Rage against the machine?

The effects of robots on labour demand in the Norwegian manufacturing industry from 2003 to 2016

Ola Helmich Borchgrevink Pedersen

ECON4091 – Master's Thesis and Oral Exam (30 credits)

Master thesis submitted for the degree of

Master in Economics

Department of Economics

Faculty of Social Sciences

University of Oslo

May 2022

Rage against the machine?

The effects of robots on labour demand in the Norwegian manufacturing industry from 2003 to 2016

Ola Helmich Borchgrevink Pedersen

Rage against the machine? The effects of robots on labour demand in the Norwegian manufacturing industry from 2003 to 2016

Ola Helmich Borchgrevink Pedersen

http://www.duo.uio.no/

Print: Graphic center, University of Oslo

Abstract

This paper investigates the effects of robots on labour demand at the firm level in the Norwegian manufacturing industry from 2003 to 2016. Labour demand is proxied by within-firm wage shares of different worker groups and access to robot technology is proxied by imports of industrial robots. The worker groups are based on their level of completed education; compulsory school (CS), upper secondary school (US) and higher education (HE). I use both dynamic and static two-way fixed effects (TWFE) regressions to examine the effects of robots on these worker groups' wage share. Firstly, the results show that robot firms generally have a higher wage share of HE workers and a lower share of CS and US workers *before* robot adoption. Secondly, through the time period, there is a skill biased and *robot-independent* trend of increasing wage share for HE workers in all firms, although this trend is slightly stronger in robot firms. The opposite is true for the wage shares of CS and US workers, which have a negative trend. Thirdly, *after* the firm adopts robots the wage share for US workers increases and the wage share for HE workers decreases. CS workers' wage share is unaffected. The results are confirmed by robustness tests. Viewed together with the findings of Barth et al. (2020), who examines the same firms and workers, these results suggests that robots lead to relatively higher wages for HE workers, while US workers are relatively more employed. The latter effect seems to dominate as the total effect, examined in this paper, is a relative increased demand for US workers and reduced relative demand for HE workers after Norwegian manufacturing firms adopt robots.

Acknowledgements

Without the outstanding support and guidance by eminent supervisors, this thesis would have been difficult to write. In particular, discussions and input from my main supervisor Erling Barth from ISF have been invaluable. I would also like to thank my co-supervisors from University of Oslo, Jo Thori Lind and Karl Ove Moene, for interesting discussions, constructive feedback and friendly support throughout.

This degree in general and this thesis in particular would not have been the same without my good friend and study buddy Fanny Cecile Berg. I would also like to thank other co-students and professors from whom I have learned a great deal and gotten inspiration from over the last two years. Lastly, I would like to thank my family, good friends and my brilliant girlfriend who all have been of great support and a source of encouragement when writing this thesis.

In addition to supervisors, I would like to thank Fanny Cecile Berg, Nora Jungeilges Heyerdahl, Marit Gran Andreassen, Johannes Terjeson Bangum, Halvor Sætre, Håkon Dragesæt Tomter, Rolf Helmich Pedersen and Christine Borchgrevink for comments and ideas.

Any errors are my own. Data processing and estimation are all done in RStudio. For questions or access to R-code, feel free to reach out to me on email: olahbpedersen@hotmail.com.

This thesis have been supported by access to data, a grant and research facilities from Institute for Social Research (Institutt for Samfunnsforskning, ISF). This thesis has been written as part of "Changing Health And Skills Requirements in the Labour Market", a collaborative project between ISF, Harvard University and National Bureau of Economic Research (NBER) (project code #280307 from The Research Council of Norway). I would like to thank ISF in general, the ARV (Work and Welfare) working group, my funny and clever office buddy Ebba Henrekson and all the other amazing people working at ISF for providing me with this great opportunity to working with and learning from you.

Contents

Figures and tables

List of Figures

List of Tables

1 Introduction

I fear old technology, not new!
- Stefan Löfven[1]

Who will win and who will lose from the rapid advances in technology? Two trends in recent decades have increased the motivation to understand this question. Firstly, there has been exceptional improvements in technology and increased use of it in the production process. Development of new technologies, and the distresses of their consequences, are indeed not new (recall for example The Luddites, the infamous 19$^{\text{th}}$ century machine breakers (Hobsbawm, 1952)). What has changed since then is the confidence in human abilities to stay competitive (D. H. Autor, 2015; Furman, 2019). When the first machines came, humans could invest in education to improve their most important comparative advantage, namely cognitive skills and knowledge. Now, whispers of artificial intelligence (AI) (Barrat, 2013), even super-intelligence (Bostrom, 2014), and robots dancing as good as most of us, have lead some to believe that humans soon will lack sufficient competitive advantages and thus, sooner or later, be obsolete for purposes of production. Although such extreme consequences are disputed and far from certain,[2] recent studies have predicted a loss of up 50% of jobs in the US economy and 33 % in the Norwegian economy due to automation in the coming decades (Frey & Osborne, 2017; Pajarinen et al., 2015). Surveys also find increasing concerns in the population about automation, although this varies across countries (Ingelsrud & Steen, 2021; Smith & Anderson, 2017). The second trend motivating this question is the increases in inequality documented in most countries in recent decades (Aaberge et al., 2021; Piketty, 2013). Although the causes of this trend are complex, there is a consensus that technology has at least partly participated to this development (Acemoglu, 2002; Korinek & Stiglitz, 2019). Trends increasing inequality, whether it is globalisation or technology or something else, are tough to subdue, but there is room for policy to level the playing field (Goolsbee, 2019). The sense that policy have failed in this manner in the last couple of decades have pushed for increased understanding of the effects of technology and how we might respond.

This thesis attempts to contribute to the former of these aspirations; increased understanding of how robots affect labour demand. Although many recent studies have focused on the classic division of income between *capital and labour* (in particular the new evidence of a falling labour share of income in advanced economies in recent decades (D. Autor et al., 2020b; Elsby et al., 2013; Karabarbounis & Neiman, 2014)), this paper will follow recent literature on how the labour share is divided between different groups of workers (Acemoglu & Autor, 2011a; Acemoglu & Restrepo, 2019a; Asplund et al., 2011; Graetz & Michaels, 2018; Hirvonen et al., 2021; Humlum, 2019; Koch et al., 2021; Koren et al.,

[1] Former leader of the largest manufacturing industry union in Sweden (Industrifacket Metall) and former prime minister of Sweden (2014-2021, Social democrats): Original quote "Jeg frykter ikke fremtiden og ny teknologi, men gammel", quoted by Jonas Gahr Støre in Malvik (2016).

[2] In a recent survey of leading economic experts, when asked whether new robots and AI will lead to more unemployment (in the medium run), the most common answer was "Agree", both in the US and European panel. Many were uncertain and disagreed, so there was no clear consensus. However, in both panels there is a consensus that the benefits for this technology is sufficient to compensate those that loose their jobs (IGM, 2019).

2020). In particular, I will investigate the following question: *how is the relative demand for different groups of labour, based on their level of education, affected by robots at the firm level in the Norwegian Manufacturing industry?* To investigate this question I use a Difference-in-Differences (DiD) method allowing for heterogeneous treatment time to assess the changes in the worker groups' wage share (proxy for relative demand) following robot imports (proxy for access to robot technology).[3] As such, this paper complements a recent paper by Barth et al. (2020) who studies the same firms and workers as this thesis. They investigate the effect of robots on the demand for different worker groups (identically grouped by education) by looking at what happens to within-firm wage differentials. They find that the wage for workers with higher education (HE) increase relative to workers with upper secondary (US) school or compulsory school (CS). Although this suggests increased demand for higher education workers, this measure does not capture what happens to employment for the different worker groups following introduction of robots. Using the wage share as the outcome variable allows us to capture both the effects on wages and employment to understand changes in demand.

This thesis is structured as follows. Firstly, recent literature is examined in Chapter 2, both the developments in robot technology and its usage, and its effects when introduced in manufacturing firms. Chapter 3 presents two stylized models to understand how introduction of robots at the firm level might influence the demand for labour. The first model is a classic CES model, where robots enter the production function by augmenting the other input factors. The second model is the increasingly popular task-based framework developed by Acemoglu and Restrepo (2019a). Here, robots enter the production function by substituting the workers with the lowest skills. The over all outcome is determined by the increased need for workers as the firm expands and introduces new tasks needed for production. I also briefly discuss their lack of focus on labour market institutions, as these could impact both the firms decision to adopt robots and how their workers are affected (Moene & Wallerstein, 1997). Chapter 4 introduces the data that is going to be analyzed and then, through various descriptive statistics, set the scene to better understand the context at hand. I will look at developments in robot import and their value, what manufacturing robot firms typically look like (sectors, size and frequency of imports), as well as changes in the wage share and employment share over time. Chapter 5 introduces the empirical approach used to estimate the effects of robots on the wage share. Before presenting the models used in the estimation, I present these models in a potential outcome framework to understand how I can obtain internally valid estimates with the TWFE estimator. Before concluding in the last chapter, Chapter 6 presents the results. I find that there is a general robot-independent trend for both robot firms and non-robot firms toward higher wage share for HE workers and a lower wage share for CS and US workers, and that all these trends are stronger in robot firms. However, when introducing robots, the wage share of HE workers drop and the wage share US workers increase, suggesting a robot bias for the latter worker group. Viewed together with the results from Barth et al. (2020), this suggests that when robots are

[3] The wage share for group g in firm i with n worker groups is given by $S_{gi} = \frac{w_{gi}L_{gi}}{\sum_{g=1}^{n}(w_{gi}L_{gi})}$. Other papers have used the same measure to proxy labour demand for other types of changes in the production function, like Caroli and Van Reenen (2001) who look at the impact of organisational change.

introduced to Norwegian manufacturing firms, HE workers get relatively higher wages, while US workers are relatively more employed, and that the latter effect dominates. Thus, the total effect of robots is a relative increase in demand for US workers.

2 Literature review

2.1 Trends in automation and robots

Through the last couple of centuries, several technological paradigm shifts have transformed our societies. With recent development in information and communication technology (ICT) and robots, several authors have argued that we are presently at the beginning of another such paradigm shift (Brynjolfsson & McAfee, 2014; BSG, 2015; Ford, 2016; Frey & Osborne, 2017; McKinsey, 2017). Most famously, by assessing the probability of automation for 702 detailed occupations, Frey and Osborne (2017) found that new technology makes it possible to automate around half of all jobs in the US economy in the coming decades. Using these probabilities of automatisation on Norwegian and Finish jobs, Pajarinen et al. (2015) finds that around 1/3 of both Norwegian and Finish jobs are likely to be automated during the same time span. In Norway, they estimate that 44% of low education jobs are likely to be automated, substantially higher than the estimate for high education jobs (14%). They also find that a higher share of manufacturing jobs are likely to be automated (51%). These predictions complement much previous research, in Norway and elsewhere, on increased wage inequality and job polarization following automatisation, especially on low and medium skilled workers (Asplund et al., 2011; D. H. Autor et al., 2003; Goos & Manning, 2007; Michaels et al., 2014).

Automation is a broad term, defined as a process of replacing labour by other input factors for production. Of the many was to automate generally, robots is the most relevant cause in the manufacturing industry. There are two types of robots used for production, industrial robots and service robots, and only the former is the focus of this thesis. The first industrial robot, *Unimate*, was made in 1958 by Joseph Engelberger - the Father of Robotics - and his companion George Devol to be used on the assembly line for General Motors. Since then, industrial robots have change radically in quality and quantity. The International Federation of Robotics (IFR) surveys firms worldwide and is a source for much of the empirical work in recent years on the effects of robotisation. They document a five fold increase in use of robotics in Europe, the US and Asia between 1993 and 2015 (IFR, 2016). Before a sudden plunge prior to the financial crisis in 2008, the average annual number of robots sold was about 115,000 units. In 2010 it was back up to around 120,000 before peaking at over 400,000 new installed units in 2018. The latest numbers suggest around 400,000 new units in 2020 adding up to a total operational stock of industrial robots of more than 3 million units (IFR, 2021). Historically, Japan, South Korea, the US and Germany have been the largest markets for industrial robots (Cheng et al., 2019), only recently passed by China. Through the last couple of years, the automotive industry has been the sector with the highest implementation of robots, but was passed by electronics in 2020 (IFR, 2021). The drivers behind the increased pace of robotisation is cheaper, better and more adoptable robots, more access to high skills workers and high manufacturing labour costs. Most recent reports expect robotisation to continue at an increased rate as robots become even better and cheaper relative to labour (BSG, 2015; McKinsey, 2019).

Barth et al. (2020) document a steady increase in import of industrial robots in all sectors in Norway, from around 100 robots imported annually in 1999 to around 300 in 20016. Although the IFR data for Norway give slightly lower numbers, the trend follows a similar pattern. This thesis (see Section 4.2) documents a similar development for the manufacturing sector. In 2003, 19 robots worth 19 million NOK were imported. This figure rose to 41 robots worth around 110 million in 2016, doubling the number of robots and increasing the value six fold.

2.2 Effects of robots on employment, wages and skill

The effect of robots have mainly focused on the manufacturing industry, which is also the focus of this thesis. The literature is divided between papers focusing on the aggregate economy (region) and papers focusing on the firm level. The data on robots are usually from IFR, survey data or import data, the latter is used in this thesis. The methodological approach is usually some DiD approach and event studies, but several other methods have also been applied.

Two influential papers have studied the effects of adoption of industrial robots on the aggregate level. In their seminal paper, Graetz and Michaels (2018) study the effect of industrial robots on several economic output variables in seventeen developed countries from 1993 to 2007, using IFR data. They found that industrial robots increased both labour and total factor productivity, increase in growth and wages, and reduction of factor prices. Total employment does not decrease, but demand for low-skilled labour fall. Acemoglu and Restrepo (2020a) uses the same data source on the US labour market and identify an aggregate negative effect on employment. The effect is strongest in manufacturing sectors with high penetration of robots and for workers with routine jobs like machinists, assemblers, material handlers, and welders. Surprisingly, this is true for workers with less than high school, high school degree, but also for workers with only some college, and college or professional degree. Moreover, there is no positive effect for workers with a masters or doctoral degree. An important point in regard to the broader automatisation debate is that they only find these negative effects from industrial robots, not for IT technology or other capital deepening, suggesting the automating effect of industrial robots stand out in the manufacturing industry. In addition, Aghion et al. (2019) finds the same negative effect on employment on a district level in France from 1994-2014, also based on IFR data. Dauth et al. (2021) follows the same strategy as Acemoglu and Restrepo (2020a) for the German Economy from 1994 to 2014, and find that employment was not negatively impacted by robots on an aggregate level. They argue that this finding veils the underlying effects of displacement and reallocation effects, of which I will return to in the next section.

In recent years, several papers have also investigated the effects of robots at the firm level in several countries. Most studies find that robotisation increases output and productivity, but reduce the labour share (Acemoglu et al., 2020; Dauth et al., 2021; Humlum, 2019). On aggregate employment, some papers find that firms adopting robots increase overall employment, for example in Finland between 1994–2018 (Hirvonen et al., 2021), the Canadian Economy from 2000 to 2015 (Dixon et al., 2021) and in the French manufacturing sector between 2010 to 2015 (Acemoglu et al., 2020). However, some papers also find no effects on over all employment (Dauth et al., 2021), or even negative effects in some cases (e.g. in all non-financial firms in Netherlands between 20000-2016 (Bessen et al., 2019)).

Most interesting for our purposes, however, is to investigate previous findings on demand for different types of labour following robot adoption at the firm level. Several studies do not find any such bias. Hirvonen et al. (2021) show that increased employment in Finish firms apply for all skill levels and thus do not change the skill composition. Although finding the opposite employment effect, i.e. increased chance of displacement, Bessen et al. (2019) find that the effect of robots do not differ across skill levels. In the German context, Dauth et al. (2021) also find that the effects of robots are not biased across skill levels, but within, depending on whether workers stay at the robot adopting firm or leave for education or other firms. However, other studies find that robots are indeed biased, usually toward high skilled workers. Humlum (2019) show that robotisation increases demand for high skilled tech workers while decreasing demand for low skilled production workers, similar to what Acemoglu et al. (2020) find in French manufacturing firms. In the Canadian setting, Dixon et al. (2021) finds evidence of polarization as demand for both high and and low skilled workers increased following robotisation, while demand for middle skilled workers fall. Using the same data as this thesis on the Norwegian manufacturing industry, Barth et al. (2020) find that hourly wages increase for HE workers, decrease for CS workers and do not affect US workers. Looking at occupation rather than skill level, they find that manager wages increase more than those in STEM or professional occupations, while blue collar workers wages decreases. Although they did not look at the relative effects on employment, the effect on wages suggest a relative increase in the demand for HE workers, implying that robots are indeed skill biased in the Norwegian context.

2.3 Mechanisms behind the findings

As I will examine in the next chapter, theory suggest several ways in which the observed outcomes can be explained. In most papers, interpretation of the mechanisms behind the observed outcomes are usually based on some assumed theoretical model rather than empirically investigated. However, it is worth mentioning those that follow the latter strategy. They typically apply the task-based framework, as this can give guidance for the expectation of our own inquiry.

Dauth et al. (2021) apply the task-based framework to specifically investigate how robots affect workers (see Chapter 3 for description of this model). In the first instance, they find that exposure to robot technology reduce employment for manufacturing firms, illustrating the displacement effect from the task-based framework. They suggest this effect first and foremost lead to less demand for new labour in the automating industries, burdening newly educated workers entering the labour force. However, this effect is countered by increased employment in other local service sectors, in particular firms that is used by other businesses, and that production workers change occupation within their firm after robot adoption. They argue this is a replacement effect, following firm expansion and assignment of new tasks to incumbent workers. Acemoglu et al. (2020) suggests that productivity effects of robots are larger than the displacement effect within French firms. However, total employment decrease as the increased productivity effect in robot adopting firms leads to expansion at the expense of competitors which reduce employment to the extent that total employment is reduced.

Although most papers use the task-based framework, some papers argue that what they observed cannot be fully explained by the conventional models. In particular, Hirvonen et al. (2021) argue that firms typically get robots to perform new tasks to produce new types of output, not replace workers in tasks they previously performed themselves, which is assumed in the task-based framework. Based on their observations, they develop a new theoretical framework to fit their results, building on work from Dixit and Stiglitz (1977) and Melitz (2003). To fit their observations, their model distinguishes between process technological change (the common way of understanding technological change, associated with *mass production*) and product technological change. The latter is what they claim to find for the Finish manufacturing industry between 1994–2018 and is typically associated with *flexible specialization* manufacturing. Looking at survey data on strategy priorities in Canadian firms, Dixon et al. (2021) also find that having a strategy on decreasing labour costs do not predict whether a firm adopts robot technology, but that improvement of product quality does. This also suggest the conventional models might not always be appropriate to explain what happens following robotisation. However, this contrasts what McKinsey (2019) finds in their recent report on robot adoption. They claim that lowering costs is the most important factor motivating firms to adopt robots, although this varies across sectors.

