0.0021	0.0019
	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)
$\ln TFPR$	−0.070***	−0.069***	−0.025***	−0.023***	−0.021***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
N	215,862	215,862	166,920	166,920	163,705
Year FE:		✓		✓	✓
Firm FE:			✓	✓	✓
Remove outliers:					✓

*Note: Estimates statistically significant at levels of 1%: ***, 5%: **, 10%: *. Robust standard errors in brackets, clustered at the level of the fixed effects included. Revenue elasticity and log TFPR estimated following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function. Regressions weighted by value-added at the firm level. Outliers are the top and bottom 1% of firms by revenue elasticity.*

3.6 Discussion

I describe the theory to help understand scale economies, which describe the relationship between firm inputs and firm outputs. I distinguish between returns to scale in output and revenue, which is simply the sum of output and revenue elasticities respectively. With a fixed cost in the net output function, I show that both output and revenue elasticity is declining in firm productivity, in a framework with constant markups across firms.

I extend the theoretical framework to allow firms to invest in software which scales down costs, and an associated fixed cost (De Ridder 2019; Altomonte, Favoino, Morlacco, and Sonno 2021). With this extension, firms solve the standard two-stage problem to minimise costs and maximise profits. I show that revenue elasticity is greater in an environment with computer software, and the relationship between scale economies and productivity is moderated by software.

This paper's contribution to understanding revenue elasticity in the UK is threefold: (1) I present updated estimates of revenue elasticity across a services-dominated economy from 1998 - 2014, and report results at various levels of aggregation, (2) I present evidence of a rise in revenue elasticity over time, (3) I investigate the relationship between revenue elasticity and computer software.

I find that most industries have revenue elasticity close to unity, but there are some sectors with revenue elasticity as low as 0.53 and up to 1.52. I report a rise in revenue elasticity over time. I show a strong negative relationship between revenue elasticity and productivity within firms over time. However, across industries, the relationship is positive. These results are consistent with my theoretical framework. Firstly, the output-denominated fixed cost leads to an endogenous revenue elasticity that depends on firm size. Secondly, a standard productivity cut-off exists when fixed costs vary across industries. Finally, I present evidence that revenue elasticity has risen more rapidly where more computer software has been adopted. This relationship is stronger in the latter half of the sample, perhaps as software has become even more capable of allowing firms to scale up more rapidly.

A Appendix - Chapter 1

A.1 Appendix - Tables & Figures

Table A.1: Robots Per Thousand Workers by Industry

	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

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

Table A.2: Robots & Patents: Summary Statistics

	Mean	Stdev	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

Note: Summary statistics which describe the robot and automation variables data across the 348 Local Authorities in England and Wales.

Table A.3: *Employment Summary 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	0.012	−0.037	−0.019	−0.062
(1991 - 2011)	(0.02)	(0.03)	(0.02)	(0.03)
Observations	348	348	348	348

Note: Summary statistics which 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.

Table A.4: *Employment Summary 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

Note: Summary statistics which 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.

Table A.5: Income Summary Statistics

	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

Note: Summary statistics which 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.

Table A.6: Earnings Inequality Summary Statistics

	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	0.015	−0.018	0.003
(1997 - 2011)	(0.056)	(0.083)	(0.087)
Observations	242	332	242

Note: Summary statistics which 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.

Table A.7: Control Variables Summary Statistics

	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

Note: Summary statistics which describe the covariates data across the 348 Local Authorities in England and Wales.

Table A.8: 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 A.9: 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 A.10: 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 A.11: *p-values for 2SLS tests for all dependent variables*

Dependent Variable	Weak Instruments	Sargan	Wu-Hausman
Δ E/Pop	0.00	0.06	0.58
Δ FTE/Pop	0.00	0.21	0.36
Δ E/WA	0.00	0.07	0.54
Δ FTE/WA	0.00	0.29	0.37
Δ mean Full-Time hours	0.00	0.50	0.73
Δ mean Part-Time hours	0.00	0.84	0.15
Δ median Full-Time pay	0.00	0.49	0.44
Δ median Part-Time pay	0.00	0.26	0.94
Δ mean Full-Time pay	0.00	0.92	0.36
Δ mean Part-Time pay	0.00	0.13	0.75
Δ 80/50 ratio	0.00	0.92	0.99
Δ 50/10 ratio	0.00	0.20	0.56
Δ 80/10 ratio	0.00	0.24	0.47

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 A.12: *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	0.41	0.58
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	0.00	0.09
Low Tech Manuf.	0.00	0.00	0.01
High Tech Manuf. (incl. Auto)	0.00	0.08	0.00
Utilities	0.00	0.15	0.26
Construction	0.00	0.93	0.89
Services	0.00	0.34	0.00
Education	0.00	0.51	0.82
Agriculture	0.00	0.03	0.03
Mining	0.00	0.70	0.81

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 A.13: 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.

Figure A.1: Impact of Exposure to Robots on Industry Employment (1991 - 2011), removing regions most-exposed to Robots

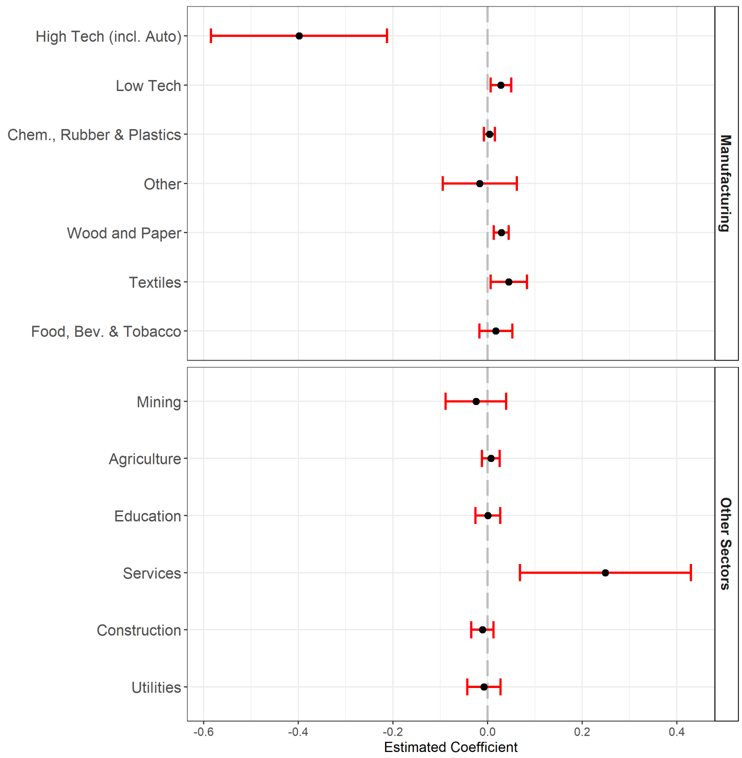

Note: Estimated 2SLS coefficients and 95% confidence interval for Exposure to Robots on change in employment shares ratio (1991-2011) across a set of industries, conditional on the full set of controls used in baseline estimation. Instruments are Lagged EU Exposure to Robots and Exposure to Automation-Related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

Figure A.2: Impact of including Chinese Trade control

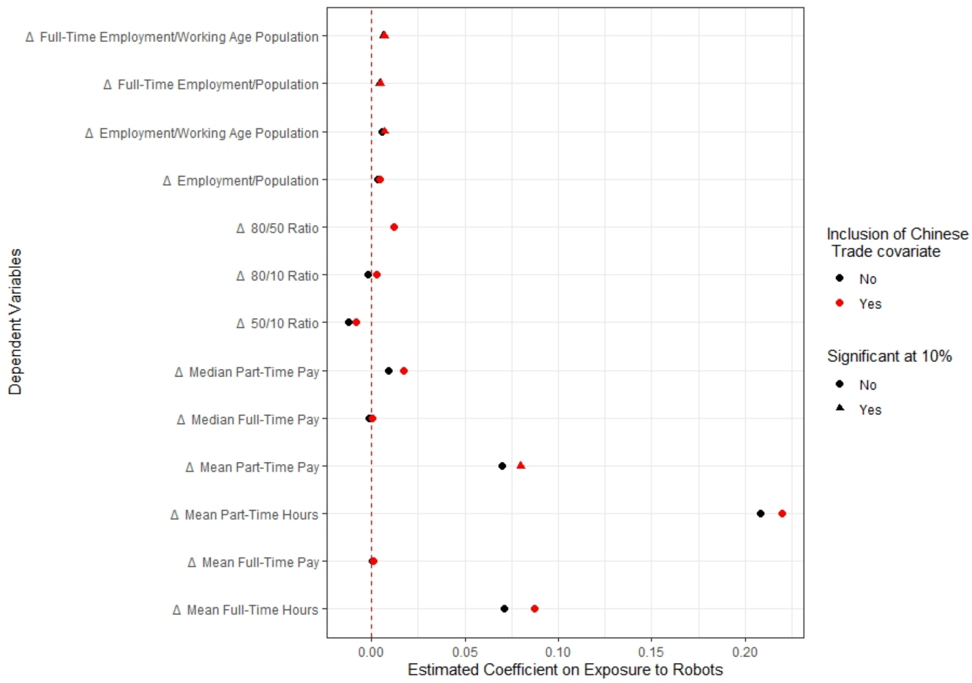

Note: Estimated 2SLS coefficients on Exposure to Robots on models with all controls, when including and excluding covariate which controls for Chinese Trade. Instruments are Lagged EU Exposure to Robots and Exposure to Automation-Related Patents.

A.2 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 (2020) 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 (2020).

The general equilibrium impact is given by:

$$d\ln L_c = -\frac{1+\eta}{1+\epsilon}\sum_{i\in\mathcal{I}}\ell_{ci}\frac{dM_i}{1-M_i} + \frac{1+\eta}{1+\epsilon}\pi_c\sum_{i\in\mathcal{I}}\ell_{ci}\frac{s_{icL}}{s_{cL}}\frac{dM_i}{1-M_i} \tag{19}$$

$$d\ln W_c = -\eta\sum_{i\in\mathcal{I}}\ell_{ci}\frac{dM_i}{1-M_i} + (1+\eta)\pi_c\sum_{i\in\mathcal{I}}\ell_{ci}\frac{s_{icL}}{s_{cL}}\frac{dM_i}{1-M_i} \tag{20}$$

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 (19) 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 (20) 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:

$$d\ln L_c^d = -\sum_{i\in\mathcal{I}} \ell_{ci} \frac{dM_i}{1-M_i} - \lambda \sum_{i\in\mathcal{I}} \ell_{ci} d\ln P_{Xci} + (\lambda - \sigma) \sum_{i\in\mathcal{I}} \ell_{ci} d\ln P_{Yi} + d\ln Y \tag{21}$$

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 (21)) is different from autarky (equation (1.1)). We still have the same displacement effect in equation (21) as in equation (1.1). 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.

A.3 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.

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.[17] 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).[18]
- 2011 WF01EW - commuting flows across (merged) UK Local Authorities.[19]
- 1991 SWS Set C (2001 geog.) - commuting between 2001-defined Standard Table wards.

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

[18] 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.

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

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 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.[20]

Clustering of Wards

I follow Tolbert and Sizer (1996) in performing hierarchical clustering on wards.[21] 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

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

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

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.

Table A.14: Results from Hierarchical Clustering on Commuting Data

	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

Note: 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.

A.4 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.

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 A.2 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.

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.[22]

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 elements. Table A.10 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 A.10). 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.

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 A.9, for OLS and 2SLS estimation with the full set of controls.

The impact on the extensive margin of employment is stark: all OLS estimates

[22]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.

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.

A.5 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.[23]

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).

[23]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.

Figure A.3: Impact of Exposure to Robots on Occupation Employment (1991 - 2011)

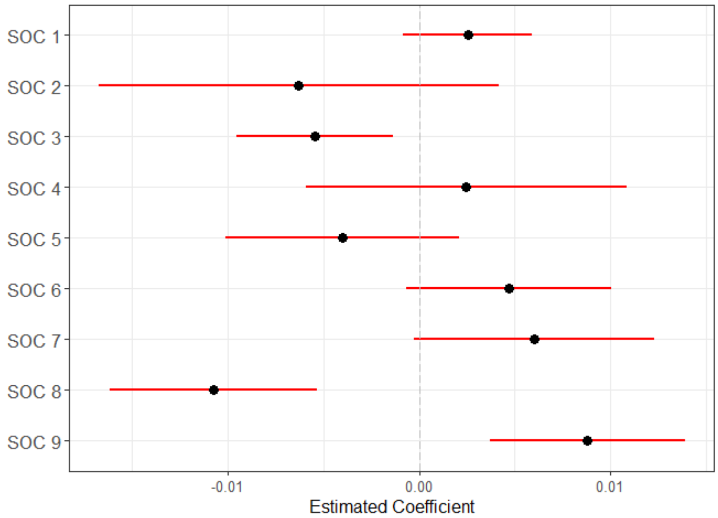

Note: The points are estimated 2SLS coefficients of Exposure to Robots on the change in employment shares (1991 - 2011) across a set of occupations. The lines indicate the 95% confidence intervals. The models have the full set of controls used in baseline estimation. SOC 1: Managers and Senior Officials. 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.

A.6 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 A.15 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 A.17.

Table A.15: Industry Mapping for Robots

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

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

Table A.16: Industry Mapping for Trade

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	-

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

Table A.17: Occupational Mapping

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*	+/−

Note: Mapping Standard Occupation Classifications (SOC) to Autor and Dorn (2013) Routine Task Intensity

A.7 Appendix - Patent Search

On the WIPO IP Portal, the PATENTSCOPE database search is used.[24] 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, Raffo, and Andrew Keisner (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."[25] 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.[26] The next step was to match the NACE Rev. 2 to ISIC Rev. 4 using the Eurostat correspondence table.[27] It was now straightforward to link to the industries in set $\mathcal{I}$ using Table A.15.

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

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

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

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

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

variable could be computed by summing patents across industries in a given year:

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

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

B Appendix - Chapter 2

B.1 Appendix - Survey Details

This section includes further details on the Bank of Italy's "Survey of Industrial and Service Firms".[28]

Survey Design

The survey collects annual data on firm employment, investment, turnover, debt, and other structural information. It also has further detail on specific issues, which may not be asked annually, such as strategies, governance, technological and organisational factors. Up until 1998, only manufacturing firms with more than 50 employees were covered, but this was expanded in 1999 to include extractive and energy firms. Firms with more than 20 workers were included from 2002.

The survey population is divided into strata and firms are chosen randomly from each for the sample (one-stage stratified sample). Each strata are defined by economic activity, firm size, and region of head office. Employment is measured as the average number of workers during the year. Firms included in the previous year of the survey are always contacted, but may not be kept in the sample if they change activity class or fall below the threshold number of employees (they are replaced by firms of the same industry and size).

The data undergoes extensive quality checks: answers must fit within the range for the question; the panel data must be consistent; outliers must be checked. Statistical methods such as 'selective editing' are used to summarise the impact of outliers, to reduce the requirement to re-contact firms.

[28] https://www.bancaditalia.it/pubblicazioni/metodi-e-fonti-note/metodi-note-2017/en_survey_methodology_invind.pdf?language_id=1

Advanced Technologies Questions

The automation technology questions are contained here in full. In 2015, firms were asked if they use or intend to adopt the following advanced technologies:

1. Mobile broadband and the cloud (e.g. wireless technology, apps, smartphones, tablets, high-speed broadband, and cloud management software)

2. Artificial intelligence and big data (e.g. collection and use of large data sets that, with the application of specific algorithms for machine learning, can provide support to decision making; possible applications are: in remote access diagnostics, defining algorithms for financial investments, patent-related or legal searches)

3. The internet of things (e.g. the use of technologies that, by means of advanced sensors, allow apparatus to be used in the production and commercial processes promoting their integration)

4. Industrial robotics using artificial intelligence (advanced robotics)

5. 3D printing

In 2017, firms were asked if they use or intend to adopt the following advanced technologies:

1. Cloud computing

2. E-commerce

3. Big data (e.g. the collection and use of large quantities of data which, also through machine learning algorithms, can assist decision-making; possible applications: distance diagnosis, financial trading algorithms, patent and legal research)

4. Internet of things (e.g. the use of technologies that, by means of advanced sensors,enable communication between the various devices used in production and business processes, facilitating their integration)

5. Artificial intelligence

6. Industrial robotics using artificial intelligence (advanced robotics)

7. 3D printing

In 2019, firms were asked if they use or intend to adopt the following advanced technologies:

1. Cloud computing

2. Big data

3. Artificial intelligence

4. Advanced robotics

5. 3D printing

Firm Exit

There is some attrition in the panel, as firms leave and are replaced. "The firms observed in the previous edition of the survey are always contacted again if they are still part of the target population, while those no longer wishing to take part are replaced with others in the same branch of activity and size class."[29] The main reasons given for leaving the survey are change of activity and staff cutbacks to below the entry threshold. These numbers are somewhat higher than estimates of the Italian exit rate of 4 -8 % (Carree, Santarelli, and Verheul 2008).

[29] https://www.bancaditalia.it/pubblicazioni/metodi-e-fonti-note/metodi-note-2017/en_survey_methodology_invind.pdf?language_id=1

Table B.1: Italian Survey Data: Firm Exit Proportions

Year	% in subsequent year	% in preceding year
2012	0.82	-
2013	0.84	0.82
2014	0.86	0.83
2015	0.84	0.84
2016	0.86	0.87
2017	0.85	0.83
2018	-	0.88

Note: the second column contains the proportion of firms appearing in that year which are present in the panel the following year. The third contains the proportion of firms appearing in that year which were present in the panel the previous year.

B.2 Appendix - Firm Automation Adoption

Advanced Technology Adoption Firm Shares 2015

Table B.2: Number of Firms Adopting Technologies (2015)

	Mobile & Cloud	AI & Big Data	IoT	Robots	3D Printing
Yes	2,524	440	642	485	360
No	934	2,953	2,742	2,912	3,037
NA	937	1,002	1,011	998	998
% Yes (all)	*0.57*	*0.10*	*0.15*	*0.11*	*0.08*
% Yes (answered)	*0.73*	*0.13*	*0.19*	*0.14*	*0.11*

Note: Answers from firms on use of advanced technologies in 2015, along with proportion of firms adopting each technology (of all firms, and of firms that were asked).

Table B.3: Share of Firms Adopting Individual Technologies (2015)

	Mobile & Cloud	AI & Big Data	IoT	Robots	3D Printing
All firms	0.56	0.09	0.12	0.06	0.05
Answered firms	0.73	0.11	0.16	0.08	0.07

Note: The weighted proportions of firms that use advanced technologies in 2015.

Table B.4: Share of Firms Adopting Technologies (2015)

	Any Tech.	Digital Tech.	Physical Tech.
All firms	0.23	0.16	0.10
Answered firms	0.30	0.21	0.13

Note: The weighted proportions of firms that use combinations of advanced technologies in 2015.

Advanced Technology Adoption Firm Shares 2017

Table B.5: Number of Firms Adopting Technologies (2017)

	Cloud	AI	Big Data	IoT	Robots	3D Printing	E-Comm.
Yes	849	226	416	733	477	375	871
No	2,976	3,591	3,404	3,092	3,349	3,443	2,974
NA	566	574	571	566	565	573	546
% Yes (all)	*0.19*	*0.05*	*0.10*	*0.17*	*0.11*	*0.09*	*0.20*
% Yes (answered)	*0.22*	*0.06*	*0.11*	*0.19*	*0.12*	*0.10*	*0.23*

Note: answers from firms on use of advanced technologies in 2017, along with proportion of firms adopting each technology (of all firms, and of firms that were asked).

Table B.6: Share of Firms Adopting Individual Technologies (2019)

	Cloud	AI	Big Data	IoT	Robots	3D Printing	E-Comm.
All firms	0.18	0.03	0.06	0.11	0.07	0.05	0.19
Answered firms	0.21	0.03	0.07	0.13	0.08	0.06	0.22

Note: weighted proportions of firms that use advanced technologies in 2017.

Table B.7: Share of Firms Adopting Technologies (2017)

	Any Tech.	Digital Tech.	Physical Tech.
All firms	0.28	0.25	0.10
Answered firms	0.33	0.29	0.11

Note: weighted proportions of firms that use combinations of advanced technologies in 2017.

Advanced Technology Adoption Firm Shares 2019

Table B.8: Number of Firms Adopting Technologies (2019)

	Cloud	**AI**	**Big Data**	**Robots**	**3D Printing**
Yes	633	149	293	317	192
No	1,440	1,916	1,778	1,751	1,877
NA	1,116	1,124	1,118	1,121	1,120
% Yes (all)	*0.20*	*0.05*	*0.09*	*0.10*	*0.06*
% Yes (answered)	*0.31*	*0.07*	*0.14*	*0.15*	*0.09*

Note: answers from firms on use of advanced technologies in 2019, along with proportion of firms adopting each technology (of all firms, and of firms that were asked).

Table B.9: Share of Firms Adopting Individual Technologies (2019)

	Cloud	**AI**	**Big Data**	**Robots**	**3D Printing**
All firms	0.18	0.03	0.07	0.04	0.03
Answered firms	0.28	0.04	0.11	0.07	0.05

Note: weighted proportions of firms that use advanced technologies in 2019.

Table B.10: Share of Firms Adopting Technologies (2019)

	Any Tech.	**Digital Tech.**	**Physical Tech.**
All firms	0.24	0.20	0.06
Answered firms	0.37	0.32	0.10

Note: weighted proportions of firms that use combinations of advanced technologies in 2019.

Pairwise Technology Adoption

Table B.11: Adoption of Multiple Technologies (2015)

	Mobile & Cloud	**AI & Big Data**	**IoT**	**Robots**	**3D Printing**
Mobile & Cloud	1.00	0.39	0.44	0.37	0.35
AI & Big Data	-	1.00	0.83	0.81	0.82
IoT	-	-	1.00	0.80	0.79
Robots	-	-	-	1.00	0.84
3D Printing	-	-	-	-	1.00

Note: proportion of technology adopters adopting another technology in 2015.

Table B.12: Adoption of Multiple Technologies (2017)

	Cloud	**AI**	**Big Data**	**IoT**	**Robots**	**3D Printing**	**E-Comm.**
Cloud	1.00	0.79	0.81	0.78	0.76	0.77	0.76
AI	-	1.00	0.90	0.83	0.88	0.88	0.77
Big Data	-	-	1.00	0.84	0.84	0.85	0.77
IoT	-	-	-	1.00	0.81	0.81	0.72
Robots	-	-	-	-	1.00	0.86	0.72
3D Printing	-	-	-	-	-	1.00	0.75
E-Comm.	-	-	-	-	-	-	1.00

Note: proportion of technology adopters adopting another technology in 2017.

B.3 Appendix - Analysis of Technology Adoption

The following section contains detailed analysis on adopters of automation technologies contained in the Bank of Italy's "Survey of Industrial and Service Firms".

There are five distinct automation technologies in the survey: Mobile & Cloud Computing, Artificial Intelligence & Big Data, Internet of Things (IoT), Industrial Robotics, 3D Printing. Firms are asked about adoption of each of these technologies in 2015, 2017, and 2019 (except IoT). In this section I look at each technology separately and investigate the firm characteristics associated with adoption. Further figures and tables are included at the end of this section.

Mobile and Cloud Computing

In 2015, the survey combined Mobile Broadband and Cloud Computing together, whereas the latter was isolated in 2017 and 2019. Mobile Broadband is quite a widely-defined technology, and the survey notes this by giving specific examples encompassing wireless technology, high-speed internet and mobile devices (see Appendix B.4). Cloud Computing is not narrowly defined in the survey, and can encompass a range of technologies such as software (e.g. Google Docs), platforms (e.g. Microsoft Azure), and data storage (e.g. Amazon S3) (Bayrak, Conley, and Wilkie 2011).

I find that a large proportion of Italian firms use Mobile and Cloud Computing: 73% use this technology in 2015. For Cloud Computing specifically, it is unsurprisingly much lower, at 21% in 2017. This rises to 28% by 2019, and is the most-adopted advanced technology in this survey. It is more common in larger firms, but adoption doesn't seem to vary with firm age.[30] The strong positive relationship between use and size is even more stark when just looking at Cloud Computing in 2017 and 2019, and this relationship holds across all age groups.

The sectoral distribution of Cloud Computing is relatively equal, as can be seen in

[30]This actually contradicts one of the few firm-level analyses of Cloud adoption, which suggests that firm size does not matter (Benlian, Hess, and Buxmann 2009). However, this study is limited to small- and medium-sized firms, with just one hundred firms, based in a substantially different economic environment in India.

Figure B.1: Cloud Computing Adoption by Firm Size

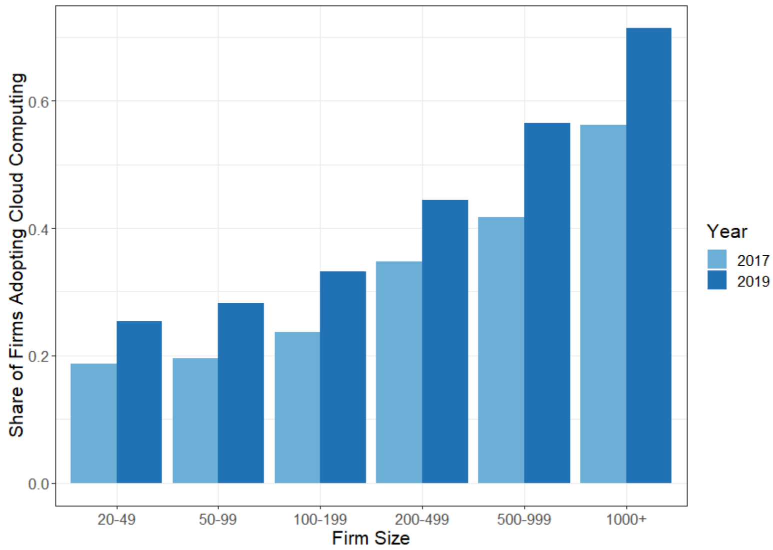

Figure B.2. For example, while approximately 21% of firms use Cloud Computing in 2017, the sectors in which it is most commonly adopted - Real Estate and Transport & Communications - are at around 30%. Cloud adoption sits between 15 - 20% for most other sectors, with low uptake only in Hotels & Restaurants. Cloud Computing is the only technology in this survey more commonly adopted in Services firms compared to other industries.

Figure B.2: Cloud Computing Adoption by Industry (2017)

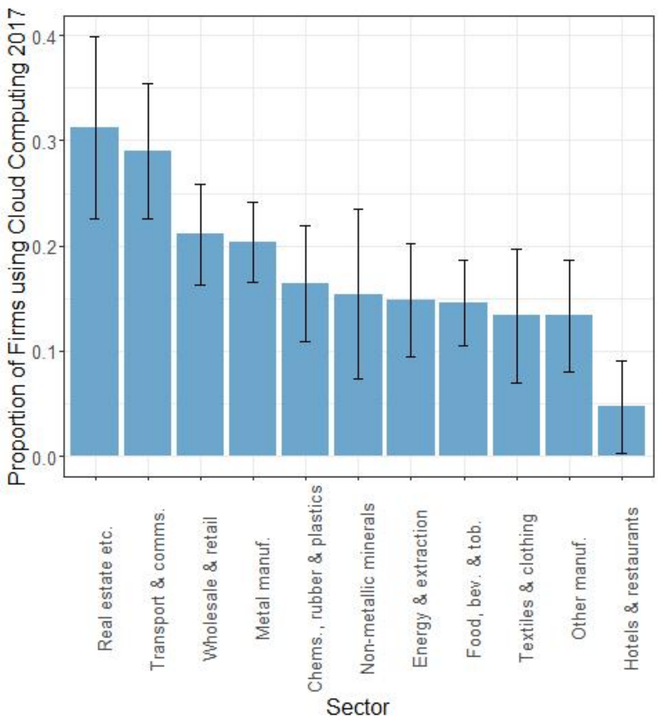

Artificial Intelligence and Big Data

Artificial Intelligence (AI) is "the capability of a machine to imitate intelligent human behavior" (Merriam Webster, 2020) and has been subject to a lot of discussion within the economics profession. It is sometimes used interchangeably with 'automation'. For example, Aghion, B. Jones, and C. Jones (2017) study the impact of AI on long-run growth, extending the Zeira (1998) model (which itself considers 'machines' that are more akin to robots) and they regularly employ the language of tasks being 'automated' or not. However, AI may not *always* simply be an automation technology, and may need to be combined with robots, big data, cloud computing, or other technologies, in order to automate a task. It seems more reasonable to think of the common tasks of AI: decision-making, prediction, anomaly detection, among others (Babina, Fedyk, He, and Hodson 2020).

Big Data is the availability of real time, larger scale, more varied, unstructured data (Einav and Levin 2014). It is often reasonable to combine AI with Big Data as the former typically depends on the latter.

Firms that use AI & Big Data are younger, larger, have higher turnover per worker, and higher fixed investment. On average, firms adopting this technology in 2015 had 63 workers and were 33 years old, compared to 49 workers and a mean age over 35 for non-adopters. The size and age gap between adopters and non-adopters of AI & Big Data rose in the subsequent years: by 2019, the average age gap rose to almost four years, and the sizes were 66 to 49. Importantly, these differences are statistically significant.[31]

There has been a gradual rise in firms adopting this technology over time, rising from 13% in 2015 to 17% by 2019. When firms are binned by employment, the difference in adoption is stark, as can be seen in Figure B.3.

For firms with under 200 workers, between 10 - 20% of them use AI & Big Data. It sits between 20 - 40% of firms that employ over 200 workers. There has also been a

[31] At the 1% level, computed with Welch's t-test and the Welch-Sattherwaite equation for degrees of freedom.

Figure B.3: AI & Big Data Adoption by Firm Size

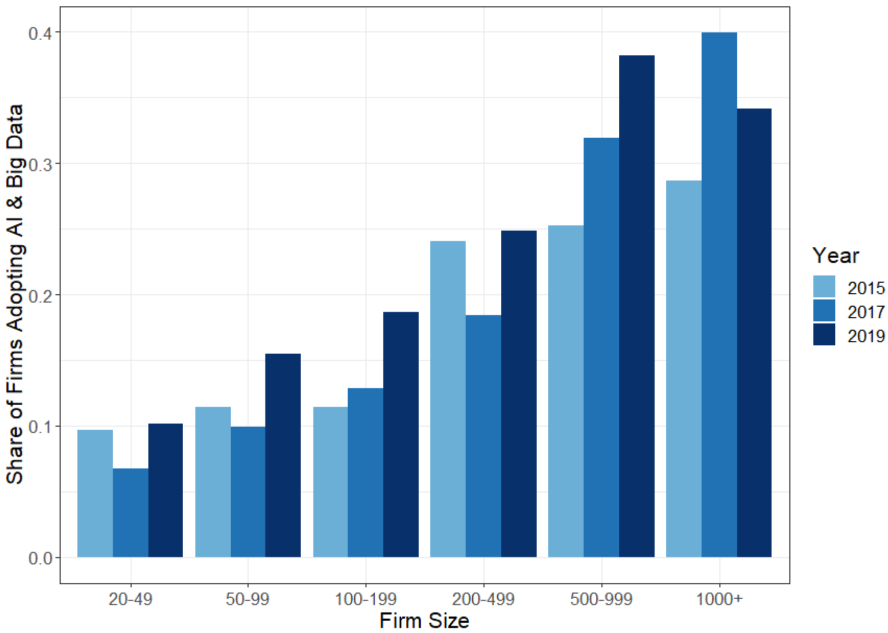

rise in firm adoption across the size distribution between 2015 and 2019, albeit with a slight dip in 2017 for some firms. Comparing the age distribution in the same way presents a less clear pattern, although there is some evidence that younger firms tend to adopt AI & Big Data.

The fact that a larger proportion of big firms use AI & Big Data also holds when the data is split by age and size. Figure B.4 shows the proportion of AI firms in age × size bins - it doesn't seem that AI differs by firm age, but it does vary by employment. This can be seen more clearly in Figures B.5, which show AI & Big Data use by employment (age), split into lines by age (employment).

AI & Big Data firms also have a much lower proportion of workers in blue collar jobs, with just 43% compared to 55% in firms that don't use this technology. In fact, of all technologies in this dataset, AI & Big Data firms have the lowest share of manual labourers by some margin (the next lowest is 50.1% for firms using Internet of Things). This could speak to the complementarity of high-skilled workers with this advanced technology.

Adoption of AI & Big Data is far more common in Services, such as Transport, Retail and Communications, ranging from 12 - 21%. On the flip side, adoption is below

Figure B.4: Age-size Distribution of Firms using AI & Big Data

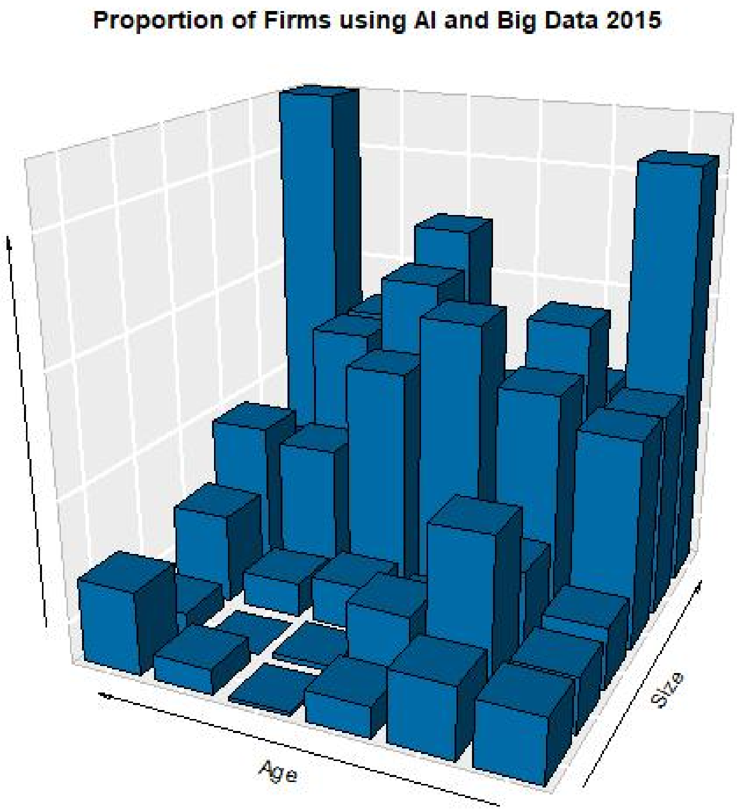

Note: each bar is ratio of firms using AI & Big Data in each age × size cell.

