Minimum Wage and Skills: Evidence from Job Vacancy Data

Authors:
Elodie Andrieu[x]
King's College London
Malgorzata Kuczera[y]
Cambridge University

Date:
July 2023

The Productivity Institute
Working Paper No.034

[x] King's Business School
[y] Education Department

Key words
minimum wage, job vacancy, labour hiring

Authors' contacts:
elodie.andrieu@kcl.ac.uk

Acknowledgements
I am particularly grateful to Malgorzata Kuczera for providing the data and helping with the construction of the education variable and regular feedback. I am grateful to the Gatsby for their financial assistance in the production of this report. The views and opinions expressed in this report are those of the author and do not necessarily state or reflect those of the Gatsby Charitable Foundation. I thank Mary O'Mahony and Kerry L. Papps for useful comments, as well as Seyhun Sakalli, John Morrow and Tommaso Ciarli for helpful suggestions. The paper also benefited from discussions at the OECD Brownbag, ESCOE conference, the Low Pay Commission seminar, Online Job Posting Workshop, Monash-Warwick-Zurich Text-As-Data Workshop, and internal King's College London events.

Suggested citation

E. Andrieu, M. Kuczera (2023) *Minimum Wage and Skills: Evidence from Job Vacancy Data* Working Paper No. 034, The Productivity Institute.

The Productivity Institute is an organisation that works across academia, business and policy to better understand, measure and enable productivity across the UK. It is funded by the Economic and Social Research Council (grant number ES/V002740/1).

The Productivity Institute is headquartered at Alliance Manchester Business School, The University of Manchester, Booth Street West, Manchester, M15 6PB. More information can be found on The Productivity Institute's website. Contact us at theproductivityinstitute@manchester.ac.uk

Abstract

Low-wage occupations tend to be populated by workers with low levels of education. An increase in the minimum wage, while designed to protect workers in the lower part of the wage distribution, might result in unintended consequences for those same workers. In this paper, we study firms' reaction to higher minimum wages, exploiting a change to the minimum-wage policy in the UK in 2016. We document how an increase in the minimum wage affects the labour hiring for different education and technical skill levels of workers. The results show that an increase in the minimum wage compressed both the demand for low educated workers and the demand for workers with low levels of technical skills (tech workers) for graduates in low and middle skilled occupations. Using a difference-in-differences framework, we find that a large and unexpected change to the minimum wage led to a 11 percentage point decrease in the proportion of non-graduate vacancies and a 15 percentage point decline in the share of low-tech ads. There is evidence for labour-labour substitution at the low-end of the skill distribution and labour-technology substitution for more educated workers as a way to compensate for labour costs increases.

1 Introduction

The minimum wage policy has been subject to an on-going debate among academics and policymakers in the UK and abroad. It was initially put in place to avoid exploitation of workers. Nowadays, in developed countries, it aims to protect workers against unduly low pay (ILO definition). The 2021 Nobel prize winner David Card and others in their work from the 1990's (Card and Krueger (1993)), challenge the classical theory that wage increases cost jobs. From a firm's point of view, two opposite channels are consistently discussed. An increase in the minimum wage leads to higher labour costs and therefore lower employment. Conversely, it can increase firms' productivity through investment in workers and in innovative production methods (technology), rather than cutting jobs. As pointed out by Clemens (2021), research has long-focused on wage and employment outcomes but "what non-employment margins might firms adjust in response to minimum wage increases?". Low-wage occupations tend to be populated by workers with low levels of education. An increase in the minimum wage, while designed to protect workers in the lower part of the wage distribution, might result in unintended consequences for those same workers.

This paper studies empirically how an increase in the minimum wage shapes firms' hiring decisions regarding different levels of education and technical skills of workers, exploiting 2016 changes in UK minimum-wage policy. We develop an empirical strategy that allows us to use detailed online job postings data to quantify the impact of this minimum wage increase on the labour hiring decisions of firms.