The observation that the employment effects differ from country to county, and particular between US and Europe, has motivated researchers to understand the role of labour unions. Haapanala et al. (2022) make a cross country comparison and find that, in countries with higher union density, robots are likely to benefit insiders, in particular high skilled labour, old employees and tertiary workers. They suggest that this because labour unions primary concern is the employment and income security of the 'insiders', rather than a concern for all workers in the economy. When facing robots, unions might agree to adopting robots if they ensure employment of insiders though training and new tasks. Although this findings is worth paying attention to, the exclusion of non-union workers might be less precises description of the Norwegian context where the large unions represent a broad range of worker interests. Nevertheless, (Dauth et al., 2021) also suggest that labour unions played a crucial role in the German context. They find that the displacement effect following robotisation is greater in areas with lower labour union density.

3 Models to understand technological changes

The aim of this chapter is to provide some intuition as to how robotisation might influence labour demand by looking at some commonly used models. After presenting the wage share, the CES production function is outlined. In this model, robots enters the production function by *augmenting* the other production inputs (differently skilled labour and capital). Although this can provide us with useful intuitions for changes in relative productivity following robotisation, several authors have pointed out its weaknesses when trying to understand changes like automatisation (Acemoglu & Autor, 2011a; Acemoglu & Restrepo, 2019b). Based on this criticism, the task-based framework has been adopted and developed in recent time. After presenting this framework, some concluding remarks are given to highlight how these frameworks can be used for interpretation of the results in this thesis.

Before moving on, it is worth noting some limitations to the models presented here. Firstly, all our models cannot explain *how* certain technological changes might come about. Instead, changes in technology are treated as completely exogenous. This thesis will follow this convenient trend, although one might be justified in questioning this assumption. Secondly, these models assume a fully competitive economy with costs minimizing firms. This ignores important labour market institutions such as unions. As discussed in Chapter 2, several papers highlights the importance of these institutions and they might be particularly relevant in the Norwegian context with a relatively high union density and centralized bargaining. In addition to be a mediating factor between robots and the wage share, it could also be that the level of unions could influence the decision to adopt robots of not. Nevertheless, they are useful to understand some core mechanisms between technology and labour demand.

3.1 Understanding the wage share

The wage share is used as a proxy for the relative demand of different labour groups in several papers related to changes in the production function, for example organisational change (Caroli & Van Reenen, 2001). Thus, it can also proxy the bias of such changes. The wage share for worker group g in firm i with n worker groups is given by

$$S_{gi} = \frac{w_{gi}L_{gi}}{\sum_{g=1}^{n}(w_{gi}L_{gi})} \tag{1}$$

The numerator is the total wage costs for worker group g and the denominator is the total wage costs for all workers in firm i. Crucial for interpretation is to realize that the wage share is a relative measure. If demand for some worker group increases (higher wage, w_g, and/or more workers employed, higher L_g) it does not necessarily imply that the wage share increase if the total wage costs for the firm increase at the same rate, meaning that demand other worker groups also increase. It could even be the case that the wage share for a worker group decreases even though they are in higher demand, if the demand for other workers increase relatively more. Thus, the wage share for worker group g increases only if the the wage costs for g *increases relatively more or decreases relatively less* compared to other worker groups. If the opposite is true, then the wage share decreases.

3.2 Effects of technology in a CES neoclassical framework

The commonly used CES framework treats improvements in technology as something augmenting input factors. To start with, let's consider a neoclassical production function, where technology is treated as an aggregate measure in a somewhat simplified way. Usually it takes the form of Hicks-neutral technological change, $Y = \tilde{F}(K, L; A) = AF(K, L)$, where the exogenous technology parameter represents the total factor productivity and thus assumes that all input factors' productivity are improving at the same rate. An alternative to this is to assign all productivity improvements to one factor. The Solow-neutral technological process takes the following form $Y = F(A_K K, L)$, i.e. technological change is always assumed to be improve the productivity of capital (the technology is capital-augmenting). Contrary to this, a Harrod-neutral production function, $Y = F(K, A_L L)$, assumes that all technological improvements are improving the productivity of labour (labour-augmenting). A more comprehensive model that allows to combine all these forms of technological change is the constant elasticity of substitution (CES) production functions (Jehle & Reny, 2011). To better fit this model to our purposes, let's imagine that the inputs are high skilled (L_H) and low skilled (L_L) labour so that $X = F(L_L, L_H)$ and A, A_L and A_H are the three forms of technological change. Then, the CES production function is given by

$$F(L_L, L_H) = A\left[\alpha(A_H L_H)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)(A_L L_L)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{2}$$

The CES production function above allows us to include all the three aforementioned forms of technological change in our production function. It also introduces a new term, the elasticity of substitution between the input factors, σ, which is a measure on how easily we can change between high and low skilled labour to produce a given output (see appendix for definition and further explanation of σ). In other words, it implies a certain curvature of the production isoquant curve.

For our purposes, it is useful to understand how technology might influence the relative marginal product (productivity) of the different labour groups. In a competitive context, this is assumed to determine the relative wage, which is a way to express the relative demand for workers. If some technological innovation change the relative marginal product of a worker group, then the technology is said to be *biased* toward this group. From the production function above we can derive the relative return of high and low skilled labour $\frac{MP_{L_H}}{MP_{L_L}}$ (see derivation in the appendix).

$$\frac{w_H}{w_L} = \frac{MP_{L_H}}{MP_{L_L}} = \frac{\alpha}{(1-\alpha)}\left(\frac{A_H}{A_L}\right)^{\frac{\sigma-1}{\sigma}}\left(\frac{L_L}{L_H}\right)^{\frac{1}{\sigma}} \tag{3}$$

Equation 3 shows us that the bias of the technology depends on how a certain technology augments the input factors (i.e. which of A, A_L and A_H are increased following the introduction of robots) and the substitution between high and low skilled labour (σ). To see this, let's consider the effects of technology augmenting high skilled labour (A_H) (the opposite will be true for increased A_L). If $\sigma < 1$, the power is negative and increasing A_H will correspond to a decrease in $\frac{w_H}{w_L}$, implying that low and high skilled labour are complements. This implies a bias toward low skilled labour. The opposite is true for $\sigma > 1$, which implies that low skilled and high skilled labour are substitutes and $\frac{w_H}{w_L}$ will increase, implying that the technological change is biased toward high skilled labour. If $\sigma = 1$ the production function collapses into a Cobb-Douglas production function and the relative return will not be affected by changes in either A_H or A_L as these both will correspond to a Hicks neutral technological change. In addition, it is worth noting that the bias of the technology are inversely related to the corresponding factor input ratio $\frac{L_H}{L_L}$. This represents a substitution effect; if the relative input of high skilled labour is high, the relative return (wage) of high skilled labour will decrease relative to low skilled labour. This effect will vary in magnitude depending on σ, but will always work in this direction. Secondly, we can see that the technological term representing Hicks-neutral technological process (total factor productivity), A, does not influence the relative factor returns. This is because it will influence the returns in the exact same way, and thus not influence their ratio.

Moreover, assuming cost minimizing and price taking behaviour of the firm, Biørn (2008) shows that we can derive the corresponding wage shares for the respective worker groups to see how they directly relate to technological changes. The wage share for worker group g is given by

$$\frac{w_g L_g}{c} = A_g^{\sigma} \left(\frac{w_g}{q} \right)^{1-\sigma} \tag{4}$$

where c is total wage costs $(w_L L_L + w_H L_H)$ and q is an aggregate price factor [4]. From equation 4 we can see that the wage share depend on the type of technological improvements (A_L or A_H) and the value of the elasticity of substitution σ as well as the wage. If A_g increases following robotisation, so will its factor cost share do as well. However, it is far from clear which of the labour groups are augmented following robotisation. Indeed, such unclear predictions from the CES framework has been the motivation to develop the task-based framework.

3.3 Effects of technology in a task-based framework

A core issue of the CES production function is that changes following technological improvements, in particular automatisation, is hard to isolate. This is because such technological improvements do not necessarily improve a particular input factor in all ways, as is assumed in the CES framework, but rather automate certain task needed to be done for production. Zeira (1998) recognized this and first introduced a task-based framework of production. This framework views tasks as the basic building block of production. Input factors such as differently skilled labour, robots and capital can provide task services to perform these tasks. Thus, to understand the effect of robots, this framework is particularly useful and is therefore applied in many recent papers (Acemoglu et al., 2020; Barth et al., 2020; Dauth et al., 2021; Graetz & Michaels, 2018; Humlum, 2019). Let's consider the version used by Acemoglu and Restrepo (2018), as they have pioneered this framework. The production function is given by

$$logY = \int_{N}^{N-1} lny(x)dx \tag{5}$$

[4] Aggregate price factor $= q = \left[A_L^{\sigma} w_L^{1-\sigma} + A_H^{\sigma} w_H^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$

The aggregate output Y is produced by combining task services from labour and machines to perform tasks x in the following interval $x \in [N, N-1]$. Thus, $y(x)$ is the output produced by performing task x. Two types of input can perform these task, either labour $l(x)$ or machines $m(x)$, and these are assumed to be given and fixed at the aggregate level (i.e. capital (K, including robots) and labour (L) are fixed). The core assumption of this model is that any task x can be performed by labour with productivity $a_L(x)$. Machines have productivity $a_L(x)$ and can only perform tasks up to some threshold I, exogenously given by the available level of technology. However, for these task, machines and labour are perfect substitutes. We assume that $\frac{a_{Li}}{a_{Ki}}$ is increasing in i, meaning that labour's competitiveness increases with the tasks complexity. Moreover, we assume that at the threshold, I, the productivity of labour relative to capital is smaller than the price of labour relative to the price of capital, $\frac{a_{Li}}{a_{Ki}} < \frac{w}{R}$. This means that, at the threshold, labour is more expensive than capital so the firm will automate this task. This implies that task below I are automated and thus output for task x is given by

$$y(x) = \begin{cases} \mathrm{a}_L(x)l(x) + a_M(x)m(x) & \text{if } x \in [N-1, I] \\ \mathrm{a}_L(x)l(x) & \text{if } x \in (I, N] \end{cases} \tag{6}$$

In equilibrium, wages will equal the marginal productivity of workers, and so labour demand is given by $W = (N-I)\frac{Y}{L}$. From equation 6, we can assess what is crucial in our model, namely the effect of robot adoption on the labour demand (i.e some exogenous improvements in robot technology, increasing the threshold I). Taking derivatives of Equation 6 gives

$$\frac{dlnw}{dI} = \underbrace{\frac{dln(1-I)}{dI}}_{Displacement\ Effect < 0} + \underbrace{\frac{dln(Y/L)}{dI}}_{Productivity\ Effect > 0} \tag{7}$$

Equation 7 show two main effects on demand following robotisation. The first effect of automation, the displacement effect (DE), implies a decrease in labour demand (always < 0) as the robots substitute labour in all tasks below the newly increased I. However, the productivity effect (PE) can counter this movement (always > 0) if a firm expands and thus demand more labour. The total effect will be determined by the difference in productivity at the marginally automated task. If capital is much more productive than labour in performing this task ($\frac{a_{LI}}{w} \ll \frac{a_{KI}}{R}$), then automation will have great benefits and will lead to a substantial PE, potentially offsetting DE and increase demand for labour. If the automation technology is only marginally more productive than labour (known as "so-so" technology), the PE will be lower and the total effect is likely to be negative.

In addition to these effects, Acemoglu and Restrepo (2019a) show that the consequences for labour demand also depend on the invention of new tasks needed for production (increase in N). By assuming that all new tasks initially only can be performed by labour, they show that these new tasks will both lead to a reinstatement effect of labour and productivity effect which both leads to higher demand for labour.

3.4 Summing up: what we can learn from theory

The purpose of outlining these models has been to give some intuition about the mechanisms affecting labour demand following technological change, in particular robotisation. In our econometric analysis, we will be able to estimate how the wage share (relative demand) changes for a certain worker group when a firm introduces robots. Although this results in it self does not unveil the causal mechanisms behind the results, we can apply the outlined theoretical frameworks to try to explain what might have happened. The CES framework suggests that changes in the wage share depend on how the different types of workers are augmented by robots and to what extent the workers are complements or substitutes with each other. The task-based framework think of robots as a perfect substitute for low skilled labour, which implies an displacement of these workers. However, this effect could be outweighed if robots improve productivity so firms hire more workers (productivity effect) or if the firm introduces new tasks for the workers to perform. Although these models provide useful intuition, they also ignore several aspects of the labour market, for example institutions such as unions and centralized bargaining. As discussed in Chapter 2, previous literature have suggested these institutions play a key role when firms adopt robots. This is important to bare when moving forward as they are likely to be particularly important in the Norwegian context.

4 Data and descriptive statistics

4.1 Data

The data used in this analysis is a long panel data set with firm-level variables from the Norwegian manufacturing industry from 2003 to 2016. The data is based on two major sources. The first data set give us data on our outcome and treatment variable. The second provides data on control variables.

The first data set contains data on the dependent and independent (treatment) variables of interest, namely the wage share for various worker groups (based on education) and access to robot technology, respectively. The data is obtained from Statistics Norway (SSB)[5] from the years 2003 to 2018 and contains data on all Norwegian manufacturing firms.[6] The register data from SSB on wage costs for different worker groups comes from the payroll statement, annually reported to The Norwegian Tax Authority ("Skatteetaten") from all Norwegian firms. The wage data show the annual wage costs and the number of employees by firm in four different worker groups, based on their level of completed education; compulsory school (CS), upper secondary school (US), some level of completed higher education (HE) and unknown level of education (UE). The latter group is marginal and is likely present due to registration and classification issues, and is likely to mostly consist of migrant workers. In a manufacturing context, CS workers are likely to be older workers who have worked in the firm in many years. US workers are likely to have High School vocational training and be blue collar workers. Based on this wage and employment data I calculate the annual total wage expenditures and number of employees for each firm, as well as employment and wage share for the respective worker groups for all firms. The latter share will be the dependent variable in this analysis.

Like other recent studies, this thesis use data on the import of industrial robots as a proxy for access to robot technology, which is the treatment variable in this thesis (Acemoglu et al., 2020; Aghion et al., 2019; Barth et al., 2020; Dixit & Stiglitz, 1977; Humlum, 2019). The import data is based on international trade data from Norway Trade Statistics Register and COMTRADE data from the United Nations Trade Statistics Database. Imported goods with the Harmonised System code "84795000" indicate that the good has been classified as an industrial robot. IFR defines an industrial robot as an automatically controlled, re-programmable, and multipurpose machine used in industrial automation applications (ISO-8373, 2021). This also implies that machines are not robots if they are not possible to reprogram to perform various different tasks or need humans to operate. The import data tells us if a robot was imported, the value of the imported robot(s), what firm imported the robot and what time it crossed the border. When referring to "robots" in this thesis, I always mean industrial robots.

[5] I got access to the data from Institute for Social Research (ISF) as a part of their student grant.

[6] As opposed to *corporate groups*. A corporate group is a conglomerate of firms operating under administrative and financial control of a parent company, and can consist of both manufacturing and non-manufacturing firms. In the SSB robot and wage data, the basis are manufacturing firms that are either not a part of a corporate group or are operating under a manufacturing corporate group. In addition, *non-manufacturing* firms from *manufacturing* corporate groups are excluded and *manufacturing* firms form ı corporate groups are included.

Several things are important to note on this first data set from SSB. Firstly, I assume that once firms are treated, they will stay treated throughout the time period. This means that I assume that no robots are destroyed and that new rounds of robot import for a firm do not change the treatment variable. Secondly, there are several reasons to believe that imported robots are not homogeneous. Any point in time, a certain set of various robots with various quality are available for import and it is possible to import any robot technology in different *quantities*. Import of any such combination will all count as being treated. Moreover, the set of available robots for import are likely to differ in quality *across time* in the period examined. Together with the possibility of several imports rounds for the same firm, these aspects of robot technology are a source of heterogeneity in treatment, which in turn can be a source of heterogeneity in treatment *effect*. I will discuss this issue further in Chapter 5. Thirdly, I assume that all industrial robots used in manufacturing firms are in fact imported, not produced in Norway. This assumption is quite plausible considering the early period studied. Most robots in the world are produced in Japan, South Korea, The US, Germany and increasingly in China (Dixon et al., 2021; Humlum, 2019). Barth et al (2020), using the same data source and same time period as this paper, shows that the import data matches the trend estimates done by IFR for Norwegian robots access, suggesting import data to be reliable source in this time period. Fourthly, I assume that the robot importers are in fact the users of the robots. If this is not the case, there might be a contaminating problem because non-adopters could be in the treatment group and robot adopters could be in the control group. Although some robot distributes have popped up in recent years, this problem is likely to be small in the time period I consider. The similarity of the data at hand and the IFR data also suggest that this potential issue is not a major concern. Moreover, important note is that, although the data set begins in 2003, the data captures whether firms had access to robot technology before 2003, avoiding a potential problem of including robot firms in the control group.[7] Lastly, throughout this thesis, when referring to ”robot adopters”, ”robot firms” or ”treated firms”, I mean firm that has gotten treatment (access to robot technology at some point between 2003 and 2016). When the opposite is true, firms are referred to as ”non-adopters”, ”non-robot firms” or ”control group”.

A last consideration should be given to treatment timing. There are three potential points in time to choose as treatment timing. Firstly, when the firm decides to buy a robot. Secondly, when it is utilized for production. And thirdly, when it was imported (crosses the border in customs). As the latter is one only data point we have access to, this is what we use as treatment time in this thesis. This has also been the most normal approach in similar papers (Acemoglu et al., 2020; Aghion et al., 2019; Barth et al., 2020; Caroli & Van Reenen, 2001; Dixit & Stiglitz, 1977; Humlum, 2019). However, there are possible implications of this that needs to be considered. The effect on the wage share seems likely both to be affected at the time of deciding to buy a robot and when robots are utilized for production. The former because firms might *expect* a need for certain worker groups as it anticipate robots to come in the near future. The latter because firms might be surprised by the effects robots have on worker productivity and

[7] 69 firms had imported robots before 2003.

Table 1: Summary statistics

	Descriptive statistics	Regression
Total number of firms in sample	9372	6428
Total number of robot adopting firms in sample	300	274
Share robot adopting firms	3.2 %	4.3 %

Table notes: The ”Descriptive statistics” sample include data on all manufacturing firms. As control variables (from Kapitaldatabasen) lack for manufacturing firms in a non-manufacturing corporate groups, these firms are excluded from the ”Regression” sample, which is used in the event studies and the regression analysis.

thus respond by changing their demand for workers after robots have begun operating. It is not certain which of these periods are actually appropriate to use as treatment timing, and it could even be that this is heterogeneous across labour groups. I will discuss possible implications of this in the next chapter.

The second data set comes from Kapitaldatabasen (Capital database) also from SSB and contains firm-level data from 2002 to 2016 for all firms that are a part of a manufacturing corporate group, including non-manufacturing firms. The data set contains several important firm characteristics like employment, sector, capital endowments, operating revenue, value-added, etc. Some of these variables will work as important controls (further discussed in Chapter 5). I will use the variables that are highlighted by the literature, in particular annual capital stock, annual value-added, employment and operating revenues (Acemoglu & Autor, 2011a; Barth et al., 2020; Humlum, 2019). As manufacturing firms in a non-manufacturing corporate group are not in the data, these firms will lack control variables and are not included in the regression analysis. However, these firms are included in the descriptive statistics so that we get a best possible understanding of the contexts in the manufacturing sector generally. The implications of this is likely to be small as we go from 280 to 300 robot firms, and these additional 20 firms are likely to be small (see Table 1).