Figure B.5: Adoption of AI & Big Data by Firm Age and Size (2015)

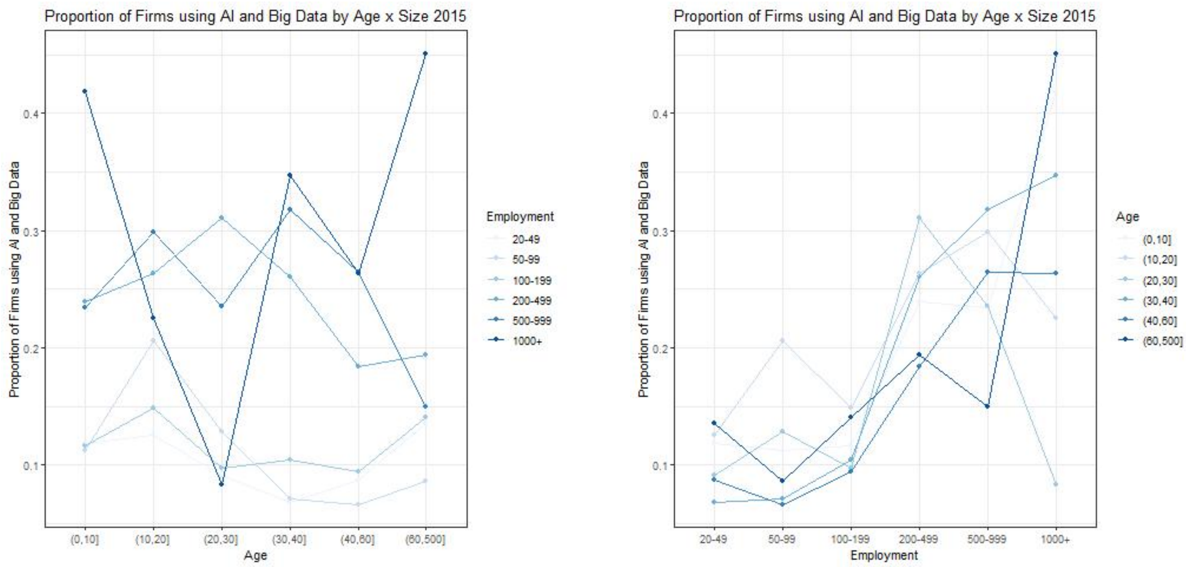

5% in Food, Drink, and Textiles manufacturing firms. Adoption across all industries can be found in Figure B.7.

Firms using AI & Big Data also have an unusual relationship with exporting. It is

Figure B.7: Industry Adoption of AI & Big Data (2015)

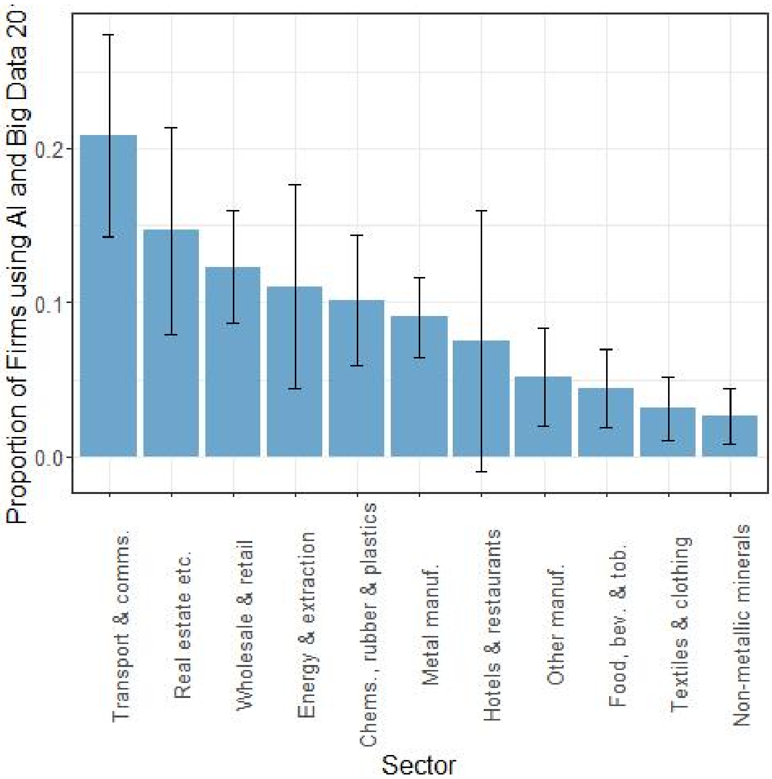

only this technology for which firms export significantly *less*, whether we compare by average share of sales in exports (see Table B.13) or looking at the distribution, as in Figure B.8.

Figure B.8: AI & Big Data Adoption by Export Intensity

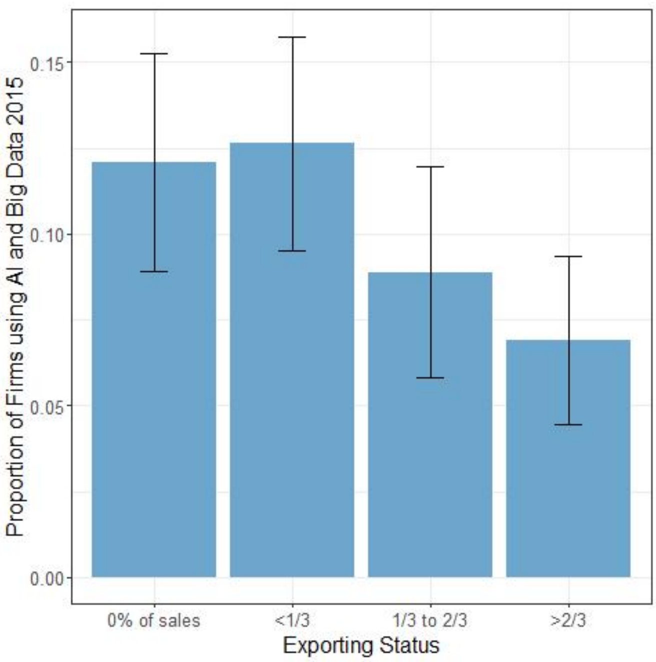

Internet of Things

The Internet of Things (IoT) is a set of technologies which connect machines and the internet. They are part of an infrastructure that connects sensors and actuators to computer systems via networks, and these systems can track the performance and actions of machines through the internet (Manyika, Chiu, Bisson, Woetzel, Bughin, and Aharon 2015). There is limited economic research on the impact of these technologies, although Edquist, Goodridge, and Haskel (2019) find significant relationships between TFP growth and the number of IoT connections per person. This country-level analysis likely suffers from endogeneity concerns, as the authors note, so can really only be interpreted in terms of high-level correlations. The firm-level survey - and panel component - allows me to investigate IoT with finer detail (across industries, firm age etc.) and hopefully alleviate concerns of endogeneity.

Firms using IoT technologies are indistinguishable in their age profile from those that do not adopt, both in the average age (see Table B.13) and age distribution (see Figure B.9).

Figure B.9: Internet of Things Adoption across Age and Size Distributions (2015)

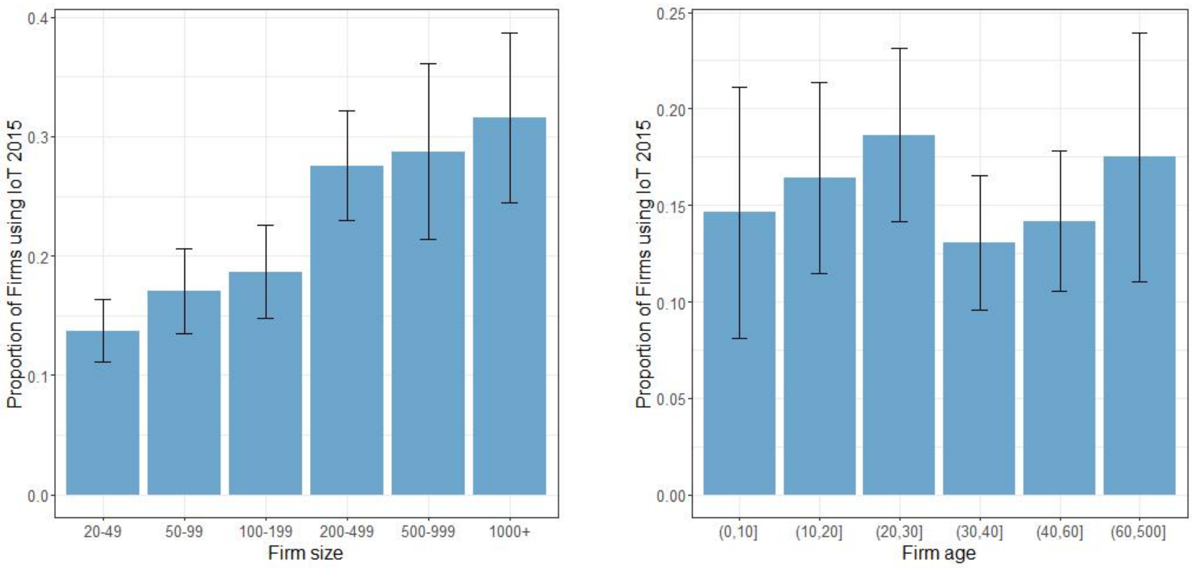

However, once again, it seems that larger firms are more likely to adopt this advanced technology: the average employment of IoT firms is almost 62 workers, compared to 48 for non-adopters. This pattern is consistent in 2015 and 2017, with the larger firms more likely to adopt IoT technology across all age groups, as in Figure B.10.

Figure B.10: Internet of Things Adoption by Firm Size and Age

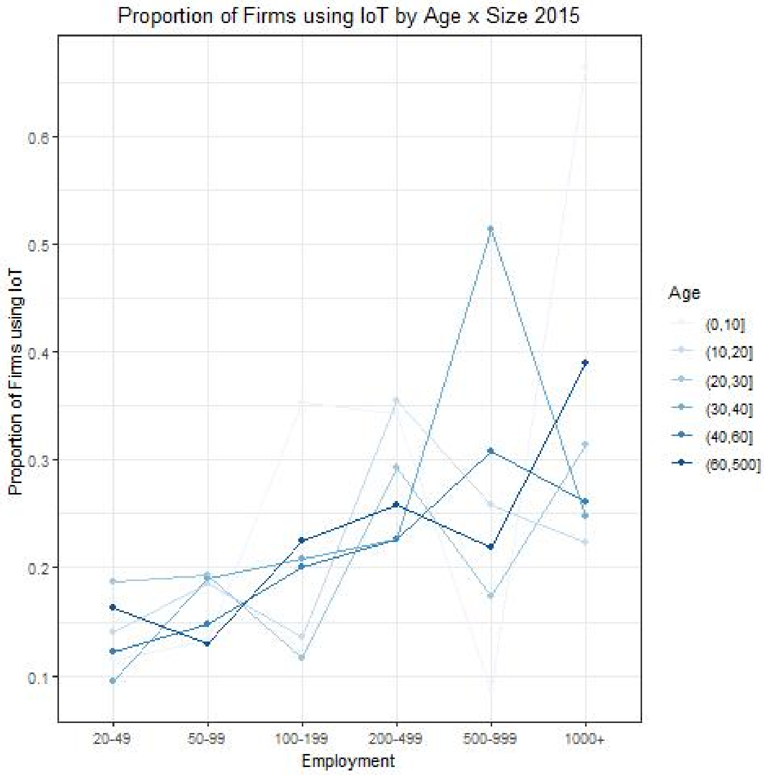

Furthermore, there are some noticeable differences in the characteristics of firms using IoT between 2015 and 2017, which may not be surprising for a novel technology. In 2015, there doesn't seem to be any pattern in IoT adoption by either exporting behaviour or macro industry. By 2017, it seems clear that a very small proportion of non-exporters use IoT technology compared to exporters. This can be clearly seen in Figure B.11. Furthermore, there is growth in IoT use among firms in manufacturing and mining, with little change in adoption for businesses in services.

Finally, there is differing IoT adoption by region, with a higher uptake for firms in the North-East. This likely reflects sectoral differences, with around 20% of those in Metal Manufacturing and Transport & Communication adopting the technology, compared to under 10% in Textiles Manufacturing and Hospitality. IoT firms have a significantly lower proportion of workers in blue-collar occupations, and exports contribute a smaller percentage of total sales, compared to non-adopters.

Figure B.11: Adoption of Internet of Things by Export Intensity

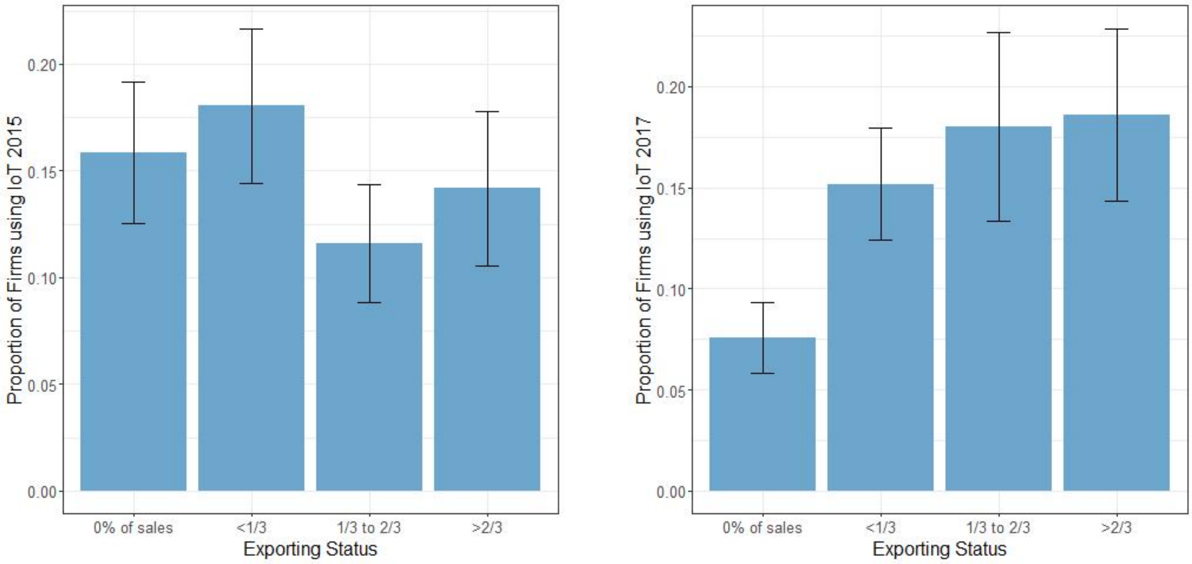

Industrial Robotics

Research on the adoption and impact of industrial robots is widespread relative to the other technologies in this paper. The findings are mixed, with evidence of displacement of jobs in the U.S., while there was no overall fall in employment in Germany, but job composition was affected (Acemoglu and Restrepo 2020; Dauth, Findeisen, Südekum, and Wößner 2017). However, these studies use quite aggregated robots data, which may mask more interesting behaviour occurring where the robots are purchased: the firms.

More recently, there has been firm-level research on the impact of robots (Acemoglu, Lelarge, and Restrepo 2020; Koch, Manuylov, and Smolka 2019; Cheng, Jia, D. Li, and H. Li 2019). There are some consistent findings, for example that robot-adopting firms tend to be larger, and that adopters tend to grow, while their non-adopting rivals shrink (Acemoglu, Lelarge, and Restrepo 2020; Koch, Manuylov, and Smolka 2019).

I find that adopters are significantly larger and older than firms that don't use robots. This pattern can be seen in Figure B.12, as the share of firms using robots is tilted towards the combination of older and larger firms.

In fact, the difference in both the average age, and the average size, between adopters

Figure B.12: Age-size Distribution of Firms using Robotics

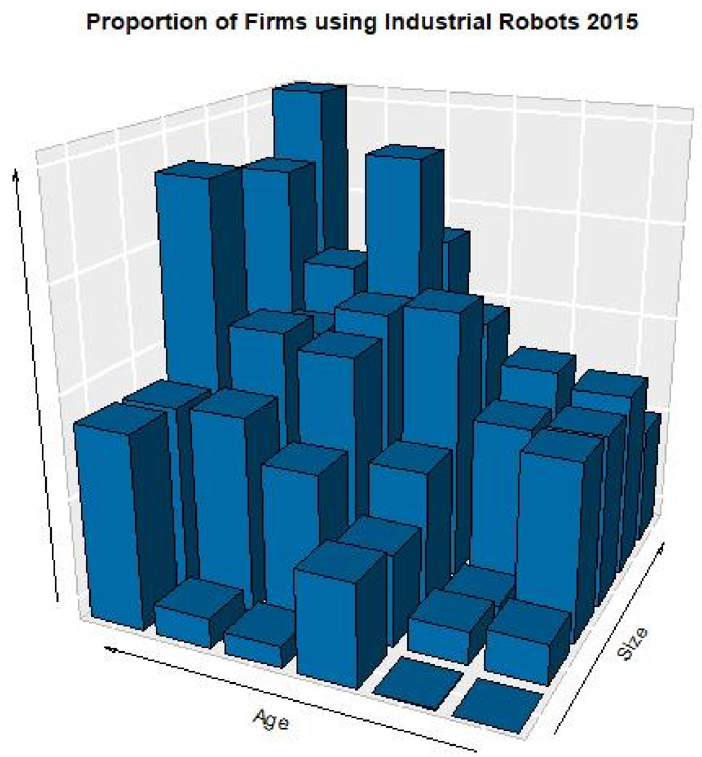

Note: each bar is ratio of firms using Industrial Robots in each age × size cell.

and non-adopters is the largest out of the technologies considered in this paper. Firms using industrial robots have a mean age of over 41, compared to under 35 for non-adopters. In addition, firms adopting industrial robotics employ over 68 workers on average, compared to just 48 employees for non-adopters.

Firms using industrial robots pay more, on average, and have a higher proportion of blue-collar workers, compared to non-adopters (see Table B.13). In fact, industrial robot adopters have 64% of their workforce in blue-collar occupations, which is by far the highest of any of the technologies in this survey (the next highest is firms using 3D printing, with 56%).

The sectoral distribution of robot adoption is, unsurprisingly, extremely unequal: robots are used by around 25% of Metal Manufacturing firms, but by barely any businesses in Services. This distribution is plotted in Figure B.13.

It is rare that non-exporting firms use robots. For example, in 2017 adoption was under 2.5% for domestic-only firms, compared to around 15% adoption for firms with

Figure B.13: Adoption of Industrial Robotics by Industry (2017)

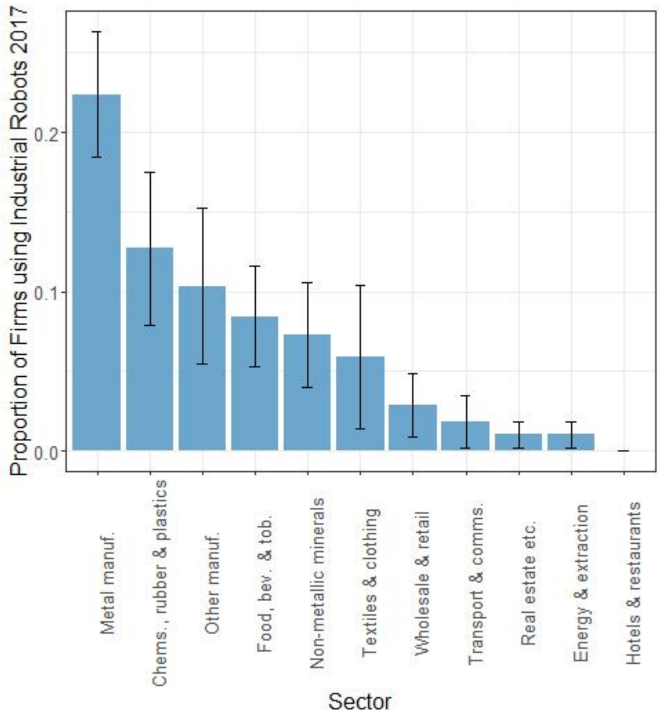

over a third of sales coming from exports. This is consistent with other research on the relationship between robot adoption and exporting (Stapleton and Webb 2020). This relationship can be seen in Figure B.14, which plots the share of firms using robots by their export share.

Figure B.14: Industrial Robotics Adoption by Export Intensity

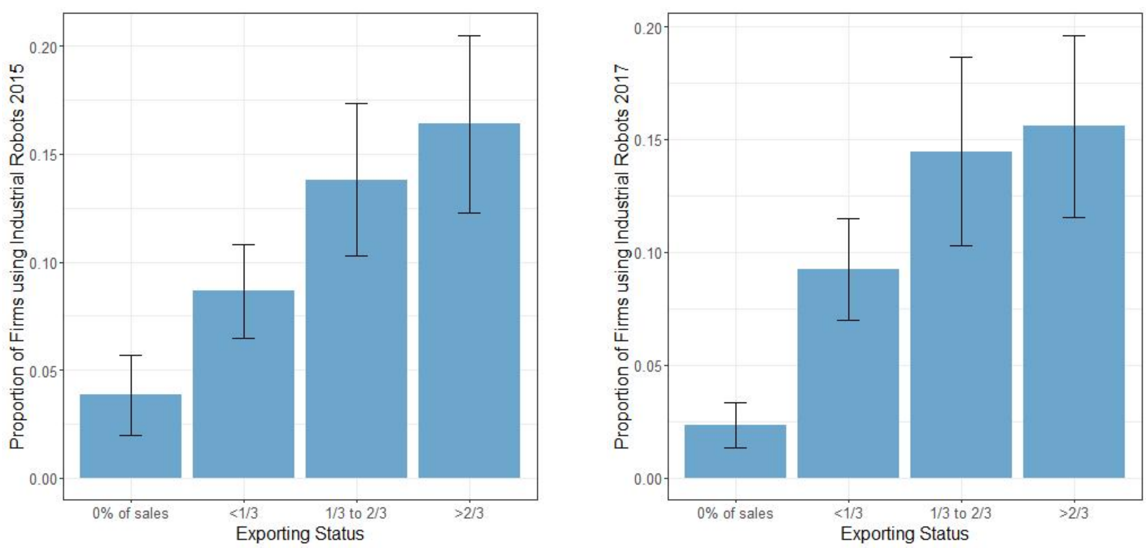

However, the relationship between robots, exporting, and size is not simple. Consider Figure B.15, which plots the proportion of robot adopters across the firm-size distribution for four levels of exporting behaviour (from non-exporting to a high share of sales in exporting). It is clear that robot adoption rises with firm size for exporting-focused firms (exporting more than a third of sales), but this relationship disappears for firms that export up to one third of sales.[32] This means that robot adoption rises with firm size, but only for heavy exporters.

Figure B.15: Industrial Robots Adoption by Employment and Exporting Status (2015)

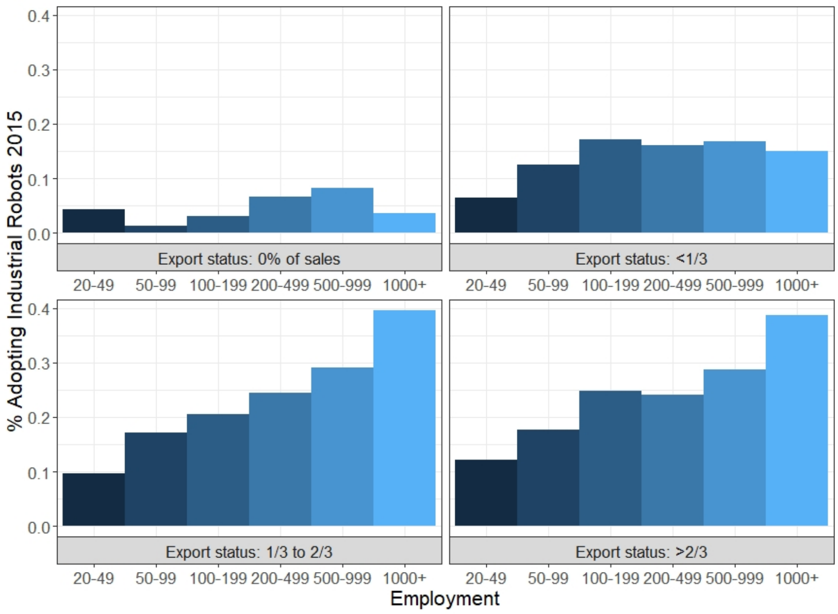

Note: each bar indicates the share of firms adopting Industrial Robots within that employment bin. Each panel is split by firms with different export shares.

Firms that use robots have more employees, but they employ workers for fewer hours on average B.13. The difference is half an hour per worker per week, but (1) this is the only technology for which adopters have lower average hours, and (2) this scales up to almost 1,800 worker-hours for firms using robots each year (or, approximately one worker). A back-of-the-envelope calculation suggests these lower hours save the firm €31,500 per year.[33] This could be consistent with robots being used to automate some worker-tasks, which reduce hours but not the number of employees.

[32] Around 31.5% of firms are in this bucket.

[33] 1,800 worker-hours divided by the average number of weekly hours (31.56), divided by 52 weeks gives 1.08, which is the equivalent number of full-time employees. At an annual wage of just over €29,150, the associated wage for 1.08 workers is 29,150 × 1.08 ≈ €31,400.

3D Printing

3D printing technology uses a digital depiction of an item alongside "additive manufacturing" techniques to create physical objects (Bechtold 2015). This data allows me to highlight facts about firm adoption of 3D printing that, to my knowledge, have not yet been investigated. In 2015 I find that adopting firms are older and larger. However, by 2017 and 2019, firms using 3D printing are still larger, but the age distribution is much flatter. It seems that fast-growing young firms adopted this technology in recent years. This is suggestive evidence that 3D printing allows young firms to grow very quickly. Plots of these patterns can be found in Figure B.16.

Similarly to robotics, it is rare than non-exporting firms use 3D printing. More specifically, under 5% of non-exporting firms adopt this technology in all years of the survey, compared to over 12% of 'heavy' exporters (those exporting over two-thirds of sales). Notably, adoption of 3D printing becomes more closely related to firm exporting behaviour over time: the proportion of non-exporters using this technology falls, while the share of 'heavy' exporters rises.

As with industrial robots, 3D printing is much more common in Manufacturing than other sectors, and it is especially concentrated in Metal and Rubber, Chemical & Plastics (see Figures B.18).

Firms using 3D printing technology are, on average, around four years older than those that do not (see Table B.13). The average firm size is around 64 workers for adopters, compared to just 49 for non-adopters. Surprisingly, fixed investment per worker is significantly lower for firms using this advanced technology, although this could be because adopters are larger, 3D printing reduces investment, or such firms are different in another way.

Figure B.16: 3D Printing Adoption by Firm Age and Size

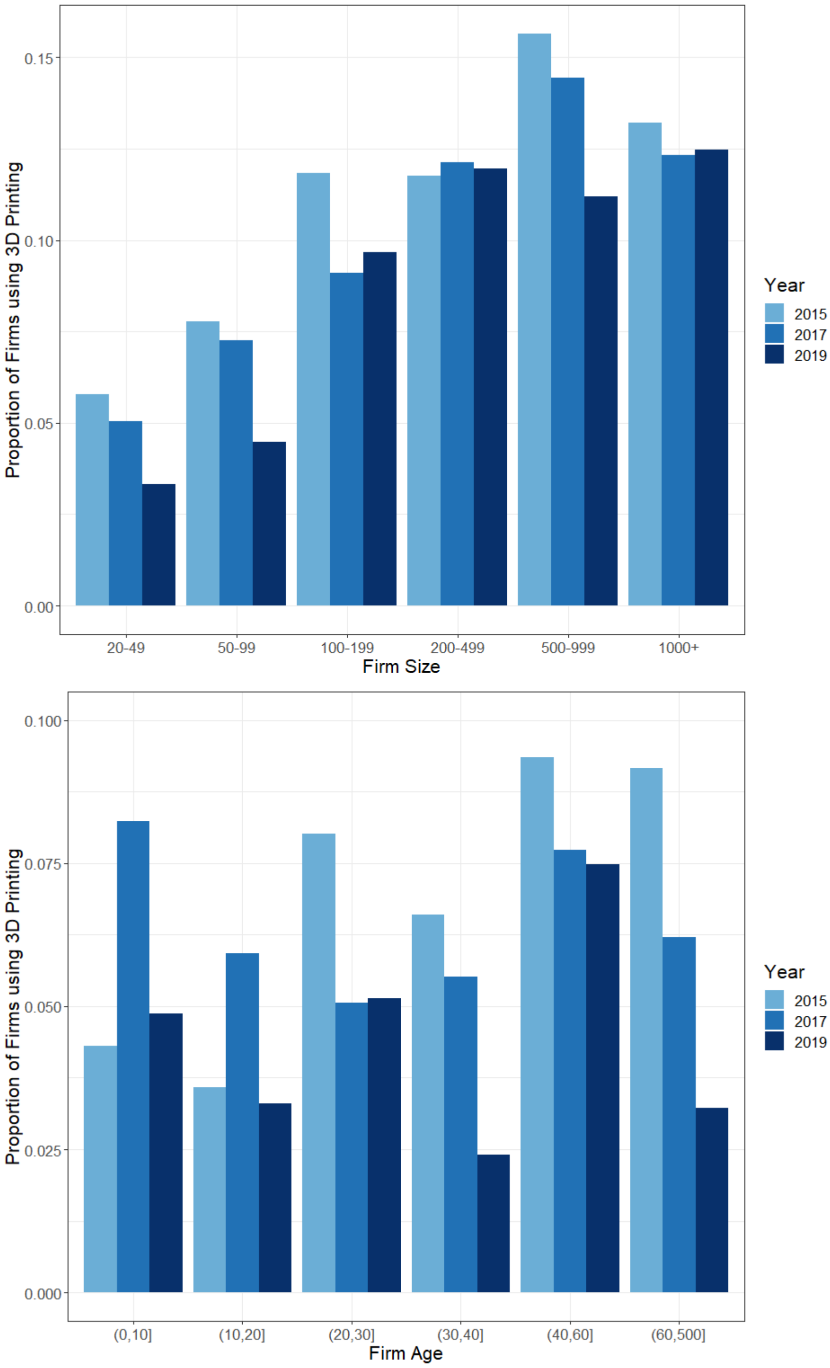

Figure B.17: 3D Printing Adoption by Export Intensity

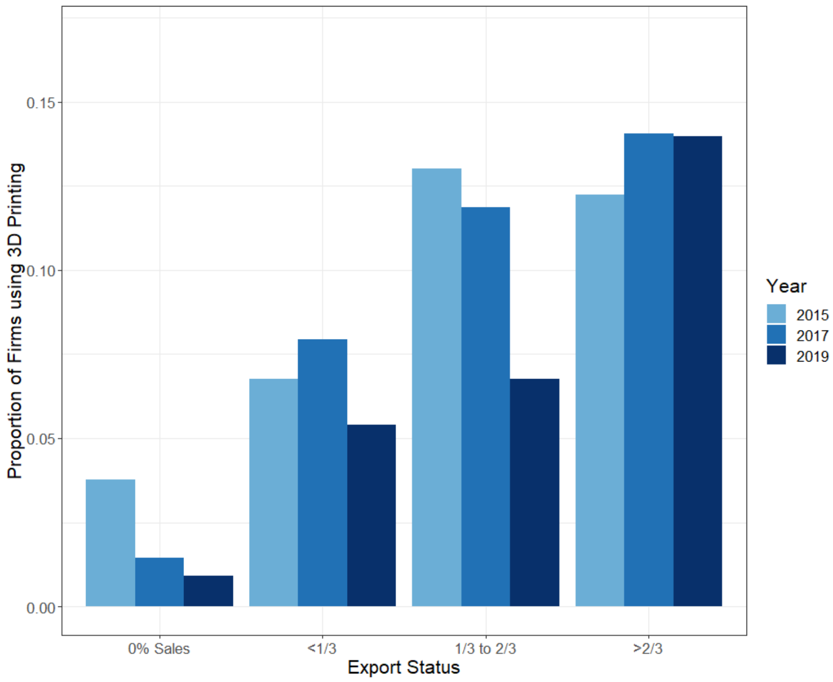

Figure B.18: Adoption of 3D Printing across Industries

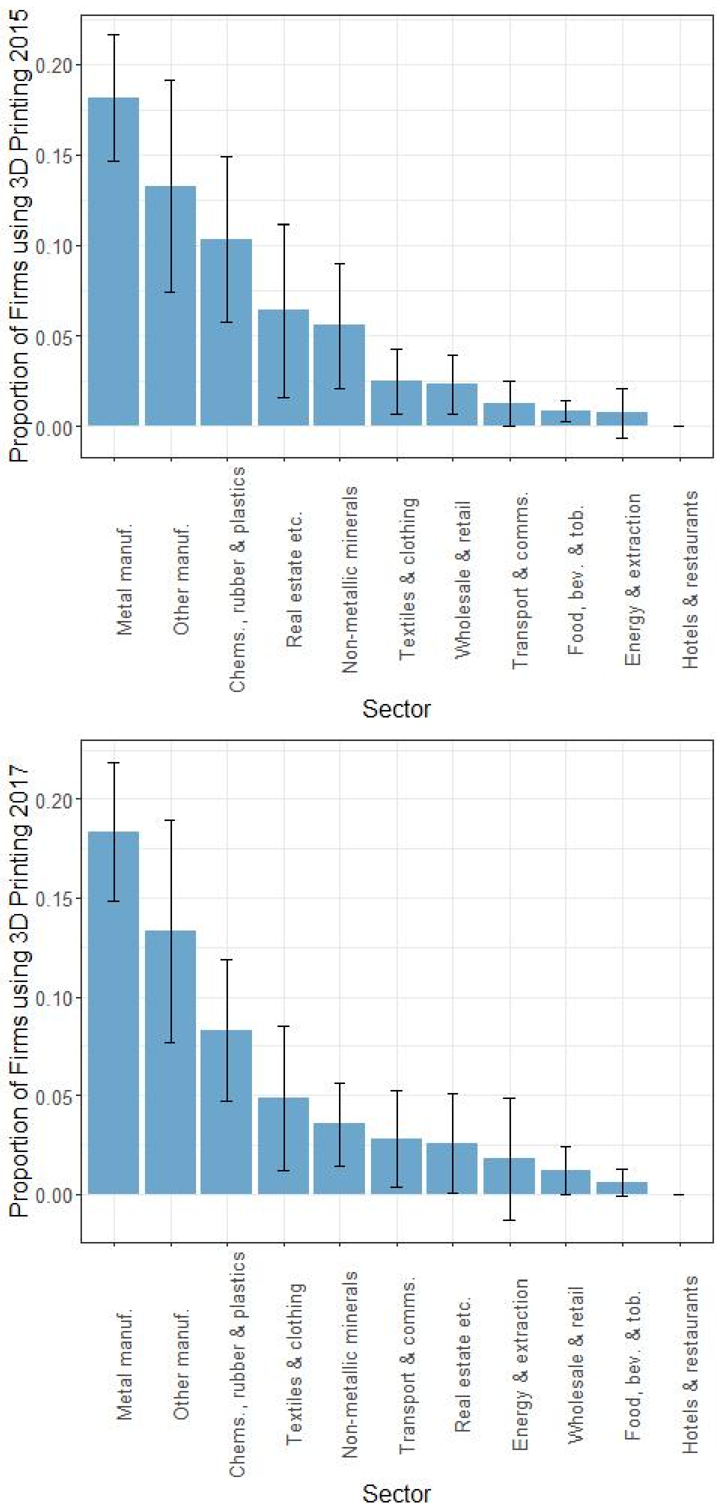

Table B.13: Advanced Technology Adopters vs Non-Adopters: Summary Statistics 2015

	Mobile & Cloud		*AI & Big Data*	
	Adopters	**Non-Adopters**	**Adopters**	**Non-Adopters**
Age	34.9	**35.7**	33.0	**35.4**
Log(employment)	**3.94**	3.83	**4.14**	3.89
Log(wage per worker) (€)	**10.28**	10.28	**10.35**	10.27
Hours per worker (weekly)	**31.46**	31.40	**31.65**	31.49
Turnover per worker (thou., €)	**459**	345	**499**	423
Fixed investment per worker (thou., €)	**10.9**	8.12	**14.5**	9.61
% blue-collar	0.51	**0.61**	0.43	**0.55**
% export sales	0.09	0.11	0.06	**0.10**
Terminations per worker	**0.18**	0.14	0.17	0.17
Hirings per worker	**0.19**	0.13	0.18	0.17
Number of firms	2524	934	440	2953
	Internet of Things		*Industrial Robotics*	
	Adopters	**Non-Adopters**	**Adopters**	**Non-Adopters**
Age	35.1	35.1	**41.1**	34.7
Log(employment)	**4.13**	3.87	**4.22**	3.88
Log(wage per worker) (€)	**10.32**	10.27	**10.33**	10.28
Hours per worker (weekly)	**31.81**	31.46	31.05	**31.56**
Turnover per worker (thou., €)	397	**439**	331	**442**
Fixed investment per worker (thou., €)	**12.4**	9.78	8.22	**10.4**
% blue-collar	0.51	**0.55**	**0.64**	0.53
% export sales	0.04	**0.11**	0.05	**0.10**
Terminations per worker	0.16	0.17	0.11	**0.17**
Hirings per worker	0.18	0.18	0.11	**0.18**
Number of firms	642	2742	485	2912
	3D Printing			
	Adopters	**Non-Adopters**		
Age	**38.5**	34.9		
Log(employment)	**4.16**	3.89		
Log(wage per worker) (€)	**10.33**	10.28		
Hours per worker (weekly)	**31.65**	31.52		
Turnover per worker (thou., €)	273	**444**		
Fixed investment per worker (thou., €)	5.71	**10.5**		
% blue-collar	0.56	0.54		
% export sales	0.11	0.10		
Terminations per worker	0.08	**0.17**		
Hirings per worker	0.11	**0.18**		
Number of firms	360	3037		

Note: Summary statistics from 2015 for firms that do and don't use advanced technologies. All values (other than number of observations) are weighted means. Bold values are the larger of the two, if there is a significant difference between adopters and non-adopters at the 1% level, computed with Welch's t-test and the Welch-Sattherwaite equation for degrees of freedom.