We study the minimum wage policy and its impact on labour demand composition - characteristics of job attributes - in the context of the UK local labour markets studies. In April 2016, the UK increased the National Minimum Wage (NMW) for workers aged 25 and above by 7.5%, from £6.70 to £7.20.[1] The UK is a country where approximately 10% of the labour force is affected by the recent minimum wage increase.[2] Within the low and middle skilled occupations, where wages are lower, more than 15% of online job postings offer a wage below the minimum wage in 2014 and more than 20% of people employed are affected by the change in policy using Labour Force Survey (LFS) data.[3]

Using online job postings data from Burning Glass Technologies, we document several facts about how the increase in the minimum wage alters the characteristics of online job postings. We examine firms' (local labour markets') adjustment by increasing productivity of its workforce in response to the minimum wage increases through two different channels. The first one is labour-labour substitution, we document that shares of online vacancies with low levels of education requirements fall after the increase in the minimum wage. The second one is technology-labour substitution; we show that the ads with low requirements for technical skills in low-middle skilled occupations decreases following minimum wage hikes. The intuition is that firms may respond to increased labour costs by replacing low-wage jobs with technology or other capital investments and workers who are complementary to these technologies.

[1] See Table A1 for a detailed evolution of the NMW since 2014.
[2] Computed using LFS data.
[3] We use 2014 in order to not take into account firms' expectations.

We employ different strategies to remove concerns that our findings arise from omitted factors or through reverse causality. We control for a range of fixed effects such as Travel To Work Areas (TTWA) by season-year fixed effects to capture TTWA trends, as previous work suggests that results can be sensitive to those. We also confirm the results using a difference-in-differences framework that exploits the pre-policy variation of the exposure to the shock in different local labour markets. We rely on the considerable variation in the exposure to the minimum wage across TTWA areas (Dustmann et al. (2022); Ahlfeldt et al. (2018); Adams-Prassl et al. (2020)) and use the initial share of online job postings offering a wage below £7 an hour in local labour markets (intensity), under the assumption that localities that have a large share of low-wage jobs should be more affected by the minimum wage change.

We document two facts that we observe in the data. First, in local labour markets with a high share of low-wage jobs (highly exposed to the policy change), firms typically rely on employing low educated workers. Second, we show that in areas where the share of ads below the minimum wage is high, the level of technology is quite low.

Our results suggest that minimum wage policies may be important in shaping labour market job characteristics. Our empirical analysis shows that, among occupations highly exposed to the minimum wage, i.e. middle and low skilled occupations, there is a shift away from lower levels of education within occupations, after controlling for location-time trends. Also, there is a decrease in the demand for ads with low levels of technical skills. We confirm that the technological transformation is not structural as we do not find the same downward shift in demand for technical skills within higher skilled occupations, those occupations less affected by the policy change. Our findings support the idea that employers also play a role in shaping labor market inequalities as highlighted in Card et al. (2013) and Song et al. (2018), yet using different mechanisms.

Our paper relates to the extensive literature around the minimum wage and its impact on wages and employment. The seminal work by Card and Krueger (1993) questioned an old and assumed theory that minimum wage increases were reducing employment. Using quasi-experiments in the US, their study finds no adverse effects on employment. Yet, many studies came after their work and are synthesised in Neumark et al. (2008) where they conclude that effects are heterogeneous across skill groups: the minimum wage reduces employment for less skilled workers in many countries. However, more recent studies in different countries (UK: Stewart (2004) and Manning (2012), Germany: Dustmann et al. (2022) and the US: Black et al. (2016)) seem to agree that minimum wage increases wages for the low-wage occupations without reducing employment. So far, there is very little evidence that the UK minimum wage has had an adverse effect on employment at an aggregated level.

Our paper builds on a sparse but growing literature looking at minimum wage increases and its effect on other outcomes such as non-wage compensation and other job characteristics.[4] Those can be for example health insurance, training possibilities, education and skill requirements. Recent studies documented different

[4] For a recent review of existing margins see Clemens (2021)

margins using online job vacancies, for example, Adams-Prassl et al. (2020) find that the change in the minimum wage in the UK increased the demand for flexible hour contracts (Datta et al. (2019)). Similarly, using online job vacancies in the UK, Delaney and Papps (2022) study the effect of the minimum wage on the number of posted ads, the quality of the advertised jobs, the wage offered for a *same* vacancy and the Beveridge curve. Our study complements the above by looking at different dimensions on the hiring behaviour, in particular on the education and tech requirements of ads. Moreover, using online job openings, Horton (2017) finds that firms shifted from low-productivity workers towards higher-productivity workers. Similarly to our study, Clemens et al. (2021) use BGT data for the US and they conclude that increasing the minimum wage has shifted firms demand towards older workers and higher levels of education in low-wage occupations. Dickens et al. (2015) document the gender inequalities of the minimum wage increase and they conclude that it reduced the part-time employment of women, which is a population that is highly exposed to minimum wages.