4.2 Descriptive statistics

This section presents and discusses some descriptive statistics and relates them to the earlier literature. After a brief presentation of summary statistics, I present the developments in robot imports and how lumpy they are. Thereafter, characteristics of robot firms are outlined, before looking at the developments in the outcome variable, the wage share.

The sample used for descriptive statistics and the regressions are different as we lack control variables for manufacturing firms in a non-manufacturing corporate group. As discussed in section 4.1 around 3000 firms drops out due to lack of such data, however only 20 of these are robot firms. As evident from Table 1, of the plentiful manufacturing firms in Norway, only few of them are actually importing robots. In both samples, the number of robot adopters are between 3% and 4.5%.

Figure 1 displays how the value of imported robots from these firms are distributed across time in the Norwegian manufacturing industry. The total and the average value (for importing firms) of the imported robots each year are displayed on the left and right axis respectively. Throughout the period, the lines roughly move together and presents a clear trend toward increased import of industrial robots. In 2016, the total value of annual imports was just above 110 million NOK, corresponding to around a 6 fold increase from 19 million in 2003. The average value of robot imports started at 1 million NOK in 2003 and almost tripled to 2.7 million in 2016. This implies that the number of imports of an industrial robots doubled in the same time period; in 2003 19 manufacturing firms imported robots, in 2016, 41 did. In 2010, both the total and the average value drops to around half of their value in 2009. This is likely due to the economic context following the financial crisis in 2008. One might have expected this fall to come in 2009. An explanation for the lagged response might be that it takes some time from ordering the robot, possibly with many firm-specific details, through production, until it crosses the border and is registered as imports in our data. The accumulated (non-depreciated) total value of robots in manufacturing firms increased by almost 400% from 174 million NOK in 2003 to around 680 million in 2016 (see Appendix C for figure of development of accumulated value).

Figure 1: Value of imported robots in the manufacturing sector (2003 - 2016)

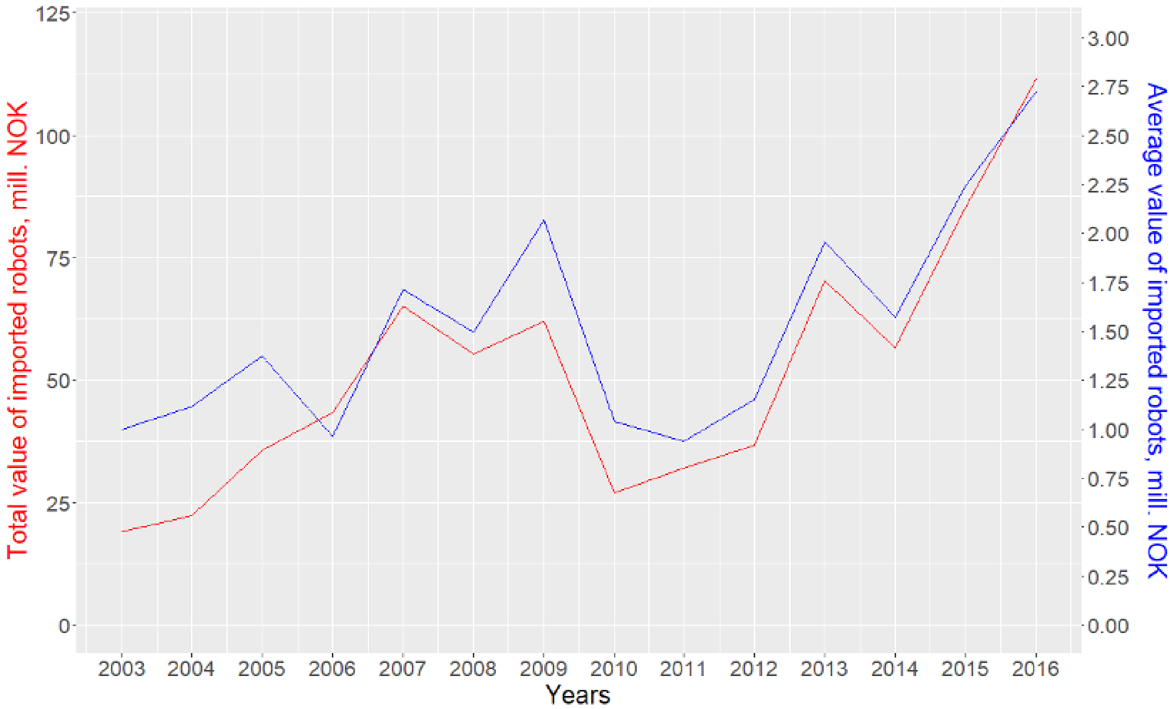

Figure notes: Total import (red line) and average value per import of industrial robots (blue line) between 2003 and 2016, both in mill. NOK. The average is calculated from robot firms only. Source: SSB.

Figure 2 illustrates the lumpiness of robot imports. Previous literature have highlighted the lumpy nature of investments similar to robots (Cooper et al., 1999; Doms & Dunne, 1998). In the robot literature, Humlum (2019) find that 70.6 % of danish manufacturing firms between 1995 and 2015 invested in robots only in a single year and Bessen et al. (2019) identifies automation based on the assumption of lumpy investments. In our case, almost half of robot adopters only imported robots once in the time period. This is significantly less than in the Danish case which is surprising considering our shorter time frame. There are some outliers importing almost all years in the time period, which are likely to be large firms. These results are interesting because they tell us that robot purchases are typically a substantial investments for firms. If the bar is high for buying a robot, investments might demand substantial planning prior to the decision to buy and the expected returns are likely to be high. Moreover, the information about lumpiness is importation information for methodological considerations. As I will return to later, heterogeneity in treatment effects can pose problems and be present for various reasons, and difference in import frequency might be one of them.

Figure 2: Lumpiness of robot imports (2003 - 2016)

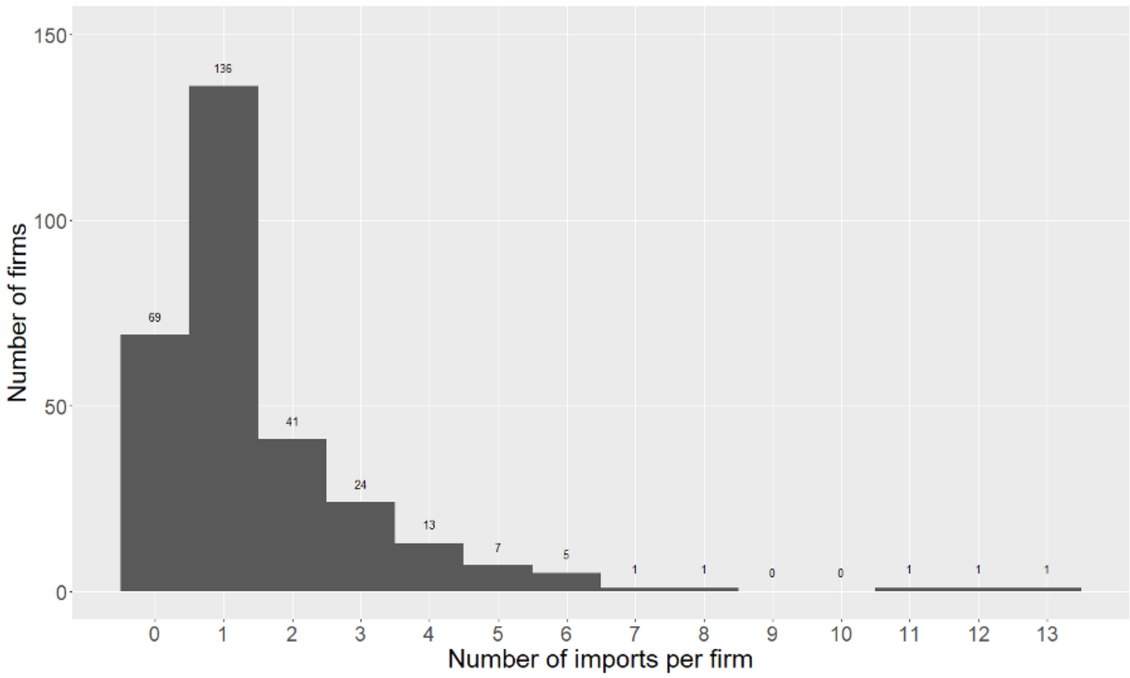

Figure notes: This figure displays the lumpiness of robot imports, i.e how many times a firm imported a robot between 2003 and 2016. 0 imports means that the firm imported a robot before 2003. Total number of firms are 300. 45% of robot firms imported robots only once, 14% twice and 8% three times. Source: SSB.

In addition to the aggregate trends of robot imports, it is useful to better understand what characterises the robot firms. Figure 3 displays the share of the total value of imported robots for all manufacturing sectors (using sector categories from the Standard Industrial Classification 2002 (SIC 2002)).[8] Some sectors stand out, but the introduction of robots is present in the majority of the sectors. Production of automotive (parts) makes up around 25% of the import, followed by the production of other means of transport and production of radio, television and communication equipment at around 15% each. 13

[8]This classification was used by SSB from 2002 to 2009. I use this classification rather than Standard Industrial Classification 2007 (SIC 2007) as this was provided from SSB and several issues occurred when attempting to convert this data to SIC 2007.

of the remaining 19 other sectors did also import robots, with varying shares under 10% of the total import. The results correspond roughly to the findings of Acemoglu and Restrepo (2020b), who studies the development of robot technology in the manufacturing sector in the US (2004-2007) and EU countries (1993-2007). Like in our results automotive stands out, followed by electronics, chemicals, metals, machines and plastic. The dominant position of transport-related manufacturing is also known in other development countries (Dauth et al., 2021; IFR, 2016). The two most dominating sectors in Norway might be linked to manufacturing clusters like Raufoss, known for, amongst other things, employing robots to produce specific parts for the automotive industry (Stensvold, 2016).

Figure 3: Share of total value of imported robots, by manufacturing sector (2003 - 2016)

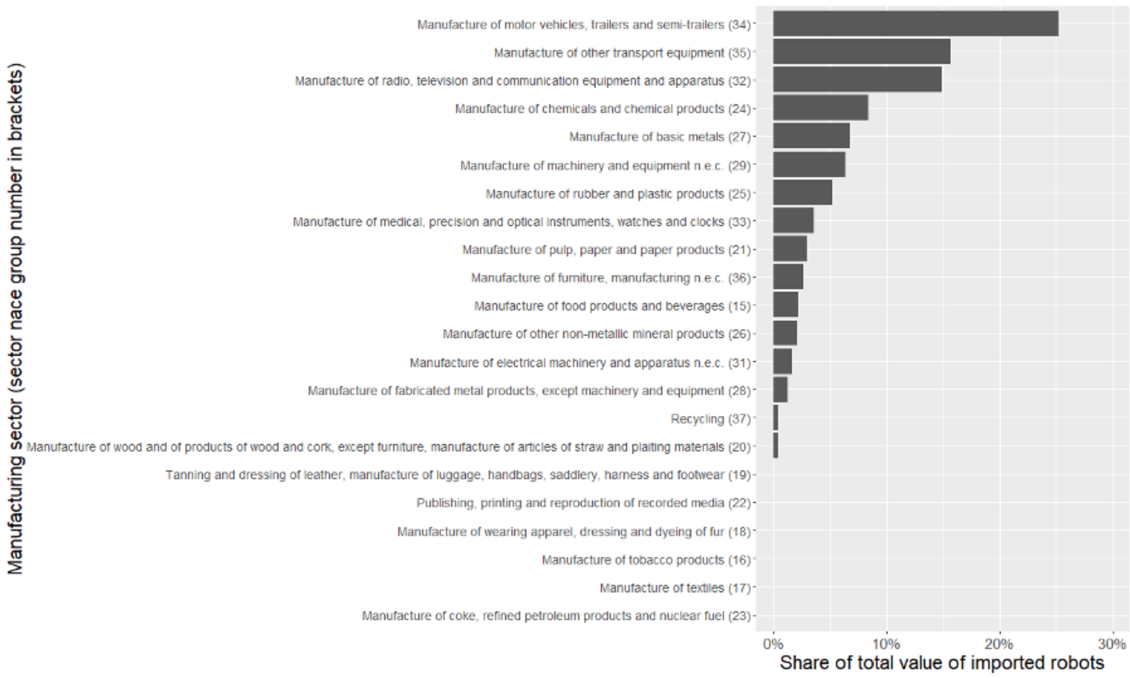

Figure notes: Share of total value of imported robots from 2003 to 2016 by manufacturing sectors (Standard Industrial Classification 2002, SSB). The list of manufacturing sectors can also be found in Appendix C. Source: SSB.

Figure 4 confirms previously discussed literature on the size of firms robot adopters. All firms in the data are arranged by their size (proxied by their average annual operating revenues) and placed into their respective percentiles. The bar shows the share of firms in each percentile group that got a robot in the time period. There is a clear trend: of the larger firms, there is a substantially higher share that imported a robot. Conclusions on the causal direction must be done with caution. It is not clear whether the firms were already large before adopting a robot, or if adopting a robot expanded the firm after it was employed. The same finding has been done in other recent papers and might be related to the so-called superstar firm phenomenon; globalisation and technological change give more productive firms a rising market share (superstar firms) with higher mark-ups and a lower labour share of value-added (D. Autor et al., 2020a). Acemoglu et al. (2020) find the same pattern for manufacturing firms in France between 2010 and 2015. However, the trend is clearer in Norway. While “Top 0.1”, “Top 1” and “1

- 10" in France had an import share of 18%, 13% and 7% respectively, the corresponding values for Norway is 75%, 35% and 18%, illustrating the higher rate of robot adoption in large Norwegian firms. Humlum (2019) argues that the size of the firms is the key feature that separates robot adopters from non-adopters. Koch et al. (2021) find that bigger firms are more likely to introduce robots in production when looking at the Spanish economy.

Figure 4: Share of firms that imported a robot, by percentile (2003 - 2016)

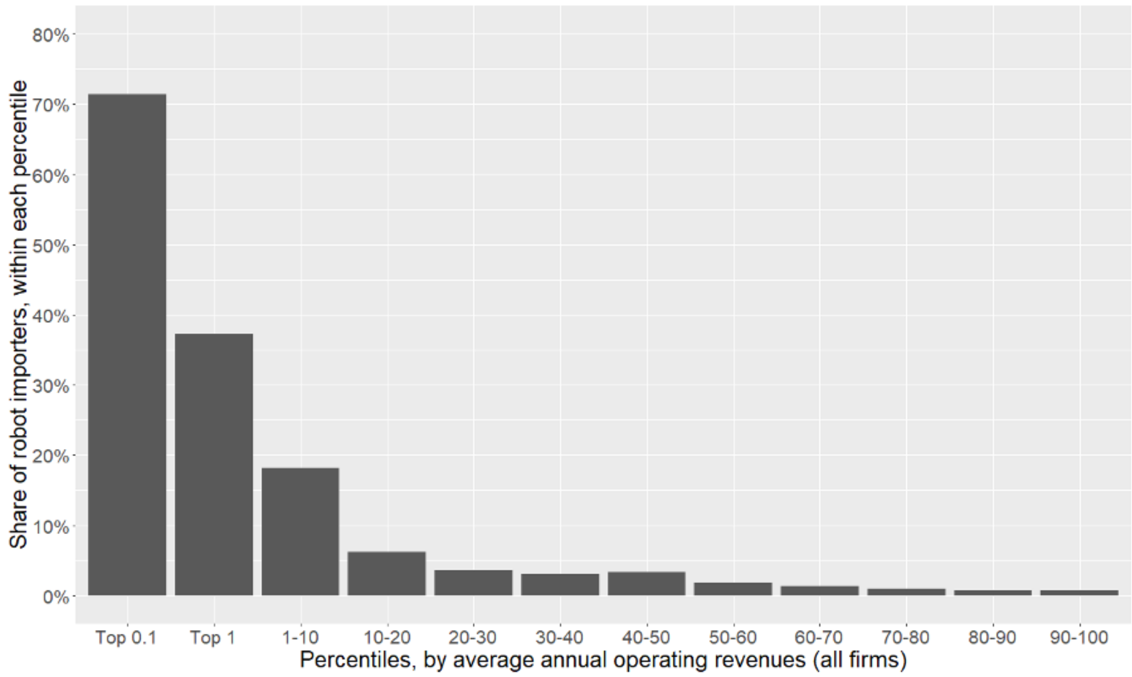

Figure notes: All firms (including non-robot firms) are arranged by their size (proxied by their average annual operating revenues) and placed into their respective percentiles. Top 0.1 are excluded from Top 1, which again are excluded from 1-10, so each firm is only in one percentile group. The bar shows the share of firms in each percentile group that got a robot between 2003 and 2016. Source: SSB.

Lastly, Figure 5, 6, 7 and Table 2 helps us understand the development in the outcome variable in question. Figure 5 displays the different distributions of the total wage costs for different labour groups (i.e the wage shares) between robot firms and non-robot firms. Figure 6 displays the same, only the employment share instead of the wage share. Although this latter variable is not used in our regression, it is useful to better understand if it is the changes in employment or hourly wages that drive the observed changes in the wage share. Both Figure 5 and 6 show a similar trend. Firms that got a robot use a higher share of their wage costs and have a higher employment share of HE workers, and lower share of CS and US workers. This is true for all years, and this pattern is increasing with time for both groups. This is similar to what is found in other countries, e.g. Spain (Koch et al., 2021).