Table B.14: Automating Firms vs Traditional Firms: Summary Statistics

	Employment	*log(Hours)*	*log(Wage)*	*Age*
'Digital' Automating Firms	59.6	7.41	10.32	34.1
Non-Adopters	47.6	7.40	10.27	35.4
'Physical' Automating Firms	63.8	7.40	10.31	38.7
Non-Adopters	48.1	7.40	10.28	34.6

Note: 'Digital' automation technologies are Mobile & Cloud, AI & Big Data, and the Internet of Things. 'Physical' automation technologies are Industrial Robotics, and 3D Printing. Note that hours are average hours per worker, over the year.

B.4 Appendix - Firm Dynamics Analysis

This section contains figures and tables from firm growth dynamics analysis on the data in this paper.

Figure B.19: Firm Distributions

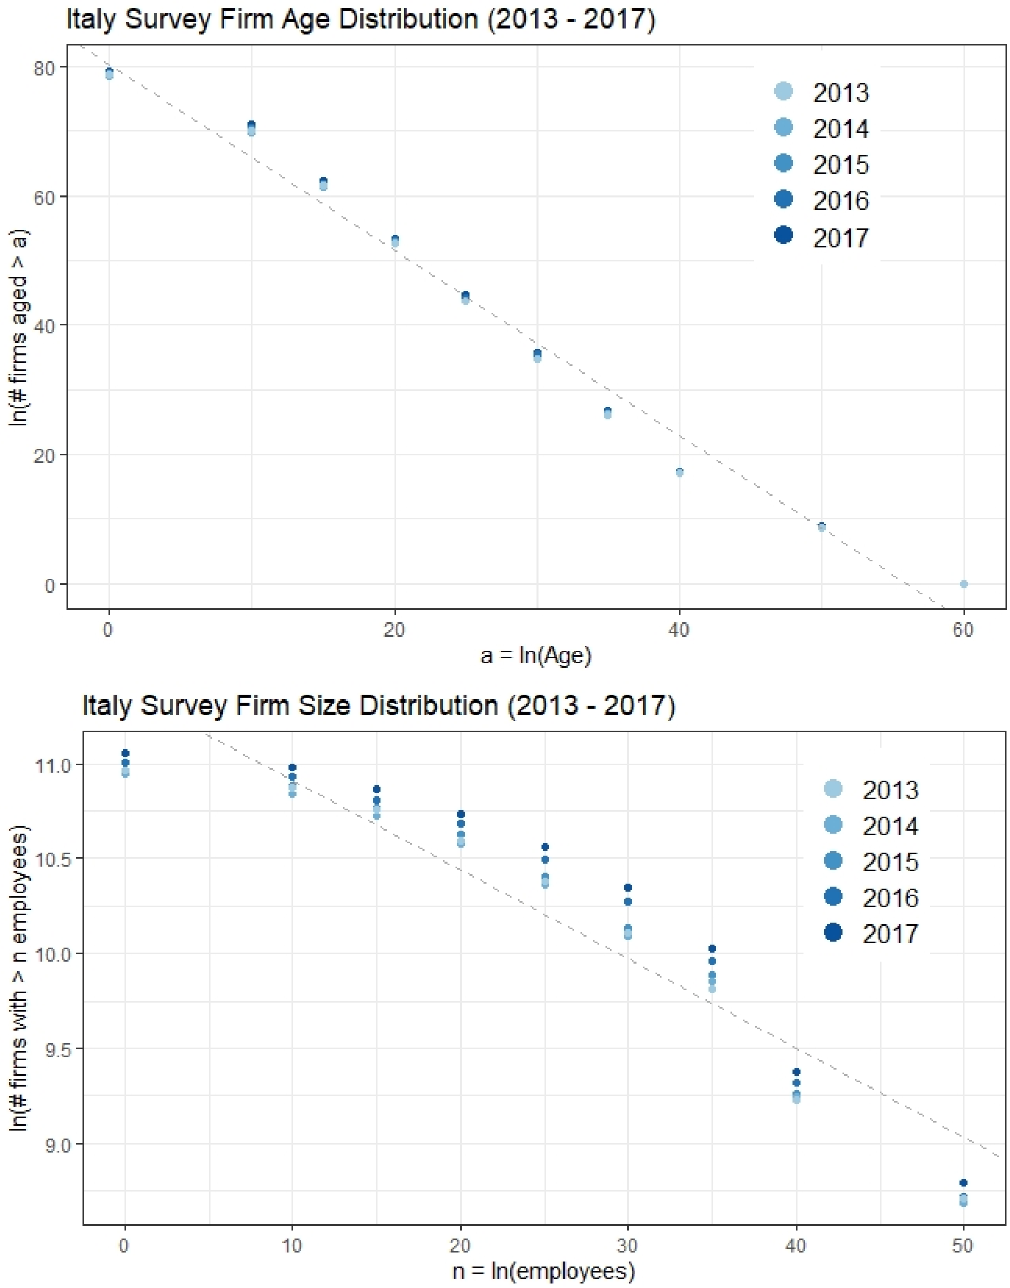

This table contains estimated coefficients from standard firm growth regressions, on log employment and age, with sector, year, and firm fixed-effects. The regressions are run over all firms in the panel that appear in more than two years between 2010 -

Figure B.20: Net Employment Growth Rates

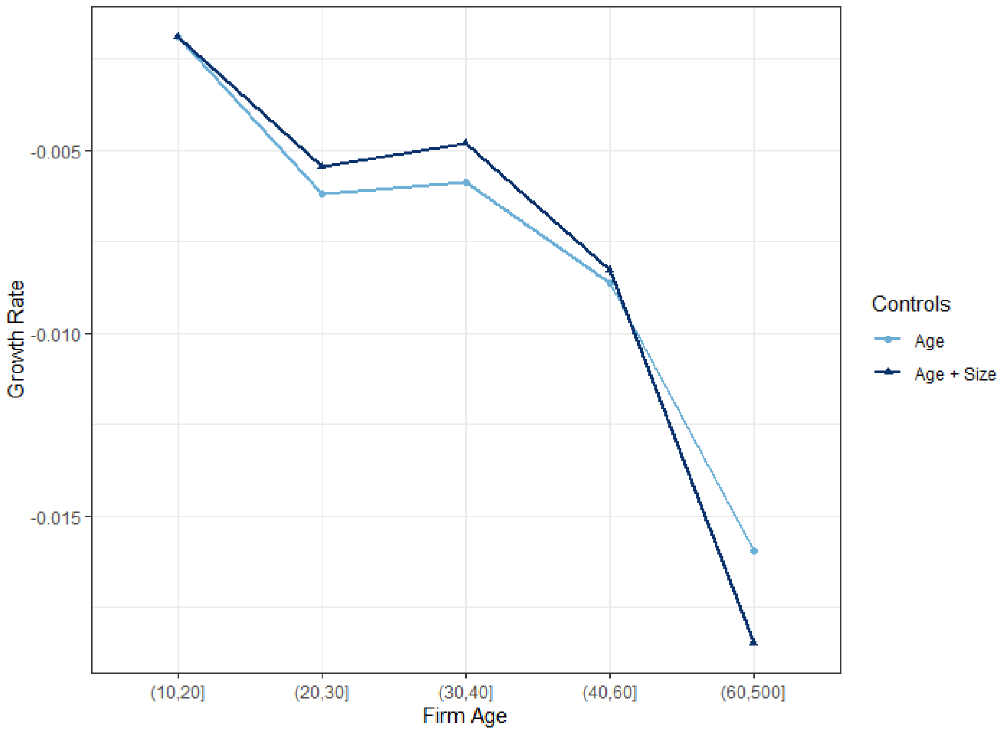

2018.

Table B.15: Estimated Coefficients from Firm Growth Regressions 2010 - 2018

	Dependent variable: *Employment Growth*							
log(Emp.)	0.014***	0.015***	0.014***	0.015***	0.366***	0.366***	0.365***	0.364***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.09)	(0.09)	(0.09)	(0.09)
Age		−0.0003***		−0.0003***		0.0005		−0.0006
		(0.00007)		(0.00007)		(0.0004)		(0.0004)
Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE			✓	✓			✓	✓
Firm FE					✓	✓	✓	✓
N	20450	20446	20450	20446	20450	20446	20450	20446

*Note: Estimates are significant at levels of 0.1%: ***, 1%: **, 5% *.*

Figure B.21: Net Growth Rates, 2010 - 2018

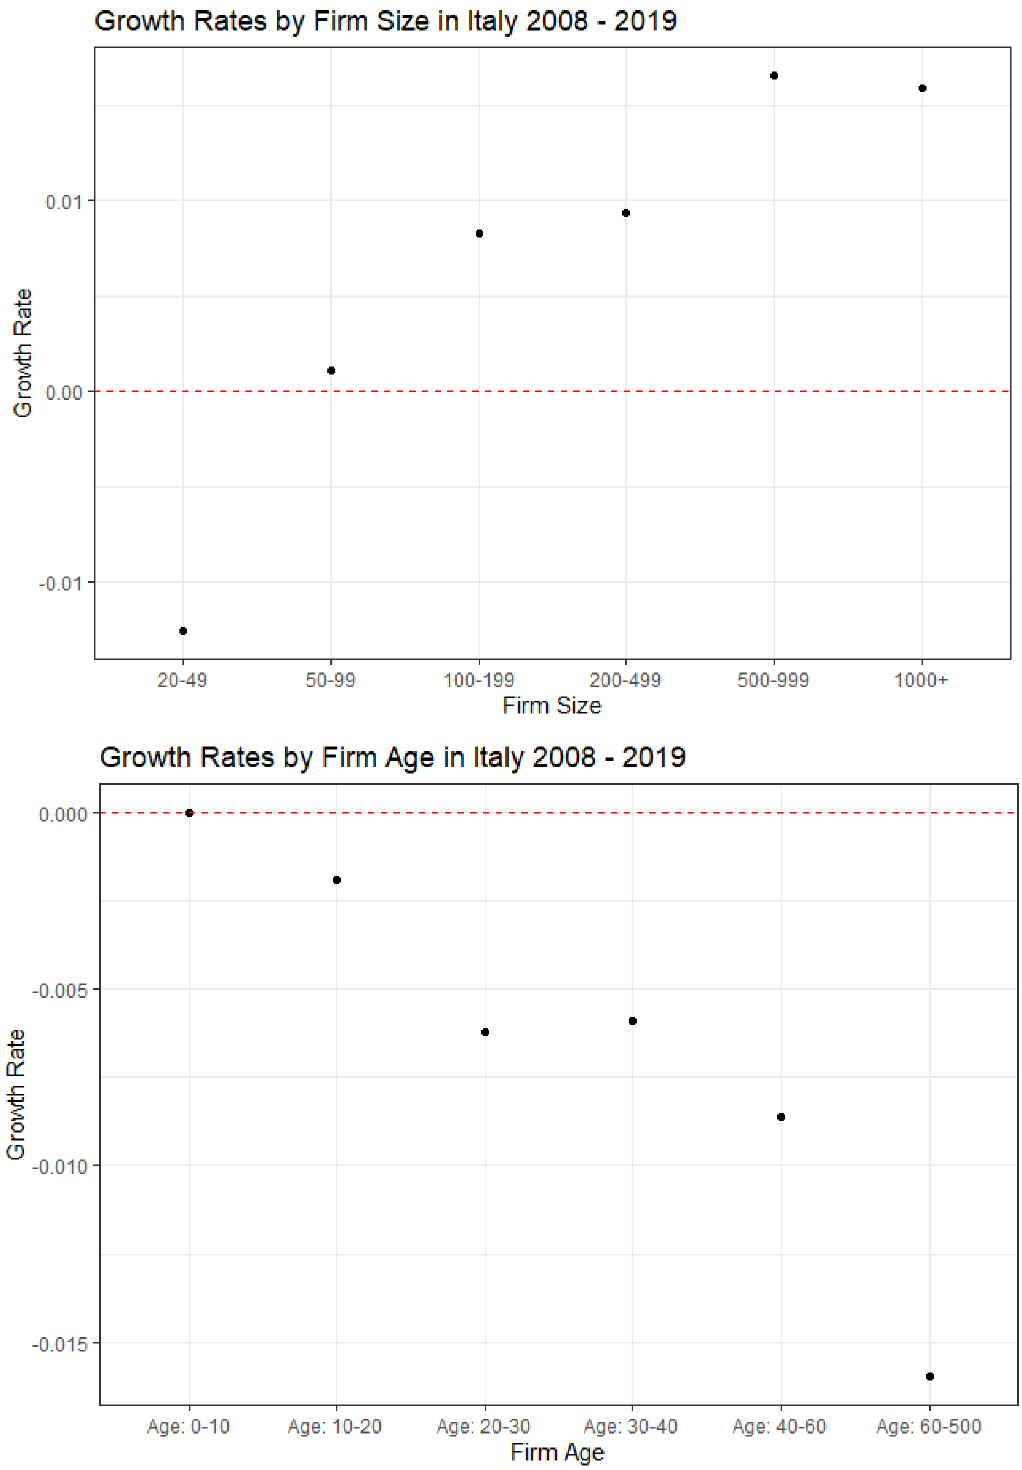

Table B.16: Average Growth Rates by Technology, for Adopters and Non-Adopters

	Cloud Computing	*AI & Big Data*	*IoT*	*Industrial Robotics*	*3D Printing*
Adopters	1.0065	1.0133	1.0133	1.0122	1.0171
Non-Adopters	0.9989	1.0036	1.0033	1.0038	1.0037

Firms are compiled in a panel from 2010 - 2018 and labelled as adopters or non-adopters of each technology. Each number here represents the average employment growth rate of firms in these groups over this time period.

Figure B.22: Average Growth Rates (2010 - 2018) for Automation Adopters compared to Non-Adopters, Across Age Distribution

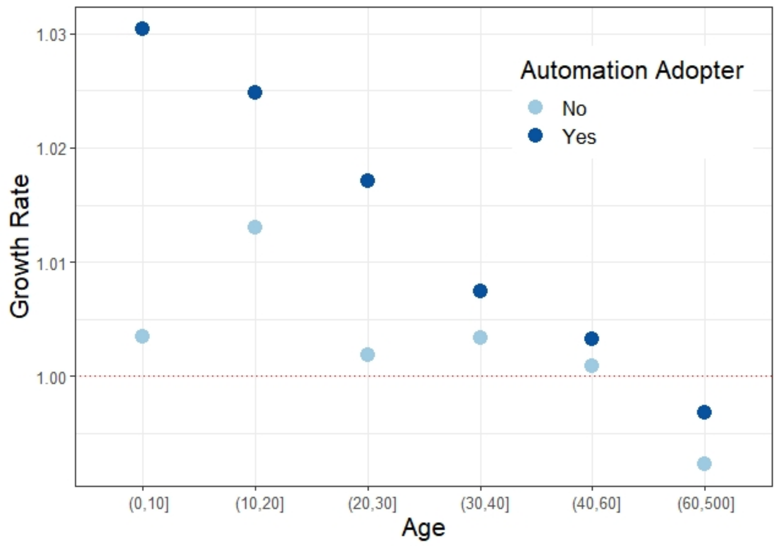

Figure B.23: Average Growth Rates (2010 - 2018) for Automation Adopters compared to Non-Adopters, Across Size Distribution

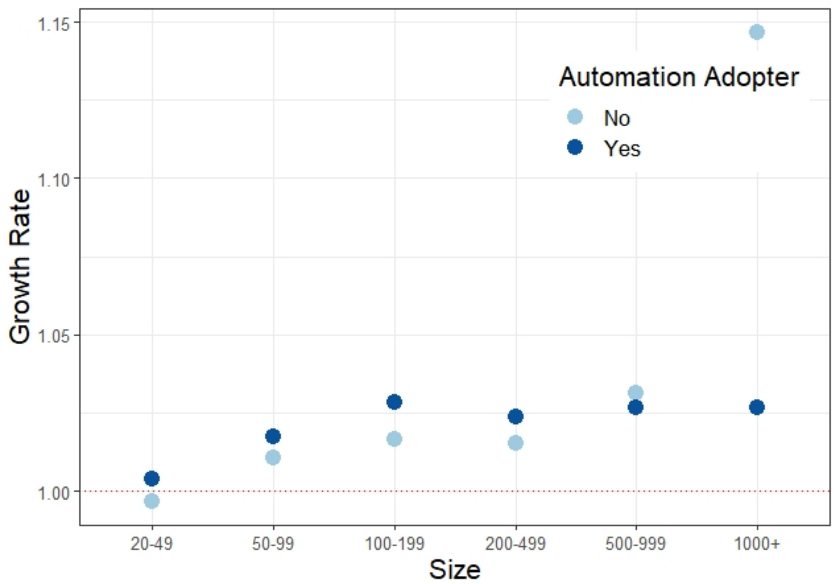

Figure B.24: Average Growth Rates (2010 - 2018) for Automation Adopters compared to Non-Adopters

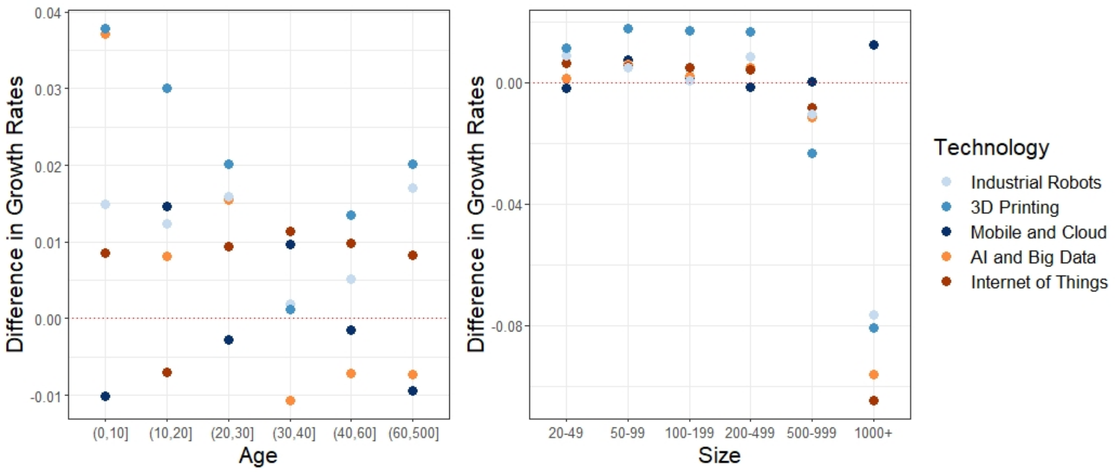

Figure B.25: Average Growth Rates (2010 - 2018) for Automation Adopters vs Non-Adopters, By Size and Age

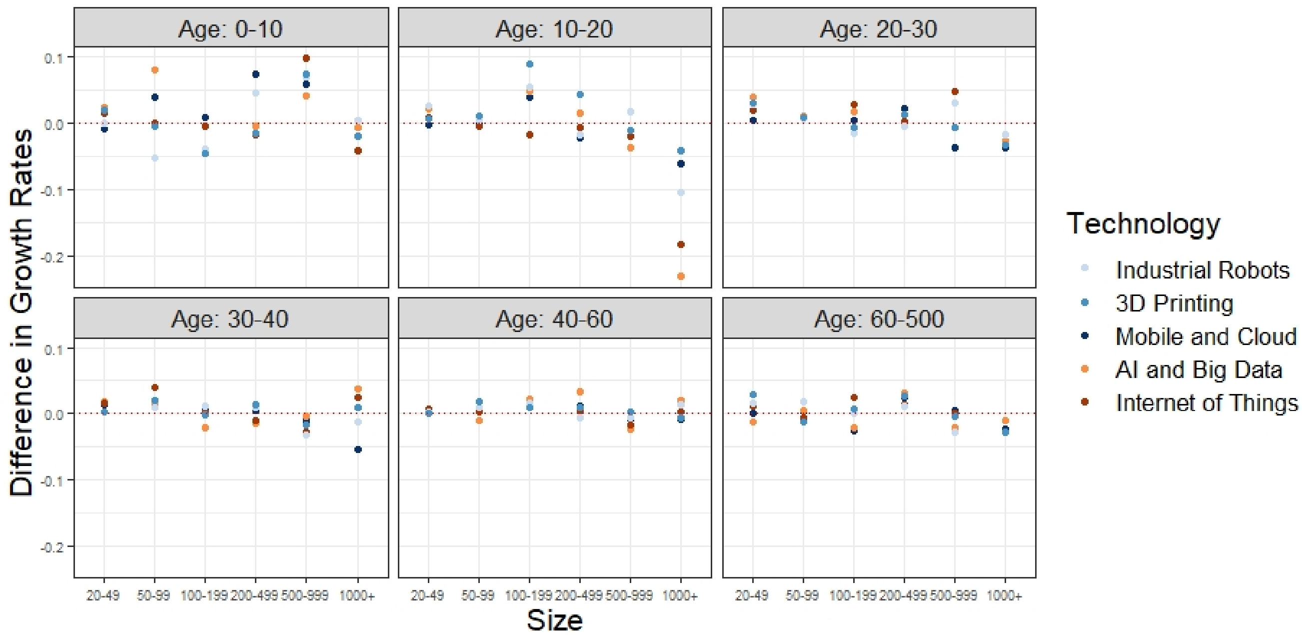

B.5 Appendix - Further Figures

Table B.17: Estimated Coefficients from Advanced Tech. Investment Regressions

Dependent variable: *Share of Investment in Advanced Tech.*						
	2016 ($N = 3756$)			2017 ($N = 3926$)		
log(Emp.)	0.279***	0.278***	0.254***	0.337***	0.329***	0.299***
	(0.025)	(0.025)	(0.026)	(0.028)	(0.028)	(0.028)
Age		−0.0000004	0.0002		0.0034**	0.0026*
		(0.001)	(0.001)		(0.001)	(0.001)
	2018 ($N = 3715$)			2019 ($N = 2075$)		
log(Emp.)	0.294***	0.295***	0.261***	0.275***	0.274***	0.258***
	(0.030)	(0.030)	(0.030)	(0.039)	(0.039)	(0.040)
Age		−0.0004	−0.002		0.0005	−0.001
		(0.001)	(0.001)		(0.001)	(0.002)
Fixed Effects			✓			✓

*Note: Estimates are significant at levels of 0.1%: ***, 1%: **, 5% *. Fixed effects are sector and region.*

Table B.18: Estimated Coefficients from Technology Adoption Regressions 2015

Dependent variable: *Technology Adoption*										
	Mobile & Cloud		*AI & Big Data*		*Internet of Things*		*Industrial Robotics*		*3D Printing*	
log(Emp.)	0.033**	0.034**	0.039***	0.041***	0.052***	0.052***	0.040***	0.038***	0.027***	0.025***
	(0.010)	(0.011)	(0.009)	(0.009)	(0.009)	(0.009)	(0.006)	(0.006)	(0.006)	(0.006)
Age		−0.0005		−0.0007		−0.0002		0.0009**		0.0004
		(0.0005)		(0.0004)		(0.0004)		(0.0003)		(0.0003)

*Note: Estimates are significant at levels of 0.1%: ***, 1%: **, 5% *.*

Table B.19: Estimated Coefficients from Technology Adoption Regressions 2017

Dependent variable: *Technology Adoption*										
	Cloud Computing		*AI & Big Data*		*Internet of Things*		*Industrial Robotics*		*3D Printing*	
log(Emp.)	0.074**	0.073**	0.064***	0.065***	0.076***	0.076***	0.036***	0.036***	0.033***	0.033***
	(0.010)	(0.010)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.004)	(0.004)
Age		0.0005		−0.0005*		−0.00007		0.0004**		−0.00006
		(0.0005)		(0.0004)		(0.0004)		(0.0003)		(0.0003)

*Note: Estimates are significant at levels of 0.1%: ***, 1%: **, 5% *.*

Table B.20: *Estimated Coefficients from Technology Adoption Regressions 2019*

	Dependent variable: *Technology Adoption*							
	Cloud Computing		*AI & Big Data*		*Industrial Robotics*		*3D Printing*	
log(Emp.)	0.084***	0.086***	0.060***	0.062***	0.054***	0.053***	0.031***	0.031***
	(0.016)	(0.016)	(0.013)	(0.013)	(0.007)	(0.007)	(0.006)	(0.006)
Age		−0.0008		−0.0011		0.0007		0.00001
		(0.0008)		(0.0008)		(0.0005)		(0.0002)

*Note: Estimates are significant at levels of 0.1%: ***, 1%: **, 5% *.*

Figure B.26: *Adoption of Automation Technologies by Firm Size and Age (2015)*

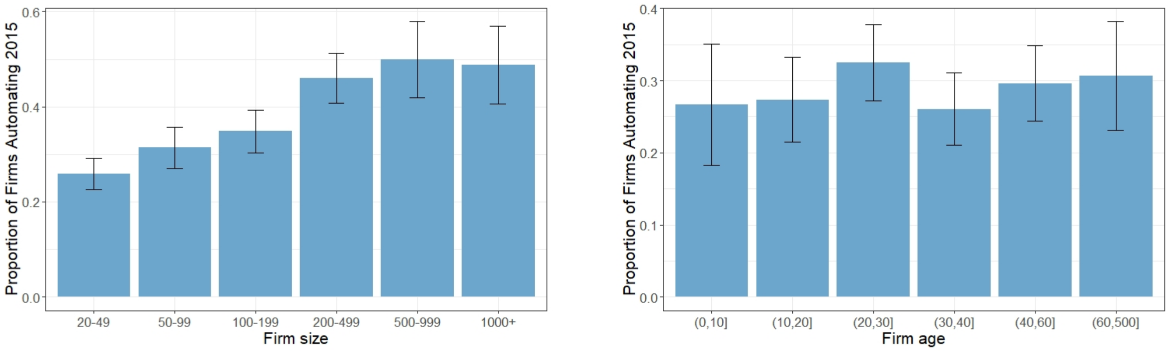

Figure B.27: *Adoption of Automation Technologies by Firm Size and Age (2017)*

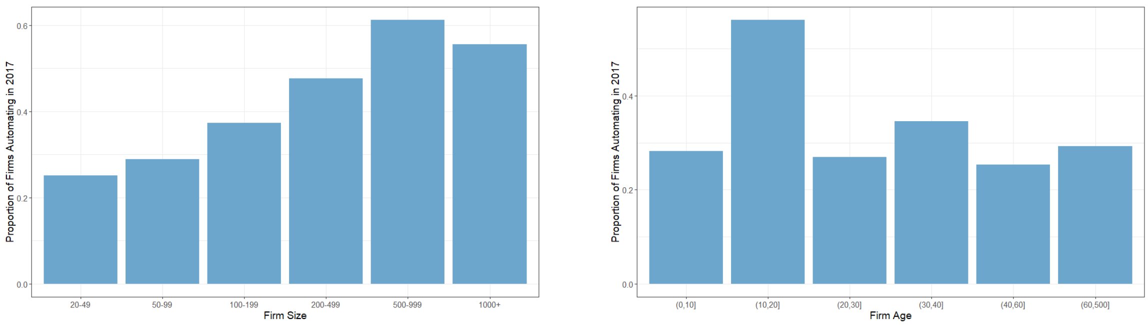

Figure B.28: *Adoption of Automation Technologies by Firm Size and Age (2019)*

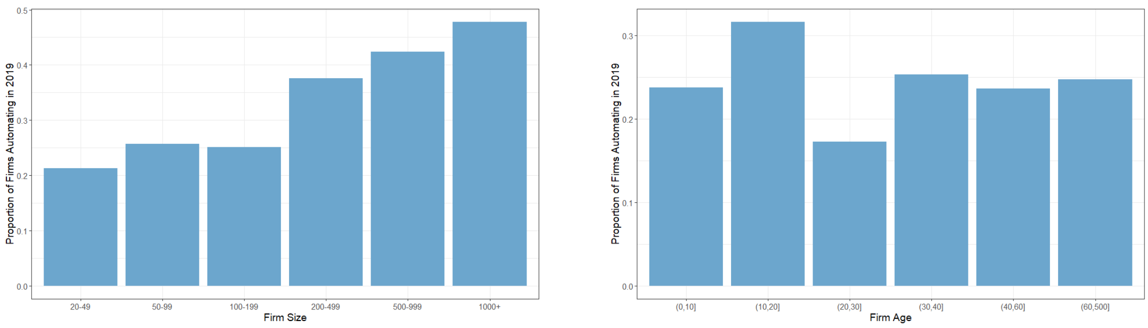

Figure B.29: Firm Age × Size Distributions of Advanced Technology Spending 2016 - 2017