Our study builds on recent literature looking at firm responses to an increase in the minimum wage. One of the strategies for firms to compensate for the increase in labour costs is to increase their productivity. Ku (2020) finds that the increased productivity among the workers at the bottom of the wage distribution offsets about half of the expected rise in the firm's wage bill. This can be achieved either by hiring more productive workers (typically workers with higher ability) or investing in technology. Evidence for labour-labour substitution within low-skill groups can also be found in Neumark (2001).

We also contribute to the growing literature using online job postings data. (Jardim et al. (2018); Gopalan et al. (2021)) found that firms adjust their hiring to a larger extent than their firing. Therefore, using online job vacancies data (BGT) is valuable to document local labour markets responses to a change in the minimum wage. Initially based on US data and increasingly used in analysis of the UK data, this strand of the literature looked at various dimensions of the labour market, such as skill demand (Grinis (2019); Deming and Kahn (2018); Hershbein and Kahn (2018)), labour market concentration (Azar et al. (2020)) and flexible work arrangements (Adams-Prassl et al. (2020)). We add to this by looking at the education requirements and the tech level of jobs after an increase in the minimum wage. Clemens et al. (2021) explore a similar question for the US using online job vacancies. While they focus on high-school diploma, we use the job vacancy text to infer the education level of the job at higher levels. Moreover, instead of using the skill classification for 10 general skills as defined by Deming and Kahn (2018) we distil thousands of keywords using Machine Learning (ML) techniques to get a tech level for each job.

The remainder of this paper proceeds as follows. In Section 2, we describe our data, outline the classification of education arrangements, and describe the program used to identify the technical skills level of vacancies. In Section 3, we chart the variation in education requirements from firms across occupations, pay, geographical locations, and time. In Section 4, we analyse changes in firms' demand for education requirements and technical skills in response to the 2016 increase in the national minimum wage. Section 5 concludes.

2 Conceptual framework

There are many different channels through which minimum wages affect the economy. The mechanisms we describe below are still widely debated in the literature. A study by the ECB detail 6 different channels, namely, layoffs, cuts in hiring, price rises, cuts in non-labour costs, wage rises for employees earning above the minimum wage, and improvements in productivity.[5] The results of the survey show that productivity adjustment is one of the most popular reactions from firms. This paper will discuss the change in productivity using two different measures: the change in education requirements from employers and the change in tech skills.

Until the influential research by Card and Kueger (Card (1992); Katz and Krueger (1992); Card and Krueger (1993)), it was assumed that when labour costs increased, a decrease in employment was the clear implication of the theory of the firm and profit maximization under competition. Researchers, but also policymakers, have for a long time focused on wages and employment outcomes when looking at the effects of the change in the minimum wage (Clemens (2021)). It is only recently that studies have started to look into changes within jobs characteristics after those policy changes. Recent studies have shed light on different margins using online job vacancies, for example, Adams-Prassl et al. (2020) find that the change in the minimum wage in the UK has increased the demand for flexible hour contracts (Datta et al. (2019) also look at contract types). Moreover, using online job openings, Horton (2017) find that firms shifted from low-productivity workers towards higher-productivity workers. Similarly to our study, Clemens et al. (2021) use BGT data for the US and they conclude that increasing the minimum wage has shifted firms demand towards older workers and higher levels of education in low-wage occupations. Dickens et al. (2015) document the gender inequalities of the minimum wage increase and they conclude that it reduced the employment of part-time women, which is a population that is highly exposed to the minimum wage increase. In our paper, we look at the change in effort requirement measured by the education and skill content of jobs ads.