Figure 7 sums up Figure 5 and 6 by displaying the long changes (from 2003 to 2016) in the wage and employment shares for all worker groups in both robot and non-robot firms. Based on Figure 7, Table

Figure 5: Share of total wage costs for different worker groups (2003 - 2016)

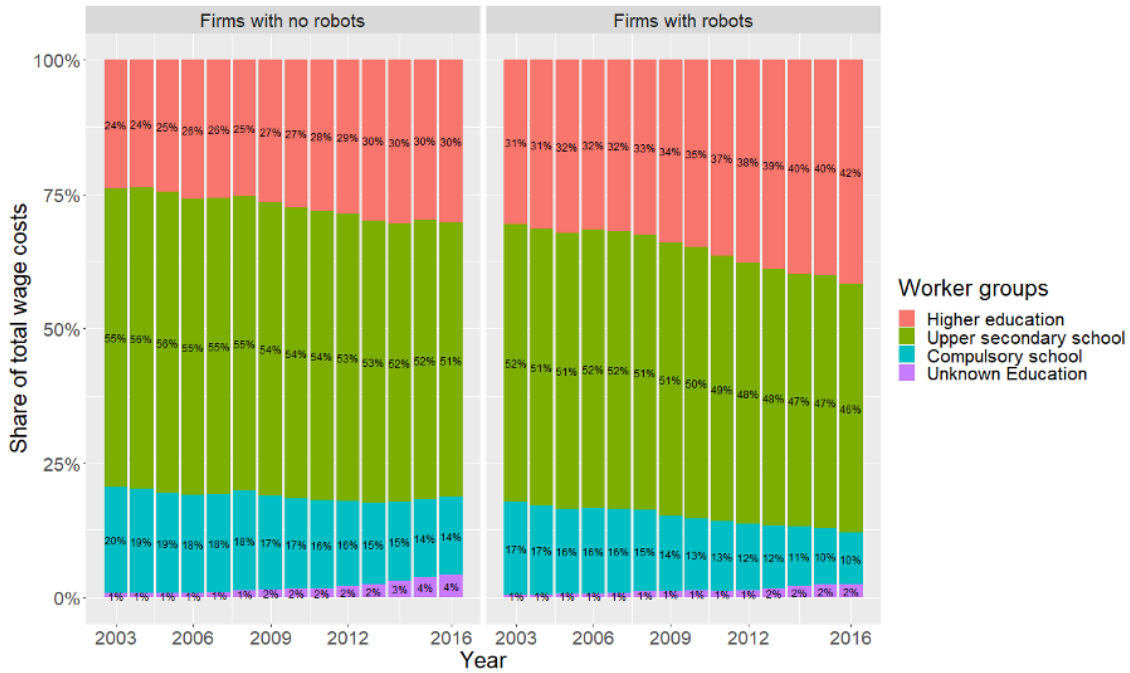

Figure notes: Firms are split into two categories. "Firms with no robots" consists of 9072 firms that did not import robots between 2003 and 2016. The 300 firms that did are in the "Firms with robots" category. Each bar show the annual wage share for each worker group. Source: SSB.

2 displays the long DiD in the wage and employment share between robot firms and non-robot firms for all worker groups.[9] This is a naive estimate of the effect of robots, as it implies no selection bias, i.e. that the robot firms and non-robot firms are on average similar and thus would be affected similarly from treatment. It shows that robots have the same effect (in terms of sign) on both the wage share and employment share, suggesting that much of the observed changes in the wage share might also drive the changes by employment. HE workers are likely to increase both their employment and wage share, and CS and US workers are likely to decrease their wage share. Again, it is important to note that this estimate is likely to suffer from selection bias.

In sum, there is a clear increasing trend in the use of robots by Norwegian manufacturing firms, both in numbers and value, although the latter constitutes the most substantial increase. Robot investment in Norway are lumpy and the most dominant robot importing manufacturing sector is related to the production of automotive (parts). There is also a clear tendency that robot adopters in the Norwegian manufacturing industry are larger firms. Robot firms have more HE workers and less CS and US workers to start with, compared to non robot firms. The former worker group also increase more for robot firms and the latter two decreases more for robot firms.

[9] The long DiD for share j for each worker group g is given by $(S^R_{2016,j} - S^R_{2003,j}) - (S^C_{2016,j} - S^C_{2003,j})$ where $S^R_{2016,j}$ is wage or employment share for worker group g in robots firms in 2016 and $S^C_{2003,j}$ is wage or employment share for worker group g in non-robots firms (control group) in 2003.

Figure 6: Share of total employment for different worker groups (2003 - 2016)

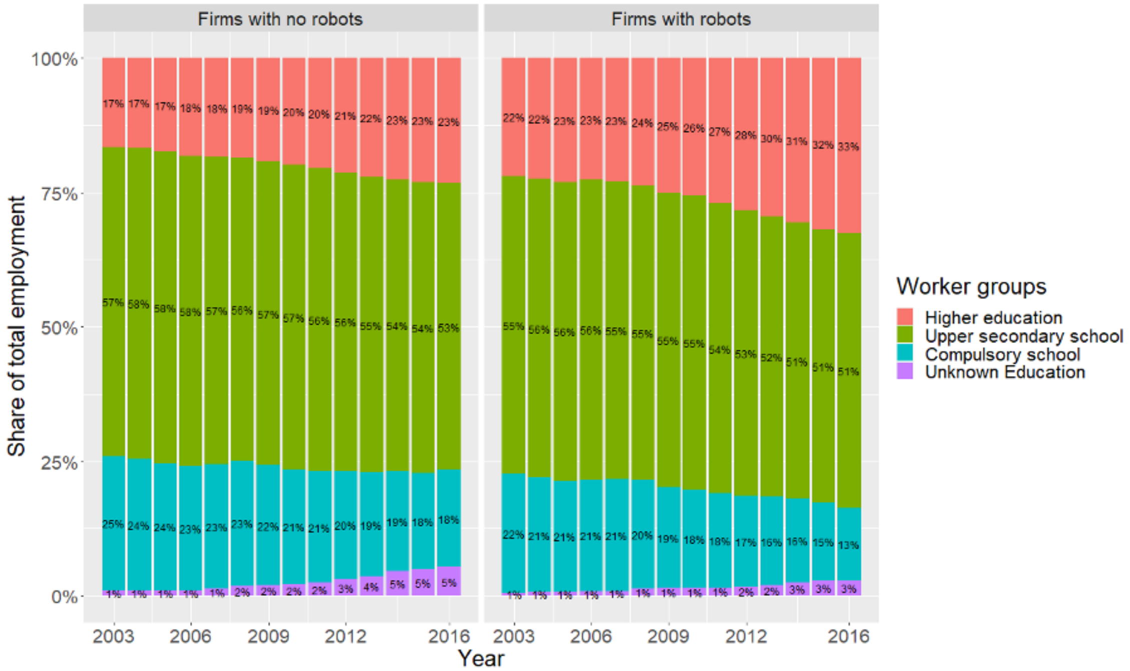

Figure notes: Firms are split into two categories. "Firms with no robots" consists of 9072 firms that did not import robots between 2003 and 2016. The 300 firms that did are in the "Firms with robots" category. Each bar show the annual employment share for each worker group. Source: SSB.

Figure 7: Changes in cost and employment shares (2003 - 2016)

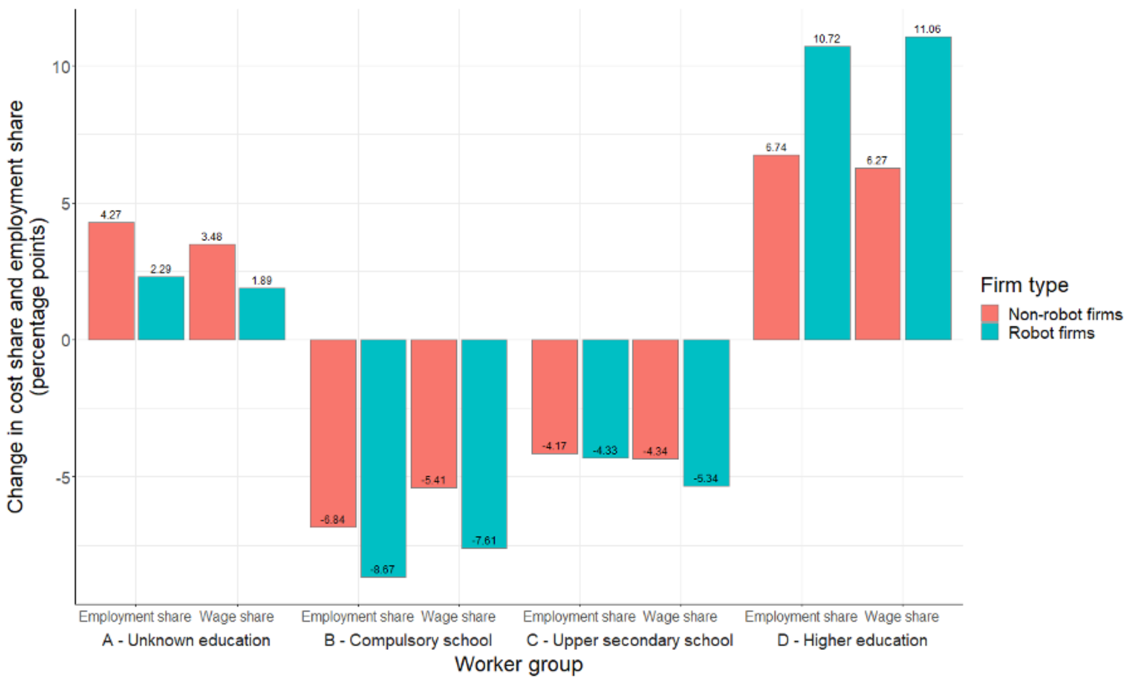

Figure notes: Changes in the wage share and the employment share (percentage points) for all worker groups in both robot firms and non-robot firms between 2003 and 2016. In other words, $(S^g c_{2016} - S^g c_{2003})$, where g is the worker group and c is the two categories ("Firms with robot") and ("Firms with no robot"). Based on numbers in Figure 5 and Figure 6. Source: SSB.

Table 2: Long Difference-in-Differences in wage and employment shares (2003 - 2016)

Worker group (g) / Share (j)	Wage share	Employment share
Unknown Education (UE) workers	$(1.89) - (3.48) = -1.59$	$(2.29) - (4.27) = -1.98$
Compulsory School (CS) workers	$(-7.61) - (-5.41) = -2.2$	$(-8.67) - (-6.84) = -1.83$
Upper Secondary School (US) workers	$(-5.34) - (-4.34) = -1$	$(-4.33) - (-4.17) = -0.16$
Higher Education (US) workers	$(11.06) - (6.27) = 4.79$	$(10.72) - (6.74) = 3.98$

Table notes: The long DiD for share j for each worker group g is given by $(S^R_{2016,j} - S^R_{2003,j}) - (S^C_{2016,j} - S^C_{2003,j})$ where $S^R_{2016,j}$ is wage or employment share for worker group g in robots firms in 2016 and $S^C_{2003,j}$ is wage or employment share for worker group g in non-robots firms (control group) in 2003. If the is not selection bias (which is highly unlikely), this number would be the causal effect of robots on the respective wage share.

5 Empirical Approach

This chapter outlines the empirical approach used to estimate the effects of robots on the wage shares. I will use a Difference-in-Differences (DiD) estimator appropriate for contexts with heterogeneous treatment timing, namely the two-way fixed effects (TWFE) regression, both a static version (for regression analysis) and dynamic (event studies). When performing such estimates, two things are crucial. Firstly, it is important to understand how our method can provide estimates of *causal* effects. Secondly, it is crucial to understand what might hinder us in this endeavour. To better understand this, I start this chapter by outlining what is known as the *potential outcome framework*. Based on this, I discuss how the DiD approach, both the basic set up and the TWFE model, can be used to estimate the causal estimates - and what might cause problems. I conclude the chapter by presenting the models I will estimate. These discussions are based on the overviews of these methods in Cunningham (2021) and (Roth et al., 2022).

5.1 Causal mechanisms and potential outcomes

The aim of this analysis is to obtain an internally valid estimate of the *causal* effect of access to robot technology on wage shares of different worker groups. Although there is still a philosophical discussion regarding the definition of causality, social scientist have adopted what is known as the *counterfactual approach to causality* to investigate such relations. To understand this view of causality, we can imagine the same manufacturing firm in two identical scenarios. The only difference is that in one of the scenarios the firm adopts a robot. According to the counterfactual approach, the causal effect of the robot on some variable of interest, in our case the wage share, is the difference in the wage share between these two scenarios. The fundamental problem of causal inference is that it is not possible that both the scenarios can be observed, as both scenarios cannot be realized (Holland, 1986). Fortunately, methods have been developed to overcome this issue so that we can obtain unbiased *estimates* of causal effects. In particular, economist have, based on the counterfactual notion of causality, developed what is known as the *potential outcome framework*. In this section I will outline this framework, as this will be key to understand the methods used in this thesis.

The potential outcome framework bridges this more philosophical understanding of causality to econometric methods, allowing us to model causal mechanisms that are possible to estimate (Cameron & Trivedi, 2005). For the purpose of explanation, let us again use the context in question, the effect of robot technology on the wage shares for different worker groups. Before treatment (access to robot technology), the wage share for worker group g in firm i have two hypothetical outcomes, either Y^0_{gi} if the firm does not have access to robot technology, or Y^1_{gi} if it does. However, only one of these outcomes will actually occur. This is the realized outcome and is denoted as Y_{gi}. The connection between the potential outcomes Y^0_{gi} and Y^1_{gi} and the actual outcome is given by the *switching equation*

$$S_{gi} = R_i S^1_{gi} + (1 - R_i) S^0_{gi} \tag{8}$$

where $R_i = 1$ and thus $S_{gi} = S^1_{gi}$ if the firm gets access to robot technology and $R_i = 0$ and thus $S_{gi} = S^0_{gi}$

if it does not. It also follows that the causal effect of access to robot technology on the wage share for group g in firm i is the difference in these outcomes, given by $\delta_i = S^1_{gi} - S^0_{gi}$. Note that this causal effect is entity specific for this particular firm. However, in most cases, certainly in this thesis, we are concerned about the *average causal effect* across all firms in a given population (Norwegian manufacturing industry from 2003 to 2016). The corresponding population notion is called the *Average Treatment Effect (ATE).* We take the leap from individual effects to the ATE by looking at expected outcomes at the population level. The ATE is given by

$$\begin{aligned} ATE &= E[\delta_i] \\ &= E[S^1_{gi} - S^0_{gi}] \\ &= E[S^1_{gi}] - E[S^0_{gi}] \end{aligned} \tag{9}$$

In addition to this concept, two other definitions that are similar to the ATE are useful. To understand this, recall that for most cases, as in the Norwegian manufacturing industry, some entities (firms) get treatment (access to robot technology), while some does not. The first notion is concerned with the former of these groups, and is called the *Average Treatment Effect for the Treated group (ATT)* and is defined as

$$\begin{aligned} ATT &= E[\delta_i | R_i = 1] \\ &= E[S^1_{gi} - S^0_{gi} | R_i = 1] \\ &= E[S^1_{gi} | R_i = 1] - E[S^0_{gi} | R_i = 1] \end{aligned} \tag{10}$$

This concept captures the true causal effect of access to robot technology on the wage share for the worker group g for the firms that actually got the treatment. As with the ATE, ATT is unknowable because we cannot observe both S^1_{gi} and S^0_{gi} for the treated group, only the former is realized and thus we have $D_i = 1$. Given some assumptions, which I will explain shortly, it could be the case that ATT = ATE. The last concept we need to define is the opposite of ATT, namely *Average Treatment Effect for the Untreated group (ATU)* and is defined as

$$\begin{aligned} ATU &= E[\delta_i | R_i = 0] \\ &= E[S^1_{gi} - S^0_{gi} | R_i = 0] \\ &= E[S^1_{gi} | R_i = 0] - E[S^0_{gi} | R_i = 0] \end{aligned} \tag{11}$$

Note that all these casual effects are hypothetical as they demand that we know the outcome of several counterfactual scenarios. In fact, we only know $E[S^1_{gi} | D_i = 1]$ (the expected value of the wage share for the firms that got robots) and $E[S^0_{gi} | D_i = 0]$ (the expected value of the wage share for the firms that did not get robots), which can be found by estimation. The difference of these estimations is called the *Simple Difference in Outcomes (SDO)* and can be written as

$$SDO = E[S^1_{gi} | R_i = 1] - E[S^0_{gi} | R_i = 0] = \frac{1}{N_T} \sum_{i=1}^{n} (s_i | r_i = 1) - \frac{1}{N_C} \sum_{i=1}^{n} (s_i | r_i = 0) \tag{12}$$

where the two latter expressions are estimations of the two realised outcomes, $E[S^1_{gi}|R_i = 1]$ and $E[S^0_{gi}|R_i = 0]$. The problem is that we can estimate SDO, but we aim to find the ATE of robots. Are these two the same? The equation below displays the connection between SDO and ATE (Cunningham, 2021)

$$\underbrace{\frac{1}{N_T}\sum_{i=1}^{n}(s_i|d_i=1)-\frac{1}{N_C}\sum_{i=1}^{n}(s_i|d_i=0)}_{Simple\ Difference\ in\ Outcomes\ (SDO)} = \underbrace{E[S^1]-E[S^0]}_{Average\ Treatment\ Effect\ (ATE)} +$$

$$\underbrace{E[S^0|D_i=1]-E[S^0|D_i=0]}_{Selection\ Bias\ (SB)} + \underbrace{(1-\pi)(ATT-ATU)}_{Heterogeneous\ Treatment\ Effect\ Bias\ (HTEB)}$$

The above equation show that we can estimate the true causal effect of adopting robots at the population level if we manage to avoid *Selection Bias (SB)* and *Heterogeneous Treatment Effect Bias (HTEB)*. As is clear from the equation above, SB occurs when the potential outcomes of the treatment group and the control group are different so that $ATT \neq ATU$, i.e if the expected value of the wage share for robot adopters had not been the same as the expected value of the wage share for non-adopters, had they not adopted robots. This difference can occur if it is not random which firms buy robots, so that the wage share would have differed independent of the the treatment. This is also known as *selection bias*.

The second term, HTEB, describes a bias that occur if the treatment *effect* is different, on average, for entities in the control group and the treatment group, i.e. that $ATT \neq ATU$. In our case this would mean that, if the firms that did not actually got access to robot technology nevertheless adopted robots, the average effect of this would be different from the average effect it has on the firms that actually got robots. In other words, the effect of access to robot technology is heterogeneous. In HTEB, π is the share of the sample that got treatment (the share of firms that got robots).

If we can avoid SB and HTEB, then our *estimate* of access to robot technology on the wage share is an internally valid estimate of the true causal effect in our population. This implies that

- The estimate is **unbiased**. An estimate is unbiased if the expected value of our estimate is equal to the true population parameter: $SDO = ATE$
- The estimate is **consistent**. An estimate is consistent if it converges in probability to the true value of the parameter as the sample size tends to infinity: $SDO \xrightarrow{p} ATE$
- **Statistical inference of the estimate is valid**. This is the case when the hypothesis tests have the desired size and confidence intervals should have the desired coverage probability.

5.2 The Difference-in-Differences set up

Based on the potential outcome framework, this section first outlines the basic DiD set up with two time periods and homogeneous treatment timing. After some discussions that relates to the nature of our data, a modified version of the DiD method, the TWFE model, is presented.