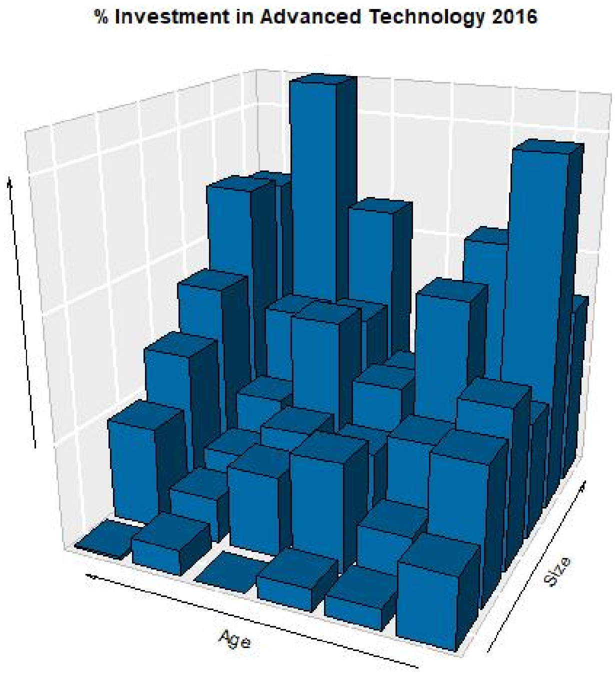

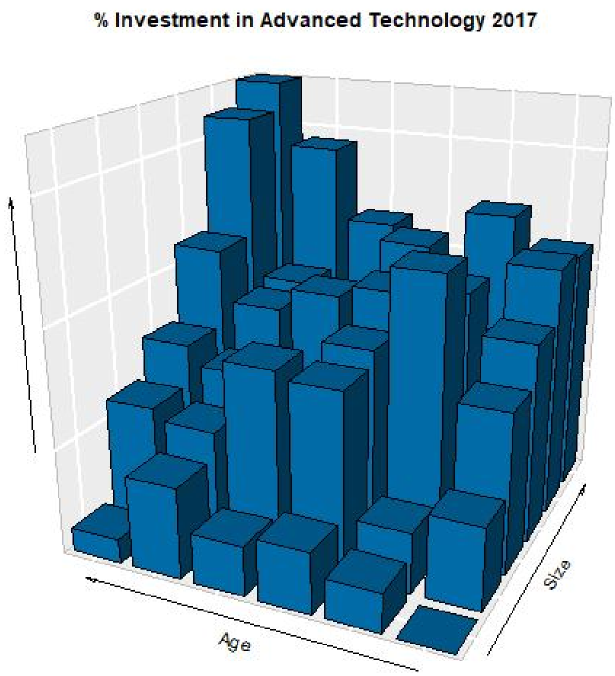

Figure B.30: Firm Age × Size Distributions of Advanced Technology Spending 2018 - 2019

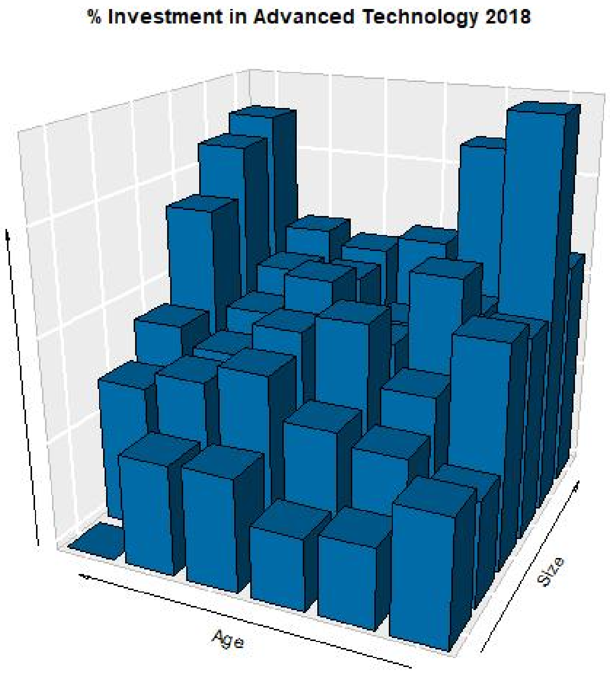

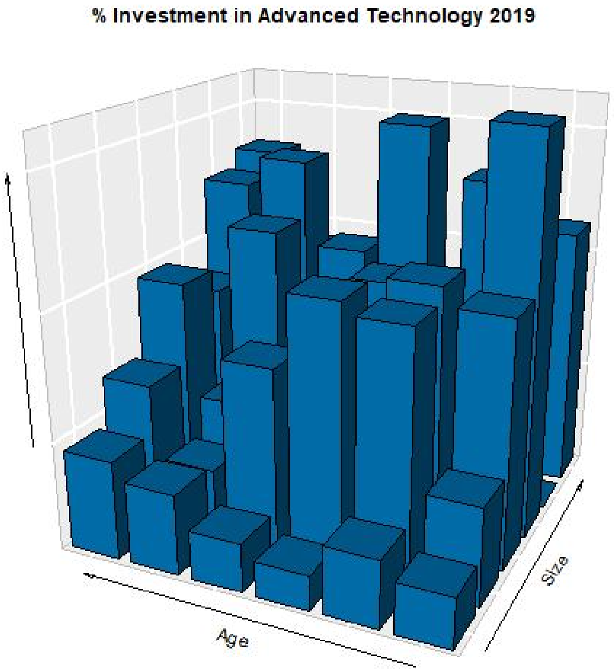

Figure B.31: Sectoral Distribution of Advanced Technology Spending 2016 - 2017

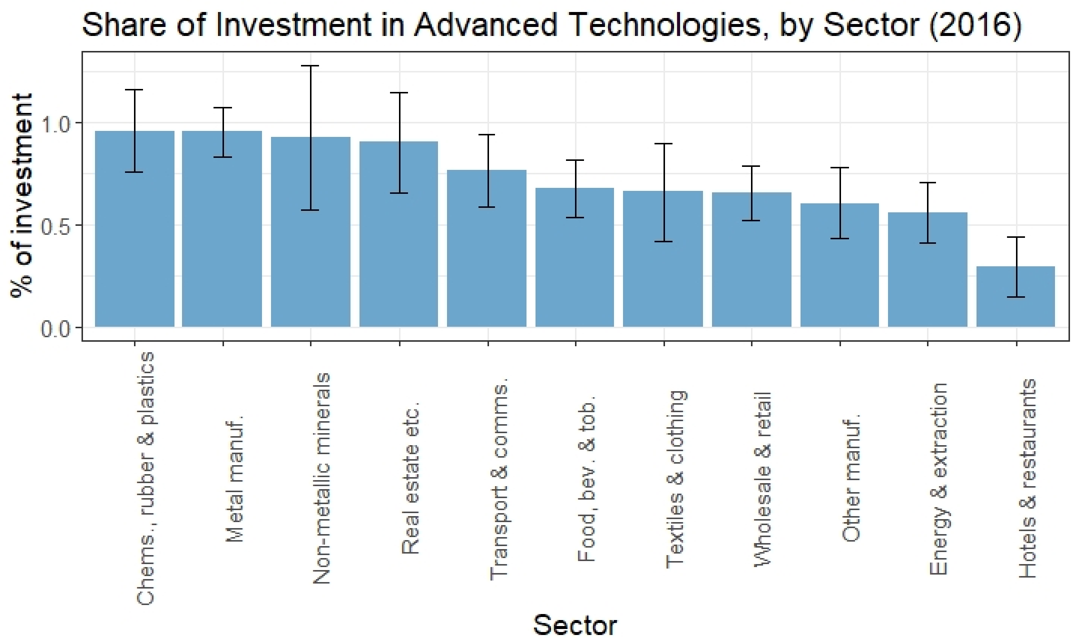

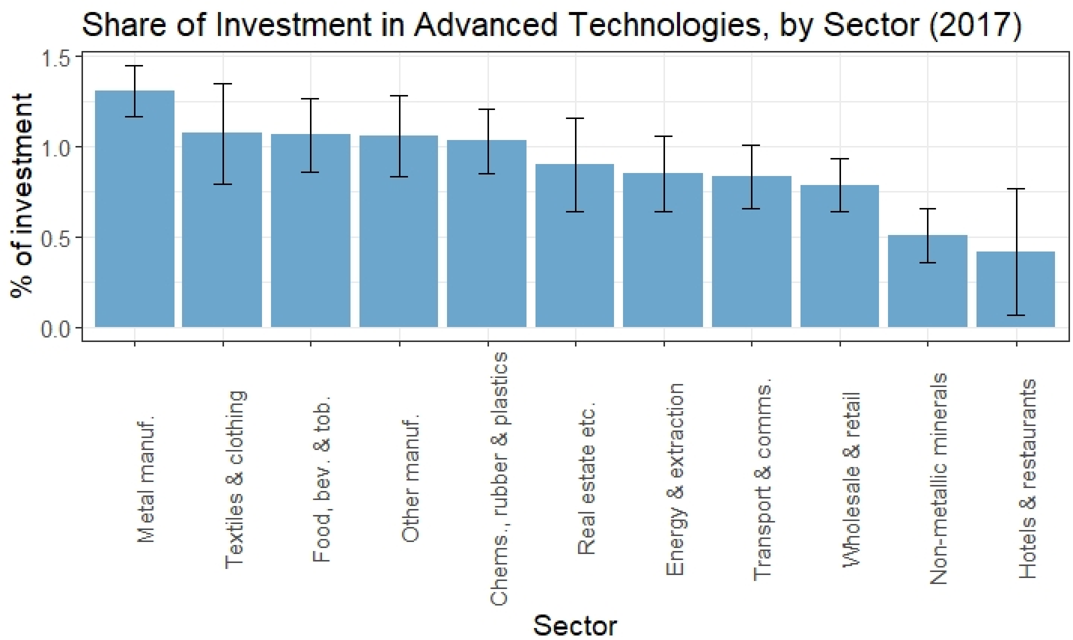

Figure B.32: Sectoral Distribution of Advanced Technology Spending 2018 - 2019

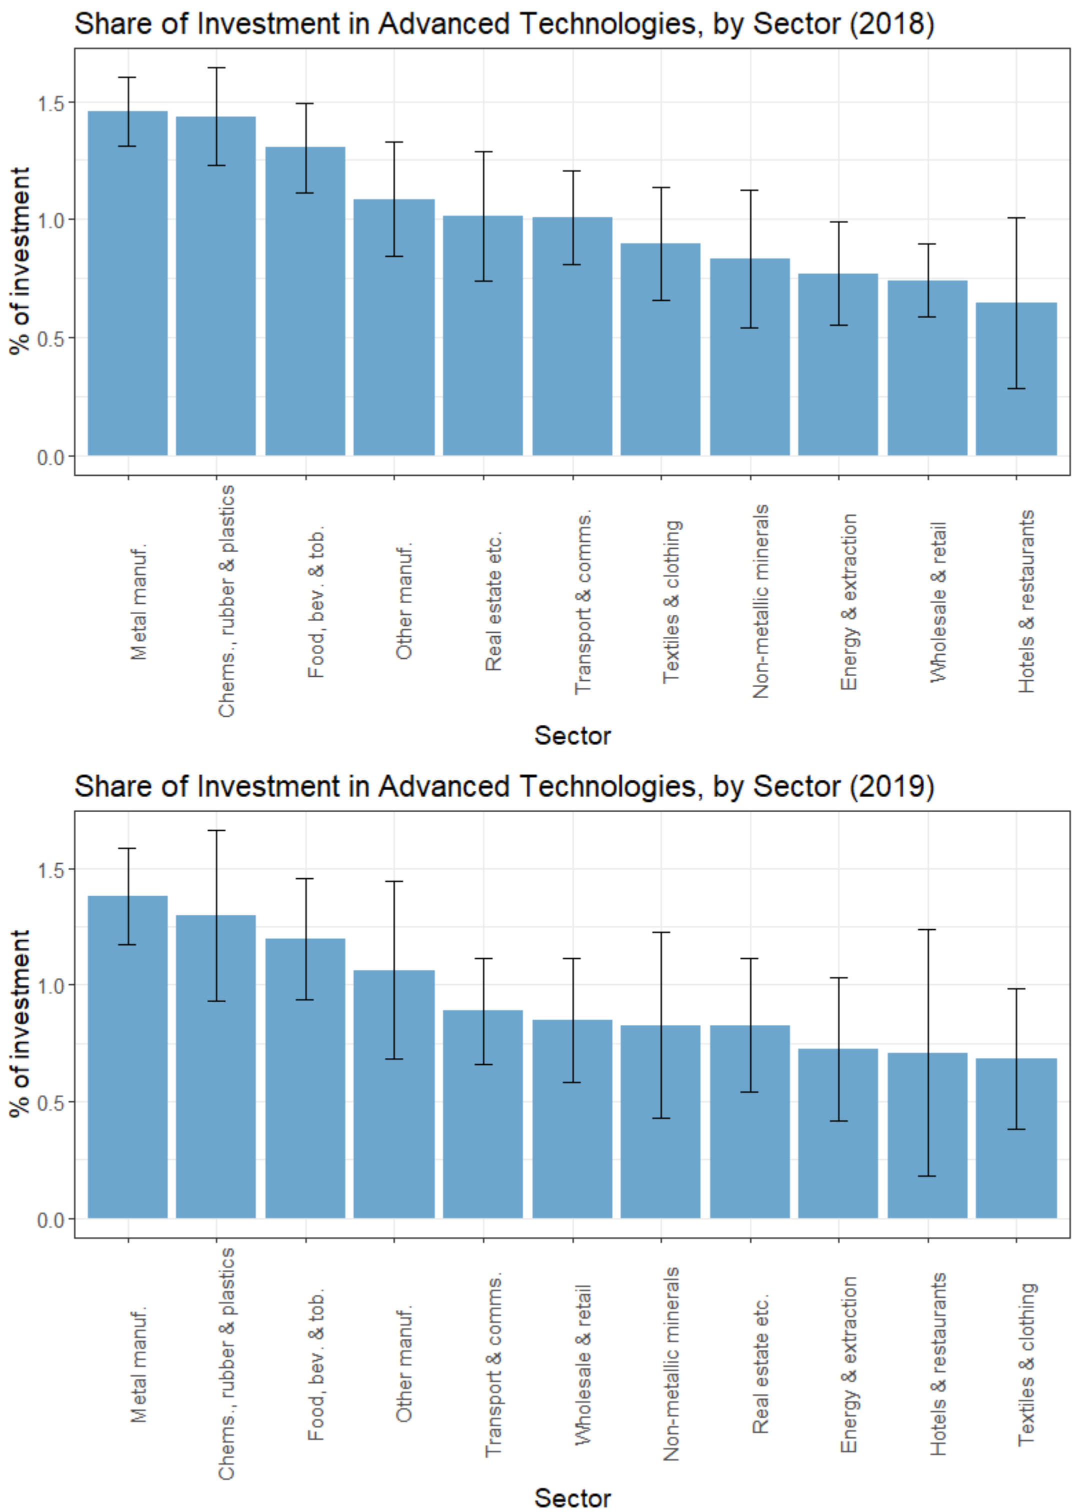

Figure B.33: Technology Switching Behaviour, 2015 - 2019

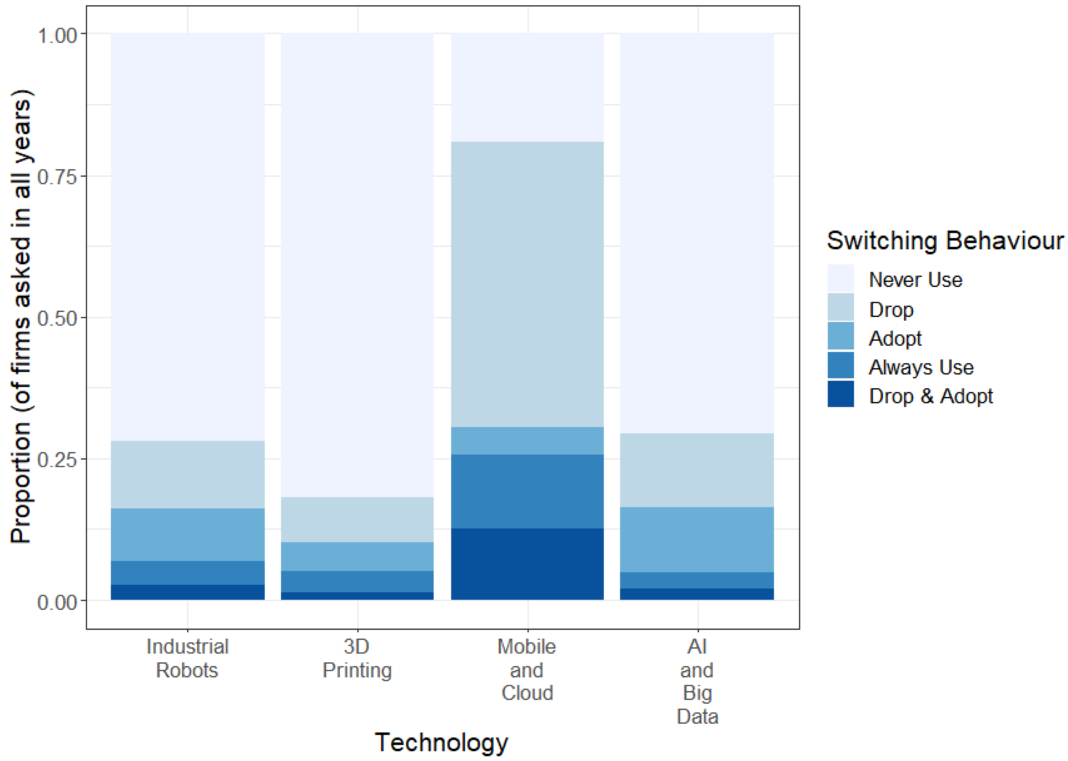

Note: Consider firms surveyed in 2015, 2017, and 2019. Never Use: firm does not adopt technology in any period. Drop: firm initially uses technology in 2015, but doesn't in subsequent periods. Adopt: firm doesn't initially use technology in 2015, but does adopt in a subsequent period. Always Use: firm uses technology in all periods. Drop & Adopt: firm initially uses technology in 2015, doesn't in a subsequent period, but then adopts it again in the future.

Figure B.34: Investment Share on Advanced Technology doesn't vary with Firm Age.

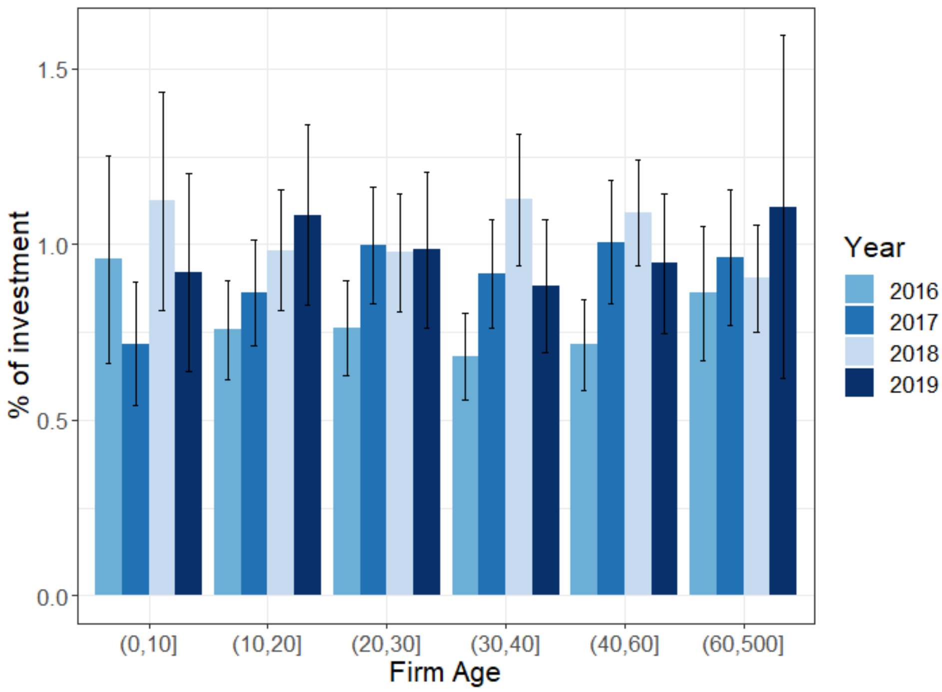

Figure B.35: Estimates from Event Study Regression

(a) Standard event study design on employment, without controlling for pre-trends (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

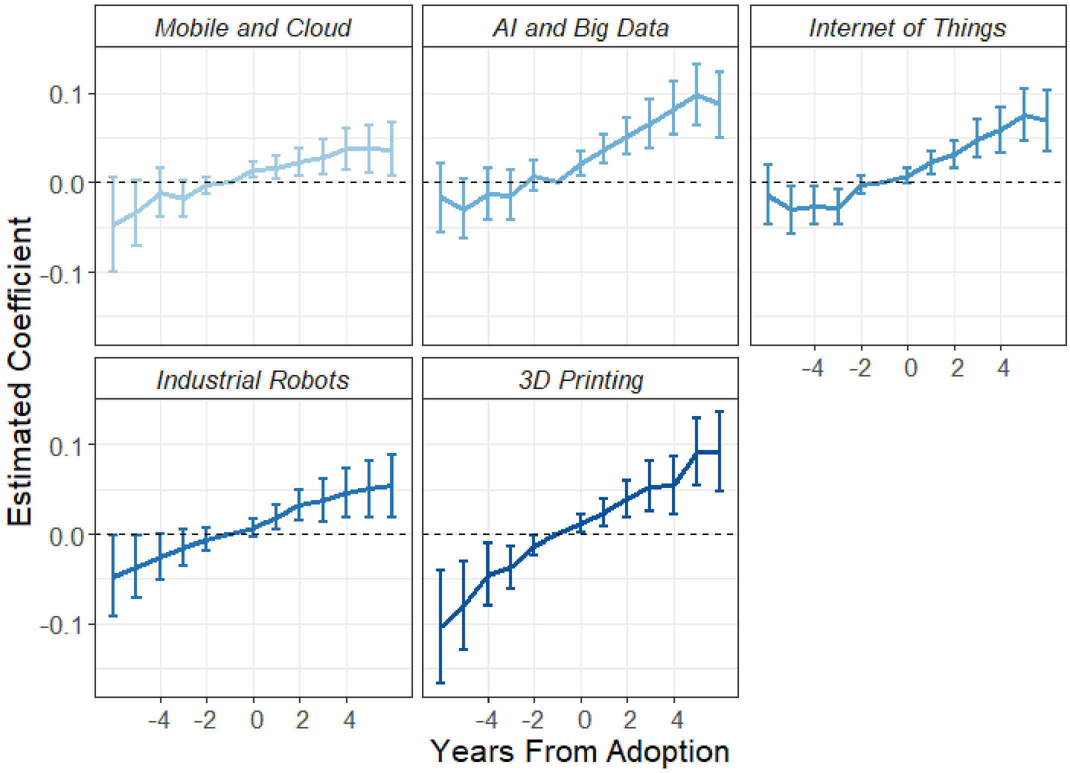

(b) Standard event study design on hours worked, without controlling for pre-trends (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

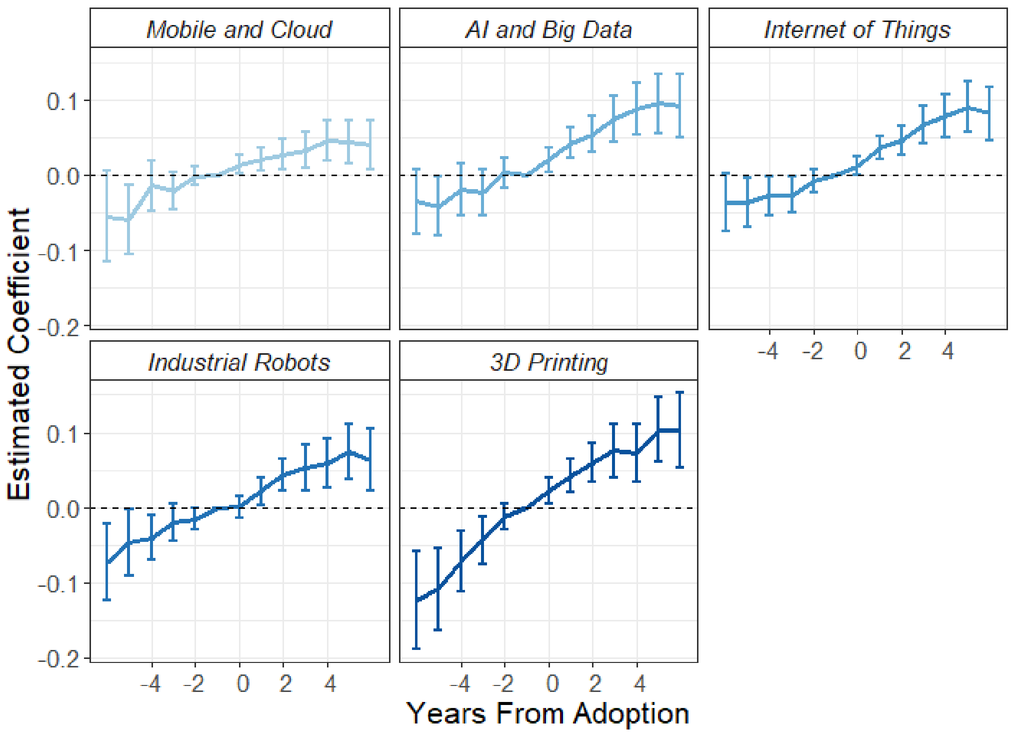

Figure B.36: Estimates from Event Study Regression

(a) Standard event study design on turnover, without controlling for pre-trends (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

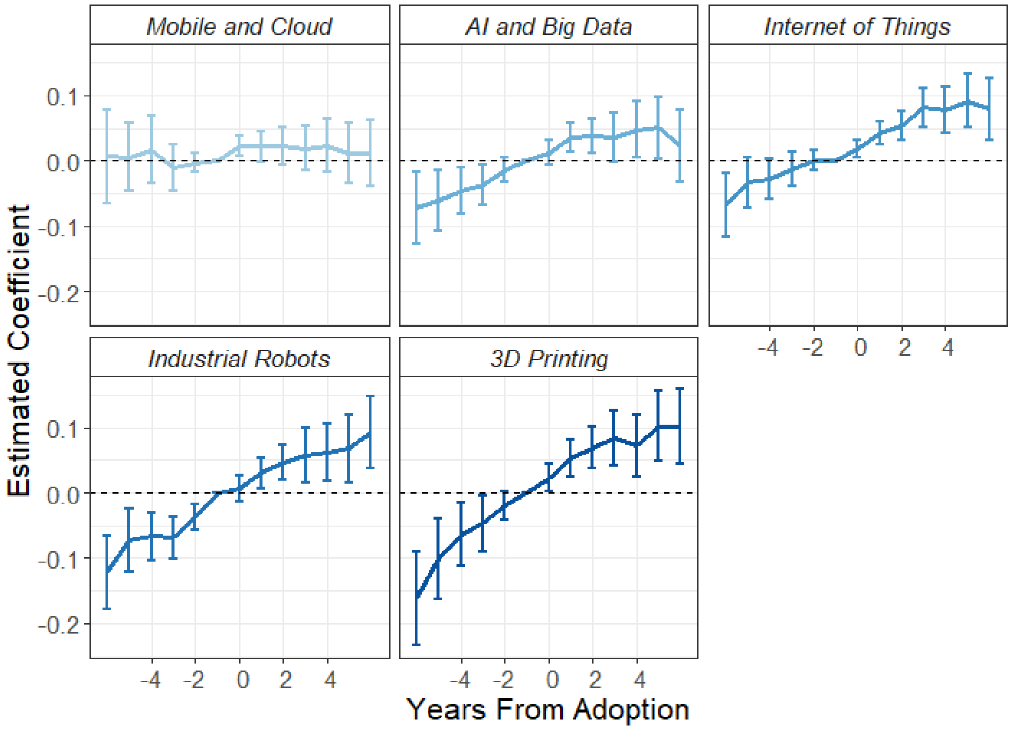

(b) Standard event study design on wages, without controlling for pre-trends (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

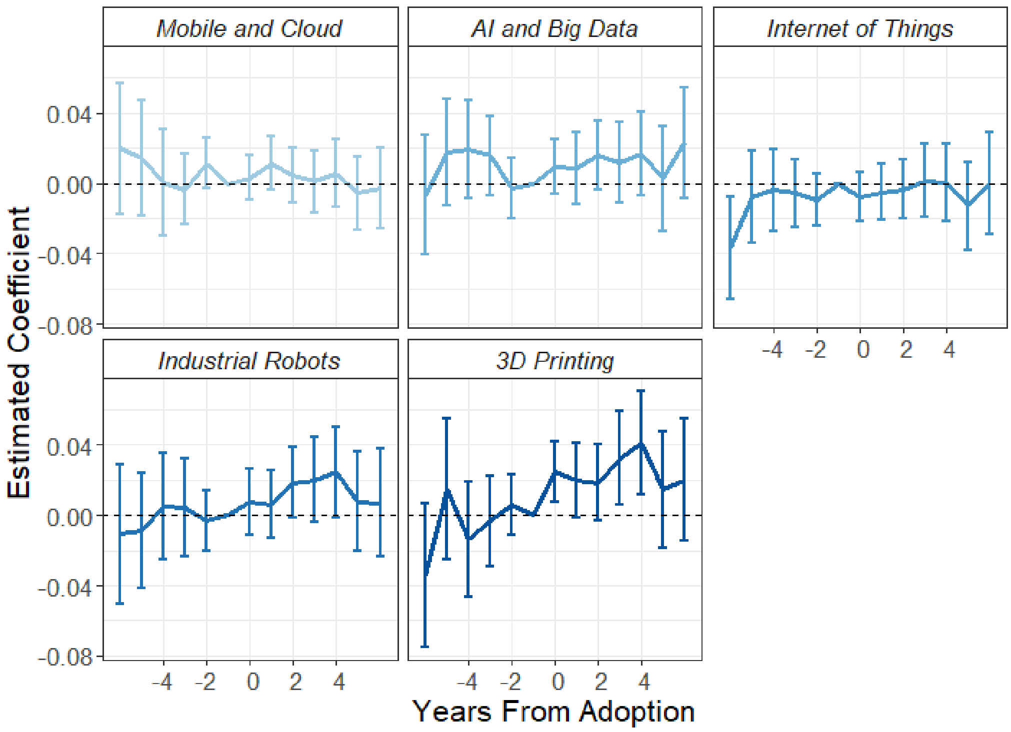

Figure B.37: Estimates from Event Study Regression

(a) Event study design on employment with pre-trends linearized (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

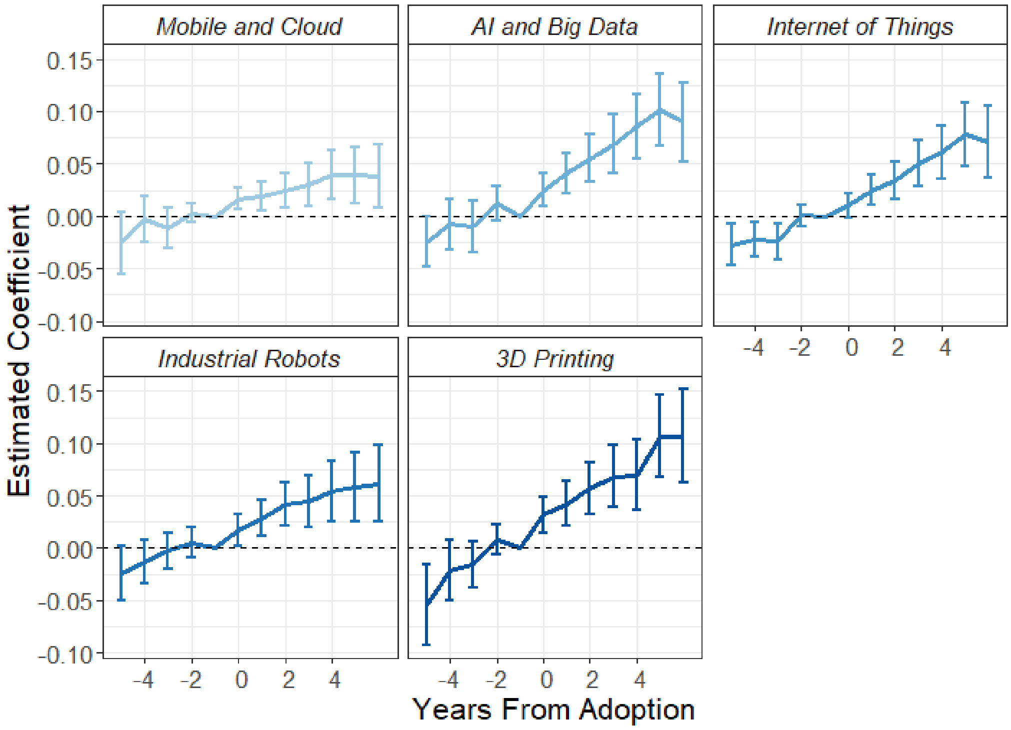

(b) Event study design on hours worked with pre-trends linearized (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

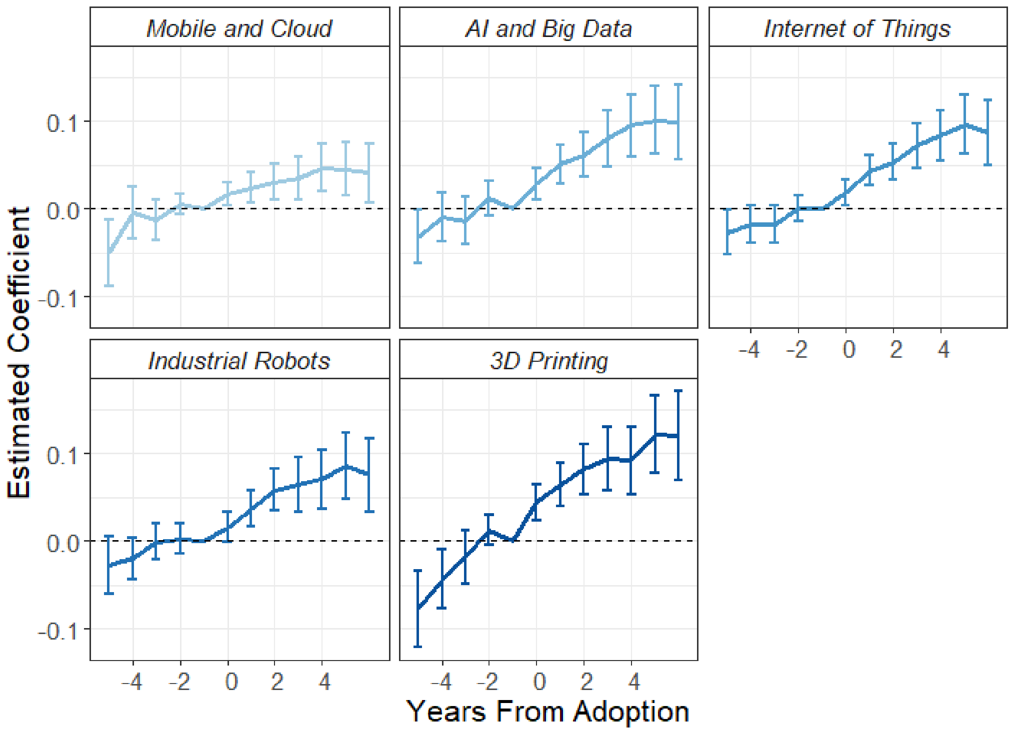

Figure B.38: Estimates from Event Study Regression

(a) Event study design on turnover with pre-trends linearized (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

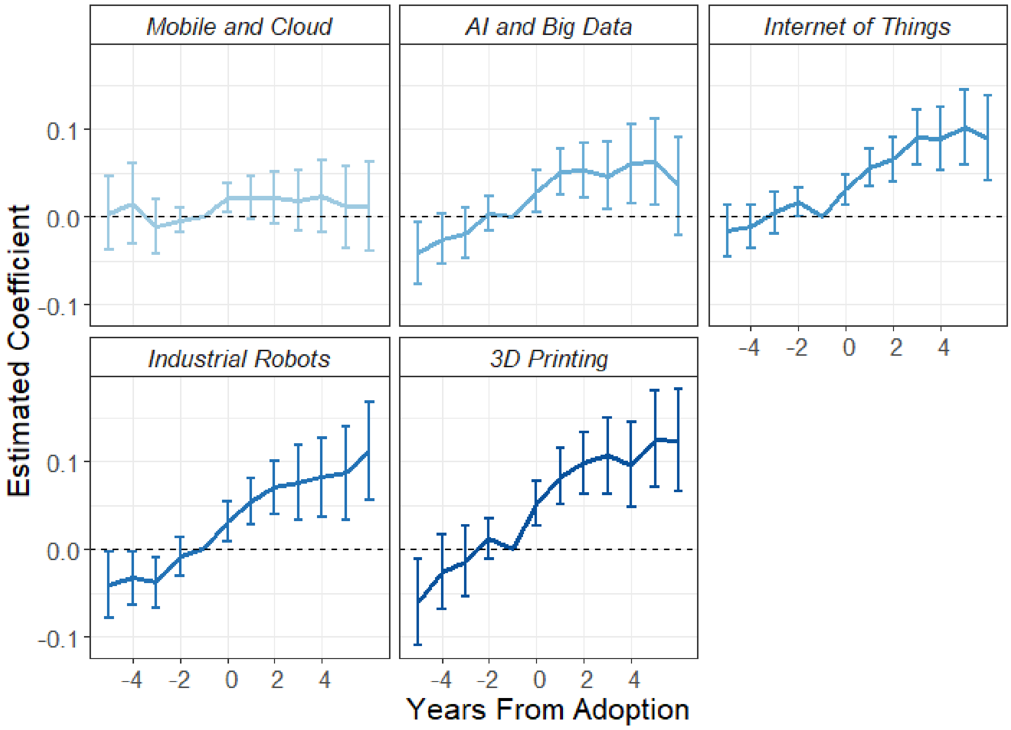

(b) Event study design on wages with with pre-trends linearized (regression run on 3,197 - 3,271 firms and 24,036 - 24,571 firm-year observations.)