We describe the mechanisms through which a minimum wage increase can generate labour-labour substitution and labour-technology substitution. Labour-labour substitution is a substitution between different categories of labour. Here, we concentrate on the relative increase in the education level of workers and skill requirements in job postings in local labour markets and especially in low-skilled occupations. We emphasise on those occupations as they are the ones highly exposed to the NMW change. Technology-labour substitution is replacing workers with machines or technological production processes with the aim of increasing productivity and reducing the unit cost of production. The link between innovation and employment has been studied: firms use technologies to improve efficiency and generate labour-saving costs.

One important channel by which firms may adjust in the face of minimum wage increases is through substitution towards higher-skilled labour in order to increase the productivity of their workers. Doing so will allow firms to increase output to counterbalance the increase in labour costs. By hiring a higher educated

[5] Not mentioned but firms can also cut profits.

worker, firms expect to increase their productivity and obtain higher skills or workers with higher ability. We document the change in hiring in the different education levels. A second mechanism that we look at is the technology and labour substitution. The idea is that the increase in labour costs will affect firms' decisions on investment. As we do not have data for the technological level of local labour markets, we use workers as a proxy for that. The intuition is that employers looking for workers with tech skills are expected to implement innovative processes and ICT. We analyse the change in demand for tech skills in job vacancies.

In summary, our study illustrates two primary channels that could increase the education requirements from firms for workers employed in minimum wage jobs: shifts in employment towards more tech workers and the substitution of low-skilled workers for high-skilled workers. Our methodology which we explain further relies on the idea that when the NMW increases, it is more likely to impact areas that employ workers earning low wages. The more the area is exposed to the increase in the minimum wage, the more likely firms in that area may exit, reduce employment or offset the extra cost by increasing its productivity due to the increase in labour costs.

We need to keep in mind some shortcomings of our analysis. The effects could be partly biased due to the change in labour supply. Minimum wage increases may have effects not only on the hiring side of the labour market but also on the supply of workers (Acemoglu (2001); Flinn (2006); Adams et al. (2018)). Moreover, we do not look at the impact on the extensive margin, where it is possible that the non-productive firms with financial difficulties exit due to the increase in labour cost. Dustmann et al. (2022) find that the quality of German firms increases as low-wage workers move from smaller and low-paying firms to larger, higher-paying firms. Similar evidence for the US by Luca and Luca (2019) and Jha and Rodriguez-Lopez (2021) both document an exit of low-productive firms after a minimum wage increase for the restaurant sector. Mayneris et al. (2018) exploit the 2004 minimum-wage reform in China and they find that after experiencing an increase in labour costs, firms also became more productive. However, the conclusion for job seekers remains the same: an increase in the level of education in order to get a job. Another concern could be the effect of Brexit as the referendum was voted in June 2016. We provide an extensive explanation in Appendix D in the Appendix, where we conclude that our results here should not be impacted by Brexit effects. If anything, Brexit induced a decline in the posting of adverts for higher-skilled jobs, that would shift the composition of adverts in favour of lower-skilled jobs. Hence, our results are under-estimated if Brexit acts as a confounder. Finally, we do not consider that employers can replace workers without a degree with slightly younger, and cheaper, workers who are still below the relevant age threshold for the minimum wage change. Younger workers (age 18-24) represent less than 10% of the total workforce. However, Giupponi and Machin (2018) document no spillover effects on employment for workers under 25 in the UK context. Many reasons can explain this. First along with the mechanisms described above, firms hire more productive workers, more experience and education to face labour costs increase, which is not necessarily the case for the younger workers. Also, firms need to anticipate that the cost of its young workforce will jump (+20-25%) when employees reach the threshold age.

3 Data and measurement

Online job postings data offer near real-time information on labour demand and help shed light on labour market dynamics and reallocation patterns, thus complementing traditional sources of information, such as surveys collected by National Statistical Offices. With the increasing use of online platforms, employers often post online a job advert containing information about the position they want to fill but also characteristics the candidate should have. The real-time information has an added value over traditional data and statistics, as traditional data are backward-looking and may be unable to capture ongoing phenomena.