5.2.1 The basic DiD set up

In the basic DiD framework, known as the 2x2 DiD there is one group, R, where all the entities (firms) in the group gets, at the exact same time, some homogeneous treatment (access to robot technology) and some control group, C, that does not (Goodman-Bacon, 2021). We also assume two time periods, one period before the firms were given robots and one after. As it is possible to observe the average of some outcome ($\bar{S}$, the wage share) for each group, we can estimate the ATT for group R by simply taking the difference between the change in the outcome for the treated group and the change in outcome for the control group

$$\hat{\delta}_{RC}^{2x2} = \left(\bar{s}_R^{post(R)} - \bar{s}_R^{pre(R)}\right) - \left(\bar{s}_C^{post(R)} - \bar{s}_C^{pre(R)}\right) \tag{13}$$

However, as noted in the previous section, there might be some issues that bias our results so that this estimate do not reflect the ATT, which we aim to estimate. If we apply the conditional expectations from the potential framework we get that

$$\hat{\delta}_{RC}^{2x2} = \underbrace{E[S_R^1|post] - E[S_R^0|post]}_{ATT} + \underbrace{\left[E[S_R^0|post] - E[S_R^0|pre]\right] - \left[E[S_C^0|post] - E[S_C^0|pre]\right]}_{Non-parallel\ trend\ bias} \tag{14}$$

The first term of this equation, the ATT, shows the difference in outcome between the scenario with robots (which is realized) and the scenario without (counterfactual) for the group of firms that got access to robot technology. This is what we aim to estimate. The second term is concerned with the parallel-trend assumption (the DiD version of the selection bias term in Equation 12) and consists of two parts (Cunningham, 2021). The first term is the change in outcome over time for the group of firms that actually got robots, had it not received treatment. The second term is the same change, only for the control group that did not actually received robots. In total, the Non-parallel trend bias term show the difference between these two terms, and equals 0 if they are the same. If this is true, it means that the outcome variable, the wage share, had changed in the exact same way for both the robot firms and the control group, had they both not received treatment. In other words, they would have had *parallel trends* after treatment. If this is the case we get that $\hat{\delta}_{RC}^{2x2} = ATT$. This highlights the crucial important role of the parallel trend assumption as the main identifying assumption in the DiD framework (Cunningham, 2021).

This framework highlights the potential power of the DiD approach in identifying causal effects. However, as this outlined framework assume only one treatment time and two treatment group, it is not fit for our case. Instead, we must develop the framework one step further with the TWFE approach.

5.2.2 Two-way fixed effects

When treatment is not given at the same time, the TWFE model has been the most commonly used DiD estimator (Goodman-Bacon, 2021). In our basic DiD model, it is quite straight forward to estimate the ATT because there are only two differences to compare, namely the change in the treatment group that got robots and the change in the control group that did not. However, this changes when the treatment happens at different times, which is the case in our data (firms adopt robots at various different times).

The idea of the TWFE model is to calculate an aggregate coefficient based on all potential 2x2 DiD coefficients we can obtain from the data. These coefficients are aggregated by weights that are both based on group size and variance. To simplify the idea, let's image that one group of firms (k) got access to robot technology at time t_k, another group of firms (l) that got access to robot technology at a later time t_l and another group of firms that never got treatment. We call the time period before group k got treatment t_k as "pre", the period between t_k and t_l as "mid" and the period after t_l as "post". Note that all firms are now put into *groups* depending on he time of treatment. In this case, there would be three general estimates of the ATT.

$$\hat{\delta}_{kU}^{2x2} = \left(\bar{y}_k^{post(k)} - \bar{y}_k^{pre(k)}\right) - \left(\bar{y}_U^{post(k)} - \bar{y}_U^{pre(k)}\right) \tag{15}$$

$$\hat{\delta}_{kl}^{2x2} = \left(\bar{y}_k^{mid(k,l)} - \bar{y}_k^{pre(k)}\right) - \left(\bar{y}_l^{mid(k,l)} - \bar{y}_l^{pre(k)}\right) \tag{16}$$

$$\hat{\delta}_{lk}^{2x2} = \left(\bar{y}_l^{post(l)} - \bar{y}_l^{mid(k,l)}\right) - \left(\bar{y}_k^{post(l)} - \bar{y}_k^{mid(k,l)}\right) \tag{17}$$

Equation 15 show the ATT estimate of any treated group compared to the untreated group (k or l), Equation 16 is the early treated group compared to the later treatment group and Equation 17 is the late treatment group compared to the already treated group. The TWFE ATT estimate is just an aggregated and weighted average across all these possible 2X2 ATT estimates, and is formally defined as

$$\hat{\delta}^{DD} = \sum_{k \neq U} s_{kU} \hat{\delta}_{kU}^{2x2} + \sum_{k \neq U} \sum_{l > k} s_{kl} \left[\mu_{kl} \hat{\delta}_{kl}^{2x2,k} + (1 - \mu_{kl}) \hat{\delta}_{kl}^{2x2,l}\right] \tag{18}$$

where the weights s_{kU}, s_{kl} and μ_{kl} are calculated from the within-group variance. To get a better understanding of the intuition, it is useful to express this equation in terms of the potential outcome framework as outlined before. We obtain this by expressing all the individual ATT estimates in the equations above with corresponding potential outcome framework notation and then substituting all these back into Equation 18 (Cunningham, 2021). When this is done, we get the following equation which gives us the TWFE coefficient.

$$p\, lim \hat{\delta}_{n \to \infty}^{DD} = VW\ ATT + VW\ CT - \Delta ATT \tag{19}$$

The first two terms on the right hand side of this equation are just the TWFE version of the concepts described in the basic DiD case. The first term, $VW\ ATT$ is what we want to estimate, namely the (group) variance weighted ATT. This is just the collection of all possible ATTs weighted by the group variance. The second term, $VW\ CT$, is the variance weighted common trend assumption. This is just a collection of Non-Parallel trend assumptions written out earlier, weighted in the same way as the $VW\ ATT$. Although the intuitive idea might be that this assumption is stricter, as all the non-parallel trend assumptions are included, it is in fact a *weaker* assumption that the non-parallel trend assumption.

This is because the weights can make up for trends that are not parallel (Cunningham, 2021). However, we will focus on parallel trends when assessing our data with event studies in the next section. As this is a stricter criteria, the $VW\ CT$ will hold if the non-parallel trends bias zeros out. Lastly, the ΔATT is the ATT heterogeneity with time bias, and this is concept is slightly new in the TWFE case. In potential outcome form we can write it as

$$\Delta ATT = \sum_{k \neq U} \sum_{l > k} (1 - \mu_{kl}) \left[ATT_k(Post(l) - ATT_k(|)) \right] \tag{20}$$

This equation addresses the issue of heterogeneous treatment effects, a concern we already touched upon in Chapter 4 and in the potential outcome framework. An important thing to notice, though, is that ΔATT is not concerned with heterogeneous treatment effect across groups, as the $VW\ ATT$ will average out these differences by the weights on share and treatment variance. However, if there is treatment heterogeneity *within a group* of firms as this will bias the control group and thus the results. As shown in Roth et al. (2022), this can in fact lead to ”forbidden comparisons” that could bias our results. In our data, this can be a problem. Firstly, as discussed in Chapter 4, there are several ways in which the treatment might be heterogeneous, including difference in quantity and quality both across firms and across time, as well as several rounds of robot imports (see Figure 2). Moreover, we know that the majority of firms get robots more than once from Figure 2 in Chapter 4. All these sources of heterogeneous treatment, in addition to different response to the exact same treatment, might lead to heterogeneous treatment *effects* and thus bias our results. Although recent literature have highlighted this problem, there is yet no consensus on how to correct for it (Cunningham, 2021). Our model will not try to resolve this issue, so caution should be taken when we interpret the results. I will return to this in the Chapter 6, Results.

Lastly, in order to provide unbiased results, the DiD framework relies on what is commonly called the staggered no-anticipation assumption (Malani & Reif, 2015). This assumption holds if the outcome (the wage share for a worker group) in some time period v before treatment time v_0 does not depend on what time period they will treated in the future (Roth et al., 2022). Formally, this is given by

$$S_{gi}^{R} = S_{gi}^{C} \ for\ all\ v < v_0 \tag{21}$$

where S_{gi}^{R} and S_{gi}^{C} are the wage shares for the treated group and the control group respectively. Read more intuitively, it says that the wage share of the treatment group should not be different to the control group, before treatment is given. As discussed in Chapter 4, this might be an issue in our case. I have assumed that treatment time is when the robots crossed the boarder. However, it could be that case that the effect of robots on the wage share happens before this, when the firm decides to order the robot, or after, when it has been put to production. It could also be the case that this is different for the various worker groups. For example, it could be the case that firms need HE workers prior to robot purchase to find the right type and to understand how to leverage it most effectively. It could also be that firms need more US and CS workers after robots are employed as they see these are now more productive when they can work with robots. This uncertainty of knowledge about treatment timing could pose a problem for our estimates. I will also highlight this issue in Chapter 6, Results.

5.3 The final model

This sections builds on the theoretical framework described above to present the models used to estimate the causal effect of access to robot technology on the demand for labour. I will first present the event studies regression, also called the dynamic TWFE model, which will be used to investigate both the parallel trend assumption and give us an idea of the effect of robots. Thereafter, I will outline the model used in the regression analysis, namely the (static) TWFE model.

5.3.1 Event studies model

As discussed, violation of the parallel trend assumption leads to non-parallel trend bias in our estimates. The parallel trend assumption holds if our control group (firms that not had access to robot technology) and the treated group (firms that had access to robot technology) have similar developments for the variable of interests (the wage share) *both* before and after the treatment period. Recall, however, that we cannot observe the potential outcome of not being treated for our treatment group, i.e. $E[S^0_{gi}|R_i = 1]$. Thus, we cannot compare the post treatment trends and thus we cannot really test the parallel trend assumption. Although this is certainly a problem, we *can* investigate pre treatment trends for both the control group and the treatment group. And if these trends are parallel, it might be suggestive of the post treatment trends as well. Although this certainly is a bold leap, this is the closest we get and thus is the most frequently used method to test the parallel trend assumption.

To investigate the pre trends, economists have used a method called the dynamic TWFE model, or more commonly known as event studies. Formally, this is done by regressing the outcome variable (wage share) on entity and time fixed effects, as well as time dummies. When treatment is given to all entities homogeneously, time dummies corresponds to years relative to the year. When treatment time is heterogeneous, as is the case with our data, time dummies are constructed for all years relative to treatment time v and equals 1 if the dummy variable corresponds to the year of the observation, relative to treatment. For example, if a firm got a robot in 2007, the observation of the cost share in 2003 and 2013 will equal one for the dummies $v - 4$ and $v + 6$, respectively. The regression model is as follows:

$$S_{git} = \alpha_i + \phi_t + \sum_{\substack{v \neq -1 \\ -\underline{T} \leq v \leq \overline{T}}} 1\left[V_{it} = v\right]\beta_v + \epsilon_{it} \tag{22}$$

where S_{git} is the outcome variable (wage share for worker group g in firm i at time t), α_i and ϕ_t are entity and time fixed effects (FE) respectively and the summation term is the vector of dummy variables for all time dummies relative to treatment time ($V_{it} = v - v_0$). Notice that the time period before treatment is excluded and thus is used as reference point so that our estimates in the event studies correspond to the estimates in the following TWFE regression where the treatment dummy equal 1 in the year of treatment (v_0). ϵ_{git} is the error term. The results of the event studies will be outlined in the next chapter.

5.3.2 The TWFE model

The TWFE model outlined in the previous section can be estimated using OLS and it will be used to estimate the causal effect of access to robots on the wage share for the various worker groups (S_{git}). Two versions will be used, differing in how they control for factors that change across time. The first model to be estimated is not really a TWFE, but a FE model with an alternative time control variable. Instead of adding time FE, a trend variable, $TREND$ is included.[10] The reason for adding this trend variable is to learn about the general robot-independent trend of the respective wage share. The trend variable can be though of as a linear fit of the time FE and, as we will see, will control for much of the same unobservables as the time FE. As such, we can view this FE regression with trend variable as a sort of quasi-TWFE regression. This gives us the following model for each worker group g:

$$S_{it} = \beta_0 + \beta_1 ROBOT_{git} + \beta_2 \mathbf{X}_{git} + \beta_3 TREND_t + \alpha_i + \epsilon_{git} \tag{23}$$

In addition, I will use the TWFE version of this model, exchanging the trend variable with time FE, ϕ_t. This model is given by:

$$S_{it} = \beta_0 + \beta_1 ROBOT_{git} + \beta_2 \mathbf{X}_{git} + \phi_t + \alpha_i + \epsilon_{git} \tag{24}$$

In both models, the other terms are the same. The effect of access to robot technology on the wage share for each worker group is the coefficients in front the $ROBOT_{git}$ dummy variable, β_1. The vector $\mathbf{X}_{git}$ is a vector of time variant firm characteristics, including capital stock, value added, employment and operating revenues. These controls are added to cope with possible selection bias and are included based on what has been emphasised in the previous literature (Acemoglu et al., 2020; Aghion et al., 2019; Barth et al., 2020; Caroli & Van Reenen, 2001; Dixit & Stiglitz, 1977; Humlum, 2019). The model will be run several times with various combination of controls. Note that this regression has swapped the time FE term with a trend variable, $TREND_t$. The sign and value of this coefficient allow us to understand the general time trend of the wage share for the respective worker groups.

[10]The trend variable is calculated by subtracting the first year, 2003, from the observation year. This gives us $TREND = year - 2003$.

6 Results

This section presents the results of the econometric analysis and aims to establish a causal relationship between access to robot technology and the wage share for various worker groups. The first section outlines the results of event studies in order to examine the parallel trend assumption and to look for suggestive results on the effect of robots on wage shares. The second section goes through the results of the TWFE analysis. These results are re-examined through robustness tests, before, lastly, the chapter is concluded by summing up findings, relating them to the existing literature and theory, as well as highlighting potential threats to internal and external validity.

6.1 Event studies

The results from the event studies suggest that the parallel trend assumption holds and gives us some suggestive results on how robots might change the wage share for the various labour groups. As described in the previous chapter, the event studies are a visual examination of the coefficients obtained from the dynamic TWFE model (consult Appendix D for tables with all regression result). As outlined in Chapter 5, the coefficients are estimated with equation 22.

Two rounds of event studies are undertaken and displayed in the panels below. Each panel displays the separate event studies for the CS, US and HE worker groups, respectively. The points in the plots display the coefficient β_v for time dummy v, relative to the time period before treatment, as given in the Equation 22. The shaded area is the 95% confidence interval for each individual coefficient. The coefficient tells us if the outcome variable (wage share) at a particular time, relative to the time period before treatment, is different for the treatment group and the control group (non-robot firms and robot firms in the period v-1). If this is the case prior to the treatment, then the treatment group and control group had pre treatment parallel trends. A secondary priority should be given to the coefficients after the treatment time. If they are significantly different from zero, the treatment seems to yield a significant change in the outcome variable for treated firms relative to non-treated firms after treatment is given.

The first plot shows the event studies when all firms are included. For all three event studies, the coefficients further from the year of treatment on both sides are those with the largest coefficient and the biggest confidence intervals. This is likely because of few observations at these points in time. Most interesting for our purposes is to examine the parallel trend assumption, i.e. if the coefficients are zero or close to zero before treatment. In all figures, this appears to be the case, although the coefficients are significantly below zero for US and above for HE for the coefficients on the ends (consult Appendix D for regression table results with all time coefficients).[11] In the post-treatment period, it is worth noting

[11] I apologies to the reader for imposing the challenged of reading very small numbers. To make room for all confidence level, these axis scales had to be used. Consistency on scales in order to more easily compare results was prioritised over readability.

the sudden break from the pre trend coefficients for US workers in the first years after treatment (significantly positive at 10%). CS workers experience a similar negative change (first two years significantly negative at 10%). All other coefficients are not significant. However, it is worth noting that there is striking heterogeneity in the post treatment coefficients for all worker groups. Although observations are few and thus uncertainty is high at the ends of the time line, this could indicate that the effects of robots change over time and in different ways for the respective worker groups, with positive impacts on the long run HE and CS wage shares, and a long run negative impact for US workers.

Figure 8: Event studies, all firms sample

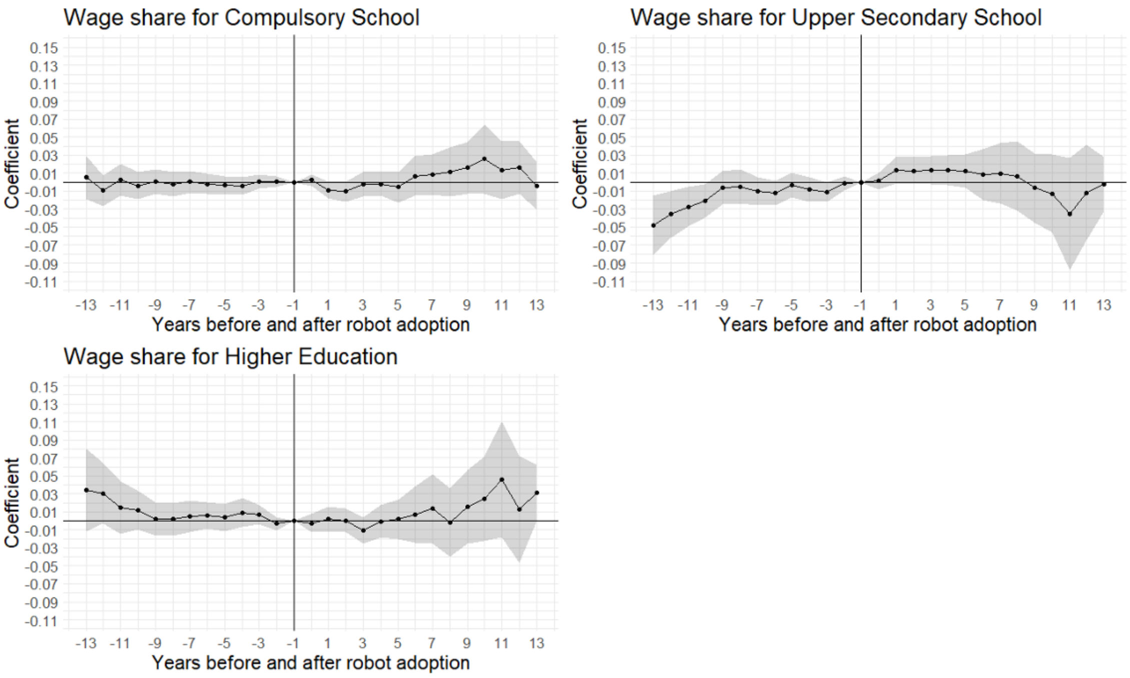

Figure notes: Event studies for all worker groups. Top left is for CS workers, top right is for US workers and bottom is for US workers. For each worker group, all years relative to the treatment time is regressed on the outcome variable following Equation 22. The year before treatment is used as reference point and thus the observations from this year are included in the control group.

To further examine the pre trend, some adjustments are made to the control group based on the observations done in the descriptive statistics in Chapter 4. The aim is to improve the similarity of the treated and control group, to potentially improve the coefficients prior to treatment. Based on the descriptive statistics, we know that there are some manufacturing sectors that do not import robots. We also know that robot adopters are typically large firms and that a very minor share of the smaller firms import robots. Thus, two adjustments are made. Firstly, all firms with industry nace code 16, 17, 18, 19, 20, 22, 23 and 37 are excluded.[12] In these sectors no firms or only very few firms had access to robot technology.

[12] 16 = Manufacture of tobacco products, 17 = Manufacture of textiles, 18 = Manufacture of wearing apparel, dressing and dyeing of fur, 19 = Tanning and dressing of leather, manufacture of luggage, handbags, saddlery, harness and footwear, 20 = Manufacture of wood and of products of wood and cork, except furniture, manufacture of articles of straw and plaiting materials, 22 = Publishing, printing and reproduction of recorded media, 23 = Manufacture of coke, refined petroleum

In addition, the smallest half of all firms (including non adopters), in terms of operating revenues, are excluded, as less than 2% of these firms imported robots. This new sample will henceforth be called the *adjusted* control group or *adjusted* firm sample. This group will also be used in the regression analysis. The results of these changes are displayed in Figure 9. Of the pre treatment coefficients, both US and HE are now closer to 0, but slightly further away for the CS group. Although the post treatment coefficient are fairly similar, the mentioned effects for the CS and US workers are not significant at 5%.