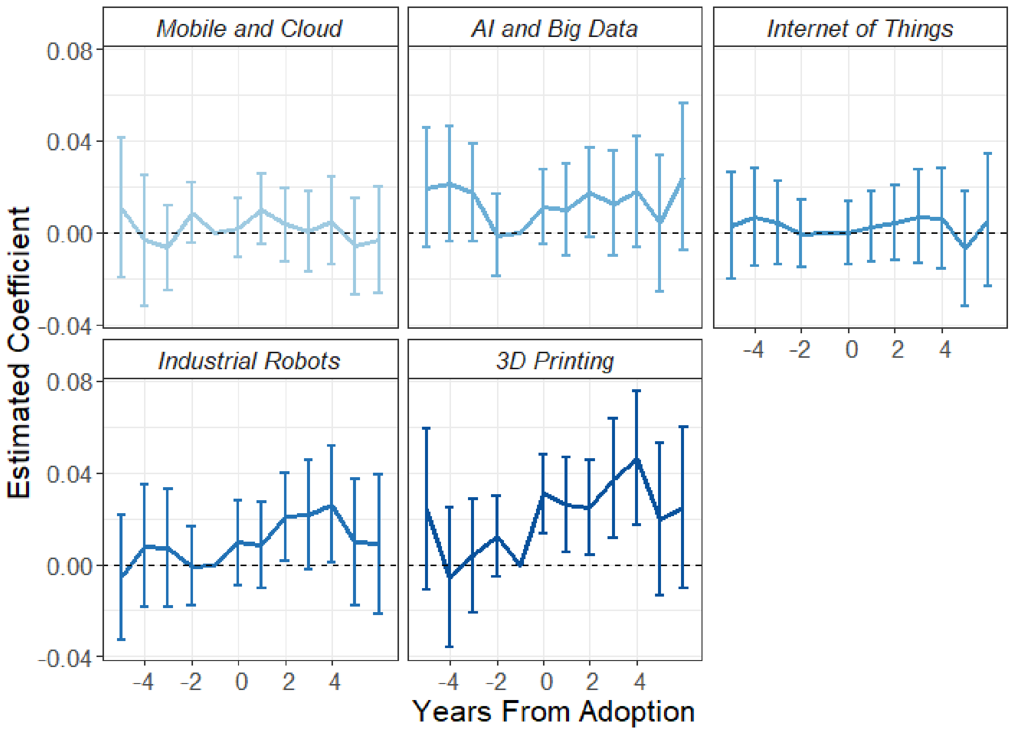

B.6 Appendix - Reweighting Strategy

Although there are clearly significant differences between firms that automate and those that do not, it may be that the sample is not representative, even though the survey itself is balanced. This is because not all of the panel is asked these automation-related questions: construction firms are excluded, and some sectors are under-represented. I compute a logit regression on the binary variable of response to the questions on technology adoption, to see which firm characteristics are related to non-response. I find consistent results across all technologies: firms in the North-East, and in sectors SS2 (textiles, clothing, and hide, leather and footwear products) and SS7 (energy and extraction) are less likely to respond to these questions. The full results are included in Table B.21.

Therefore, firms that do not answer the automation-related questions will be dropped from the subsequent panel analysis. Dealing with potentially non-random non-response requires adjusting the weights, using known information about the non-responding firms.

To this end, I update the weights supplied in the survey. I adjusts these weights depending on the scope of firms considered in any piece of analysis. Following Dargatz and Hill (1996), I use adjustment classes (sector × size) and inflate the weights of all respondents (within an adjustment class) by the class non-response adjustment factor:[34]

$$CAF_h = \frac{\sum_i w_{i,h,eligibles}}{\sum_i w_{i,h,respondents}}$$

where $w_{i,h,x}$ is the weight of firm i in adjustment class/strata h and $x \in$ [eligibles,respondents] denotes if the firm was eligible to respond to the survey, or if they actually responded. Then the non-response adjusted weight of each firm is the original weight inflated by the relevant class non-response adjustment factor: $\widetilde{w_{i,h}} = w_{i,h} \times CAF_h$.

[34] I utilise the strata in the survey for this exercise, of which there are 66 combinations of 2-digit sector and firm size.

Given that firm non-response varies systematically with sectors (see Table B.21 below), this class non-response adjustment factor should account for the selection in the survey.

Responding to the technology adoption question is regressed on industry, region, size and age of firms. The following results include just the industry and region results - the only significant results.

Table B.21: Logit Regression: Answering Technology Adoption Questions 2015

	Mobile & Cloud	*AI & Big Data*	*Internet of Things*	*Industrial Robotics*	*3D Printing*
SS2	-0.719^{***}	-0.742^{***}	-0.744^{***}	-0.643^{**}	-0.721^{***}
SS3	0.022	−0.135	−0.062	0.087	0.0215
SS4	−0.422	−0.381	−0.403	−0.290	−0.336
SS5	−0.201	−0.182	−0.168	−0.141	−0.111
SS6	-0.432^{*}	-0.455^{*}	-0.442^{*}	−0.367	−0.414
SS7	-0.960^{***}	-0.879^{***}	-0.906^{***}	-0.842^{***}	-0.873^{***}
SS8	−0.285	−0.267	−0.270	−0.218	−0.260
SS9	−0.299	−0.241	−0.254	−0.198	−0.236
SS10	−0.366	−0.341	−0.403	−0.307	−0.365
SS11	0.003	0.069	0.035	0.088	0.060
North-East	-0.506^{***}	-0.472^{***}	-0.429^{***}	-0.522^{***}	-0.495^{***}
Centre	0.139	0.173	0.183	0.114	0.139
South & Islands	0.276	0.225	0.233	0.209	0.216

*Note: Estimated coefficients from logit model, regressing the response/non-response to the technology adoption question in 2015 on industry, region, size and age of firms. Estimates are significant at levels of 1%: ***, 5%: **, 10% *. Estimates for industry are with respect to the omitted industry SS1, and for region with respect to the omitted region North-West.*

B.7 Appendix - Propensity Score Matching

This section contains details on the Propensity Score Matching (PSM) approach used in Section 2.4. For firms using each technology - along with those adopting 'any' technology - I run a logit on a wide set of covariates for 2015 technological adoption. The regressors are: age, turnover, investment, exporting behaviour, wage, hires per worker, fires per worker, blue-collar proportion of workers, alongside sector and region fixed-effects.

How do I check if these propensity scores are useful? I use them to predict technology adoption in 2017. Various measures of accuracy of this regression are contained below in Table B.22. It is clear that the propensity scores perform reasonably well across most measures, except for precision, which is the proportion of firms predicted to adopt which actually do so. This is likely due to the relatively small share of firms which do use the technologies, hence the F1 Score is preferred to account for the unbalanced data.

Table B.22: Accuracy of Propensity Scores Predicting 2015 Technology Adoption

	Dependent variable: *2015 Tech. Adoption*					
	Any Tech.	*Cloud Computing*	*AI & Big Data*	*Internet of Things*	*Industrial Robotics*	*3D Printing*
Accuracy	0.61		0.57	0.56	0.59	0.58
Sensitivity	0.65		0.76	0.65	0.83	0.89
Specificity	0.59		0.54	0.54	0.55	0.54
Precision	0.49		0.20	0.27	0.23	0.18
Recall	0.94		0.86	0.95	0.98	0.74
F1 Score	0.59		0.33	0.41	0.38	0.31
N	2547		2541	2544	2538	2554

I use the propensity scores to match adopters and non-adopters, and run regressions on how employment varies with technological adoption. The following table contains the number of firms considered for the matching algorithm, and the number matched. Although the sample size is lower than for other analysis in this paper (such as the panel models), the matching algorithm should reduce selection bias.

B.8 Appendix - Further Regression Results

In Section 2.4, I show the results of two-way fixed-effects models to estimate the effect of automation adoption on firm outcomes. Here I include two extensions: (1) I allow the coefficient to vary over time, to capture the time-varying effect of automation technology on firm outcomes, and (2) I instrument technology adoption with a set of plausibly valid and exogenous instruments.

Time-Varying Adoption

I interact the firm technology adoption dummy with the survey year, allowing estimation of time-varying coefficients. In other words, β_t provide estimates of the impact of technology adoption on the dependent variable in year t, by comparing the set of technology adopters to non-adopters in each year. Otherwise, I estimate TWFE models as before.

$$\ln Y_{it} = \alpha_i + \gamma_t + \delta X_{it} + \beta_t \mathbb{1}\text{Tech}_{it} \times \text{Year}_t + \epsilon_{it} \tag{22}$$

The results of these models are shown below, with the series of β_t in Figure B.39. There are no estimates in some years, which arises due to a lack of observations of technology adopters in that year.

The plotted TWFE estimates in Figure B.39 provide evidence that technology adopters are generally larger than non-adopters, and that this pattern is increasing over time. However, estimates from Figure B.40 shows limited impact on blue-collar employment, providing evidence of capital-skill complementarity with these advanced technologies, whereby adoption allows firms to grow by hiring more skilled workers.

Finally, evidence on the association between technology adoption and turnover per worker is presented in Figure B.41. The results cannot provide support for a relationship, other than for Industrial Robots.

Figure B.39: Estimated β_t from heterogeneous effect TWFE model on employment

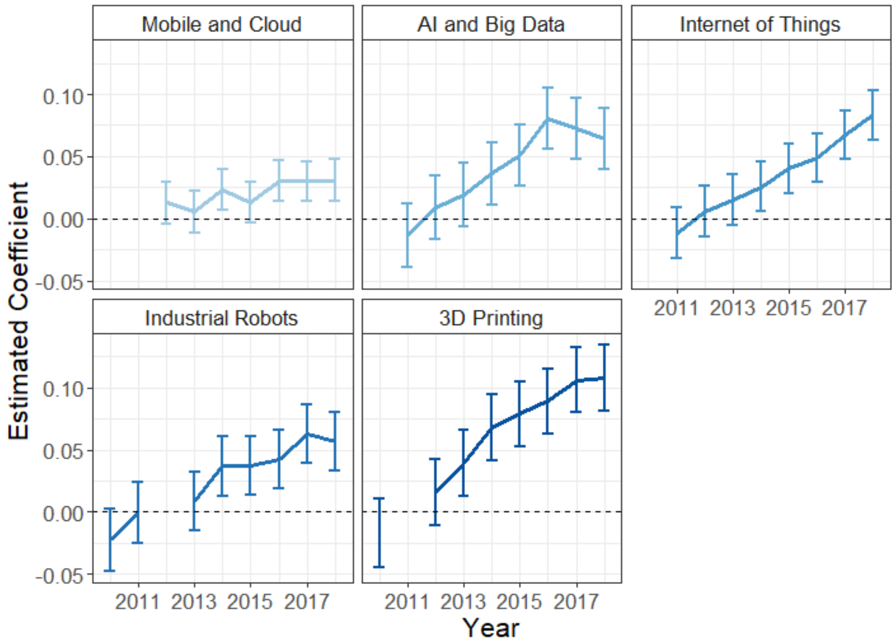

Figure B.40: Estimated β_t from TWFE model on blue-collar employment

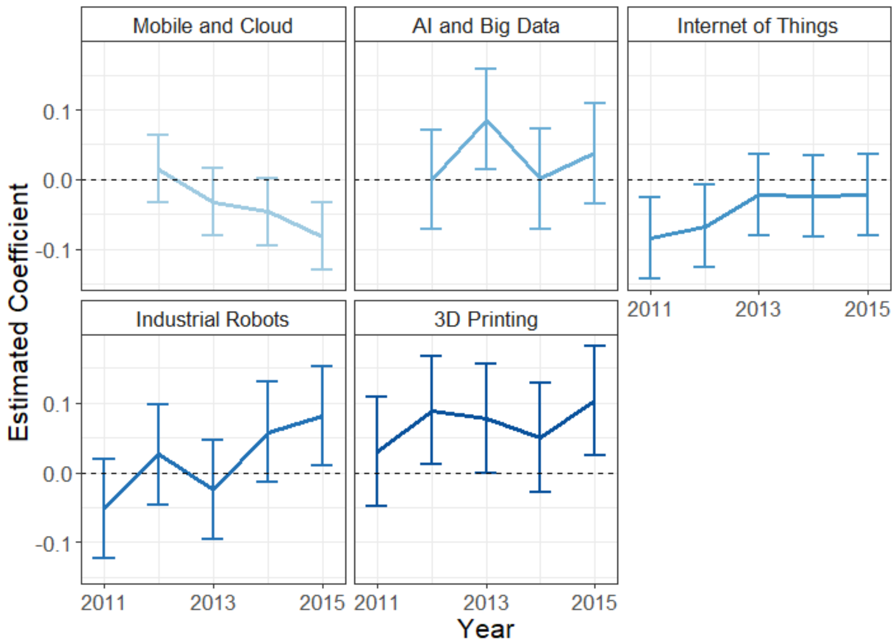

IV Analysis

In Section 2.4, I presented results from TWFE models on the relationship between adoption of automation technologies and firm outcomes: overall employment, productivity and blue-collar employment.

However, there are potential threats to identification in these specifications. It is

Figure B.41: Estimated β_t from TWFE model on turnover per worker

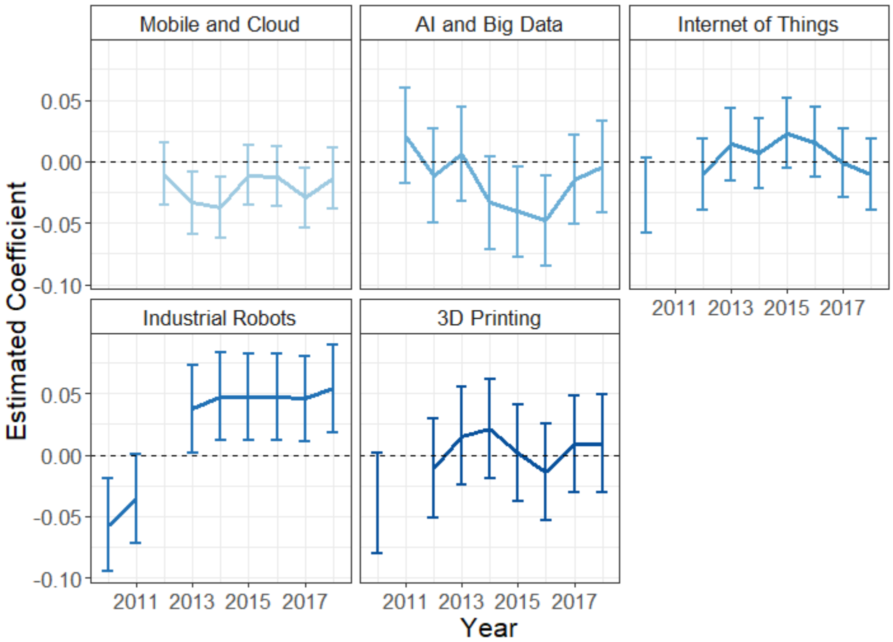

plausible that an omitted variable is correlated with both technology adoption and the various dependent variables. For example, a change in management could lead to a new strategy of aggressive growth and adoption of new technologies. Therefore, I instrument technology adoption with various sets of plausibly valid and exogenous instruments. I leverage the full panel to use a time series of instruments from 2010 to 2018.[35] The instruments used are intangible goods expenditure, expenditure on R&D, and fixed investment expenditure. The identifying assumption is that each of these series is correlated with technological adoption decisions in each year, but do not directly influence employment decisions or firm turnover. This seems possible for the former two: spending on intangible goods and R&D. However, it is reasonable to expect that fixed investment expenditure would not qualify here: to the extent that labour is chosen to equate its marginal product (MPL) to the real wage, and to the degree that investment affects the MPL. Nevertheless, it is possible that lagged values of fixed investment could satisfy the exclusion restriction.

The results for the IV regressions are presented in Table B.23, with the R&D expenditure instrument. It contains the estimated coefficients and associated tests for each

[35]They only cover 2010 - 2015 for blue-collar employment, due to data limitations.

Table B.23: Estimates of β from TWFE models, instrumented by expenditure on R&D

Technology	*Dependent Variable*	*Coefficient*	*F-stat.*	*Cond. F-stat.*	*Hausman*	*Sargan*
Cloud Computing	Employment	0.836**	2.143	0.838	4.709	2.767
	Blue-collar Emp.	−0.647	1.691	0.507	0.817	0.390
	Turnover per worker	−0.036	2.143	0.838	0.008	9.413
AI & Big Data	Employment	0.609*	4.047	1.584	2.685	8.660
	Blue-collar Emp.	−0.306	5.594	1.678	0.451	0.640
	Turnover per worker	−0.01	4.047	1.584	0.000	9.558
IoT	Employment	0.412**	6.108	2.39	3.378	2.390
	Blue-collar Emp.	0.169	7.825	2.347	0.205	0.568
	Turnover per worker	0.008	6.108	2.39	0.002	9.618
Industrial Robotics	Employment	0.516**	5.227	2.045	3.972	10.70
	Blue-collar Emp.	−0.02	6.829	2.049	0.125	0.575
	Turnover per worker	−0.064	5.227	2.045	0.237	9.684
3D Printing	Employment	0.478**	5.869	2.296	3.646	15.60
	Blue-collar Emp.	0.147	8.027	2.408	0.664	0.655
	Turnover per worker	0.033	5.869	2.296	0.005	10.42

Note: Robust standard errors clustered at firm level. Coefficients labelled by statistical significance at: *** 0.1%, ** 1%, * 5%. *The test statistics are: first-stage F-statistic; conditional F-statistic (Sanderson and Windmeijer 2016); Hasuman Test for endogeneity; Sargan statistic for over-identifying restrictions.*

technology, for the three dependent variables. I find significant, and much larger estimated coefficients for the impact of technological adoption on employment, across all technologies. Although there is evidence of OLS endogeneity in these specifications (large Hausman values), and instrument validity (large Sargan statistics), it is likely the results suffer somewhat from a weak instrument problem (F-statistics are lower than typical critical values (Stock and Yogo 2005)).

The estimates for unskilled employment and turnover per worker are all statistically insignificant in the IV setting, while evidence of OLS consistency, and invalid and weak instruments, is quite substantial.

I take this evidence to suggest that there is some endogeneity in my estimates for the impact of technological adoption on total employment, which bias the initial estimates towards zero. However, due to the likelihood my instruments are weak, these IV estimates should be treated with caution.

B.9 Appendix - Turnover per Worker

As an alternative measure of productivity, I consider turnover per worker (Abel, Tenreyro, and Thwaites 2018, e.g.).[36] I regress this proxy on technology adoption, alongside firm age and fixed-effects for sector and region. The estimated coefficients are in Table B.24. The results suggest that adoption of certain advanced technologies is associated with 9 - 17% higher in turnover per worker, and the evidence is significant for Cloud Computing, AI & Big Data, and Internet of Things - the 'Digital' automation technologies. The results are not consistent with either Robotics nor 3D Printing being correlated with higher productivity. However, it should be noted that this measure of productivity contains employment on the denominator, and as technology adoption is associated with firm size, this could mechanically hide any potential relationship between technology and turnover.[37]

Table B.24: Estimate of relationship between technology adoption and productivity

		Cloud Computing	*AI & Big Data*	*IoT*	*Industrial Robotics*	*3D Printing*
2015	Coeff	0.087^{*}	0.103	0.095^{*}	−0.019	−0.070
	SE	(0.046)	(0.100)	(0.059)	(0.074)	(0.080)
2017	Coeff	0.120^{**}	0.169^{**}	0.121^{**}	0.067	0.055
	SE	(0.056)	(0.069)	(0.054)	(0.054)	(0.068)

Note: Log turnover per worker regressed on technology adoption. Robust standard errors computed. Coefficients labelled by statistical significance at: *** 1%, ** 5%, * 10%.

[36] Abel, Tenreyro, and Thwaites (2018) show this correlates with value-added per head, and thus is a good proxy. A simple model can highlight this doesn't always hold: consider Cobb-Douglas production $y = zk^{\alpha}l^{\beta}$. Following a cost-minimisation procedure, and computing turnover per worker $\frac{py}{l}$, we get an expression that is independent of productivity z. The price-setting environment becomes important here, on which I won't elaborate. Suffice to say that this productivity proxy has its drawbacks, hence a more in-depth procedure is explored in the main paper.

[37] Regressing turnover on technology adoption, with a control for firm size, would do no better, as this would lead to issues of multicollinearity.

B.10 Appendix - Productivity Estimation

Firm-level productivity estimation is computed using control function methods. These require firm-level capital stock. This section explains the Perpetual Inventory Method (PIM) used to construct a measure of capital stock, followed by the productivity estimation approaches.

Perpetual Inventory Method

The Perpetual Inventory Method (PIM) allows construction of firm-level capital stocks when such data is unavailable, but investment data is present. The method here follows Martin (2002). The PIM is constructed using the following equation:

$$K_t = (1-\delta)K_{t-1} + i_t.$$

where K_t is the capital stock in period t, and i_t is investment in period t. However, to use this method, we need K_0 - the initial capital stock of a firm - which is not in this survey. To construct this series, each firm's K_0 is an employment-weighted share of total investment in the year they first appear in the survey. Capital stock is then constructed for all future years with the above equation. The depreciation rate is taken to be 10%.

Productivity Estimation

Consider the production function:

$$Y_{it} = Z_{it} K_{it}^{\beta_k} L_{it}^{\beta_l}.$$

where Y_{it}, K_{it}, L_{it} represent revenue, capital stock, and employment respectively, while the β's are the production elasticities to be estimated. Taking logarithms, we get:

$$y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \epsilon_{it}.$$

where $\ln z_{it} = \beta_0 + \epsilon_{it}$. Firms draw productivity Z_{it} which is unobserved by the econometrician, leading to potential omitted variable bias, as clearly the optimal firm input choices will be correlated with this variable.

We proceed with the control function method of Olley and Pakes (1996), which makes assumptions on the timing of input choices to achieve identification, and uses investment as a proxy for unobserved productivity shocks. With this approach, we split up the unobserved residual $\epsilon_{it} = \omega_{it} + \eta_{it}$, where ω_{it} is anticipated and η_{it} is an ex-post shock. Thus, inputs are correlated with ω_{it} only. The following assumptions are required:

1. **Information Sets:** firms' information sets I_{it} include current and past productivity shocks $\{\omega_{i\tau}\}_{\tau=0}^{t}$, but firms know nothing about future shocks. The ex-post shocks η_{it} are expected to be zero on average: $\mathbb{E}\{\eta_{it}|I_{it}\} = 0$.

2. **First-Order Markov Shocks:** productivity shocks follow a First-Order Markov Process, so $\omega_{it} = \mathbb{E}(\omega_{it}|\omega_{i,t-1}) + \nu_{it}$, and $\mathbb{E}\{\nu_{it}|I_{it-1}\} = 0$.

3. **Timing of Input Choices:** in the previous period $i_{i,t-1}$ determines capital in the current period k_{it}, whereas labour is chosen in the current period.

4. **Scalar Unobservable:** investment decisions $i_{it} = f_t(k_{it}, \omega_{it})$ have just one scalar unobservable ω_{it}, so there is no other across firm unobserved heterogeneity (e.g. adjustment costs, investment efficiency, input prices).

5. **Strict Monotonicity:** investment decisions are strictly monotonic in the scalar unobservable ω_{it}, so $i_{it} = f_t(k_{it}, \omega_{it})$.

Given that investment is strictly monotonic in the unobserved anticipated shock, this function can be inverted, and then substituted into the production function:

$$y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + f_t^{-1}(k_{it}, i_{it}) + \eta_{it}.$$

This inverted function is unknown, so is approximated by a polynomial in capital and investment:

$$y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \gamma_0 + \gamma_1 k_{it} + \gamma_2 i_{it} + \gamma_3 k_{it}^2 + \gamma_4 i_{it}^2 + \gamma_5 i_{it} k_{it} + \eta_{it}.$$

and estimation takes the standard two-step process.

The first-step OLS regression over the above equation yields an estimate $\widehat{\beta_l}$ and an estimate of the "composite" term $\widehat{\Phi}_{it} = \beta_0 + \widehat{\beta_k} k_{it} + \omega_{it}$. To estimate β_k, we calculate 'implied' ω_{it}'s:

$$\widehat{\omega}_{it}(\beta_k) = \widehat{\Phi}_{it} - \widehat{\beta_k} k_{it}.$$

Then, by the First-Order Markov Process of the productivity shocks, we can non-parametrically regress the implied $\widehat{\omega}_{it}(\beta_k)$'s on their lag $\widehat{\omega}_{it-1}(\beta_k)$, and the residuals $\widehat{\nu}_{it}(\beta_k)$ are the implied innovations in productivity. Finally, the sample analogue of the moment condition $\mathbb{E}\{\nu_{it} k_{it}\} = 0$ is:

$$\frac{1}{N}\frac{1}{T}\sum_i \sum_t \widehat{\nu}_{it}(\widehat{\beta_k}) k_{it} = 0.$$

and we find $\widehat{\beta_k}$ to solve this problem. The resulting estimated factor elasticities are:

Table B.25: Results from production function estimation.

	$\widehat{\beta_l}$	$\widehat{\beta_k}$
Coefficient	0.71	0.196
Standard Error	(0.057)	(0.025)

Note: Production function estimation performed with Olley and Pakes (1996) method.

C Appendix - Chapter 3

C.1 Appendix - Production Function Estimation

Consider the production function:

$$Y_{it} = e^{\eta_{it}} e^{\omega_{it}} K_{it}^{\beta_k} L_{it}^{\beta_l} M_{it}^{\beta_m}$$

where $Y_{it}, K_{it}, L_{it}, M_{it}$ represent revenue, capital stock, employment, and materials respectively, while the β's are the production elasticities. ω_{it} are ex-ante shocks. η_{it} are ex-post shocks. Taking logarithms yields a regression equation with the log of gross revenue on the left-hand side, and logs of the factor inputs on the right-hand side:

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \omega_{it} + \eta_{it}.$$

Firms draw productivity ω_{it} which is unobserved by the econometrician, leading to potential omitted variable bias, as clearly the optimal firm input choices will be correlated with this variable.

I should briefly comment on how this links to the theoretical section. The theoretical concept of net output $y = zF(k, \ell) - \phi$ can be mapped onto this production function quite simply: the function $F(\cdot)$ is Cobb-Douglas (with materials as an extra input, but the results still hold); z is unobserved idiosyncratic productivity which must be controlled for; and $y + \phi$ is the gross output measure observed in the data.

Control Function Approach

I proceed with an explanation of the control function method of Ackerberg, Caves, and Frazer (2015), which makes assumptions on the timing of input choices to achieve identification, and uses materials expenditure as a proxy for unobserved productivity shocks. The following assumptions are required:

1. **Information Sets:** firms' information sets $\mathcal{I}_{it}$ include current and past produc-

tivity shocks $\{\omega_{i\tau}\}_{\tau=0}^{t}$, but firms know nothing about future shocks. The ex-post shocks η_{it} are expected to be zero on average: $\mathbb{E}\{\eta_{it}|\mathcal{I}_{it}\} = 0$.

2. **First-Order Markov Shocks:** productivity shocks follow a First-Order Markov Process, so $\omega_{it} = \mathbb{E}(\omega_{it}|\omega_{i,t-1}) + v_{it}$, and $\mathbb{E}\{v_{it}|\mathcal{I}_{it-1}\} = 0$.

3. **Timing of Input Choices:** firms accumulate capital according to $k_{it} = \kappa(k_{it-1}, i_{it-1})$ where investment i_{it-1} is chosen in period $t-1$. Labour l_{it} is chosen at period $t, t-1$ or in between. m_{it} is either chosen at the same time, or after l_{it} is chosen.

4. **Scalar Unobservable:** investment decisions $m_{it} = h_t(k_{it}, \omega_{it}, l_{it})$ have just one scalar unobservable ω_{it}, so there is no other across firm unobserved heterogeneity (e.g. adjustment costs, investment efficiency, input prices).

5. **Strict Monotonicity:** investment decisions are strictly monotonic in the scalar unobservable ω_{it}, so $m_{it} = h_t(k_{it}, \omega_{it}, l_{it})$.

Given that investment is strictly monotonic in the unobserved anticipated shock, this function can be inverted, and then substituted into the production function:

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + h_t^{-1}(k_{it}, m_{it}, l_{it}) + \eta_{it} = \Phi_t(k_{it}, \omega_{it}, l_{it}) + \eta_{it}.$$

This inverted function is unknown, so is approximated by a polynomial. Running the first-step regression yields an estimate of the composite term Φ_t. This exploits the moment condition $\mathbb{E}(\eta_{it}|\mathcal{I}_{it}) = 0$, because the ex-post productivity shock is unanticipated to the firm. Clearly neither β_k nor β_l are identified in this stage, as both are contained in the composite term:

$$\widehat{\Phi}_t = \beta_k k_{it} + \beta_l l_{it} + \omega_{it}.$$

Production function parameters are estimated in the second stage. With the estimate of the composite term, estimates of the ex-ante productivity shock $\widehat{\omega}_{it}(\beta_k, \beta_l)$ can be computed for guesses of β_k, β_l. The implied $\widehat{\omega}_{it}(\beta_k, \beta_l)$ are non-parametrically regressed on their lag $\widehat{\omega}_{it-1}(\beta_k, \beta_l)$, and the residuals $\widehat{v}_{it}(\beta_k, \beta_l)$ are the implied innovations in productivity. The second stage moment condition is $\mathbb{E}(v_{it} + \eta_{it}|\mathcal{I}_{it-1}) = 0$. The sample analogue of the moment condition $\mathbb{E}\{v_{it}k_{it}\} = 0$ is:

$$\frac{1}{N}\frac{1}{T}\sum_i\sum_t \widehat{v_{it}}(\beta_k, \beta_l)k_{it} = 0.$$

If labour is assumed to be chosen after $t-1$, l_{it} will generally be correlated with v_{it}, so lagged labour is chosen as an additional moment condition. This procedure yields estimates $\widehat{\beta_k}, \widehat{\beta_l}$.

Revenue elasticity is computed by summing the estimated revenue elasticities $\widehat{\beta_k}$ and $\widehat{\beta_l}$. This is computed across the whole sample, but also on four 'macro sectors', and disaggregated by 2-digit industries. Other research looking at U.S. manufacturing has found the revenue elasticity to be in the range of 0.95 (e.g. Ruzic and Ho 2019; Gao and Kehrig 2020), implying diminishing returns, although industrial heterogeneity gives results both above and below this value.