We use Burning Glass Technologies data for England, which cover the near universe of online job vacancies and are increasingly used in labour economics (for example, for the UK: Grinis (2019), Adams-Prassl et al. (2020), Javorcik et al. (2020), and for the US: Deming and Kahn (2018); Hershbein and Kahn (2018)).[6] These data are provided by a US labour market analytic company daily web-scraping approximately 7,500 UK online platforms, including job boards, government job databases, company's websites, and websites of agencies specialised in recruitment. Their algorithm enables to collect all the information from millions of current job postings, which allows us to have the world's largest and most detailed database of jobs. For the UK, we have an average of 8 million annual online job adverts from 2012 to 2019 (or approximately 50 million online job postings). We use data from 2014 to 2019 for two reasons. First, between 2012 and 2014, the number of webpages that BGT scraped for job adverts increased rapidly but has been quite stable since this date. Second, the bias in the representativeness of the data is also more stable since that date (Cammeraat and Squicciarini (2021), Adams-Prassl et al. (2020)).

Using the full text of the job vacancy, BGT categorise them according to variables, such as geographical location, occupation, industry, required skills, education and experience levels. In the analysis we use both the full text of the ad and some already coded variables. Because of differences in the education system within regions in the UK, we concentrate on England only, therefore reducing the sample approximately by 15% ads over the period covered.[7][8] Our analysis is thus based on the full vacancy text of over 31 million unique online job postings between 2014 and 2019. Table 1 presents some summary statistics about the numbers of vacancies and the percentages of non-missing values in each year in the BGT dataset. About 99% of ads contain the UK SOC Code, 65% the wage information and slightly less than 100% the location at the TTWA level information.[9] We use this sample for the construction of our variables but for the analysis, we use the sample of ads that have at least the occupation and wage information provided, which leaves us with approximately 20 million ads. As is obvious from the table, the percent of ads with information on qualification from BGT is quite low. To address this we used text analysis to construct an alternative

[6] Non-exhaustive list.

[7] It reduces both because ads are not from England and because some ads do not have the country information. Table A2 in the Appendix provides descriptive statistics of the coverage of BGT data before restricting to England only.

[8] Scotland has its own qualification structure with differences in the names of qualifications making our text analysis not possible. For example, National Vocational Qualifications in England are called Scottish Vocational Qualifications.

[9] Some ads had the country information but not necessarily the TTWA information.

measure of qualification, as explained in the section "Construction of the education variable".

Table 1: Descriptive statistics, Jan. 2014 - Dec. 2019 BGT sample

Year	# ads	Job Title	TTWA	UK SOC Code	Hourly Salary	Education (BGT)	Education (our)
2014	3838043	100.00	99.90	97.96	64.46	16.28	28.41
2015	5133397	100.00	99.91	98.24	66.03	15.35	25.90
2016	5756695	100.00	99.81	98.41	67.03	14.69	25.54
2017	6305359	100.00	99.53	98.32	63.15	15.18	26.30
2018	5521480	100.00	99.80	99.58	62.70	16.51	27.27
2019	4498679	99.99	100.00	99.91	64.00	18.11	27.72

Notes: The table reports the number of observations and the share of ads with the following information non-missing: job title; Travel To Work Area (TTWA); occupation (4-digit UK SOC); salary (minimum in the ad); education levels from BGT and education level from our classification.

Source: BGT, 2014-2019.

There are a couple of main concerns when using BGT to analyse labour market dynamics. First from an aggregate point of view, even though there is an increasing use of the Internet to post job vacancies, the traditional sources such as newspaper ads, signs on shop doors etc. are still in use and are not accounted by the BGT database. However, it is unlikely that firms spread vacancies ads through only one source. Second, and more likely, there might be a disproportionate representation of industries, occupations and location, which has been found in many studies using BGT data in the UK. Third, there is no information on if the vacancy has been filled, which makes it difficult to use the number of ads in BGT data to quantify labour market demands. In the following analysis, we use BGT data to identify changes in educational requirements from employers, therefore, if biases are constant over time, we are able to correctly identify changes. Hershbein and Kahn (2018) state that to analyse differences in trends over time, as long as the distributions are relatively stable across time, then we can draw conclusions. Not all variables in BGT are complete for each posting, but we can consider them as "true" missing as Hershbein et al. (2018) have analysed a sub-sample of the original text of job postings for the US to verify that the reported missing field was not initially given. Lastly, even though BGT applies a de-duplication method[10] for ads that appear on multiple websites, it remains possible that one posting covers multiple job openings. Due to all those caveats described above, we will remain cautious in the interpretation of our results.