Figure 9: Event studies, adjusted firms sample

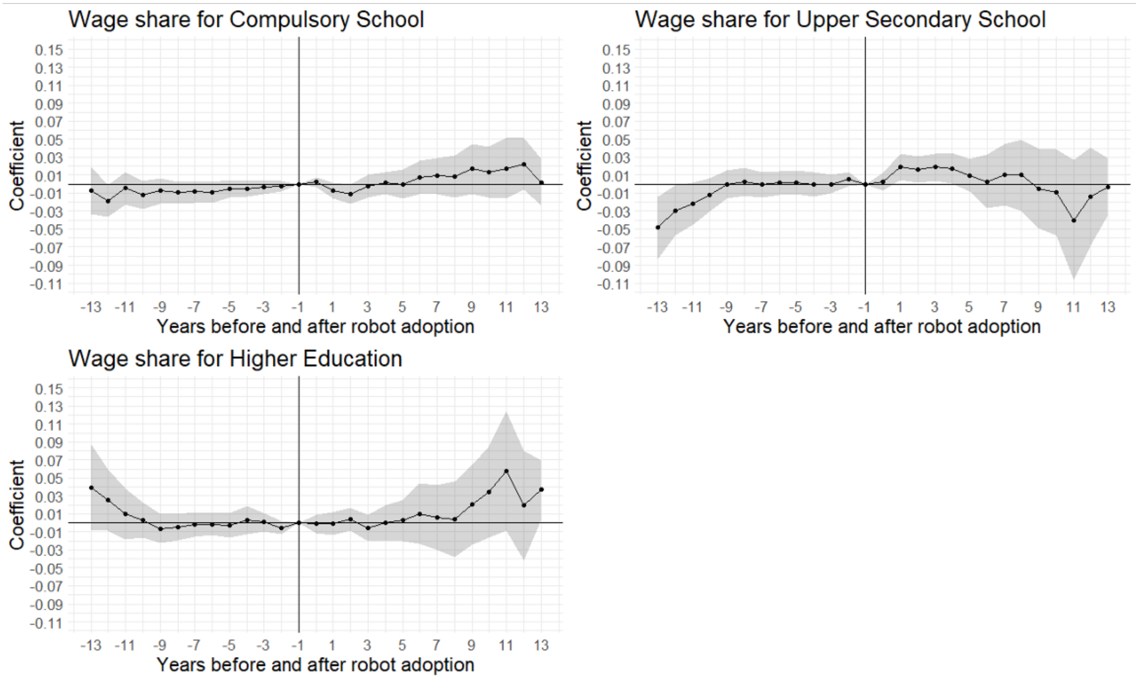

Figure notes: Event studies for all worker groups, based on the adjusted an adjusted sample where the smallest half of firms (in terms of operating revenues) and firms from certain non-adopting manufacturing sectors are excluded. Top left is for CS workers, top right is for US workers and bottom is for US workers. For each worker group, all years relative to the treatment time is regressed on the outcome variable following Equation 22. The year before treatment is used as reference point and thus the observations from this year are included in the control group.

The event studies undertaken have given support for parallel trends and have given some suggestive results on the effects of robotisation on the wage share for the respective worker groups. The parallel trends assumption seems to hold, at least for the 6-9 years prior to treatment. Adjusting the control group to only include larger firms in robot-using sectors do not significantly change this conclusion as shown in Figure 9. In addition, the event studies suggest that the wage share of CS workers decreases significantly after robot adoption, while the opposite is true for US workers, which is also the most dominant worker group in both robot firms and non-robot firms for all years (see Figure 5 in Chapter 4). The effect on workers with HE might be slightly negative, but is insignificant. The post treatment coefficients are changing over time and this might imply heterogeneous treatment effects across time.

products and nuclear fuel

Some points are worth highlighting in relation to the regression analysis in the next section. Notice that it is the year prior to the treatment that is excluded from the regression and is thus used as the reference year. This is to be consistent with the regression analysis, where the treatment dummy equals 1 from the year of import (v_0). However, as discussed in Chapter 4, it might be bold to assume that robots are taking into production the same year as it is produced. Moreover, some firms might have imported the robot late in the year, so treatment is not effective for most of the year. The fact that the r coefficients are close to 0, but changes radically the year after for CS and US workers, might be evidence for this problem. Moreover, the regression analysis presented in the next section can be thought of as a test of whether the weighted mean of the coefficients before treatment is significantly different from the weighted mean of the coefficients after the treatment. In relation to this it is worth noting that the TWFE model gives more weight to the observations closer to the treatment period as there are more of them and the variance tend to be smaller. Lastly, it is worth noticing that the confidence intervals for each coefficient is likely to be exaggerated because of few observations for each individual observation. In the regression analysis, as more observations are included in the aggregate coefficient, it is likely that the confidence intervals are smaller.

6.2 Regressions

This section goes through the results of the TWFE regression analysis. The regressions correspond to the model described in Chapter 5 (Equation 23 and 24). In each table the same regressions are reported twice, once for the control group with all manufacturing firms and another for the adjusted control group where smaller firms and certain sectors with a negligible usage of robots are excluded (as explained in the previous section).

Table 3 reports the first results based on Equation 23, including the trend variable. The regression is run separately for CS, US and HE workers in both the unadjusted and the adjusted sample. This implies that, for each group of robot adopters, the control group will be non-robot firms, firms that will adopt robots later in the time period and firms that already adopted robots. All regressions are done three times, each time with different control variables. In all regressions, a trend variable is included to identify the robot-independent time trend for the respective wage share (how many percentage points the wage share is expected to change each year if the firm does not adopt a robot). The *ROBOT* coefficient is the estimated causal effect of access to robot technology on the respective wage shares. All regressions also include entity (firm) FE, but not time FE as this will disrupt the trend coefficient.

The results show that there is a significant robot-independent trend that decreases the wage share for CS and US workers annually by around 0.4-0.5 and 0.3 percentage points respectively. The wage share for HE workers increases on average around 0.4 percentage points annually. This confirms the general trends observed in Figure 5, Figure 7 and Table 2 from Chapter 4. However, the *ROBOT* coefficients tell another story. When firms adopt robots, the effect on the wage share for CS and HE workers is insignificant. Contrary to this, the wage share for US workers have a significant increase of 1.3 - 1.4 percentage points. Seen together with the robot-independent trend, this implies that the positive effect of robots is likely to be outweighed by the general negative trend in about five years. The results are robust to changes in control variables and do not change notably between the unadjusted and adjusted sample. Although not precisely the same, these results map quite well to the predictions of the event studies; the US share increased and, although showing some slight negative trends, both the CS and HE firms are not affected by robot adoption.

Table 4 perform various regressions that include extended controls and differ only in different combinations of included time FE, firm FE and the Trend variable (i.e. both Equation 24, Equation 23 and a combination of these). Regressions in Panel A adds time FE to Equation 23. The coefficients do not change, but the trend variable is precluded and cannot inform us on the robot-independent wage share trend. In Panel B, Equation 24 is estimated, only including time and entity FE, no trend variable. These results are more or less identical to the results in Panel A, illustrating the similar function of the trend variable and time FE. The coefficients in Panel C, however, changes radically. This is because this model do not include firm FE. Ignoring firm FE changes the sign of the HE coefficient to positive and the US coefficients to negative. Nevertheless, only the CS and US negative coefficients are significant. These results are more similar to what we observed in Figure 5 and 7. What this demonstrates is that there are several unobserved variables that are constant over time for each specific worker group at each individual firm that are both correlated to the wage share and robot adoption.[13] Examples of such variables could be the local labour market, manufacturing sector, the skill of the manager or work culture. What this demonstrates is there is evidence of selection bias; the firms that adopts robots are on average different from the non-adopters. Although adding firm FE controls for some of these variables, we cannot be certain all are accounted for.

6.3 Robustness test

To investigate the credibility of the estimated coefficients, the same regressions as in Table 3 are performed again, only with a slight twist; we adjust the control group. Recall that, as discussed in Chapter 5, the TWFE regressions are a weighted average of all possible 2x2 DiD regressions. These DiD regressions are in effect comparisons between a group of firms that get treatment at a specific point in time and with several different control groups. Some regressions use firms that never adopt robots as the control group, some use firms that have not yet adopted the robot and some use the firms that adopted earlier. The weighted average of these estimates gives us the results in Table 3. The models in Table 5 are the same models as in Table 3 only that all firms that never adopt a robot are excluded. This effectively means that the only control group at any specific treatment point only include firms that either have not yet got treatment or that got treatment before. Although the event studies suggested parallel trends

[13]Note that this is the case because each regression is performed for each worker group separately. Had this not been the case, both the time and firm FE, and other control variables, would not account for much as the firm wage share would always equal to 1.

Table 3: FE regressions with trend variable. Dependent variable: Wage share (for different worker groups)

	All manufacturing firms			Adjusted sample		
	(1) Compulsory School	(2) Upper Secondary School	(3) Higher Education	(4) Compulsory School	(5) Upper Secondary School	(6) Higher Education
Panel A						
No control variables						
ROBOT	0.001	0.014^{**}	-0.007	0.004	0.013^{**}	-0.006
	(0.005)	(0.007)	(0.007)	(0.005)	(0.007)	(0.007)
TREND	-0.004^{***}	-0.003^{***}	0.004^{***}	-0.005^{***}	-0.003^{***}	0.004^{***}
	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Panel B						
Basic control variables						
ROBOT	0.001	0.014^{**}	-0.007	0.004	0.013^{**}	-0.006
	(0.005)	(0.007)	(0.007)	(0.005)	(0.007)	(0.007)
TREND	-0.004^{***}	-0.003^{***}	0.004^{***}	-0.005^{***}	-0.003^{***}	0.004^{***}
	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Panel C						
Extended control variables						
ROBOT	0.001	0.014^{**}	-0.005	0.004	0.013^{*}	-0.004
	(0.005)	(0.007)	(0.007)	(0.005)	(0.006)	(0.007)
TREND	-0.004^{***}	-0.003^{***}	0.004^{***}	-0.005^{***}	-0.003^{***}	0.004^{***}
	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Observations	61,395 / 44,158 / 44,158	71,079 / 50,215 / 50,215	49,223 / 34,668 / 34,668	31,000 / 28,780 / 28,780	33,164 / 30,627 / 30,627	27,751 / 25,577 / 25,577
Adjusted R^2	0.795 / 0.797 / 0.797	0.789 / 0.802 / 0.802	0.877 / 0.877 / 0.878	0.800 / 0.806 / 0.806	0.804 / 0.813 / 0.813	0.878 / 0.886 / 0.887

Table notes: Fixed effect regressions with trend variable (quasi TWFE). ROBOT show the estimated effect of access to robot technology on the wage share for the particular worker group. TREND show the robot-independent time trend for the wage share for the particular worker group. All regressions are performed on two samples; one including ”All manufacturing firms” (regression (1), (2), (3)) and one ”Adjusted sample” where the smallest half of firms (in terms of operating revenues) and firms from certain non-adopting manufacturing sectors are excluded (regression (4), (5), (6), see section 6.1 for further details). Basic controls in Panel B includes capital and value added. Extended controls in Panel C adds, on top of the basic controls, employment and operating revenues. In all models, entity (firm) FE are included, but no time FE so it does not compete with TREND (see Table 4 for same regressions including time FE and excluding entity FE). Observations and Adjusted R^2 are reported for all regressions in this order: No Control Variables, Basic Control Variables, Extended Control Variables. Heteroscedasticity-robust standard errors are clustered at the firm level. Significant levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 4: TWFE regressions with different fixed effects. Dependent variable: Wage share (for different worker groups)

	All manufacturing firms			Adjusted sample		
	(1) Compulsory School	(2) Upper Secondary School	(3) Higher Education	(4) Compulsory School	(5) Upper Secondary School	(6) Higher Education
Panel A						
Trend, time FE						
and entity FE						
ROBOT	0.0004	0.014^{**}	-0.005	0.003	0.013^{**}	-0.004
	(0.005)	(0.007)	(0.007)	(0.004)	(0.006)	(0.007)
TREND						
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Panel B						
Time and entity FE.						
No Trend.						
ROBOT	0.0004	0.014^{**}	-0.005	0.003	0.013^{**}	-0.004
	(0.005)	(0.007)	(0.007)	(0.004)	(0.006)	(0.007)
Panel C						
Trend. No time FE						
or entity FE						
ROBOT	-0.041^{***}	-0.022^{*}	0.019	-0.025^{***}	-0.011	0.021
	(0.008)	(0.012)	(0.013)	(0.008)	(0.011)	(0.013)
TREND	-0.004^{***}	-0.003^{***}	0.003^{***}	-0.005^{***}	-0.002^{***}	0.003^{***}
	(0.0003)	(0.0003)	(0.0004)	(0.0003)	(0.0004)	(0.0004)
Observations	44,158 / 44,158 / 44,158	50,215 / 50,215 / 50,215	34,668 / 34,668 / 34,668	28,780 / 28,780 / 28,780	30,627 / 30,627 / 30,627	25,577 / 25,577 / 25,577
Adjusted R^2	0.798 / 0.798 / 0.026	0.802 / 0.802 / 0.007	0.878 / 0.878 / 0.008	0.807 / 0.807 / 0.034	0.813 / 0.813 /0.006	0.887 / 0.886 / 0.014

Table notes: Two way fixed effect regressions. ROBOT show the estimated effect of access to robot technology on the wage share for the particular worker group. TREND show the robot-independent time trend for the wage share for the particular worker group. All regressions are performed on two samples; one including "All manufacturing firms" (regression (1), (2), (3)) and one "Adjusted sample" where the smallest half of firms (in terms of operating revenues) and firms from certain non-adopting manufacturing sectors are excluded (regression (4), (5), (6), see section 6.1 for further details). All regressions includes extended controls: capital, added value, employment and operating revenues. The regressions in Panel A includes the TREND variable, time FE and entity (firm) FE. In Panel B, the regressions includes time FE and entity (firm) FE, but not the trend variable. The regressions in Panel C do not include time FE nor entity (firm) FE, but the trend variable. Observations and Adjusted R^2 are reported for all regressions in this the order just described. Heteroscedasticity-robust standard errors are clustered at the firm level. Significant levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

before treatment, the reason for making this change is to be more certain that the control group is more similar to the treatment firms.

The results of the robustness check are reported in Table 5. Firstly, although the change in time trends are not radical, it is worth noting some changes that suggest the wage share trend for the various worker groups are different for the robot adopters compared to non-adopters. In particular, the wage share for HE workers *increase* at a higher rate for robot adopters (annual increase is 0.2 percentage point compared to Table 3). Contrary to this, the trend in the wage share for both CS and US workers decrease slightly more compared to Table 3, although this is only the case in the "All manufacturing firms" sample, not the "Adjusted sample". This is also in line with what we expect from Figure 7 in Chapter 4. In addition, there are some crucial differences for the coefficients. Although all signs are the same, which gives some strength to our previous results, the positive effect of access to robots on US workers are now not significant. The opposite is true for the negative effect on HE workers, which is now significant. The effect on CS workers has not changed. In sum, this suggests that, although robot adopting firms generally increase the demand for HE workers at a higher rate, the effect of robots is significantly negative for this group, while insignificant for the two others. This robustness check give us some new insights, but there are some reason to be cautions. Although it might be the case that these firms might be a better control group, they are in effect controls for themselves, the number of observations are also radically fewer which makes the results more uncertain. It is nevertheless comforting to see that the signs of the coefficients stays the same.

6.4 Discussion of the findings

The purpose of the event studies and the regression analysis is to understand the causal effect of access to robot technology on labour demand (wage share) for three different worker groups. Figure 5 show that robot adopters have a higher share of HE workers at the beginning of the period compared to non-adopters, while the opposite is true for CS and US workers. The trend variables in our regression analysis, together with Figure 5, 6 and 7 in Chapter 4, show that there is a general skill biased trend for all firms toward more HE workers and less CS and US workers. The robustness test show that all these trends are somewhat greater for robot firms. However, the effect of adopting robot technology contrast these trends. Although coefficients and significance levels vary slightly, the results suggests that only the relative demand for US workers increase by around 1.4 percentage points. In a manufacturing context, the US workers are likely to have High School vocational training and be blue collar workers. The demand for CS workers is not effected, but the demand for HE decreases up to 1.5 percentage points. This suggests that robot technology is positively *biased* toward US workers and negatively *biased* toward HE workers.

These findings both support and contrast some of the previous literature. The general trends toward more high skilled labour supports the recent historical developments of skilled biased technological change (Acemoglu & Autor, 2011b). Although many papers find that robot adoption is not biased (Bessen et al.,

Table 5: FE regressions with trend variable. Dependent variable: Wage share (for different worker groups). Only robot firms.

	All manufacturing firms			Adjusted sample		
	(1) Compulsory School	(2) Upper Secondary School	(3) Higher Education	(4) Compulsory School	(5) Upper Secondary School	(6) Higher Education
ROBOT	0.005	0.011	-0.015^{**}	0.006	0.011	-0.013^{*}
	(0.005)	(0.008)	(0.007)	(0.005)	(0.007)	(0.007)
TREND	-0.005^{***}	-0.002^{***}	0.006^{***}	-0.005^{***}	-0.003^{***}	0.006^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	2,916	3,111	2,835	2,696	2,793	2,676
Adjusted R^2	0.843	0.826	0.883	0.846	0.815	0.881

Table notes: Fixed effect regressions with trend variable (quasi TWFE). Only robot firms are included, implying that the control group is always robot firms that have not yet been treated or have already been treated. ROBOT show the estimated effect of access to robot technology on the wage share for the particular worker group. TREND show the robot-independent time trend for the wage share for the particular worker group. All regressions are performed on two samples; one including "All manufacturing firms" (regression (1), (2), (3)) and one "Adjusted sample" where the smallest half of firms (in terms of operating revenues) and firms from certain non-adopting manufacturing sectors are excluded (regression (4), (5), (6), see section 6.1 for further details). All regressions includes extended controls: capital, added value, employment and operating revenues. Entity (firm) FE are included, but no time FE so it does not compete with TREND (see Appendix E for same regressions including time FE). Observations and Adjusted R^2 are reported for all regressions in this order: No Control Variables, Basic Control Variables, Extended Control Variables. Heteroscedasticity-robust standard errors are clustered at the firm level. Significant levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

2019; Dauth et al., 2021; Hirvonen et al., 2021), any bias previously documented are usually in favor of high skilled labour, i.e. HE workers (Acemoglu et al., 2020; Barth et al., 2020; Dixon et al., 2021; Humlum, 2019). The fact that robot adopting firms have a higher share of HE workers *before* implementing robots while skilled blue collar (US) workers increase their wage share *after* might suggest that HE workers are needed to facilitate robot adoption, while US workers' productivity increase relatively more once robots are used in production. Particularly important for our purposes, when looking at the same robot and worker group data as this paper (only difference is that their data go back to 1999 rather than 2003), Barth et al. (2020) indeed find increased within-firm wage differentials. Following robot adoption, the wage of HE workers increase, US workers' wage is not affected while CS workers get a lower wage. Our results complements their study in some important aspects. As noted, the wage share consists of both wages *and employment*. Seen together, it seems that robot implementation leads to higher *wages* for HE workers, but increased *employment* of US workers, and that this latter effects dominates so the total effect, measured in this thesis by the wage share, is a relative increase in demand for US workers and relative decrease in demand for HE workers. This is interesting because, first and foremost, it suggests that, in certain context, the bias of robots might in fact be towards the blue collar workers. Moreover, it might show that increased demand might be expressed in different ways for different worker groups. It seems like HE workers get higher wages, while US workers are more employed.