This estimation procedure also allows us to extract firm-level productivity estimates. Productivity is computed from the estimated ω_{it}, and then aggregated over industries to investigate productivity trends.

In order to obtain revenue elasticity at the level of the firm and year, a slightly different estimation procedure is required. I need to estimate time- and firm-specific revenue elasticities $\theta^k_{it}, \theta^l_{it}$. This can be achieved by generalising the production function to translog, as in De Loecker and Warzynski (2012):

$$y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_{ll} l_{it}^2 + \beta_{kk} k_{it}^2 + \beta_{lk} l_{it} k_{it} + \eta_{it}.$$

Then follow the same procedure as used with a Cobb-Douglas production function to estimate the coefficients and productivity process. The time- and firm-specific output elasticities are easily computed:

$$\widehat{\theta^k_{it}} = \widehat{\beta_k} + 2\widehat{\beta_{kk}} k_{it} + \widehat{\beta_{lk}} l_{it}.$$

$$\widehat{\theta^l_{it}} = \widehat{\beta_l} + 2\widehat{\beta_{ll}} l_{it} + \widehat{\beta_{lk}} k_{it}.$$

Allowing heterogeneity across firms and over time in the output elasticities permits computation of distributions of revenue elasticity, and analysing changes over time.

Cost Share Approach

In this section, I describe the Gandhi, Navarro, and Rivers (2020) cost share approach to estimating production functions. The production environment is the same as before, and the productivity process is also Markov. Capital and labour are chosen at $t-1$, while materials is chosen at t. Materials is strictly increasing in ex-ante productivity shock ω_{it}.

As shown in Section 3.4.2, this method takes the first-order condition from choosing materials input to maximise profits to obtain the following relationship:

$$s_{it} = \ln\left(\frac{\partial f(l_{it}, k_{it}, m_{it})}{\partial m_{it}}\right) + \ln \mathbb{E}(\eta_{it}|\mathcal{I}_{it}) - \eta_{it}$$

where $s_{it} = \ln \frac{P_t^M M_{it}}{P_t Y_{it}}$ is the log of the materials cost share in revenue. This can be written as:

$$s_{it} = \ln D(l_{it}, k_{it}, m_{it}) - \eta_{it} \tag{23}$$

Since $\mathbb{E}(\eta_{it}|\mathcal{I}_{it}) = \mathbb{E}(\eta_{it}|l_{it}, k_{it}, m_{it}) = 0$, Equation (23) is a regression equation of the log of the materials cost share in revenue on the inputs l_{it}, k_{it}, m_{it}.

Notice that the constant $\mathbb{E}(\eta_{it}|\mathcal{I}_{it})$ can be obtained because:

$$\mathbb{E}(\eta_{it}|\mathcal{I}_{it}) = \mathbb{E}\exp(\ln D(l_{it}, k_{it}, m_{it})) - s_{it}$$

Then the materials input elasticity can be identified:

$$\left(\frac{\partial f(l_{it}, k_{it}, m_{it})}{\partial m_{it}}\right) = \frac{D(l_{it}, k_{it}, m_{it})}{\mathbb{E}(\eta_{it}|\mathcal{I}_{it})} \tag{24}$$

Gandhi, Navarro, and Rivers (2020) propose regressing the log of the materials cost share on logs of a polynomial in capital, labour, and materials.

The rest of the production function can be identified as follows. Integrating over the materials input elasticity yields an expression which represents the part of the production function related to the materials input:

$$\int \frac{\partial f\left(l_{it}, k_{it}, m_{it}\right)}{\partial m_{it}} dm_{it} = f\left(l_{it}, k_{it}, m_{it}\right) + g\left(k_{it}, l_{it}\right) \tag{25}$$

Combining this with the production function yields:

$$\mathcal{Y}_{it} = y_{it} - \eta_{it} - \int \frac{\partial f\left(l_{it}, k_{it}, m_{it}\right)}{\partial m_{it}} dm_{it} = \omega_{it} - g\left(k_{it}, l_{it}\right) \tag{26}$$

where $\mathcal{Y}_{it}$ is a function of the data: output, minus the ex-post shock, minus the materials elasticity already obtained in the first stage. Gandhi, Navarro, and Rivers (2020) suggest using a polynomial sieve estimator so that the integral has a closed-form solution.

Using the Markov process on productivity, alongside the moment condition on the 'surprise' to ex-ante productivity $\mathbb{E}(v_{it}|k_{it}, l_{it}, \mathcal{Y}_{it-1}, k_{it-1}, l_{it-1}) = 0$, we obtain:

$$\mathbb{E}\left(\mathcal{Y}_{it}|k_{it}, l_{it}, \mathcal{Y}_{it-1}, k_{it-1}, l_{it-1}\right) = -g(k_{it}, l_{it}) + d(\mathcal{Y}_{it-1} + g(k_{it-1}, l_{it-1}) \tag{27}$$

so we regress $\mathcal{Y}_{it}$ on $k_{it}, l_{it}, \mathcal{Y}_{it-1}, k_{it-1}, l_{it-1}$.

The moment conditions for this stage are:

$$\mathbb{E}\left(\eta_{it}\widehat{\mathcal{Y}}_{it-1}\right) = 0$$

$$\mathbb{E}\left(\eta_{it}k_{it}^{\tau_k}l_{it}^{\tau_l}\right) = 0$$

where the second moment condition includes polynomials of capital and labour, up to τ_k, τ_l.

C.2 Appendix - Revenue Elasticity Estimates

Table C.1: Production Function Coefficients by Macro Sector: Cobb-Douglas production function

	ACF (VA)	ACF (GO)	GNR (GO)
		Manufacturing ($N = 120,712$)	
β_l	0.789	0.399	0.297
	(0.014)	(0.061)	(0.005)
β_k	0.421	0.031	0.148
	(0.011)	(0.044)	(0.005)
β_m	-	0.759	0.590
	-	(0.121)	(0.003)
		Construction ($N = 51,784$)	
β_l	0.826	0.283	0.328
	(0.047)	(0.061)	(0.016)
β_k	0.388	0.161	0.224
	(0.027)	(0.044)	(0.007)
β_m	-	0.431	0.493
	-	(0.015)	(0.005)
		Wholesale/Trade/Transport ($N = 181,985$)	
β_l	0.669	0.269	0.198
	(0.014)	(0.575)	(0.006)
β_k	0.343	0.074	0.130
	(0.021)	(0.137)	(0.006)
β_m	-	0.510	0.688
	-	(0.034)	(0.006)
		Services ($N = 173,332$)	
β_l	0.681	0.339	0.446
	(0.005)	(0.012)	(0.005)
β_k	0.384	0.060	0.215
	(0.003)	(0.009)	(0.003)
β_m	-	0.398	0.354
	-	(0.012)	(0.002)

Note: Estimated coefficients on labour, capital and materials from Cobb-Douglas production functions, with gross output (GO) or value-added (VA). Standard errors in brackets. Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF) or Gandhi, Navarro, and Rivers (2020) (GNR).

Table C.2 contains revenue elasticity estimates by industry, at the 2-digit SIC level. The number of firms on which estimation was computed is included. If the factor elasticity on labour or capital was outside the range of $[0, 1]$, then the RTS was not computed.

Table C.2: Revenue Elasticity by 2-digit SIC

	Revenue Elasticities			
SIC	ACF (GO)	ACF (VA)	GNR (GO)	N
10	1.162	1.034	1.039	12,495
11	1.117	1.202	1.092	1,724
13	1.061	1.192	1.020	4,981
14	0.806	1.371	0.996	3,355
15	1.265	1.019	1.115	841
16	1.298	1.235	1.057	3,478
17	1.033	1.128	1.012	4,184
18	1.279	0.893	1.054	7,521
19	1.019	1.081	-	506
20	1.071	1.420	1.016	5,733
21	1.064	1.162	1.061	986
22	-	1.173	1.026	7,776
23	0.771	1.238	1.016	5,616
24	0.842	-	1.000	4,776
25	-	1.367	1.019	15,597
26	1.245	1.084	1.036	7,648
27	0.876	1.123	1.011	4,913
28	-	0.916	1.023	10,899
29	1.090	1.085	1.181	1,633
30	1.130	1.437	1.091	1,973
31	1.244	-	1.050	4,060
32	1.169	1.349	1.055	5,020
33	-	1.243	1.065	4,997
41	0.882	1.131	1.067	12,216
42	0.864	0.934	1.038	12,554
45	0.978	1.020	1.036	24,639
46	0.817	0.843	1.041	68,969
47	0.916	1.010	1.061	66,171

	Revenue Elasticities			
SIC	ACF (GO)	ACF (VA)	GNR (GO)	N
49	0.970	1.216	1.046	11,501
50	1.087	-	1.094	1,306
51	-	1.047	1.107	807
52	0.894	1.210	1.066	8,103
53	-	1.265	1.429	489
55	0.836	1.520	1.011	8,549
56	0.685	1.381	0.971	25,219
58	0.918	1.127	1.032	6,802
59	1.051	1.063	1.008	2,547
60	1.040	1.292	1.086	693
61	-	-	1.131	1,062
62	-	-	1.103	9,061
63	1.054	1.066	1.119	1,224
69	0.853	0.527	1.045	10,295
70	0.986	1.016	1.053	10,274
71	0.785	1.319	1.032	11,953
72	0.835	0.892	1.022	2,323
73	-	0.946	1.054	5,168
74	0.852	1.017	1.079	4,769
75	1.032	-	0.987	1,482
77	1.006	1.016	1.041	6,195
78	1.031	1.091	1.010	9,842
79	0.800	1.151	1.094	4,136
80	1.075	1.095	1.072	1,926
81	0.981	1.142	1.042	6,472
82	0.902	0.961	1.109	9,624
90	0.935	0.846	0.936	3,111
91	0.761	-	-	1,722
92	1.097	1.165	1.030	1,248
93	0.946	1.066	1.025	7,853
94	-	1.264	1.045	6,086
95	1.261	1.077	1.117	1,889
96	1.237	1.052	0.979	11,807

Note: Estimates of revenue elasticity across 2-digit SICs, following the Ackerberg, Caves, and Frazer (2015) and Gandhi, Navarro, and Rivers (2020) approaches with a Cobb-Douglas gross-output production function. Missing sectors have estimated coefficients on labour or capital that are negative or greater than one.

Table C.3: Ackerberg, Caves, and Frazer (2015) (GO) Production Function Coefficients by 2-digit SIC

	ACF (GO) Coefficients			
SIC	β_l	β_k	β_m	N
10	0.311	0.074	0.777	12,495
	(0.026)	(0.037)	(0.056)	
11	0.325	0.139	0.653	1,724
	(0.085)	(0.048)	(0.039)	
13	0.556	0.291	0.215	4,981
	(0.251)	(0.105)	(0.132)	
14	0.185	0.078	0.543	3,355
	(0.440)	(0.143)	(0.164)	
15	0.317	0.478	0.47	841
	(0.049)	(0.022)	(0.014)	
16	0.699	0.119	0.48	3,478
	(0.124)	(0.064)	(0.007)	
17	0.344	0.115	0.574	4,184
	(0.051)	(0.096)	(0.059)	
18	0.457	0.293	0.53	7,521
	(0.047)	(0.037)	(0.007)	
19	0.102	0.131	0.786	506
	(0.052)	(0.057)	(0.022)	
20	0.219	0.168	0.683	5,733
	(0.097)	(0.053)	(0.030)	
21	0.293	0.269	0.502	986
	(0.180)	(0.076)	(0.049)	
22	0.465	-	0.644	7,776
	(0.015)	-	(0.034)	
23	0.251	0.001	0.519	5,616
	(0.293)	(0.142)	(0.096)	
24	0.302	0.023	0.517	4,776
	(0.039)	(0.050)	(0.069)	
25	0.537	-	0.655	15,597
	(0.047)	-	(0.022)	
26	0.554	0.112	0.578	7,648
	(0.129)	(0.079)	(0.110)	
27	0.266	0.146	0.465	4,913
	(0.261)	(0.301)	(0.089)	
28	0.376	-	0.359	10,899
	(0.229)	-	(0.073)	
29	0.215	0.073	0.802	1,633
	(0.102)	(0.020)	(0.057)	
30	0.607	0.298	0.225	1,973
	(0.227)	(0.076)	(0.130)	
31	0.421	0.248	0.575	4,060
	(0.106)	(0.053)	(0.116)	
32	0.522	0.026	0.622	5,020
	(0.085)	(0.038)	(0.029)	

ACF (GO) Coefficients				
SIC	β_l	β_k	β_m	N
33	0.598	-	0.581	4,997
	(0.302)	-	(0.122)	
41	0.201	0.185	0.496	12,216
	(0.013)	(0.020)	(0.017)	
42	0.147	0.155	0.562	12,554
	(0.008)	(0.006)	(0.015)	
43	0.284	0.061	0.632	27,014
	(0.120)	(0.012)	(0.022)	
45	0.284	0.061	0.632	24,639
	(0.012)	(0.034)	(0.061)	
46	0.127	0.052	0.638	68,969
	(0.214)	(0.063)	(0.061)	
47	0.183	0.185	0.548	66,171
	(0.001)	(0.008)	(0.006)	
49	0.361	0.038	0.572	11,501
	(0.052)	(0.022)	(0.068)	
50	0.200	0.226	0.661	1,306
	(0.196)	(0.091)	(0.105)	
51	0.251	-	0.958	807
	(0.183)	-	(0.390)	
52	0.321	0.229	0.343	8,103
	(0.097)	(0.008)	(0.145)	
53	0.683	-	0.730	489
	(0.153)	-	(0.167)	
55	0.258	0.075	0.503	8,549
	(0.062)	(0.054)	(0.006)	
56	0.114	0.084	0.487	25,219
	(0.012)	(0.009)	(0.014)	
58	0.336	0.208	0.373	6,802
	(0.030)	(0.207)	(0.202)	
59	0.415	0.215	0.420	2,547
	(0.040)	(0.074)	(0.065)	
60	0.317	0.374	0.349	693
	(0.227)	(0.070)	(0.143)	
61	-	0.171	0.773	1,062
	-	(0.100)	(0.218)	
62	-	-	0.139	9,061
	-	-	(0.003)	
63	0.39	0.221	0.443	1,224
	(0.059)	(0.114)	(0.030)	
69	0.334	0.266	0.252	10,295
	(0.143)	(0.024)	(0.040)	
70	0.378	0.172	0.436	10,274
	(0.504)	(0.271)	(0.114)	

	ACF (GO) Coefficients			
SIC	β_l	β_k	β_m	N
71	0.303	0.258	0.223	11,953
	(0.146)	(0.014)	(0.030)	
72	0.276	0.219	0.339	2,323
	(0.137)	(0.053)	(0.055)	
73	0.292	-	0.428	5,168
	(0.083)	-	(0.216)	
74	0.265	0.225	0.362	4,769
	(0.113)	(0.053)	(0.148)	
75	0.617	0.253	0.162	1,482
	(0.116)	(0.079)	(0.021)	
77	0.179	0.326	0.500	6,195
	(0.054)	(0.020)	(0.116)	
78	0.429	0.059	0.544	9,842
	(0.003)	(0.012)	(0.086)	
79	0.190	0.100	0.511	4,136
	(0.012)	(0.041)	(0.072)	
80	0.431	0.227	0.417	1,926
	(0.021)	(0.029)	(0.051)	
81	0.263	0.029	0.689	6,472
	(0.132)	(0.034)	(0.144)	
82	0.285	0.179	0.438	9,624
	(0.038)	(0.022)	(0.074)	
90	0.077	0.432	0.425	3,111
	(0.057)	(0.227)	(0.022)	
91	0.129	0.184	0.449	1,722
	(0.030)	(0.031)	(0.056)	
92	0.399	0.099	0.598	1,248
	(0.111)	(0.038)	(0.009)	
93	0.249	0.266	0.431	7,853
	(0.030)	(0.031)	(0.136)	
94	-	0.476	0.458	6,086
	-	(0.028)	(0.046)	
95	0.997	0.144	0.130	1,889
	(0.352)	(0.148)	(0.123)	
96	0.552	0.447	0.239	11,807
	(0.049)	(0.068)	(0.101)	

Note: Estimated coefficients from production function across 2-digit SICs, following Ackerberg, Caves, and Frazer (2015) with a Cobb-Douglas gross-output production function. Standard errors in brackets. Missing values are point estimates that are less than zero or greater than one.

Table C.4: Ackerberg, Caves, and Frazer (2015) (VA) Production Function Coefficients by 2-digit SIC

	ACF (VA) Coefficients		
SIC	β_l	β_k	N
10	0.694	0.340	12,495
	(0.025)	(0.029)	
11	0.731	0.471	1,724
	(0.293)	(0.229)	
13	0.848	0.344	4,981
	(0.038)	(0.038)	
14	0.990	0.381	3,355
	(0.388)	(0.405)	
15	0.640	0.380	841
	(0.062)	(0.076)	
16	0.793	0.442	3,478
	(0.045)	(0.042)	
17	0.604	0.524	4,184
	(0.137)	(0.132)	
18	0.544	0.349	7,521
	(0.098)	(0.077)	
19	0.694	0.387	506
	(0.681)	(0.751)	
20	0.942	0.478	5,733
	(0.701)	(1.000)	
21	0.813	0.350	986
	(0.033)	(0.036)	
22	0.796	0.376	7,776
	(0.008)	(0.009)	
23	0.974	0.264	5,616
	(0.021)	(0.022)	
24	0.797	-	4,776
	(0.126)	-	
25	0.589	0.779	15,597
	(0.027)	(0.046)	
26	0.812	0.272	7,648
	(0.017)	(0.016)	
27	0.812	0.311	4,913
	(0.048)	(0.051)	
28	0.593	0.324	10,899
	(0.025)	(0.023)	
29	0.987	0.097	1,633
	(0.209)	(0.182)	
30	0.965	0.472	1,973
	(0.056)	(0.052)	
32	0.891	-	5,020
	(0.034)	-	

	ACF (VA) Coefficients		
SIC	β_l	β_k	N
33	0.771	0.578	4,997
	(0.041)	(0.061)	
41	0.989	0.254	12,216
	(0.010)	(0.014)	
42	0.663	0.468	12,554
	(0.094)	(0.083)	
43	-	0.340	27,014
	-	(0.093)	
45	0.736	0.284	24,639
	(0.133)	(0.192)	
46	0.584	0.259	68,969
	(0.013)	(0.038)	
47	0.656	0.354	66,171
	(0.101)	(0.055)	
49	0.735	0.481	11,501
	(0.129)	(0.107)	
50	-	0.183	1,306
	-	(0.200)	
51	0.666	0.381	807
	(0.169)	(0.520)	
52	0.816	0.394	8,103
	(0.083)	(0.083)	
53	0.729	0.536	489
	(0.208)	(0.156)	
55	0.978	0.541	8,549
	(0.043)	(0.032)	
56	0.834	0.547	25,219
	(0.111)	(0.182)	
58	0.734	0.392	6,802
	(0.039)	(0.060)	
59	0.412	0.65	2,547
	(0.307)	(0.190)	
60	0.865	0.427	693
	(0.329)	(0.117)	
61	-	0.091	1,062
	-	(1.710)	
62	-	0.588	9,061
	-	(0.094)	
63	0.601	0.465	1,224
	(0.114)	(0.074)	

	ACF (VA) Coefficients		
SIC	β_l	β_k	N
69	0.501	0.026	10,295
	(0.056)	(0.113)	
70	0.49	0.526	10,274
	(0.099)	(0.060)	
71	0.795	0.524	11,953
	(0.069)	(0.177)	
72	0.707	0.185	2,323
	(0.167)	(0.047)	
73	0.255	0.691	5,168
	(0.414)	(0.173)	
74	0.515	0.502	4,769
	(1.258)	(0.179)	
75	-	0.373	1,482
	-	(0.056)	
77	0.513	0.504	6,195
	(0.019)	(0.028)	
78	0.807	0.283	9,842
	(0.209)	(0.108)	
79	0.907	0.244	4,136
	(0.244)	(0.450)	
80	0.729	0.365	1,926
	(0.025)	(0.013)	
81	0.792	0.35	6,472
	(0.149)	(0.085)	
82	0.578	0.383	9,624
	(0.109)	(0.038)	
90	0.214	0.632	3,111
	(0.355)	(0.052)	
91	0.501	-	1,722
	(0.572)	-	
92	0.596	0.568	1,248
	(0.124)	(0.178)	
93	0.649	0.417	7,853
	(0.057)	(0.031)	
94	0.793	0.471	6,086
	(0.145)	(0.089)	
95	0.861	0.216	1,889
	(0.117)	(0.245)	
96	0.593	0.460	11,807
	(0.161)	(0.107)	

Note: Estimated coefficients from production function across 2-digit SICs, following Ackerberg, Caves, and Frazer (2015) with a Cobb-Douglas value-added production function. Standard errors in brackets. Missing values are point estimates that are less than zero or greater than one.

Table C.5: Gandhi, Navarro, and Rivers (2020) Production Function Coefficients by 2-digit SIC

GNR (GO) Coefficients				
SIC	β_l	β_k	β_m	N
10	0.238	0.081	0.719	12,495
	(0.006)	(0.008)	(0.003)	
11	0.348	0.160	0.585	1,724
	(0.031)	(0.062)	(0.007)	
13	0.315	0.032	0.672	4,981
	(0.007)	(0.009)	(0.005)	
14	0.279	0.050	0.667	3,355
	(0.024)	(0.022)	(0.012)	
15	0.334	0.391	0.390	841
	(0.068)	(0.041)	(0.005)	
16	0.285	0.180	0.592	3,478
	(0.043)	(0.040)	(0.009)	
17	0.243	0.041	0.728	4,184
	(0.003)	(0.011)	(0.006)	
18	0.298	0.221	0.534	7,521
	(0.017)	(0.006)	(0.005)	
19	0.188	-	0.705	506
	(0.204)	-	(0.032)	
20	0.231	0.177	0.608	5,733
	(0.012)	(0.005)	(0.003)	
21	0.281	0.217	0.563	986
	(0.263)	(0.092)	(0.035)	
22	0.270	0.115	0.641	7,776
	(0.018)	(0.006)	(0.010)	
23	0.280	0.043	0.693	5,616
	(0.009)	(0.012)	(0.004)	
24	0.291	0.013	0.697	4,776
	(0.05)	(0.054)	(0.01)	
25	0.346	0.133	0.541	15,597
	(0.027)	(0.024)	(0.004)	
26	0.324	0.066	0.646	7,648
	(0.027)	(0.035)	(0.009)	
27	0.315	0.062	0.634	4,913
	(0.009)	(0.002)	(0.007)	
28	0.362	0.050	0.612	10,899
	(0.015)	(0.004)	(0.007)	
29	0.501	0.096	0.584	1,633
	(0.128)	(0.093)	(0.006)	
30	0.434	0.219	0.438	1,973
	(0.064)	(0.060)	(0.012)	
31	0.315	0.222	0.514	4,060
	(0.026)	(0.017)	(0.002)	

GNR (GO) Coefficients				
SIC	β_l	β_k	β_m	N
32	0.339	0.213	0.502	5,020
	(0.056)	(0.02)	(0.006)	
33	0.442	0.082	0.542	4,997
	(0.027)	(0.022)	(0.011)	
41	0.433	0.220	0.414	12,216
	(0.040)	(0.008)	(0.010)	
42	0.236	0.172	0.630	12,554
	(0.044)	(0.025)	(0.008)	
43	0.336	0.221	0.502	27,014
	(0.012)	(0.007)	(0.005)	
45	0.228	0.044	0.764	24,639
	(0.026)	(0.009)	(0.006)	
46	0.193	0.145	0.704	68,969
	(0.002)	(0.024)	(0.02)	
47	0.219	0.101	0.741	66,171
	(0.026)	(0.015)	(0.000)	
49	0.277	0.238	0.531	11,501
	(0.019)	(0.028)	(0.005)	
50	0.238	0.306	0.55	1,306
	(0.055)	(0.077)	(0.018)	
51	0.301	0.224	0.582	807
	(0.043)	(0.054)	(0.018)	
52	0.351	0.226	0.489	8,103
	(0.008)	(0.011)	(0.007)	
53	0.954	0.043	0.432	489
	(0.101)	(0.024)	(0.008)	
55	0.279	0.229	0.503	8,549
	(0.042)	(0.020)	(0.002)	
56	0.213	0.136	0.622	25,219
	(0.006)	(0.009)	(0.004)	
58	0.350	0.229	0.453	6,802
	(0.031)	(0.024)	(0.007)	
59	0.207	0.387	0.415	2,547
	(0.069)	(0.007)	(0.005)	
60	0.290	0.335	0.460	693
	(0.071)	(0.091)	(0.007)	
61	0.462	0.278	0.391	1,062
	(0.042)	(0.011)	(0.015)	
62	0.450	0.334	0.318	9,061
	(0.052)	(0.004)	(0.007)	
63	0.450	0.339	0.330	1,224
	(0.031)	(0.015)	(0.010)	

GNR (GO) Coefficients				
SIC	β_l	β_k	β_m	N
69	0.392	0.375	0.278	10,295
	(0.019)	(0.01)	(0.005)	
70	0.376	0.396	0.280	10,274
	(0.050)	(0.015)	(0.010)	
71	0.349	0.341	0.342	11,953
	(0.017)	(0.007)	(0.001)	
72	0.400	0.129	0.492	2,323
	(0.105)	(0.033)	(0.023)	
73	0.259	0.376	0.419	5,168
	(0.011)	(0.005)	(0.004)	
74	0.380	0.348	0.352	4,769
	(0.036)	(0.014)	(0.006)	
75	0.411	0.218	0.358	1,482
	(0.016)	(0.016)	(0.014)	
77	0.338	0.350	0.353	6,195
	(0.052)	(0.030)	(0.009)	
78	0.565	0.296	0.148	9,842
	(0.017)	(0.012)	(0.003)	
79	0.317	0.139	0.638	4,136
	(0.031)	(0.054)	(0.016)	
80	0.532	0.336	0.203	1,926
	(0.015)	(0.026)	(0.005)	
81	0.534	0.272	0.236	6,472
	(0.007)	(0.009)	(0.003)	
82	0.498	0.267	0.344	9,624
	(0.022)	(0.009)	(0.007)	
90	0.080	0.413	0.444	3,111
	(0.048)	(0.019)	(0.019)	
91	-	0.194	0.780	1,722
	-	(0.061)	(0.028)	
92	0.337	0.234	0.458	1,248
	(0.048)	(0.046)	(0.020)	
93	0.242	0.322	0.460	7,853
	(0.002)	(0.013)	(0.005)	
94	0.052	0.394	0.598	6,086
	(0.043)	(0.042)	(0.006)	
95	0.463	0.192	0.461	1,889
	(0.040)	(0.007)	(0.011)	
96	0.239	0.392	0.348	11,807
	(0.015)	(0.004)	(0.003)	

Note: Estimated coefficients from production function across 2-digit SICs, following Gandhi, Navarro, and Rivers (2020) with a Cobb-Douglas gross-output production function. Standard errors in brackets. Missing values are point estimates that are less than zero or greater than one.

Table C.6: Overidentified Aggregate Production Functions

	Wooldridge (GO)	Wooldridge (VA)
β_l	0.282	0.651
β_k	0.110	0.255
β_m	0.370	-
Revenue Elasticity	0.762	0.907
Hansen test statistic	1155.34	2113.38
N	170,495	153,467

Note: Estimated coefficients from gross-output (GO) or value-added (VA) Cobb-Douglas production functions. Estimation approach follows Wooldridge (2009). Moment conditions include one additional lag in instruments.

Table C.7: Overidentified Macro Sector Production Functions

	Wooldridge (GO)	Wooldridge (VA)
	Manufacturing	
β_l	0.234	0.597
β_k	0.137	0.291
β_m	0.377	-
Revenue Elasticity	0.749	0.888
Hansen test statistic	780.82	1247.47
N	56,296	57,948
	Construction	
β_l	0.171	0.446
β_k	0.159	0.333
β_m	0.500	-
Revenue Elasticity	0.830	0.778
Hansen test statistic	114.56	131.78
N	12,059	11,740
	Wholesale/Trade/Transport	
β_l	0.134	0.577
β_k	0.140	0.279
β_m	0.444	-
Revenue Elasticity	0.718	0.856
Hansen test statistic	228.86	385.66
N	50,223	38,538
	Services	
β_l	0.378	0.602
β_k	0.101	0.277
β_m	0.352	-
Revenue Elasticity	0.831	0.879
Hansen test statistic	578.08	1038.45
N	49,317	43,001

Note: Estimated coefficients from gross-output (GO) or value-added (VA) Cobb-Douglas production functions. Estimation approach follows Wooldridge (2009). Moment conditions include one additional lag in instruments.

Table C.8: Revenue Elasticity from Overidentified Production Funcitons by 2-digit SIC

	Overidentified 2-digit SIC Production Functions					
SIC	Rev. Elas. (GO)	Rev. Elas. (VA)	Hansen (GO)	Hansen (VA)	N (GO)	N (VA)
10	0.793	0.877	145.83	195.61	6,835	6,705
11	0.842	1.126	40.95	41.27	1,079	1,051
13	0.903	0.902	73.7	76.86	2,440	2,384
14	1.040	1.217	42.35	73.26	1,247	1,229
15	0.909	1.391	28.83	38.07	391	381
16	0.607	0.917	59.64	51.97	1,196	1,183
17	0.666	0.880	74.01	85.58	2,305	2,269
18	0.780	1.065	61.81	56.16	2,616	2,607
19	0.807	0.617	28.95	21.76	280	273
20	0.865	0.908	54.77	165.66	3,222	3,849
21	0.847	1.078	27.59	20.31	591	554
22	0.791	0.898	108.05	105.7	3,402	3,354
23	0.787	0.860	103.04	92.47	2,876	2,825
24	0.794	0.834	86.77	92.44	2,601	2,514
25	0.816	1.014	190.44	218.31	5,763	5,690
26	0.863	1.018	61.01	127.48	3,543	3,445
27	0.604	0.891	91.41	93.19	2,286	2,231
28	0.909	1.007	141.02	121.77	4,800	4,685
29	0.763	0.995	49.65	101.27	780	2,508
30	0.652	0.876	38.74	46.48	1,041	1,197
31	0.808	0.955	48.79	60.78	1,773	1,745
32	0.869	1.007	40.37	72.09	1,801	1,759
33	-	0.873	58.12	83.52	2,202	2,185
41	0.806	0.644	48.82	66.48	2,780	2,599
42	0.935	0.668	69.03	78.92	3,395	3,331
43	0.853	0.922	71.19	56.42	5,442	5,381
45	0.772	0.884	82.23	159.94	6,740	6,351
46	0.843	0.811	100.07	143.84	21,310	20,126
47	0.878	0.941	92.48	79.28	12,865	3,211
49	0.825	0.940	59.38	97.71	4,087	3,882
50	0.833	0.789	21.26	22.28	550	491
51	-	1.003	25.18	42.15	408	397
52	0.749	0.996	80.5	122.51	3,304	3,187
53	-	1.209	28.98	51.75	164	506
55	1.035	0.922	111.91	48.71	2,938	650
56	0.934	0.915	151.59	97.32	5,188	1,513
58	0.940	1.129	62.09	88.36	2,462	2,415
59	0.694	0.973	28.85	40.71	861	790
60	1.132	1.329	22.1	28.79	255	234
61	0.750	1.021	37.65	64.95	297	763
62	0.948	1.105	26	44.29	2,055	1,962
63	0.738	0.752	24.62	34.68	478	474

Overidentified 2-digit SIC Production Functions						
SIC	Rev. Elas. (GO)	Rev. Elas. (VA)	Hansen (GO)	Hansen (VA)	N (GO)	N (VA)
69	0.826	0.977	88.88	77.07	3,261	4,824
70	0.807	0.851	48.37	54.34	1,905	1,813
71	0.823	0.997	64.09	36.79	3,252	3,190
72	-	0.816	46.05	30.14	1,181	951
73	0.803	0.899	51.8	78.02	1,542	1,512
74	0.782	0.834	46.38	46.22	881	862
75	0.990	1.039	32.23	38.73	475	460
77	0.827	0.869	58.75	55.97	2,017	1,981
78	0.922	0.910	105.23	189.2	3,686	3,674
79	0.684	1.214	30.3	40.44	1,318	1,194
80	0.831	0.798	26.9	34.72	745	743
81	0.908	0.809	57.22	102.3	2,531	2,520
82	0.762	0.801	46.1	46.45	2,116	2,037
90	-	0.305	26.74	30.29	617	464
91	-	0.602	39.48	44.82	865	491
92	0.783	0.831	46.68	56.4	551	536
93	0.851	0.797	68.86	119.65	2,857	2,675
94	0.778	0.716	45.53	67.57	1,822	1,430
95	0.618	1.103	38.58	36.2	530	396
96	0.640	0.832	57.02	50.56	1,537	1,489

Note: Estimated revenue elasticity from overidentified production function estimation across 2-digit SICs, with a Cobb-Douglas gross-output (GO) or value-added (VA) production function. Estimation approach follows Wooldridge (2009). Moment conditions include one additional lag in instruments. Missing values where estimated input elasticities were less than zero or greater than one.

Table C.9: Revenue Elasticity Summary Statistics: translog production function

Mean	Stdev	Min	Max	N
0.940	0.344	0.003	1.993	313,282

Summary statistics of revenue elasticity estimates from a value-added translog production function following Ackerberg, Caves, and Frazer (2015). Firms with labour or capital elasticities below zero or greater than one are dropped.

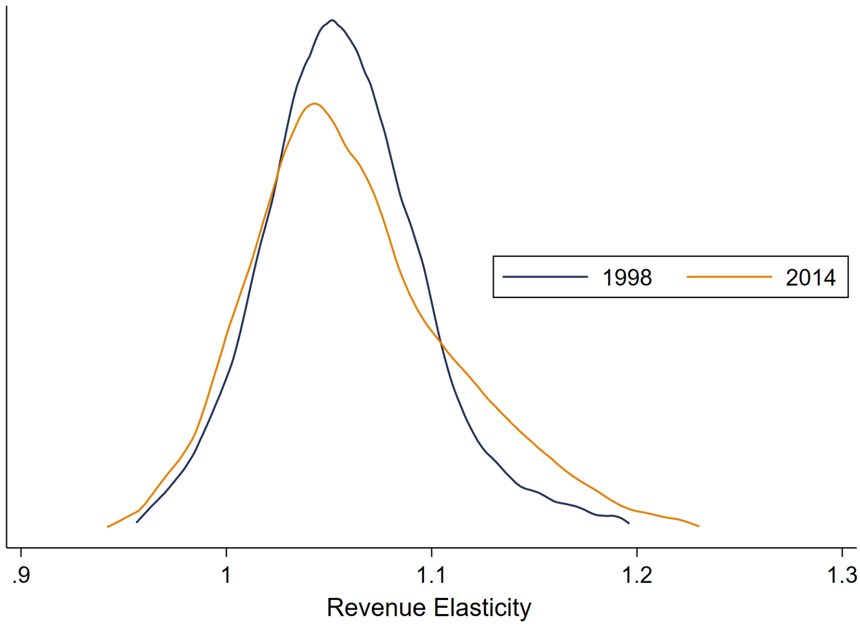

Figure C.1: Distribution of firm-level revenue elasticity estimates, following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function.

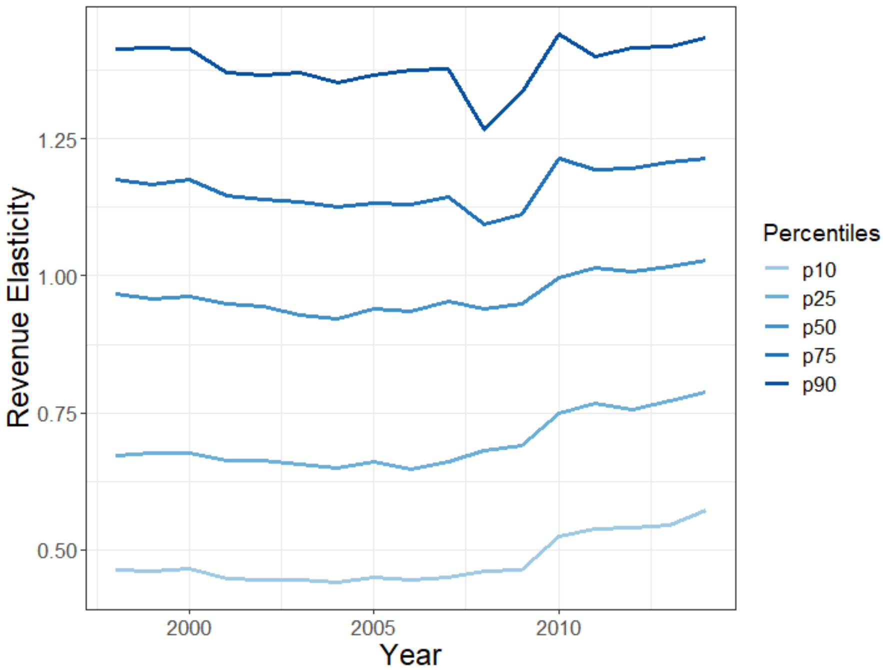

Figure C.2: Percentiles of revenue elasticity over time, following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function.

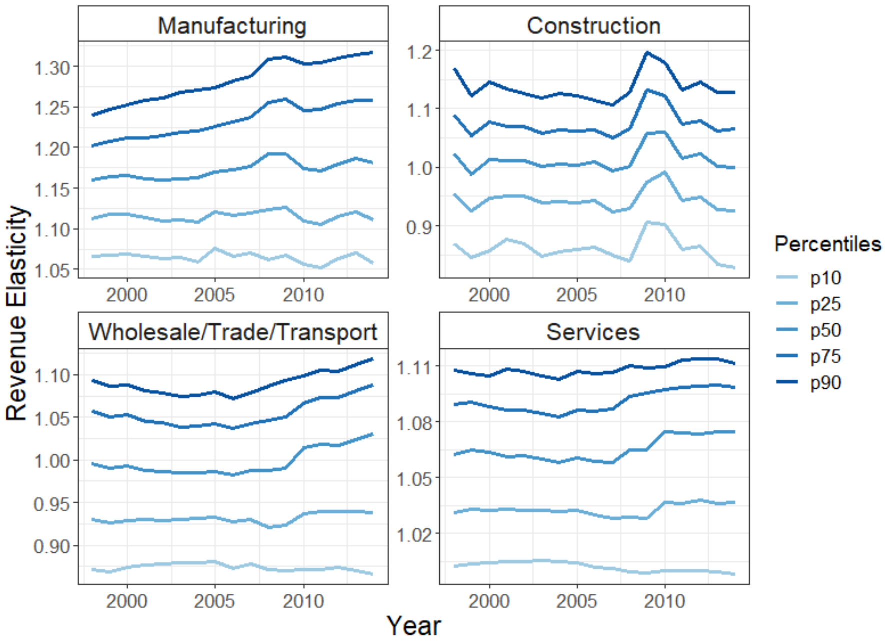

Figure C.3: Percentiles of revenue elasticity by macro sector over time, following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function.

Table C.10: Changing Coefficients: Cobb-Douglas production function (GNR)

	GNR (GO)			
	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
β_l	0.265	0.313	0.367	0.390
	(0.003)	(0.009)	(0.007)	(0.007)
β_k	0.121	0.106	0.194	0.245
	(0.003)	(0.010)	(0.023)	(0.008)
β_m	0.609	0.610	0.471	0.391
	(0.005)	(0.007)	(0.012)	(0.006)
N	153,874	144,465	108,619	120,855

Note: Estimated coefficients from a gross-output Cobb-Douglas production function. Standard errors in brackets. Estimation approach follows Gandhi, Navarro, and Rivers (2020) (GNR).

Table C.11: Changing Coefficients: Cobb-Douglas production function (ACF)

	ACF (GO)			
	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
β_l	0.148	0.372	0.218	0.591
	(0.020)	(0.058)	(0.040)	(0.027)
β_k	0.029	0.022	0.038	0.442
	(0.030)	(0.023)	(0.104)	(0.234)
β_m	0.536	0.677	0.720	0.278
	(0.025)	(0.102)	(0.030)	(0.022)
N	153,874	144,465	108,619	120,855

Note: Estimated coefficients from a gross-output Cobb-Douglas production function. Standard errors in brackets. Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF).

Table C.12: Changing Coefficients by Macro Sector: Cobb-Douglas production function (GNR)

	GNR (GO)			
	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
		Manufacturing		
β_l	0.164	0.229	0.322	0.434
	(0.024)	(0.014)	(0.011)	(0.008)
β_k	0.163	0.125	0.173	0.231
	(0.014)	(0.026)	(0.004)	(0.008)
β_m	0.669	0.662	0.558	0.470
	(0.002)	(0.003)	(0.004)	(0.003)
N	39,876	34,678	24,011	22,147
		Construction		
β_l	0.249	0.245	0.444	0.391
	(0.009)	(0.027)	(0.003)	(0.012)
β_k	0.149	0.123	0.205	0.306
	(0.005)	(0.007)	(0.004)	(0.024)
β_m	0.619	0.644	0.443	0.365
	(0.006)	(0.002)	(0.006)	(0.002)
N	13,484	13,416	10,210	14,674
		Wholesale/Trade/Transport		
β_l	0.148	0.173	0.222	0.294
	(0.003)	(0.011)	(0.012)	(0.010)
β_k	0.073	0.088	0.194	0.217
	(0.010)	(0.007)	(0.012)	(0.001)
β_m	0.788	0.813	0.619	0.542
	(0.003)	(0.005)	(0.018)	(0.005)
N	53,814	50,631	37,906	39,634
		Services		
β_l	0.389	0.457	0.482	0.445
	(0.004)	(0.006)	(0.007)	(0.010)
β_k	0.182	0.132	0.234	0.290
	(0.005)	(0.003)	(0.006)	(0.007)
β_m	0.412	0.421	0.326	0.271
	(0.000)	(0.006)	(0.002)	(0.004)
N	46,700	45,740	36,492	44,400

Note: Estimated coefficients from a Cobb-Douglas gross-output production function. Standard errors in brackets. Estimation approach follows Gandhi, Navarro, and Rivers (2020) (GNR).

Table C.13: Changing Revenue Elasticities by Macro Sector: Cobb-Douglas production function (ACF)

	ACF (VA)			
	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
	Manufacturing			
Revenue Elasticity	1.061	1.135	1.158	1.200
N	39,876	34,678	24,011	22,147
	Construction			
Revenue Elasticity	-	0.909	1.173	-
N	13,484	13,416	10,210	14,674
	Wholesale/Trade/Transport			
Revenue Elasticity	0.810	1.005	1.047	0.999
N	53,814	50,631	37,906	39,634
	Services			
Revenue Elasticity	1.039	1.060	1.036	1.110
N	46,700	45,740	36,492	44,400

Note: Revenue elasticity from a value-added Cobb-Douglas production functions. Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF).

Table C.14: Changing Coefficients by Macro Sector: Cobb-Douglas production function (ACF)

	ACF (VA)			
	1998 - 2001	2002 - 2005	2006 - 2009	2010 - 2014
		Manufacturing		
β_l	0.364	0.573	0.732	0.701
	(0.004)	(0.006)	(0.036)	(0.041)
β_k	0.697	0.562	0.426	0.498
	(0.002)	(0.008)	(0.033)	(0.014)
N	39,876	34,678	24,011	22,147
		Construction		
β_l	0.464	0.699	0.949	-
	(0.221)	(0.040)	(0.066)	-
β_k	-	0.211	0.224	-
	-	(0.108)	(0.011)	-
N	13,484	13,416	10,210	14,674
		Wholesale/Trade/Transport		
β_l	0.242	0.308	0.395	0.419
	(0.071)	(0.018)	(0.024)	(0.447)
β_k	0.568	0.698	0.652	0.580
	(0.086)	(0.038)	(0.029)	(0.235)
N	53,814	50,631	37,906	39,634
		Services		
β_l	0.501	0.664	0.650	0.725
	(0.011)	(0.004)	(0.010)	(0.047)
β_k	0.537	0.397	0.386	0.384
	(0.003)	(0.010)	(0.004)	(0.034)
N	46,700	45,740	36,492	44,400

Note: Estimated coefficients from a value-added Cobb-Douglas production functions. Standard errors in brackets. Estimation approach follows Ackerberg, Caves, and Frazer (2015) (ACF). Missing values are estimates below zero or greater than one.

C.3 Appendix - Productivity Estimates

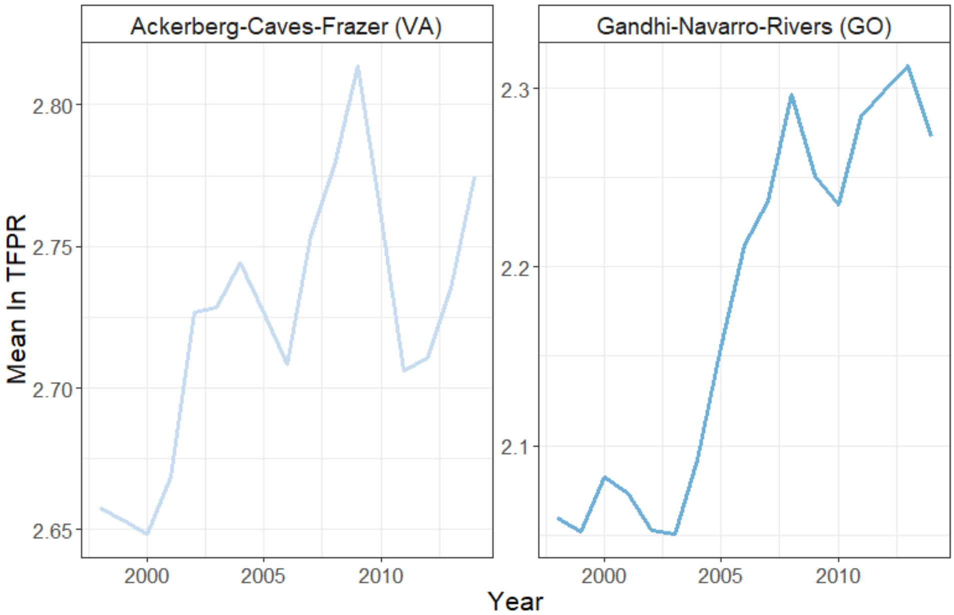

Figure C.4: Mean log TFPR from Cobb-Douglas production function estimation, estimated at the aggregate level by year.

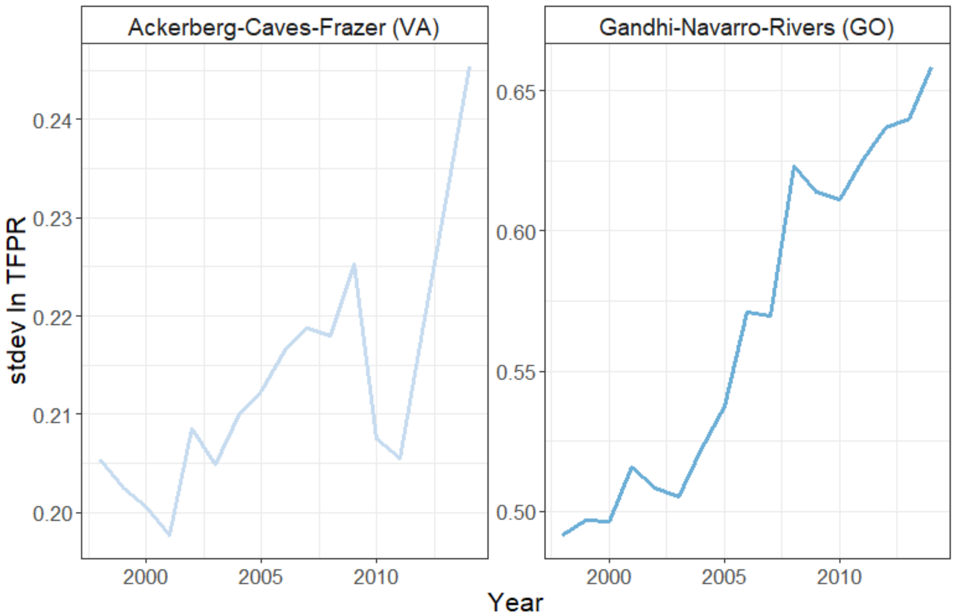

Figure C.5: Standard deviation of log TFPR from Cobb-Douglas production function estimation, estimated at the aggregate level by year.

Figure C.6: Mean log TFPR by macro sector, following Ackerberg, Caves, and Frazer (2015) with value-added Cobb-Douglas production function.

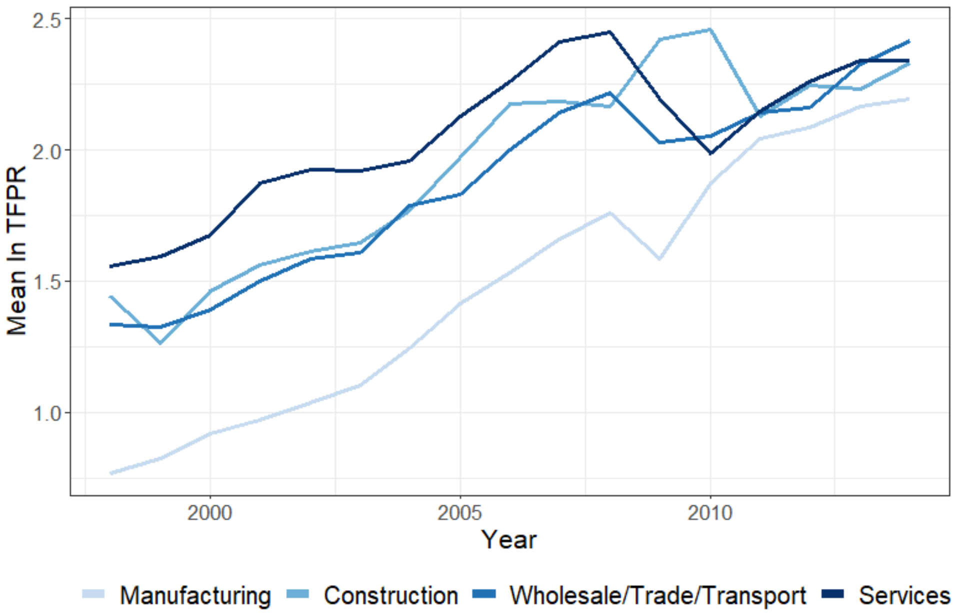

Figure C.7: Mean log TFPR by macro sector, following Ackerberg, Caves, and Frazer (2015) with value-added translog production function.

C.4 Appendix - Further Figures & Tables

Table C.15: Summary Statistics from the ARDx, 1998 - 2014

		Economy-Wide		
	y	k	l	m
Mean	7.6	6.0	3.2	7.0
Stdev	2.4	2.2	1.9	2.6
N		527,800		
		Manufacturing		
Mean	8.5	7.4	4.0	8.0
Stdev	2.0	1.9	1.6	2.3
N		125,700		
		Construction		
Mean	7.0	4.9	2.6	6.3
Stdev	2.3	2.1	1.8	2.6
N		51,800		
		Trade, Wholesale & Transport		
Mean	7.7	5.8	2.9	7.3
Stdev	2.4	2.2	1.8	2.5
N		182,800		
		Services		
Mean	7.0	5.7	3.1	6.0
Stdev	2.3	2.2	2	2.5
N		179,000		

Summary statistics on panel of firms used for production function estimation. Output, labour, capital, and materials are measured in logs.

Table C.16: Regression: Revenue Elasticity on Productivity at Industry Level

	Dependent variable: Revenue Elasticity		
	ACF (VA)	ACF (GO)	GNR (GO)
Mean log TFPR	0.335***	0.347***	0.102
	(0.031)	(0.025)	(0.068)
N	59	54	60

*Note: Revenue elasticity estimated with gross-output (GO) or value-added Cobb-Douglas production functions following Ackerberg, Caves, and Frazer (2015) (ACF) or Gandhi, Navarro, and Rivers (2020) (GNR). Estimates statistically significant at levels of 1%: ***, 5%: **, 10%: *. Robust standard errors in brackets.*

Figure C.8: Missing UK Firm Data

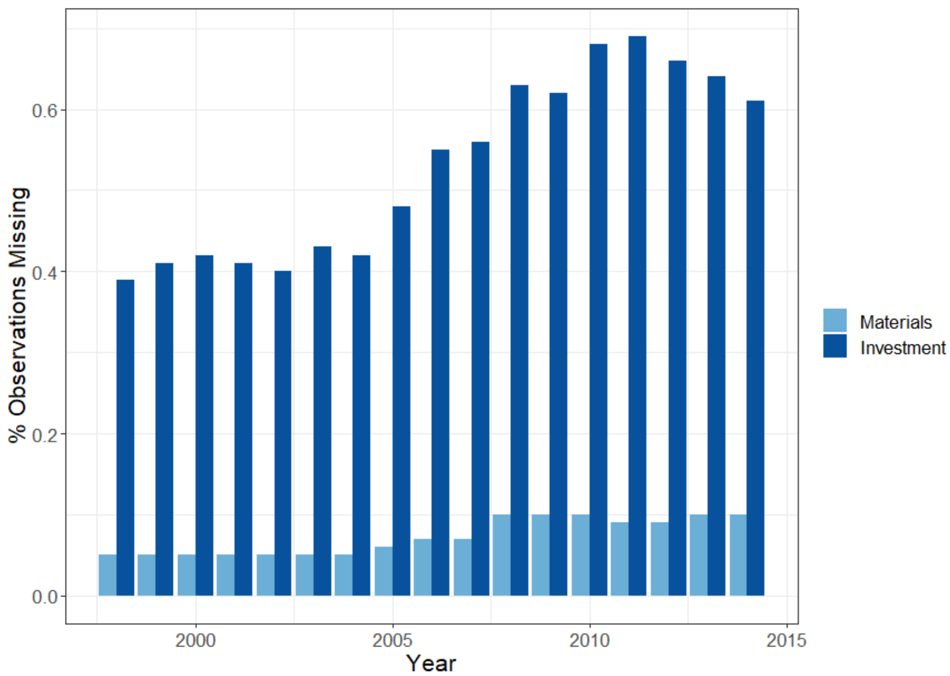

Note: for each year, the share of firm observations with missing data for materials inputs or investment.

Table C.17: Computer software share in total investment across 2-digit SICs, averaged over 1998 to 2014

SIC	Share
10	0.039
11	0.044
13	0.056
14	0.120
15	0.056
16	0.055
17	0.047
18	0.067
19	0.033
20	0.051
21	0.059
22	0.040
23	0.042
24	0.041
25	0.056
26	0.093
27	0.066
28	0.084
29	0.047
30	0.071
31	0.069
32	0.088
33	0.084
41	0.052
42	0.054
43	0.072
45	0.069
46	0.101
47	0.080
49	0.036
50	0.073
51	0.078
52	0.109
53	0.088
55	0.048
56	0.043
58	0.189
59	0.152
60	0.173
61	0.146
62	0.211
63	0.206

SIC	Share
69	0.260
70	0.197
71	0.206
72	0.107
73	0.252
74	0.181
75	0.139
77	0.029
78	0.242
79	0.266
80	0.152
81	0.087
82	0.198
90	0.162

Table C.18: Regression: Revenue Elasticity on Software Investment Share, Split by Sub-Period

	Dependent variable: Revenue Elasticity				
	1998 - 2006				
Software Investment Share	0.023***	0.024***	0.005**	0.003	0.002
	(0.008)	(0.008)	(0.002)	(0.002)	(0.002)
N	132,469	132,469	92,278	92,278	90,706
	2007 - 2014				
Software Investment Share	0.035***	0.036***	0.003**	0.004***	0.004***
	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
N	83,393	83,393	62,961	62,961	61,440
Year FE:		✓		✓	✓
Firm FE:			✓	✓	✓
Remove outliers:					✓

*Note: Estimates statistically significant at levels of 1%: ***, 5%: **, 10%: *. Robust standard errors in brackets, clustered at the level of the fixed effects included. Revenue elasticity and log TFPR estimated following Ackerberg, Caves, and Frazer (2015) with a translog value-added production function. Regressions weighted by value-added at the firm level. Outliers are the top and bottom 1% of firms by revenue elasticity.*

C.5 Appendix - Standard Model

Cost Relationships

From cost minimisation, variable costs are $\mathcal{C} = wL + rK = z\lambda y\left(\frac{\partial y}{\partial L}\frac{L}{y} + \frac{\partial y}{\partial K}\frac{K}{y}\right)$. Euler's homogeneous function theorem states that a function multiplied by its degree of homogeneity will be equal to the sum of partial derivatives multiplied by the arguments:

$$\nu F(X_1, ..., X_N) = \sum_{i=1}^{N} \frac{\partial F(X_1, ..., X_N)}{\partial X_i} X_i$$

Applying this theorem to our minimised variable costs yields:

$$\begin{aligned}
\mathcal{C} &= z\lambda y\left(\frac{\partial y}{\partial L}\frac{L}{y} + \frac{\partial y}{\partial K}\frac{K}{y}\right) \\
&= z\lambda\left(\frac{\partial y}{\partial L}L + \frac{\partial y}{\partial K}K\right) \\
&= \lambda\nu z F(K, L) \\
&= \lambda\nu(y + \phi)
\end{aligned}$$

Notice that λ is marginal cost $\frac{\partial \mathcal{C}}{\partial y}$, so we have a partial differential equation:

$$\mathcal{C}(y; w, r) = \frac{\partial \mathcal{C}(y; w, r)}{\partial y}\nu(y + \phi),$$

which can be integrated in multiplicatively-separable form where $G(w, r, 1)$ is an arbitrary function independent of output y, as in Savagar (2021):

$$\mathcal{C}(y; w, r) = \left(\frac{y + \phi}{z}\right)^{\frac{1}{\nu}} G(w, r, 1).$$

Therefore, the marginal cost is:

$$\lambda = \frac{\partial \mathcal{C}(y; w, r)}{\partial y} = \frac{1}{\nu z^{\frac{1}{\nu}}}\left(\frac{y + \phi}{z}\right)^{\frac{1}{\nu} - 1} G(w, r, 1).$$

Returns to Scale: Markups and Profit Share

Consider the profit function $\pi = \mathcal{P}(y)y - \mathcal{C}(y;w,r)$ and note that the markup $\mu = \frac{p}{MC}$. Some simple algebra yields another equation for returns to scale:

$$
\begin{aligned}
\pi &= \mathcal{P}(y)y - \mathcal{C}(y;w,r) \\
\frac{\pi}{\mathcal{P}(y)y} &= 1 - \frac{\mathcal{C}(y;w,r)}{\mathcal{P}(y)y} \\
s_\pi &= 1 - \frac{\lambda}{\lambda}\frac{\text{AC}}{\mathcal{P}(y)} \\
s_\pi &= 1 - \mu^{-1}\frac{\text{AC}}{\text{MC}} \\
\frac{\mathcal{AC}}{\mathcal{MC}} &= \mu(1 - s_\pi)
\end{aligned}
$$

C.6 Appendix - Fixed Cost in the Cost Function

Consider an alternative set-up where the fixed cost ϕ is still output-denominated, but is included in the cost function rather than taken from gross output.

Output is given:

$$Y = zF(k, \ell)$$

Profit maximisation with monopolistic competition and CES demand yields price over marginal cost equal to a constant markup:

$$\frac{p}{\lambda} = \mu$$

The cost function is:

$$\mathcal{C} = w\ell + rk + p\phi$$

By cost minimisation and applying Euler's homogeneous function theorem, the cost function can be written:

$$\mathcal{C} = \nu \lambda y + p\phi$$

To obtain returns to scale, we divide average costs by marginal costs λ:

$$\text{RTS} = \frac{\mathcal{AC}}{\mathcal{MC}} = \nu + \frac{p}{\lambda}\frac{\phi}{y} = \nu + \mu s_{\phi}.$$

C.7 Appendix - Model with Adjustment Costs

Following Appendix A.2 of Bond, Hashemi, Kaplan, and Zoch (2021), I introduce an adjustment cost on capital k_{it}. A firm faces such a cost when the capital input choice deviates from $\overline{k_{it}}$ which can be interpreted as the previous period decision. Consider that the cost is given by a smooth convex function $\eta(k_{it})$, which satisfies the condition $\eta\left(\overline{k_{it}}\right) = \frac{d\eta\left(\overline{k_{it}}\right)}{dk_{it}} = 0$.

The firm's cost minimisation problem is:

$$\mathcal{C}_{it} := \min_{k_{it},\ell_{it}} w_{it}\ell_{it} + r_{it}k_{it} + r_{it}\eta(k_{it}) \quad \text{s.t.} \quad y_{it} \geq z_{it}F(k_{it},\ell_{it}) - \phi.$$

where the adjustment cost on capital is normalised by the rental rate. The first-order conditions are:

$$w_{it} = \lambda_{it}z_{it}\frac{\partial F}{\partial \ell_{it}}$$

$$r_{it}\left(1 + \frac{d\eta(k_{it})}{dk_{it}}\right) = \lambda_{it}z_{it}\frac{\partial F}{\partial k_{it}}$$

From here, I obtain expressions for returns to scale and revenue elasticity. I compute the optimised cost function and use Euler's Homogeneous Function Theorem to substitute the net output function.

$$\begin{aligned}
\mathcal{C}_{it} &= w_{it}\ell_{it} + r_{it}k_{it} + r_{it}\eta(k_{it}) \\
&= \lambda_{it}z_{it}\frac{\partial F}{\partial \ell_{it}}\ell_{it} + \lambda_{it}z_{it}\frac{\partial F}{\partial k_{it}}k_{it} - r_{it}\frac{d\eta(k_{it})}{dk_{it}}k_{it} + r_{it}\eta(k_{it}) \\
&= \nu\lambda_{it}(y_{it} + \phi) + r_{it}\left(\eta(k_{it}) - \frac{d\eta(k_{it})}{dk_{it}}k_{it}\right) \\
&= \nu\lambda_{it}(y_{it} + \phi) + r_{it}\eta(k_{it})\left(1 - \frac{d\eta(k_{it})}{dk_{it}}\frac{k_{it}}{\eta(k_{it})}\right)
\end{aligned}$$

Returns to scale is:

$$\text{RTS} = \nu(1 + s_{\phi}) + \frac{r_{it}\eta(k_{it})}{\lambda_{it}y_{it}}\left(1 - \varepsilon_{\eta(k),k}\right)$$

where $\varepsilon_{\eta(k),k}$ is the elasticity of the adjustment cost function to capital. Revenue elasticity takes the form:

$$\text{Revenue Elasticity} = \frac{\nu}{\mu}(1 + s_{\phi}) + \frac{r_{it}\eta(k_{it})}{p_{it}y_{it}}\left(1 - \varepsilon_{\eta(k),k}\right)$$

which differs from the no-adjustment-cost baseline $\frac{\nu}{\mu}(1 + s_{\phi})$ by the adjustment cost share in revenue, multiplied by one minus the adjustment cost elasticity.

This elasticity can be greater or less than one, and depends on whether the firm's capital stock is growing or shrinking. Consider the convex adjustment costs $\frac{\gamma}{2}\left(\frac{I}{K}\right)^2 K$ of the form in Cooper and Haltiwanger (2006), where $I = K - \overline{K}$ is investment, K is capital, and $\overline{K}$ is previous period capital. The elasticity $\varepsilon_{\eta(k),k}$ takes the form $\frac{K^2-\overline{K}^2}{(K-\overline{K})^2}$, which exceeds one if this period's capital stock exceeds the previous period's ($K > \overline{K}$). Thus if firm's capital stock is growing, returns to scale and revenue elasticity is lower than in the absence of adjustment costs on capital.

An alternative way to see this phenomenon is to sum the output and revenue elasticities straight from rearranged first-order conditions. Consider just capital:

$$\begin{aligned}
r_{it}\left(1 + \frac{d\eta\left(k_{it}\right)}{dk_{it}}\right) &= \lambda_{it}z_{it}\frac{\partial F}{\partial k_{it}} \\
\frac{r_{it}k_{it}}{p_{it}y_{it}}\left(1 + \frac{d\eta\left(k_{it}\right)}{dk_{it}}\right) &= \frac{1}{\mu}z_{it}\frac{\partial F}{\partial k_{it}}\frac{k_{it}}{y_{it}} \\
\alpha_{it}^{k}\mu\left(1 + \frac{d\eta\left(k_{it}\right)}{dk_{it}}\right) &= \varepsilon_{y,k}
\end{aligned}$$

where α_{it}^{x} is the cost share in revenue for input x.

Returns to scale is equal to the sum of output elasticities:

$$\text{RTS} = \mu\left[\alpha_{it}^{k}\left(1 + \frac{d\eta\left(k_{it}\right)}{dk_{it}}\right) + \alpha_{it}^{\ell}\right]$$

And revenue elasticity is the above equation divided by the markup:

$$\text{Revenue Elasticity} = \alpha_{it}^{k}\left(1 + \frac{d\eta\left(k_{it}\right)}{dk_{it}}\right) + \alpha_{it}^{\ell}$$

In a world without adjustment costs, revenue elasticity is simply the sum of cost shares in revenue. When adjustment costs are present, there is an extra wedge due to $\left(1 + \frac{d\eta(k_{it})}{dk_{it}}\right)$. Whether this is greater or less than the case without adjustment costs on capital depends on the slope of the adjustment cost function, which is determined by whether the capital stock is growing or shrinking.

C.8 Appendix - Model with Software Extension

Demand Side

Following Melitz (2003), Cobb-Douglas preferences are defined over J sectors:

$$U = \sum_j \beta_j \ln Y_j.$$

with $\sum_j \beta_j = 1$.

In each sector j, there's a continuum of horizontally-differentiated varieties, and preferences for these are CES:

$$Y_j = \left(\int_{\omega \in \Omega_j} y_j(\omega)^{(\sigma_j - 1)/\sigma_j} d\omega \right)^{\sigma_j/(\sigma_j - 1)}.$$

where $\sigma_j > 1$.

If M is aggregate income, then Cobb-Douglas preferences over sectors yields expenditure $E_j = \beta_j M$. The demand for each variety in each sector is thus:

$$y_j(\omega) = A_j p_j(\omega)^{-\sigma_j}.$$

where $A_j = E_j P_j^{\sigma_j - 1}$, and the price index:

$$P_j = \left(\int_{\omega \in \Omega_j} p_j(\omega)^{(1-\sigma_j)} d\omega \right)^{1/(1-\sigma_j)}.$$

A_j is an index of market demand that scales each firm's residual demand, and is determined by sectoral expenditure and the CES price index. Given a continuum of firms, each firm takes this as given. Dropping the sector and variety notation for simplicity, we can analyse firm behaviour within each sector in this environment.

The demand function for each firm $y(z)$ can be rewritten by combining the equation A_j above, and noting that expenditure equals revenue (so $E = R = py$) in each sector: $y(z) = \left(\frac{p(z)}{P}\right)^{-\sigma} \frac{E}{P} = \left(\frac{p(z)}{P}\right)^{-\sigma} Y$.