Theory can help us understand what mechanisms leads to the observed results. Adopting the technology augmenting framework, the mostly likely interpretation is that robots are augmenting *all* worker groups, but that US workers are more augmented compared to other worker groups and thus their relative demand increases. This might be because it is US workers that operate the robots and thus their relative productivity increase. An alternative explanation could be that only US workers are augmented by robots, leading to the same results. However, this is unlikely as previous findings have found a positive effect of robots on HE workers and these workers are usually needed to implement the robots. Lastly, it could be the case that robots are in fact not augmenting CS and US workers, only HE workers, and thus reduce demand in the first instance. However, if US workers are complements with HE workers, demand for US workers might increase in the second instance, leading to a higher wage share. Although this could indeed be the case, the former of these explanations seems more plausible.

Adopting the task-based framework, robots are implemented and displaces low skilled labour, which might be both CS and US workers. This will in the first instance lead to a lower wage share, which is not what we observe for the US workers. This implies that other effects are dominating the displacement effect. In particular, if these robots are substantially more productive at the margin compared to the workers they displace, it is likely that the firm can expand and thus employ more workers (productivity effect). In addition, new tasks dependent on US workers might be created, also increasing demand. This could for example be to operate the robots, a new task introduced following robotisation. Which of these effects are at work cannot be determined with the results at hand, only that they together outweigh the displacement effect and thus increase the relative demand for US workers.

It could also be the case that it is not lower costs that motivate robot adoption, but performance of new tasks to produce new types of output, as observed in Finland between 1994 and 2018 (Hirvonen et al., 2021). If this is the case, then our results could suggest that US workers are a complement to robots, increasing in demand relatively more than other worker groups.

As previously discussed in Chapter 2 and 3, labour unions could play an important role both in the decision to adopt robots and in mediating its effects on the wage share. The bargaining power of strong unions might force firms to make sure new workers are trained to perform new tasks or employed otherwise. Thus, the high degree of unionisation in Norway generally and the manufacturing sector specifically could be a reason why we observe the positive effect of robots on the wage share for US workers.

6.5 Internal and external validity

The aim of any statistical analysis is to provide internally valid estimates of the causal mechanism in question, meaning that they are unbiased and consistent. Although several adjustments have been done to ensure this, some issues is worth highlighting that could cause issues to the internal validity of our results.

Firstly, it is important to note that there are only 280 robot firms in the regression analysis, which is not a substantial number. This is particularly true for the coefficients in the event studies case, evident from the increasing confidence intervals as we move away from the treatment time and observations become more scarce.

The most eminent threat to econometric analysis is usually that we do not sufficiently control for unobserved variables, thus leaving us with omitted variable bias. This means that our control group, conditional on control variables and fixed effects, is not appropriate as a proxy for the counterfactual. In the DiD framework, event studies is the most common way to examine if this is a problem by looking at pre-treatment parallel trends. The event studies in Section 6.1 suggest that there are parallel pre trends. Although this provides some comfort, we cannot be certain that there are no omitted variables affecting the *post*-treatment development of the wage share. Including control variables and FEs helps, but are far for a guarantee from omitted variable bias. Some papers use the long difference estimator to avoid potential issues with omitted variable bias (Acemoglu et al., 2020; Caroli & Van Reenen, 2001). Although this could help avoiding the issue to some degree, this approach might still suffer from this bias.

Another potential issue with our TWFE estimate could be that our model includes ”forbidden comparisons” that could bias our results (Roth et al., 2022). As pointed out in the methodology chapter (Equation 20), this could happen if within-group treatment heterogeneity is present, i.e. that the treatment effect would vary across time for a group of firms that got the treatment in the same time period. As discussed broadly in Chapter 4, this is likely to be the case with our data. Firstly, differences in robot quality and quantity across firms and time suggest differences in treatment. So does the fact that more than half of the robot firms import in more than one time period. These differences in treatment are likely to transfer in *heterogeneity in treatment effect*. Unfortunately, lack of data does not allow to us to more precisely discriminate between different types of robot technology. And even if we could, it could still be the case that, even with homogeneous treatment, the treatment *effects* could be heterogeneous, both across and within groups. The potential outcome framework outlined in Chapter 5 helps clarify the threat from this bias. However, this awareness is fairly recent (Roth et al., 2022). Roth et al. (2022) have suggested that the best way to cope with this issue is to manually exclude any possible forbidden comparisons by explicitly specifying all 2x2 DiD regressions that should be undertaken, as well as the corresponding weights used to aggregate the result to a single coefficient.

The last major threat to internal validity is anticipatory effects. As discussed in Chapter 4, this could indeed be a problem in our case. We have assumed treatment to start the same year as the robots were imported. However, two other points in time might be more relevant for the effect on the wage share, namely when robots are ordered (before import) and when robots are put into production (after import). As we have seen, robot firms typically have higher shares of HE workers prior to robot adoption. It seems plausible that these workers are employed before the robots come in order to facilitate. It might also be that CS and US workers are more affected after robots are put into place, as their productivity might increase. A way to be more certain about the results in regard to this issue could be to perform the same analysis again, only setting treatment time to start a year or two before, as for example done in Hirvonen et al. (2021).

A note should also be made on external validity, i.e. the degree to which these results applies to other contexts in time and space. The best way to examine the external validity is to repeat the experiment in new settings and see if they are similar. Although previous studies differ in many important respects, several other studies have found different effects on the bias of technological change. This suggests that the we need to be cautious before applying these results to other contexts. For example, the results might be more relevant in similar economic sectors and with similar labour market institutions. Some might be tempted to export these results to the more fundamental discussion about the future of work and inequality. Although these results have a place in this broader debate, they cannot help predict with much certainty how the economic future of automatisation might look like.

7 Conclusion

A growing literature tries to figure out who might win and who might lose from robotisation. This thesis contributes to this endeavour by investigating the effects of robotisation, in particular with-firm changes in labour demand in Norwegian Manufacturing firms between 2003 and 2016. Using a TWFE regression analysis, I have found that the wage share for US workers increased and the wage share for HE workers decreased after manufacturing firms in Norway adopted robot technology. Viewed together with the results from Barth et al. (2020), who examines the same firms and workers, this suggests that when robots are introduced to Norwegian manufacturing firms HE workers get relatively higher wages. At the same time, US workers are relatively more employed, and that the latter effect dominates. Thus, the total effect of robots is an increased relative demand for US workers and a reduced relative demand for HE workers. These effects following robotisation contrast a general skill biased and robot independent trend of falling wage shares for CS and US workers, and increasing wage share for HE workers. Two models might help us understand these results. Applying the factor augmenting framework, this result might suggest that robot technology is biased toward US workers as their relative productivity increases and thus they are employed at a higher rate. If we rather use the task-based model, our results can be explained by a combination of new tasks and increased productivity (and thus more labour demand for workers) outweighing the initial displacement effect of robots. Although these models provide useful intuitions that can explain parts of our observations, it is also crucial to highlight the particular Norwegian context with strong labour market institutions such as unions.

Although the results are robust to various control groups, several issues might threat the internal validity of the results. Omitted variable bias can never be ruled out, even though parallel trends seems to be present before robot adoption and control variables are included. More worryingly, our analysis might suffer from heterogeneous treatment effect bias due to both robots quality and quantity heterogeneity. In addition, we might have reasons to suspect anticipatory effects on the wage share before treatment. Considering the external validity, these results might have some value to related industries and labour markets with similar institutions. Although the results presented here have its place in the broader ongoing debate about the future of work, they cannot be used to predict any of these fundamental questions in any confident way.

In a context of ever faster technological improvements, on top of increased income and wealth inequality, it is imperative to understand effects of automation technologies such as robots. Although the efficiency outcomes are certainly important, it is particularly crucial to pay attention to the various outcomes for different groups of people. Indeed, it could be the case that there is an overall gain from robotisation, but that certain groups lose out. The recent political turmoil and protest movements in many Western countries serve as an example of what might happen if the effects of rapid changes in our economies are not well understood and met with appropriate policy. With this in mind, the results of this thesis might bring some comfort. In the Norwegian manufacturing context, it is skilled blue collar workers that in fact are the winners. Although their wage share is generally decreasing, it increases when their firm adopts robots. This might be good news, but should not lead to the conclusion that all future waves of automation automatically will bring the same outcomes. Rather, it is likely that this outcome will depend on policy regarding education and training, as well labour market institutions. This calls for more empirical work, but also better theoretical frameworks that integrate factors such as unions and bargaining structures. Hopefully, these efforts might enable us to identify strategies that make automation benefit everyone.

8 Bibliography

Aaberge, R., Mogstad, M., Vestad, O. L., & Vestre, A. (2021). Økonomisk ulikhet i norge i det 21. århundre. https://www.ssb.no/inntekt-og-forbruk/inntekt-og-formue/artikler/okonomisk-ulikhet-i-norge-i-det-21.arhundre/_/attachment/inline/46945fe1-533f-45b3-9ef2-cde52936f6fc:eaf3f053ce1e5878a6fb288814b428703665122e/RAPP2021-33.pdf

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature*, *40*(1), 7–72. https://doi.org/10.1257/0022051026976

Acemoglu, D., & Autor, D. (2011a). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 197–236). The University of Chicago Press.

Acemoglu, D., & Autor, D. (2011b). Skills, tasks and technologies: Implications for employment and earnings. In D. Card & O. Ashenfelter (Eds.). Elsevier. https://doi.org/https://doi.org/10.1016/S0169-7218(11)02410-5

Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). Competing with robots: Firm-level evidence from france. *AEA Papers and Proceedings*, *110*, 383–388. https://doi.org/10.1257/pandp.20201003

Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, *108*(6), 1488–1542. https://doi.org/10.1257/aer.20160696

Acemoglu, D., & Restrepo, P. (2019a). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, *33*(2), 3–30. https://doi.org/10.1257/jep.33.2.3

Acemoglu, D., & Restrepo, P. (2019b). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.). Elsevier. https://ideas.repec.org/h/eee/labchp/5-12.html

Acemoglu, D., & Restrepo, P. (2020a). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, *128*(6), 2188–2244. https://doi.org/10.1086/705716

Acemoglu, D., & Restrepo, P. (2020b). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, *128*(6), 2188–2244. https://doi.org/10.3386/w23285

Aghion, P., Antonin, C., & Bunel, S. (2019). Artificial intelligence, growth and employment: The role of policy. *Economie et Statistique / Economics and Statistics*, *Special Issue 50th Anniversary*(510-512), 149–164. https://doi.org/10.24187/ecostat.2019.510t.1994

Asplund, R., Barth, E., Lundborg, P., & Nilsen, K. M. (2011). Polarization of the nordic labour markets. *Finnish Economic Papers, Finnish Economic Association*, *24*(2), 87–110.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020a). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, *135*(2), 645–709. https://doi.org/10.1093/qje/qjaa004

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020b). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, *135*(2), 645–709.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, *29*(3), 3–30. https://doi.org/10.1257/jep.29.3.3

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill contentof recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–333.

Barrat, J. (2013). *Our final invention: Artificial intelligence and the end of the human era*. Thomas Dunne Books.

Barth, E., Roed, M., Schøne, P., & Umblijs, J. (2020). How robots change within-firm wage inequality.

Bessen, J., Goos, M., & Salomons, A. (2019). What happens to workers at firms that automate? *Boston Univ. School of Law, Law and Economics Research Paper*. https://doi.org/10.2139/ssrn.3328877

Biørn, E. (2008). *Økonometriske emner*. Unipub forlag.

Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W.W. Norton.

BSG. (2015). The robotics revolution: The next great leap in manufacturing. *Boston Consulting Group*. https://circabc.europa.eu/sd/a/b3067f4e-ea5e-4864-9693-0645e5cbc053/BCG_The_Robotics_Revolution_Sep_2015_tcm80-197133.pdf

Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics: Methods and applications*. Cambridge University Press.

Caroli, E., & Van Reenen, J. (2001). Skill-biased organizational change? evidence from a panel of british and french establishments. *The Quarterly Journal of Economics*, *116*(4), 1449–1492. https://doi.org/10.1162/00335530175326562

Cheng, H., Jia, R., Li, D., & Li, H. (2019). The rise of robots in china. *Journal of Economic Perspectives*, *33*(2), 71–88. https://doi.org/10.1257/jep.33.2.71

Cooper, R., Haltiwanger, J., & Power, L. (1999). Machine replacement and the business cycle: Lumps and bumps. *American Economic Review*, *89*(4), 921–946. https://doi.org/10.1257/aer.89.4.921

Cunningham, S. (2021). *Causal inference: The mixtape*. Yale University Press.

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*, *19*(6), 3104–3153. https://doi.org/10.1093/jeea/jvab012

Dixit, A. K., & Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American Economic Review*, *63*(3), 297–308.

Dixon, J., Hong, B., & Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, *67*(9), 5301–5967. https://doi.org/10.1287/mnsc.2020.3812

Doms, M., & Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, *1*(2), 409–429. https://doi.org/https://doi.org/10.1006/redy.1998.0011

Elsby, M. W., Hobijn, B., & Şahin, A. (2013). The decline of the us labor share. *Brookings Papers on Economic Activity*, *2013*(2), 1–63.

Ford, M. (2016). *Rise of the robots: Technology and the threat of a jobless future*. Basic Books.

Frey, C. B., & Osborne, M. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*(January), 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

Furman, J. (2019). Should we be reassured if automation in the future looks like automation in the past? In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 317–328). The University of Chicago Press.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

Goolsbee, A. (2019). Public policy in an ai economy. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 309–316). The University of Chicago Press.

Goos, M., & Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics*, *89*(1), 118–133. https://doi.org/10.1162/rest.89.1.118

Graetz, G., & Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, *100*(5), 753–768. https://doi.org/10.1162/rest_a_00754

Haapanala, H., Marx, I., & Parolin, Z. (2022). Robots and unions: The moderating effect of organized labour on technological unemployment. *Economic and Industrial Democracy*, 0143831X221094078.

Hirvonen, J., Stenhammar, A., & Tuhkuri, J. (2021). New evidence on the effect of technology on employment and skill demand. *Job Market Paper*. https://economics.mit.edu/grad/tuhkuri/research

Hobsbawm, E. J. E. (1952). The machine breakers. *Past Present*, *1*(1), 57–70. https://doi.org/10.1093/past/1.1.57

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, *81*(396), 945–960. https://doi.org/10.2307/2289064

Humlum, A. (2019). Robot adoption and labor market dynamics. *Working Paper*.

IFR. (2016). World robotics report 2016. *International Federation of Robotics*.

IFR. (2021). World robotics 2021. *International Federation of Robotics*.

IGM. (2019). *Robots and artificial intelligence*. Retrieved May 7, 2022, from https://www.igmchicago.org/surveys/robots-and-artificial-intelligence-2/

Ingelsrud, M. H., & Steen, A. H. (2021). *Norsk arbeidsliv 2021 - mot en ny normal?* (Tech. rep. Arbeidslivsbarometer). YS. http://arbeidslivsbarometeret.no/wp-content/uploads/2021/08/Arbeidslivsbarometeret-2021.pdf

ISO-8373. (2021). Iso 8373. *International Organization for Standardization*. https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-3:v1:en

Jehle, G. A., & Reny, P. J. (2011). *Advanced microeconomic theory*. Prentice Hall.

Karabarbounis, L., & Neiman, B. (2014). The global decline of the labor share. *The Quarterly journal of economics*, *129*(1), 61–103.

Koch, M., Manuylov, I., & Smolka, M. (2021). Robots and firms. *The Economic Journal*, *131*(638), 2553–2584.

Koren, M., Csillag, M., & Kollo, J. (2020). Machines and machinists: Incremental technical change and wage inequality. *Quarterly Journal of Economics, Working Paper*. https://doi.org/http://koren.mk/static/pdf/machines.pdf

Korinek, A., & Stiglitz, J. E. (2019). Artifi cial intelligence and its implications for income distribution and unemployment. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 349–390). The University of Chicago Press.

Malani, A., & Reif, J. (2015). Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform. *Journal of Public Economics*, *124*, 1–17.

Malvik, J. (2016). *Jonas gahr støre om ny teknologi: her står en optimist"*. Retrieved May 5, 2022, from https://www.innomag.no/jonas-gahr-store-om-ny-teknologi-her-star-en-optimist/

McKinsey. (2017). Jobs lost, jobs gained: Workforce transitions in a time of automation. *McKinsey Global Institute*. https://www.mckinsey.com/~/media/BAB489A30B724BECB5DEDC41E9BB9FAC.ashx

McKinsey. (2019). Industrial robotics. insights into the sector's future growth dynamics. *McKinsey Global Institute*. https://www.mckinsey.com/industries/advanced-electronics/our-insights/growth-dynamics-in-industrial-robotics

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, *71*(6), 1695–1725. https://doi.org/10.1111/1468-0262.00467

Michaels, G., Natraj, A., & Van Reenen, J. (2014). Has ict polarized skill demand? evidence from eleven countries over twenty-five years. *The Review of Economics and Statistics*, *696*(1), 0-77. https://doi.org/10.1162/REST_a_00366

Moene, K. O., & Wallerstein, M. (1997). Inequality, social insurance, and redistribution. *American Political Science Review*, *95*(4), 859–874. https://doi.org/10.1017/S0003055400400067

Pajarinen, M., Rouvinen, P., & Ekeland, A. (2015). Computerization threatens one-third of finnish and norwegian employment. *https://www.etla.fi/wp-content/uploads/ETLA-Muistio-Brief-34.pdf*.

Piketty, T. (2013). Capital in the 21st century. *Cambridge, MA: President and Fellows, Harvard College*.

Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2022). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Working Paper*. https://doi.org/10.48550/arXiv.2201.01194

Smith, A., & Anderson, M. (2017). *Automation in everyday life* (tech. rep.). Pew Research Center. https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2017/10/PI_2017.10.04_Automation_FINAL.pdf

Stensvold, T. (2016). *Dropper kina - legger fabrikk for bildeler til gjøvik*. Retrieved May 5, 2022, from https://www.tu.no/artikler/dropper-kina-legger-fabrikk-for-bildeler-til-gjovik/346859

Zeira, J. (1998). Workers, machines and economic growth. *Quarterly Journal of Economics*, *113*(4), 1091–1117. https://doi.org/10.1162/003355398555847

Appendices

A First appendix: Elasticity of Substitution

$\sigma_{K,L}$ tells us to what degree labour and capital are compliments or substitutes, and thus implies a certain the curvature of the production isoquant curve. Another way to put it, $\sigma_{K,L}$ measure how easily we can change between capital and labour to produce a given output (move along the isoquant). When we move along the isoquant, both the input ratio ($\frac{K}{L}$) and marginal rates of technical substitution ($\frac{-\delta L}{-\delta K}$). The elasticity of substitution measure the relative change of these, so that we get

$$
\begin{aligned}
\sigma_{K,L} &= \frac{\frac{\Delta(K/L}{(K/L)}}{\frac{\Delta(-\delta L/-\delta K)}{(-\delta L/-\delta K)}} \\
\sigma_{K,L} &= \frac{\frac{\Delta K/L}{K/L}}{\frac{\Delta MRTS_{L,K}}{MRTS_{L,K}}} \\
\sigma_{K,L} &= \frac{\delta log(K/L)}{\delta log(MRTS)} \\
\sigma_{K,L} &= \frac{\delta log(K/L)}{\delta log(MP_L/MP_K)} \\
\sigma_{K,L} &= \frac{\delta log(K/L)}{\delta log(w/r)} \\
\sigma_{K,L} &= \frac{\%\Delta(K/L)}{\%\Delta(w/r)}
\end{aligned}
\tag{25}
$$

As we can see from the derivation, $\sigma_{K,L}$ is the response of the input ratio to the change of the input marginal product (factor price) ratio. Thus, we can examine the effects of the various technological changes by finding an expression of the relative marginal products $\frac{MP_K}{MP_L}$ (see Appendix B for derivation).

$$
\frac{r}{w} = \frac{MP_K}{MP_L} = \frac{\alpha}{(1-\alpha)} \left(\frac{L}{K}\right)^{\frac{1}{\sigma}} \left(\frac{A_K}{A_L}\right)^{\frac{\sigma-1}{\sigma}} \tag{26}
$$

If $\sigma_{K,L} = \infty$, capital and labour are perfect substitutes and the isoquant is linear. In the opposite case, when $\sigma_{K,L} = 0$, then labour and capital are perfect complements and the isoquant take the Leontief form. In intermediate cases, we can have $\sigma_{K,L} > 1$ which implies that capital and labour are gross substitutes, if $\sigma_{K,L} = 1$ then they are unit elastic (Cobb-Douglas case), and if it is $\sigma_{K,L} < 1$, which has typically been the case empirically, implies that capital and labour are gross complements. In the respective cases, the implied isoquant will become increasingly convex moving from perfect substitutes to perfect compliments.

B Second appendix: Derivation of relative marginal returns

From the CES function (Equation 2), we can find $\frac{MP_K}{MP_L}$ by taking the derivative of F(K, L) with respect to L and K and then solve.

$$CES\ function = F(K, L) = A\left[\alpha(A_K K)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)(A_L L)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{27}$$

To simplify we set $u = \frac{\sigma-1}{\sigma}$ and $v = \frac{\sigma}{\sigma-1}$, which gives.

$$F(K, L) = A\left[\alpha(A_K K)^u + (1-\alpha)(A_L L)^u\right]^v \tag{28}$$

We take the derivative of this function with respect to both the factors.

$$MP_L = \frac{\delta F(K, L)}{\delta L} = v * A\left[\alpha(A_K K)^u + (1-\alpha)(A_L L)^u\right]^{v-1} * \left[u * (1-\alpha)(A_L L)^{u-1} * A_L\right] \tag{29}$$

$$MP_K = \frac{\delta F(K, L)}{\delta K} = v * A\left[\alpha(A_K K)^u + (1-\alpha)(A_L L)^u\right]^{v-1} * \left[u * \alpha(A_L L)^{u-1} * A_K\right] \tag{30}$$

This gives us

$$\frac{MP_K}{MP_L} = \frac{v * A\left[\alpha(A_K K)^u + (1-\alpha)(A_L L)^u\right]^{v-1} * \left[u * \alpha(A_L L)^{u-1} * A_K\right]}{v * A\left[\alpha(A_K K)^u + (1-\alpha)(A_L L)^u\right]^{v-1} * \left[u * (1-\alpha)(A_L L)^{u-1} * A_L\right]} \tag{31}$$

$$\frac{MP_K}{MP_L} = \frac{\left[\alpha(A_L L)^{u-1} * A_K\right]}{\left[(1-\alpha)(A_L L)^{u-1} * A_L\right]} \tag{32}$$

$$\frac{MP_K}{MP_L} = \frac{\alpha}{(1-\alpha)}\left(\frac{K}{L}\right)^{u-1}\left(\frac{A_K}{A_L}\right)^{u-1}\frac{A_K}{A_L} \tag{33}$$

$$\frac{MP_K}{MP_L} = \frac{\alpha}{(1-\alpha)}\left(\frac{A_K}{A_L}\right)^{u}\left(\frac{K}{L}\right)^{u-1} \tag{34}$$

In the term $\left(\frac{K}{L}\right)^{u-1}$, we substitute in $\frac{\sigma-1}{\sigma}$ for u and solve

$$\left(\frac{K}{L}\right)^{u-1} = \frac{K^{u-1}}{L^{u-1}} = \frac{K^{\frac{\sigma-1}{\sigma}-\frac{\sigma}{\sigma}}}{L^{\frac{\sigma-1}{\sigma}-\frac{\sigma}{\sigma}}} = \frac{K^{-\frac{1}{\sigma}}}{L^{-\frac{1}{\sigma}}} = \frac{L^{\frac{1}{\sigma}}}{K^{\frac{1}{\sigma}}} \tag{35}$$

Lastly, substituting this back into equation (8) and substitute in $\frac{\sigma-1}{\sigma}$ for u gives us

$$\frac{MP_K}{MP_L} = \frac{\alpha}{(1-\alpha)}\left(\frac{A_K}{A_L}\right)^{\frac{\sigma-1}{\sigma}}\left(\frac{L}{K}\right)^{\frac{1}{\sigma}} \tag{36}$$

C Third appendix: Additional descriptive statistics

This figure displays the increase in the accumulated value of the robot stock in manufacturing firms.

Figure 10: Accumulated value of robots

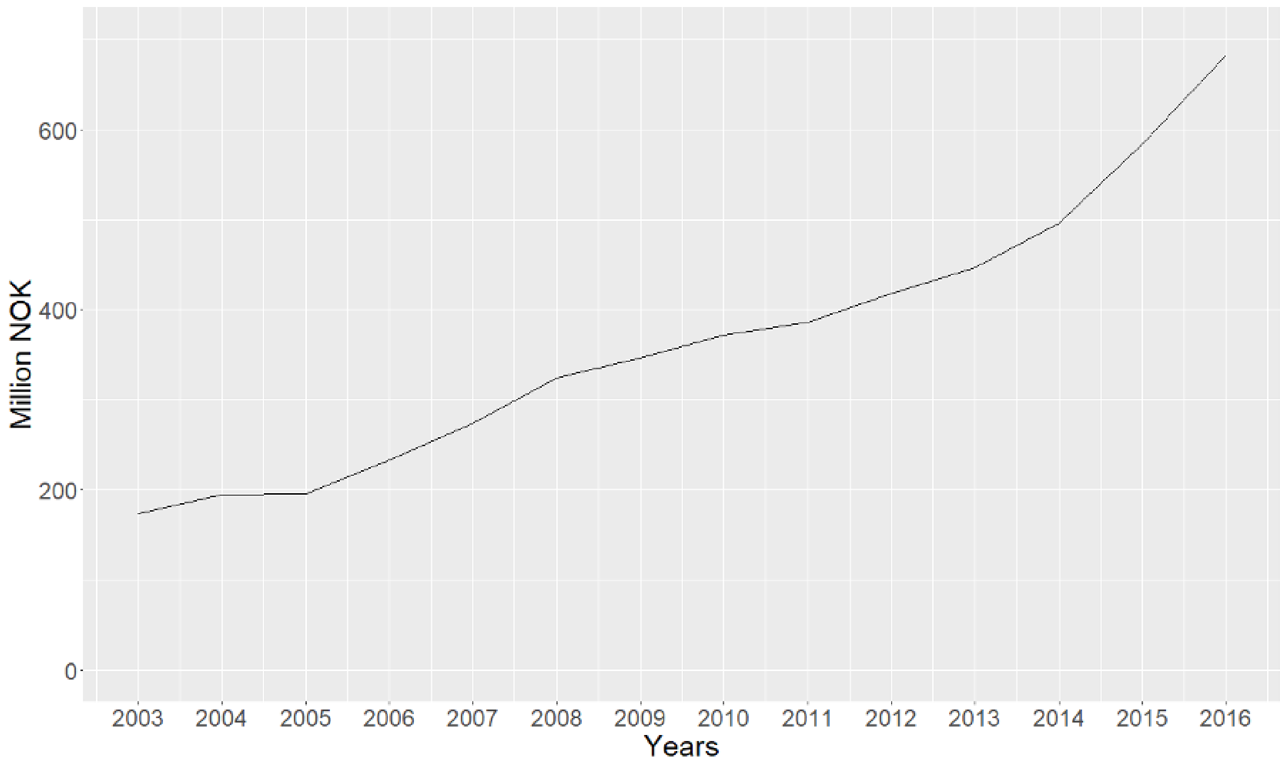

Figure notes: Accumulated value of robot imports (not depreciated). Source: SSB.

Manufacturing nace codes (Standard Industrial Classification 2002, SSB)

Manufacture of food products and beverages (15)

Manufacture of tobacco products (16)

Manufacture of textiles (17)

Manufacture of wearing apparel, dressing and dyeing of fur (18)

Tanning and dressing of leather, manufacture of luggage, handbags, saddlery, harness and footwear (19)

Manufacture of wood and of products of wood and cork, except furniture, manufacture of articles of straw and plaiting materials (20)

Manufacture of pulp, paper and paper products (21)

Publishing, printing and reproduction of recorded media (22)

Manufacture of coke, refined petroleum products and nuclear fuel (23)

Manufacture of chemicals and chemical products (24)

Manufacture of rubber and plastic products (25)

Manufacture of other non-metallic mineral products (26)

Manufacture of basic metals (27)

Manufacture of fabricated metal products, except machinery and equipment (28)

Manufacture of machinery and equipment n.e.c. (29)

Manufacture of office machinery and computers (30)

Manufacture of electrical machinery and apparatus n.e.c. (31)

Manufacture of radio, television and communication equipment and apparatus (32)

Manufacture of medical, precision and optical instruments, watches and clocks (33)

Manufacture of motor vehicles, trailers and semi-trailers (34)

Manufacture of other transport equipment (35)

Manufacture of furniture, manufacturing n.e.c. (36)

Recycling (37)

D Fourth appendix: Regression tables for event studies

These tables are the corresponding regression tables for the eventstudies in Figure 8 and Figure 9, respectively

Table 6: Dynamic TWFE regression, all firms. Compliments Figure 8

	Dependent variable:		
	Cost share		
	Compulsory School	Upper Secondary School	Higher Education
	(1)	(2)	(3)
Lead 2	0.001 (0.003)	−0.001 (0.004)	−0.003 (0.004)
Lead 3	0.001 (0.004)	−0.011* (0.005)	0.007 (0.006)
Lead 4	−0.004 (0.005)	−0.008 (0.007)	0.009 (0.008)
Lead 5	−0.004 (0.005)	−0.003 (0.007)	0.004 (0.008)
Lead 6	−0.002 (0.006)	−0.012* (0.007)	0.006 (0.007)
Lead 7	0.0002 (0.006)	−0.010 (0.008)	0.005 (0.009)
Lead 8	−0.002 (0.007)	−0.005 (0.010)	0.002 (0.009)
Lead 9	0.001 (0.007)	−0.006 (0.009)	0.002 (0.010)
Lead 10	−0.004 (0.008)	−0.021** (0.009)	0.011 (0.011)
Lead 11	0.002 (0.009)	−0.027** (0.011)	0.015 (0.015)
Lead 12	−0.009 (0.009)	−0.036*** (0.013)	0.030* (0.017)
Lead 13	0.005 (0.012)	−0.048*** (0.017)	0.034 (0.023)
Lag 0	0.002 (0.003)	0.001 (0.005)	−0.002 (0.005)
Lag 1	−0.009* (0.005)	0.013* (0.007)	0.002 (0.007)
Lag 2	−0.010* (0.006)	0.013* (0.008)	0.001 (0.007)
Lag 3	−0.002 (0.007)	0.013* (0.008)	−0.011 (0.008)
Lag 4	−0.002 (0.007)	0.013 (0.009)	−0.0005 (0.009)
Lag 5	−0.006 (0.009)	0.012 (0.010)	0.002 (0.011)
Lag 6	0.007 (0.011)	0.008 (0.015)	0.007 (0.016)
Lag 7	0.009 (0.011)	0.010 (0.017)	0.014 (0.020)
Lag 8	0.011 (0.014)	0.007 (0.020)	−0.002 (0.019)
Lag 9	0.016 (0.015)	−0.007 (0.020)	0.016 (0.021)
Lag 10	0.026 (0.020)	−0.013 (0.022)	0.024 (0.024)
Lag 11	0.013 (0.016)	−0.035 (0.032)	0.046 (0.033)
Lag 12	0.016 (0.015)	−0.012 (0.027)	0.013 (0.030)
Lag 13	−0.004 (0.013)	−0.002 (0.016)	0.031* (0.016)

Table notes: Dynamic TWFE models (Equation 22). Year before treatment is used as reference year. All firms are included. All regressions includes time FE and entity (firm FE). Significant levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 7: Dynamic TWFE regression, adjusted firms. Compliments Figure 9

	Dependent variable:		
	Cost share		
	Compulsory School	Upper Secondary School	Higher Education
	(1)	(2)	(3)
Lead 2	−0.003 (0.003)	0.005 (0.004)	−0.005 (0.004)
Lead 3	−0.003 (0.004)	0.00005 (0.005)	0.001 (0.005)
Lead 4	−0.005 (0.005)	−0.0002 (0.007)	0.003 (0.008)
Lead 5	−0.005 (0.005)	0.002 (0.007)	−0.003 (0.007)
Lead 6	−0.009 (0.006)	0.001 (0.007)	−0.001 (0.006)
Lead 7	−0.009 (0.006)	−0.001 (0.007)	−0.002 (0.007)
Lead 8	−0.009 (0.006)	0.002 (0.008)	−0.005 (0.008)
Lead 9	−0.008 (0.007)	−0.0005 (0.008)	−0.006 (0.008)
Lead 10	−0.012 (0.008)	−0.012 (0.009)	0.003 (0.010)
Lead 11	−0.005 (0.009)	−0.022* (0.012)	0.010 (0.014)
Lead 12	−0.019** (0.009)	−0.029** (0.014)	0.025 (0.017)
Lead 13	−0.007 (0.013)	−0.048*** (0.018)	0.039 (0.024)
Lag 0	0.002 (0.003)	0.002 (0.005)	−0.001 (0.005)
Lag 1	−0.007 (0.004)	0.019** (0.008)	−0.001 (0.006)
Lag 2	−0.011** (0.005)	0.016** (0.008)	0.004 (0.006)
Lag 3	−0.002 (0.006)	0.019** (0.008)	−0.006 (0.008)
Lag 4	0.001 (0.006)	0.017** (0.009)	−0.0001 (0.010)
Lag 5	−0.00000 (0.008)	0.010 (0.009)	0.003 (0.012)
Lag 6	0.007 (0.009)	0.003 (0.015)	0.010 (0.017)
Lag 7	0.009 (0.010)	0.010 (0.017)	0.006 (0.018)
Lag 8	0.008 (0.012)	0.010 (0.020)	0.004 (0.021)
Lag 9	0.017 (0.014)	−0.005 (0.023)	0.020 (0.023)
Lag 10	0.013 (0.014)	−0.009 (0.025)	0.035 (0.026)
Lag 11	0.017 (0.017)	−0.040 (0.034)	0.058* (0.034)
Lag 12	0.022 (0.015)	−0.014 (0.028)	0.019 (0.031)
Lag 13	0.002 (0.013)	−0.003 (0.016)	0.037** (0.017)

Table notes: Dynamic TWFE models (Equation 22). Year before treatment is used as reference year. Only "Adjusted sample" included (see Section 6.1) where the smallest half of firms (in terms of operating revenues) and firms from certain non-adopting manufacturing sectors are excluded. All regressions includes time FE and entity (firm FE). Significant levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

E Fifth appendix: Additional regressions, robustness test

This figure is the robustness test version of Table 4.

Table 8: TWFE regressions. Only robot firms. Dependent variable: Wage share (for different worker groups). Different fixed effects.

	All manufacturing firms			Adjusted sample		
	(1) Compulsory School	(2) Upper Secondary School	(3) Higher Education	(4) Compulsory School	(5) Upper Secondary School	(6) Higher Education
Panel A						
Trend, time FE						
and entity FE						
ROBOT	0.005	0.011	-0.015^{**}	0.005	0.011	-0.013^{*}
	(0.005)	(0.008)	(0.007)	(0.005)	(0.007)	(0.007)
TREND						
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Panel B						
Time and entity FE.						
No Trend.						
ROBOT	0.005	0.011	-0.015^{**}	0.005	0.011	-0.013^{*}
	(0.005)	(0.008)	(0.007)	(0.005)	(0.007)	(0.007)
Panel C						
Trend. No time FE						
or entity FE						
ROBOT	−0.002	0.012	−0.006	0.003	0.010	−0.007
	(0.012)	(0.015)	(0.020)	(0.011)	(0.015)	(0.020)
TREND	-0.004^{***}	−0.001	0.003^{**}	-0.004^{***}	−0.002	0.003^{**}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	2,916 / 2,916 / 2,916	3,111 / 3,111 / 3,111	2,835 / 2,835 / 2,835	2,696 / 2,696 / 2,696	2,793 / 2,793 / 2,793	2,676 / 2,676 / 2,676
Adjusted R^2	0.844 / 0.844 / 0.059	0.826 / 0.826 / 0.022	0.883 / 0.883 / 0.036	0.847 / 0.847 / 0.066	0.816 / 0.816 /0.023	0.881 / 0.881 / 0.047

Two way fixed effect regressions. ROBOT show the estimated effect of access to robot technology on the wage share for the particular worker group. TREND show the robot-independent time trend for the wage share for the particular worker group. All regressions are performed on two samples; one including "All manufacturing firms" (regression (1), (2), (3)) and one "Adjusted sample" where the smallest half of firms (in terms of operating revenues) and firms from certain non-adopting manufacturing sectors are excluded (regression (4), (5), (6), see section 6.1 for further details). All regressions included extended controls: capital, added value, employment and operating revenues. The regressions in Panel A includes the TREND variable, time FE and entity (firm) FE. In Panel B, the regressions includes time FE and entity (firm) FE, but not the trend variable. The regressions in Panel C do not include time FE nor entity (firm) FE, but the trend variable. Observations and Adjusted R^2 are reported for all regressions in this the order just described. Heteroscedasticity-robust standard errors are clustered at the firm level. Significant levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.