Software Derivative

$$
\begin{aligned}
\frac{\partial s}{\partial y} &= -\frac{1}{2-\psi}\frac{1}{\psi}\frac{1}{\nu z}\left(\frac{y+\phi}{z}\right)^{\frac{1}{\nu}-1} G(w,r,1)\left[\frac{1}{\psi}\left(\frac{y+\phi}{z}\right)^{\frac{1}{\nu}} G(w,r,1)\right]^{\frac{1}{2-\psi}-1} \\
&= -\frac{1}{2-\psi}\frac{1}{\psi}\lambda(1-s)^{\psi-1} \\
&= \frac{\lambda}{\psi(\psi-2)}(1-s)^{\psi-1}
\end{aligned}
$$

Marginal Cost

$$
\begin{aligned}
\frac{\partial \mathcal{C}}{\partial y} &= \lambda + \frac{\partial f(s)}{\partial s}\frac{\partial s}{\partial y} \\
&= \lambda + \psi\left(\frac{1}{1-s}\right)^{\psi-1} \frac{\lambda}{\psi(\psi-2)}(1-s)^{\psi-1} \\
&= \lambda\left(1+\frac{1}{\psi-2}\right) \\
&= \lambda\left(\frac{\psi-1}{\psi-2}\right)
\end{aligned}
$$

Returns to Scale

$$
\begin{aligned}
\frac{\mathcal{AC}}{\mathcal{MC}} &= \frac{\nu\lambda(1+s_{\phi}) + \frac{f(s)}{y}}{\lambda\left(\frac{\psi-1}{\psi-2}\right)} \\
&= \frac{\psi-2}{\psi-1}\left(\nu(1+s_{\phi}) + \frac{\left(\frac{1}{1-s}\right)^{\psi}-1}{\lambda y}\right) \\
&= \nu(1+s_{\phi})\frac{\psi-2}{\psi-1}\left(1+\frac{\left(\frac{1}{1-s}\right)^{\psi}-1}{\nu\lambda(y+\phi)}\right) \\
&= \nu(1+s_{\phi})\frac{\psi-2}{\psi-1}\left(1+\frac{\left(\frac{1}{1-s}\right)^{\psi}-1}{\psi\left(\frac{1}{1-s}\right)^{\psi}}\right) \\
&= \nu(1+s_{\phi})\frac{\psi-2}{\psi-1}\left(\frac{1}{\psi}\left(\psi+1-(1-s)^{\psi}\right)\right)
\end{aligned}
$$

C.9 Appendix - Alternative Model with Software

Following De Ridder (2019), consider software s as part of the production function. I allow this to be an input factor that firms choose in the cost minimisation procedure. Let the production function take the form:

$$y = \frac{1}{1-s} z k^{\alpha} \ell^{\beta}.$$

where $\alpha + \beta = \nu$. I can now compute cost minimisation over three input choices: labour, capital, and software.

Cost Minimisation:

$$\begin{aligned}
\mathcal{C} &:= \min_{k,\ell,s} w\ell + rk + f(s) \text{ s.t. } y = \frac{1}{1-s} z k^{\alpha} \ell^{\beta} \\
[k] &: r = \frac{1}{1-s} \lambda \alpha z k^{\alpha-1} \ell^{\beta} \implies k = \left(\frac{\lambda z \alpha}{r(1-s)} \ell^{\beta} \right)^{\frac{1}{1-\alpha}} \\
[\ell] &: w = \frac{1}{1-s} \lambda \beta z k^{\alpha} \ell^{\beta-1} \implies \ell = \left(\frac{\lambda z \beta}{w(1-s)} k^{\alpha} \right)^{\frac{1}{1-\beta}} \\
[s] &: f'(s) = \lambda \frac{y}{1-s}
\end{aligned}$$

From here, I can combine the FOCs for k, ℓ to get conditional factor demands:

$$\begin{aligned}
\ell &= \left(\frac{z\lambda}{1-s} \left(\frac{\alpha}{r} \right)^{\alpha} \left(\frac{\beta}{w} \right)^{1-\alpha} \right)^{\frac{1}{1-\nu}} \\
k &= \left(\frac{z\lambda}{1-s} \left(\frac{\alpha}{r} \right)^{1-\beta} \left(\frac{\beta}{w} \right)^{\beta} \right)^{\frac{1}{1-\nu}}
\end{aligned}$$

Notice also that the variable costs can be written compactly as:

$$w\ell = \beta \lambda \frac{z}{1-s} k^{\alpha} \ell^{\beta-1} \ell = \beta \lambda y.$$

$$rk = \alpha\lambda\frac{z}{1-s}k^{\alpha-1}\ell^{\beta}k = \alpha\lambda y.$$

$$\mathcal{C} = w\ell + rk + f(s) = \lambda\nu y + f(s).$$

To proceed, I require a functional form for $f(s)$, which is the software 'fixed' cost. For simplicity, consider the following function where ψ is the software cost elasticity:[38]

$$f(s) = \left(\frac{1}{1-s}\right)^{\psi} + 1.$$

Then the derivative is $f'(s) = \psi\frac{1}{(1-s)^2}\left(\frac{1}{1-s}\right)^{\psi-1}$. Combine this with the FOC for s, which is $f'(s) = \lambda\frac{y}{1-s}$, and optimal y, which is $y = \left(\frac{z}{1-s}\frac{\alpha\beta}{rw}\lambda^{\nu}\right)^{\frac{1}{1-\nu}}$, and the result is:

$$\frac{\lambda}{(1-s)}\left(\frac{z}{1-s}\frac{\alpha\beta}{rw}\lambda^{\nu}\right)^{\frac{1}{1-\nu}} = \frac{\psi}{(1-s)}\left(\frac{1}{1-s}\right)^{\psi}.$$

This can be easily rearranged to get an expression for optimal software investment:

$$s = 1 - \left(z\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}.$$

[38]This differs slightly from De Ridder (2019), but has the required properties and suffices for this exercise.

Software investment s is increasing in productivity z if $\psi(1-\nu) > 1$.[39] I can substitute optimal s back into conditional factor demands and optimal y:

$$\ell = \left(\frac{z\lambda}{\left(z\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}} \left(\frac{\alpha}{r}\right)^{\alpha} \left(\frac{\beta}{w}\right)^{1-\alpha} \right)^{\frac{1}{1-\nu}}$$

$$= z^{\frac{\psi}{\psi(1-\nu)-1}} \left(\frac{\lambda}{\left(\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}} \left(\frac{\alpha}{r}\right)^{\alpha} \left(\frac{\beta}{w}\right)^{1-\alpha} \right)^{\frac{1}{1-\nu}}$$

$$k = \left(\frac{z\lambda}{\left(z\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}} \left(\frac{\alpha}{r}\right)^{1-\beta} \left(\frac{\beta}{w}\right)^{\beta} \right)^{\frac{1}{1-\nu}}$$

$$= z^{\frac{\psi}{\psi(1-\nu)-1}} \left(\frac{\lambda}{\left(\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}} \left(\frac{\alpha}{r}\right)^{1-\beta} \left(\frac{\beta}{w}\right)^{\beta} \right)^{\frac{1}{1-\nu}}$$

$$y = \left(\frac{z}{\left(z\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}} \frac{\alpha\beta}{rw}\lambda^{\nu} \right)^{\frac{1}{1-\nu}}$$

$$= z^{\frac{\psi}{\psi(1-\nu)-1}} \left(\frac{1}{\left(\lambda\alpha\beta\psi^{\nu-1}\frac{1}{rw}\right)^{\frac{1}{1-\psi(1-\nu)}}} \frac{\alpha\beta}{rw}\lambda^{\nu} \right)^{\frac{1}{1-\nu}}$$

In terms of quantity y, the optimised software 'fixed' cost is $\frac{\lambda y}{\psi}+1$.[40] With this function for the 'fixed' cost, the cost function is:

$$C = \lambda\nu y + \frac{\lambda y}{\psi} + 1$$

[39]

$$\frac{ds}{dz} = \frac{\lambda\alpha\beta\psi^{\nu-1}}{\psi(1-\nu)-1}(1-s)^{\psi(1-\nu)}.$$

[40]Take the FOC for s and equate to the derivative of the software 'fixed' cost function:

$$\frac{\lambda y}{1-s} = \psi\frac{1}{(1-s)^2}\left(\frac{1}{1-s}\right)^{\psi-1}.$$

This can be rearranged to:

$$\frac{\lambda y}{\psi} = \frac{1}{1-s}\left(\frac{1}{1-s}\right)^{\psi-1} = \left(\frac{1}{1-s}\right)^{\psi}.$$

Then returns to scale is:

$$\begin{aligned}\text{RTS} &= \nu + \frac{1}{\psi} + \frac{1}{\lambda y} \\ &= \frac{1}{\psi} + \nu\left(1 + \frac{1}{w\ell + rk}\right)\end{aligned}$$

When software is not available, returns to scale will be equal to ν. Thus, the presence of software raises returns to scale.

Bibliography

Abel, William, Silvana Tenreyro, and Gregory Thwaites (2018). *Monopsony in the UK*. CEPR Discussion Papers 13265. C.E.P.R. Discussion Papers.

Abeliansky, Ana Lucia, Inmaculada Martínez-Zarzoso, and Klaus Prettner (2015). "The impact of 3D printing on trade and FDI". In.

Acemoglu, Daron and David Autor (June 2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *National Bureau of Economic Research, Working Papers*. Working Paper Series 4.16082.

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price (May 2014). "Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing". In: *American Economic Review* 104.5, pp. 394–99.

Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo (2020). "Competing with Robots: Firm-Level Evidence from France". In: *ASSA 2020 Papers and Proceedings*.

Acemoglu, Daron and Pascual Restrepo (2016). "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment". In: *National Bureau of Economic Research, Working Paper* 22252.

— (2018). *Artificial Intelligence, Automation and Work*. Working Paper 24196. National Bureau of Economic Research.

— (2020). "Robots and Jobs: Evidence from U.S. Labor Markets". In: *Journal of Political Economy* 128.6, pp. 2188–2244.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer (2015). "Identification Properties of Recent Production Function Estimators". In: *Econometrica* 83.6, pp. 2411–2451.

Aghion, Philippe, Céline Antonin, and Simon Bunel (2019). "Artificial Intelligence, Growth and Employment: The Role of Policy". In: *Economie et Statistique / Economics and Statistics* 510-511-5, pp. 149–164.

Aghion, Philippe, Benjamin Jones, and Charles Jones (2017). *Artificial Intelligence and Economic Growth*. Working Paper 23928. National Bureau of Economic Research.

Altomonte, Carlo, Domenico Favoino, Monica Morlacco, and Tommaso Sonno (2021). *Markups, intangible capital and heterogeneous financial frictions*. CEP Discussion Papers. Centre for Economic Performance, LSE.

Altug, Sumru G. and Alpay Filiztekin (Apr. 1999). *Estimates of the Returns to Scale for US Manufacturing*. CEPR Discussion Papers 2121. C.E.P.R. Discussion Papers.

Atkeson, Andrew and Patrick J Kehoe (2005). "Modeling and measuring organization capital". In: *Journal of political Economy* 113.5, pp. 1026–1053.

Atkeson, Andrew, Aubhik Khan, and Lee Ohanian (1996). "Are data on industry evolution and gross job turnover relevant for macroeconomics?" In: *Carnegie-Rochester Conference Series on Public Policy*. Vol. 44. Elsevier, pp. 215–250.

Autor, David (2010). "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings". In: *The Center for American Progress and The Hamilton Project*.

— (Sept. 2014). "Polanyi's Paradox and the Shape of Employment Growth". In: 20485.

— (2015). "Why Are There Still So Many Jobs? The History and Future of Workplace Automation". In: *Journal of Economic Perspectives* 29.3, pp. 3–30.

Autor, David and David Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market". In: *American Economic Review* 103.5, pp. 1553–97.

Autor, David, David Dorn, and Gordon H. Hanson (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States". In: *American Economic Review* 103.6, pp. 2121–68.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen (Feb. 2020). “The Fall of the Labor Share and the Rise of Superstar Firms*”. In: *The Quarterly Journal of Economics* 135.2, pp. 645–709.

Autor, David, Lawrence F Katz, and Melissa S Kearney (2005). “Rising Wage Inequality: The Role of Composition and Prices”. In: NBER Working Paper Series 11628.

Autor, David, Frank Levy, and Richard J. Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration”. In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333.

Autor, David and Anna Salomons (2017). “Robocalypse Now: Does Productivity Growth Threaten Employment?” In: *NBER Chapters* 14019.

Axtell, Robert L. (2001). “Zipf Distribution of U.S. Firm Sizes”. In: *Science* 293.5536, pp. 1818–1820.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson (2020). “Artificial Intelligence, Firm Growth, and Industry Concentration”. In.

Baldwin, Richard E. and Toshihiro Okubo (Nov. 2005). “Heterogeneous firms, agglomeration and economic geography: spatial selection and sorting”. In: *Journal of Economic Geography* 6.3, pp. 323–346.

Baqaee, David and Emmanuel Farhi (May 2020). *The Darwinian Returns to Scale*. Working Paper 27139. National Bureau of Economic Research.

Barseghyan, Levon and Riccardo DiCecio (2011). “Entry costs, industry structure, and cross-country income and TFP differences”. In: *Journal of Economic Theory* 146.5, pp. 1828–1851.

— (2016). “Externalities, endogenous productivity, and poverty traps”. In: *European Economic Review* 85.C, pp. 112–126.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw (2007). “How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills”. In: *The Quarterly Journal of Economics* 122.4, pp. 1721–1758.

Bartelsman, Eric, George Van Leeuwen, and Henry Nieuwenhuijsen (1998). "Adoption Of Advanced Manufacturing Technology And Firm Performance In The Netherlands". In: *Economics of Innovation and New Technology* 6.4, pp. 291–312.

Barth, Erling, Marianne Roed, Pal Schone, and Janis Umblijs (2020). *How Robots Change Within-Firm Wage Inequality*. IZA Discussion Papers 13605. Institute of Labor Economics (IZA).

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

Basso, Gaetano (2019). *The evolution of the occupational structure in Italy in the last decade*. Questioni di Economia e Finanza (Occasional Papers) 478. Bank of Italy, Economic Research and International Relations Area.

Basu, Susanto (2008). "Returns to Scale Measurement". In: *The New Palgrave Dictionary of Economics: Volume 1 – 8*. Ed. by Steven N. Durlauf and Lawrence E. Blume. London: Palgrave Macmillan UK, pp. 5559–5562.

Basu, Susanto and John Fernald (1996). "Returns to Scale in U.S. Production : Estimates and Implications". In: *Social Science Research Network* 1996, pp. 1–52.

Baumol, William J. (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis". In: *The American Economic Review* 57.3, pp. 415–426.

Bayrak, Ergin, John Conley, and Simon Wilkie (2011). "The Economics of Cloud Computing". In: 1118.

Bechtold, Stefan (2015). "3D printing and the intellectual property system". In: *WIPO Economic Research Working Paper* 28.

Bellemare, Marc F. and Casey J. Wichman (2020). "Elasticities and the Inverse Hyperbolic Sine Transformation". In: *Oxford Bulletin of Economics and Statistics* 82.1, pp. 50–61.

Benlian, Alexander, Thomas Hess, and Peter Buxmann (2009). "Drivers of SaaS-Adoption – An Empirical Study of Different Application Types". In: *Business & Information Systems Engineering* 1.357.

Bessen, James, Martin Goos, Anna Salomons, and Wiljan van den Berge (2019). "Automatic Reaction - What Happens to Workers at Firms that Automate?" In: *Boston University School of Law, Law and Economics Research Paper* 19-2.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2016). *Management as a Technology?* NBER Working Papers 22327. National Bureau of Economic Research, Inc.

Blundell, Richard and Steve Bond (2000). "GMM Estimation with persistent panel data: an application to production functions". In: *Econometric Reviews* 19.3, pp. 321–340.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch (2021). "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data". In: *Journal of Monetary Economics* 121, pp. 1–14.

Borusyak, Kirill and Xavier Jaravel (2017). "Revisiting Event Study Designs". In: *Working Paper*.

Brynjolfsson, E. and A. McAfee (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton.

Burstein, Ariel, Vasco M Carvalho, and Basile Grassi (Oct. 2020). *Bottom-up Markup Fluctuations*. Working Paper 27958. National Bureau of Economic Research.

Bustos, Paula (2011). "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms". In: *American Economic Review* 101.1, pp. 304–40.

Caliendo, Lorenzo and Esteban Rossi-Hansberg (2012). "The Impact of Trade on Organization and Productivity*". In: *The Quarterly Journal of Economics* 127.3, pp. 1393–1467.

Carbonero, Francesco, Ekkehard Ernst, and Enzo Weber (2020). *Robots Worldwide: The Impact of Automation on Employment and Trade*. Working Papers. IAB.

Carree, Martin, Enrico Santarelli, and Ingrid Verheul (2008). "Firm entry and exit in Italian provinces and the relationship with unemployment". In: *International Entrepreneurship and Management Journal* 4, pp. 171–186.

Cheng, Hong, Ruixue Jia, Dandan Li, and Hongbin Li (2019). "The Rise of Robots in China". In: *Journal of Economics Perspectives* 33(2), pp. 71–88.

Chiacchio, Francesco, Georgios Petropoulos, and David Pichler (2018). *The impact of industrial robots on EU employment and wages - A local labour market approach*. Working Papers 25186. Bruegel.

Chiavari, Andrea and Sampreet Goraya (Jan. 2021). *The Rise of Intangible Capital and the Macroeconomic Implications*. Working Paper. Department of Economics, Economics, Universitat Pompeu Fabra and Barcelona GSE.

Clementi, Gian Luca and Berardino Palazzo (2016). "Entry, exit, firm dynamics, and aggregate fluctuations". In: *American Economic Journal: Macroeconomics* 8.3, pp. 1–41.

Cogley, Timothy and James M Nason (1995). "Output Dynamics in Real-Business-Cycle Models". In: *American Economic Review* 85.3, pp. 492–511.

Cooper, Russell W. and John Haltiwanger (2006). "On the Nature of Capital Adjustment Costs". In: *The Review of Economic Studies* 73.3, pp. 611–633.

Cortes, Guido Matias, Nir Jaimovich, and Henry Siu (2016). "Disappearing Routine Jobs: Who, How, and Why?" In: *National Bureau of Economic Research, Working Paper* 22918.

Costinot, Arnaud, Andrés Rodríguez-Clare, and Iván Werning (2020). "Micro to Macro: Optimal Trade Policy With Firm Heterogeneity". In: *Econometrica* 88.6, pp. 2739–2776.

Dargatz, D.A. and G.W. Hill (1996). "Analysis of survey data". In: *Preventive Veterinary Medicine* 28, pp. 225–237.

Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Wößner (Sept. 2017). "German Robots - The Impact of Industrial Robots on Workers". In: *CEPR Discussion Papers* 12306.

Davis, Steven J., John Haltiwanger, and Scott Schuh (1996). "Small Business and Job Creation: Dissecting the Myth and Reassessing the Facts". In: *Small Business Economics* 8.4, pp. 297–315.

De Loecker, Jan and Frederic Warzynski (2012). "Markups and firm-level export status". In: *American Economic Review* 102.6, pp. 2437–71.

De Ridder, Maarten (Mar. 2019). *Market Power and Innovation in the Intangible Economy*. Discussion Papers 1907. Centre for Macroeconomics (CFM).

Dechezleprêtre, Antoine, David Hémous, Morten Olsen, and Carlo Zanella (2020). *Automating Labor: Evidence from Firm-Level Patent Data*. CEP Discussion Papers dp1679.pdf. Centre for Economic Performance, LSE.

Dinlersoz, Emin and Zoltan Wolf (2018). "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing". In: *Working Papers 18-39, Center for Economic Studies, U.S. Census Bureau*.

Dixon, Jay, Bryan Hong, and Lynn Wu (2020). *The Robot Revolution: Managerial and Employment Consequences for Firms*. Working Papers. NYU Stern School of Business.

Doraszelski, Ulrich and Jordi Jaumandreu (2018). "Measuring the Bias of Technological Change". In: *Journal of Political Economy* 126.3, pp. 1027–1084.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu (May 2021). *How Costly Are Markups?* NBER Working Papers 24800. National Bureau of Economic Research, Inc.

Edquist, Harald, Peter Goodridge, and Jonathan Haskel (2019). "The Internet of Things and economic growth in a panel of countries". In: *Economics of Innovation and New Technology* 0.0, pp. 1–22.

Einav, Liran and Jonathan Levin (2014). "The Data Revolution and Economic Analysis". In: *Innovation Policy and the Economy* 14, pp. 1–24.

Etro, F. (2009). "The Economic Impact of Cloud Computing on Business Creation, Employment and Output in Europe. An application of the Endogenous Market Structures Approach to a GPT innovation". In: *Review of Business and Economic Literature* 54.2, pp. 179–209.

Faunce, William A. (1965). "Automation and the Division of Labor". In: *Social Problems* 13.2, pp. 149–160.

Flynn, Zach, James Traina, and Amit Gandhi (2019). *Measuring Markups with Production Data*. Working Papers.

Ford, M. (2015). *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.

Fox, John and Georges Monette (1992). "Generalized Collinearity Diagnostics". In: *Journal of the American Statistical Association* 87.417, pp. 178–183.

Fujii, Hidemichi and Shunsuke Managi (2018). "Trends and priority shifts in artificial intelligence technology invention: A global patent analysis". In: *Economic Analysis and Policy* 58.C, pp. 60–69.

Gandhi, Amit, Salvador Navarro, and David A. Rivers (2020). "On the Identification of Gross Output Production Functions". In: *Journal of Political Economy* 128.8, pp. 2973–3016.

Gao, Wei and Matthias Kehrig (Feb. 2020). "Returns to Scale, Productivity and Competition: Empirical Evidence from US Manufacturing and Construction Establishments". In.

Garicano, Luis, Claire Lelarge, and John Van Reenen (2016). "Firm Size Distortions and the Productivity Distribution: Evidence from France". In: *American Economic Review* 106.11, pp. 3439–79.

Girma, S. and H. Görg (2002). "Foreign Ownership, Returns to Scale and Productivity: Evidence from UK Manufacturing Establishments". In: *CEPR Discussion Paper Series*.

Goos, Maarten and Alan Manning (2007). "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain". In: *The Review of Economics and Statistics* 89.1, pp. 118–133.

Gorodnichenko, Yuriy (Dec. 2007). *Using Firm Optimization to Evaluate and Estimate Returns to Scale*. Working Paper 13666. National Bureau of Economic Research.

Gorodnichenko, Yuriy, Debora Revoltella, Jan Svejnar, and Christoph T. Weiss (2018). *Resource misallocation in European firms: The role of constraints, firm characteristics* *and managerial decisions*. EIB Working Papers 2018/06. European Investment Bank (EIB).

Gosling, Amanda, Stephen Machin, and Costas Meghir (2000). "The Changing Distribution of Male Wages in the U.K." In: *The Review of Economic Studies* 67.4, pp. 635–666.

Graetz, Georg and Guy Michaels (Mar. 2015). "Robots at Work". In: *CEP Discussion Papers* 1335.

Griliches, Zvi (1969). "Capital-Skill Complementarity". In: *The Review of Economics and Statistics* 51.4, pp. 465–68.

Hall, Robert (1988). "The relation between price and marginal cost in US industry". In: *Journal of political Economy* 96.5, pp. 921–947.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013). "Who Creates Jobs? Small Versus Large Versus Young". In: *The Review of Economics and Statistics* 95.2, pp. 347–361.

Hamermesh, Daniel S. (1989). "Labor Demand and the Structure of Adjustment Costs". In: *The American Economic Review* 79.4, pp. 674–689.

Harris, Richard and Eunice Lau (Apr. 1998). "Verdoorn's law and increasing returns to scale in the UK regions, 1968–91: some new estimates based on the cointegration approach". In: *Oxford Economic Papers* 50.2, pp. 201–219.

Haskel, Jonathan, Peter Goodridge, and Gavin Wallis (Feb. 2015). *Accounting for the UK productivity puzzle: a decomposition and predictions*. Working Papers 21167. Imperial College, London, Imperial College Business School.

Haskel, Jonathan and Stian Westlake (2017). *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton University Press.

Hopenhayn, Hugo (1992). "Entry, Exit, and Firm Dynamics in Long Run Equilibrium". In: *Econometrica* 60.5, pp. 1127–1150.

— (2014). "Firms, misallocation, and aggregate productivity: A review". In: *Annual Review of Economics* 6.1, pp. 735–770.

Hsieh, Chang-Tai and Peter J. Klenow (2009). “Misallocation and Manufacturing TFP in China and India”. In: *The Quarterly Journal of Economics* 124.4, pp. 1403–1448.

Hubmer, Joachim and Pascual Restrepo (2021). *Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital–Labor Substitution*. Working Papers w28579. NBER.

Humlum, Anders (2019). “Robot Adoption and Labor Market Dynamics”. In.

Hwang, Kyung-In and Anthony Savagar (2020). “Product Market Power and TFP”. In: *Working Paper*.

Imai, Kosuke and In Song Kim (2019). “When Should We Use Unit Fixed Effects Regression Models for Causal Inference with Longitudinal Data?” In: *American Journal of Political Science* 63.2, pp. 467–490.

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

Jager, Angela, Cornelius Moll, Christoph Zanker, and Oliver Som (2015). *Analysis of the impact of robotic systems on employment in the European Union*. Working Papers. Fraunhofer ISI.

Jaimovich, Nir and Max Floetotto (2008). “Firm dynamics, markup variations, and the business cycle”. In: *Journal of Monetary Economics* 55.7, pp. 1238–1252.

Jaimovich, Nir and Henry Siu (2012). *Job Polarization and Jobless Recoveries*. NBER Working Papers 18334. National Bureau of Economic Research, Inc.

Jorgenson, Dale W. (2005). “Chapter 10 Accounting for Growth in the Information Age”. In: ed. by Philippe Aghion and Steven N. Durlauf. Vol. 1. Handbook of Economic Growth. Elsevier, pp. 743–815.

Jung, Jin Hwa and Dong-Geon Lim (2020). “Industrial robots, employment growth, and labor cost: A simultaneous equation analysis”. In: *Technological Forecasting and Social Change* 159, p. 120202.

Kariel, Joel (2021a). *Firms That Automate: Theory and Evidence*. Working Papers. Department of Economics, University of Oxford.

— (2021b). “Job Creators or Job Killers? Heterogeneous Effects of Industrial Robots on UK Employment”. In: *LABOUR* 35.1, pp. 52–78.

Kim, Daisoon and Anthony Savagar (2021). *Revenue Elasticities and Macroeconomic Dynamics*. Working Papers.

Koch, Michael, Ilya Manuylov, and Marcel Smolka (2019). "Robots and Firms". In: *CESifo Working Paper* 7608.

Koziel, Slawomir and Xin-She Yang (2011). *Computational Optimization, Methods and Algorithms*. Vol. 356.

Kromann, Lene and Anders Sorensen (2019). "Automation, Performance, and International Competition: A Firm-Level Comparison of Process Innovation". In: *Working Papers 03-2015, Copenhagen Business School, Department of Economics.*

Kwon, Myung Joong and Paul Stoneman (1995). "The Impact Of Technology Adoption On Firm Productivity". In: *Economics of Innovation and New Technology* 3.3-4, pp. 219–234.

Kydland, Finn E. and Edward C. Prescott (1991). "Hours and Employment Variation in Business Cycle Theory". In: *Economic Theory* 1.1, pp. 63–81.

Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard (2019). *Information Technology and Returns to Scale*. 2019 Meeting Papers 1380. Society for Economic Dynamics.

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

Levinsohn, James and Amil Petrin (2003). "Estimating production functions using inputs to control for unobservables". In: *The Review of Economic Studies* 70.2, pp. 317–341.

Lewis, Ethan (May 2011). "Immigration, Skill Mix, and Capital Skill Complementarity". In: *The Quarterly Journal of Economics* 126.2, pp. 1029–1069.

Lewis, Robert and Virginia Torczon (2002). "A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds". In: *SIAM Journal on Optimization* 12.4, pp. 1075–1089.

Liang, Kung-Yee and Scott L. Zeger (1986). "Longitudinal data analysis using generalized linear models". In: *Biometrika* 73.1, pp. 13–22.

Madić, Miloš and Miroslav Radovanović (2014). “Optimization of machining processes using pattern search algorithm”. In: *International Journal of Industrial Engineering Computations* 5, pp. 223–234.

Manaresi, Francesco (2015). *Net employment growth by firm size and age in Italy*. Questioni di Economia e Finanza (Occasional Papers) 298. Bank of Italy, Economic Research and International Relations Area.

Mann, Katja and Lukas Püttmann (2018). “Benign Effects of Automation: New Evidence from Patent Texts”. In.

Manyika, James, Michael Chiu, Peter Bisson, Jonathan Woetzel, Jacques Bughin, and Dan Aharon (2015). *The Internet of Things: Mapping the Value Beyond the Hype*. Tech. rep.

Martin, Ralf (2002). *Building the capital stock*. CeRiBA Working Paper. The Centre for Research into Business Activity.

Melitz, Marc J (Nov. 2003). “The impact of trade on intra-industry reallocations and aggregate industry productivity”. In: *Econometrica* 71.6, pp. 1695–1725.

Moretti, Enrico (2010). “Local Labor Markets”. In: *National Bureau of Economic Research, Working Paper*. Working Paper Series 15947.

O’brien, Robert (2007). “A Caution Regarding Rules of Thumb for Variance Inflation Factors”. In: *Quality & Quantity: International Journal of Methodology* 41.5, pp. 673–690.

Olley, G. Steven and Ariel Pakes (1996). “The dynamics of productivity in the telecommunications equipment industry”. In: *Econometrica* 64.6, pp. 1263–1297.

Oulton, Nicholas (1996). “Increasing Returns and Externalities in UK Manufacturing: Myth or Reality?” In: *The Journal of Industrial Economics* 44.1, pp. 99–113.

Pasold, E.W. (1975). “In search of William Lee”. In: *Textile History* 6.1, pp. 7–17.

Paulinyi, Akos (1986). “John Kay’s Flying Shuttle: Some Considerations on His Technical Capacity and Economic Impact”. In: *Textile History* 17.2, pp. 149–166.

Reeves, Phil and Dinusha Mendis (2015). *The Current Status and Impact of 3D Printing Within the Industrial Sector: An Analysis of Six Case Studies*. Tech. rep. MSU-CSE-06-2. The Intellectual Property Office.

Restuccia, Diego and Richard Rogerson (2008). "Policy distortions and aggregate productivity with heterogeneous establishments". In: *Review of Economic dynamics* 11.4, pp. 707–720.

Robotics, International Federation of (2014). *Robotics Data*. data retrieved from International Federation of Robotics, https://my.worldrobotics.org/.

Ruzic, Dimitrije and Sui-Jade Ho (2019). "Returns to Scale, Productivity Measurement, and Trends in U.S. Manufacturing Misallocation". In: *INSEAD working paper*.

Sanderson, Eleanor and Frank Windmeijer (2016). "A weak instrument F-test in linear IV models with multiple endogenous variables". In: *Journal of Econometrics* 190.2. Endogeneity Problems in Econometrics, pp. 212–221.

Savagar, Anthony (Mar. 2021). *Measured Productivity with Endogenous Markups and Economic Profits*. Discussion Papers CFM-DP2021-10. Centre for Macroeconomics (CFM).

Seamans, Robert and Manav Raj (2018a). "AI, Labor, Productivity and the Need for Firm-Level Data". In: *NBER Working Paper Series*.

— (2018b). "AI, Labor, Productivity and the Need for Firm-Level Data". In: *NBER Working Paper Series*.

Sedláček, Petr and Vincent Sterk (2017). "The Growth Potential of Startups over the Business Cycle". In: *American Economic Review* 107.10, pp. 3182–3210.

Segarra, Agustí and Mercedes Teruel (2012). "An appraisal of firm size distribution: Does sample size matter?" In: *Journal of Economic Behavior & Organization* 82.1, pp. 314–328.

Solow, Robert (1987). "We'd Better Watch Out". In: *New York Times Book Review*.

Stapleton, Katherine and Michael Webb (2020). *Automation, trade and multinational activity: Micro evidence from Spain*. CSAE Working Paper Series 2020-16. Centre for the Study of African Economies, University of Oxford.

Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley (2021). “The Nature of Firm Growth”. In: *American Economic Review* 111.2, pp. 547–79.

Stock, James and Motohiro Yogo (2005). “Testing for Weak Instruments in Linear IV Regression”. In: *Identification and Inference for Econometric Models*. Ed. by Donald W.K. Andrews. New York: Cambridge University Press, pp. 80–108.

Suedekum, Jens and Nicole Woessner (2019). *Robots & the Rise of European Superstar Firms*. European Economy - Discussion Papers 2015 - 118. Directorate General Economic and Financial Affairs (DG ECFIN), European Commission.

Susskind, Daniel (2017). “A Model of Technological Unemployment”. In: *University of Oxford, Department of Economics Discussion Paper Series* 819.

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

Syverson, Chad (2017). “Challenges to Mismeasurement Explanations for the US Productivity Slowdown”. In: *Journal of Economic Perspectives* 31.2, pp. 165–86.

Tintner, Gerhard (1944). “A Note on the Derivation of Production Functions from Farm Records”. In: *Econometrica* 12.1, pp. 26–34.

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

Triplett, Jack E. (1999). “The Solow Productivity Paradox: What do Computers do to Productivity?” In: *The Canadian Journal of Economics* 32.2, pp. 309–334.

Vannutelli, Silvia, Sergio Scicchitano, and Marco Biagetti (2021). *Routine biased technological change and wage inequality: do workers' perceptions matter?* GLO Discussion Paper Series 763. Global Labor Organization (GLO).

Weller, Christian, Robin Kleer, and Frank Piller (2015). “Economic implications of 3D printing: Market structure models in light of additive manufacturing revisited”. In: *International Journal of Production Economics* 164, pp. 43–56.

Wooldridge, Jeffrey M (2009). “On estimating firm-level production functions using proxy variables to control for unobservables”. In: *Economics Letters* 104.3, pp. 112–114.

Wunsch-Vincent, Sacha, Julio Raffo, and C Andrew Keisner (2015). “Breakthrough technologies – Robotics, innovation and intellectual property”. In: *WIPO Economic Research Working Paper* 30.

Zator, Michal (2019). “Digitization and Automation: Firm Investment and Labor Outcomes”. In.

Zeira, Joseph (1998). “Workers, Machines, and Economic Growth”. In: *The Quarterly Journal of Economics* 113.4, pp. 1091–1117.

Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz (2020). *Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey*. Working Paper 28290. National Bureau of Economic Research.

Zwick, Eric and James Mahon (2017). “Tax Policy and Heterogeneous Investment Behavior”. In: *American Economic Review* 107.1, pp. 217–48.