Grinis (2019) shows that "despite all these shortcomings, occupational and geographic distributions in the BGT data exhibit high correlations with the occupational and geographic distributions of official UK employment data (the Annual Survey of Hours and Earnings (ASHE) from the Office for National Statistics (ONS))".

[10] If two similar vacancies ads appear several times on the same or different platforms within a period of two months, one is removed as a duplicate.

Recall, this study illustrates two channels which firms might use to compensate for the labour cost increase for workers employed in minimum wage jobs: the labour-labour substitution of higher-skilled workers for the least-skilled workers and shifts in employment towards more tech workers. Therefore, in the next two subsections, we describe how we construct the two variables of interest using BGT: education requirements and technical skill level of ads.

3.1 Construction of the education variable

Our aim is to retrieve all vacancies which explicitly mention an educational requirement in the job text in our set of BGT adverts. BGT provides an educational variable that covers approximately 16% of ads. We use text analysis that relies on direct keyword matching with additional cleaning strategies. Data sources such as LFS also provide information on educational attainment of employees. Our study of online job vacancy data explores issues of the labour market demand from a different perspective, the perspective of the employer. More specifically, it looks at qualifications and skills that employers would like to see in new recruits. Information from BGT and individual survey data such as LFS are thus not entirely comparable. A major difference between the two is that job vacancies describe characteristics of an 'ideal' candidate for the job as defined by the employer, and not characteristics of the person who is already on the job. There already exist multiple job classifications in the literature, but they mainly focus on graduate jobs. For example, Green and Henseke (2016) create a graduate occupation classification based on tasks; see also, Elias and Purcell (2013). However, those indicators use surveys meaning that respondents are already on the job. Following this methodology does not allow us to identify correctly what the firm seeks as the indicators are based on the share of employees' education level. We use the education level mentioned in the job advert, yet, we do not have information on who gets the job. We believe our method partially accounts for the supply side effects that could potentially drive our results. We can imagine that firms' write in job descriptions their wish and that changes in the supply could affect the actual hiring, but not what they initially seek. For example, if a firm posts an ad with an education requirement below graduate level but a high share of applicants have a graduate degree or above, obviously firms' hiring decision can be driven by the supply of job-seekers.

We retrieve all vacancies that clearly mention the minimum educational requirement needed to get the job from the set of BGT adverts by using the text of the job vacancy and directly matching some specific keywords related to qualifications. The initial step is to extract the education specific keyword from the text of the ads and the 4 words before and 4 words after that keyword. The additional words that we extract give us some context which we use to apply cleaning and filtering strategies. The first cleaning strategy extracts the exact level of the ad using both numerical and character information around our keyword. For example, "nvq level 2" or "masters degree" will be classified as education level 2, and education level 7, respectively. Second, we flag ads that mention multiple education requirements. This step allows us to identify ads for multiple positions within the same posting. For example, if the job posting requires both a low level of education and a PhD, we drop that ad from our sample . For closer levels of education, we keep all the information but also create a variable with the lowest level mentioned. The intuition behind this is that employers stating multiple educational levels would have a minimum level for the job but would prefer a higher one. We apply other cleaning strategies using the context of the ad, such as ads mentioning "degree of" would not be classified as a degree level.[11] Finally, we suspect that employers would not mention explicitly the educational requirement for some occupations because it is straight forward that a specific qualification or a legal education requirement is needed. For example, registered nurses are required to have specific level of education such as a bachelor of science. Lawyers need to take a Legal Practice Course for which you need to have a bachelor. Therefore, we manually annotate the ads for those jobs where a degree or higher are formally required. Figure 1 provides an example of an ad and summarises the different steps of the classification program.

Figure 1: Example of an ad for a CNC technician and different steps of the classification program

"[...] currently seeking a senior cnc technician or works engineering technician to join their team in huyton you will be $\boxed{\textit{hnc}}$ *or* $\boxed{\textit{btech}}$ *trained or equivalent have experience in working within a busy works engineering department the role will be responsible for maintenance across the site you will be responsible for [...]"*

Steps:

1. direct matching of education specific keywords in the job text.

2. cleaning strategies (using context around keywords): extracting numerical values (e.g. nvq 2, nvq four etc.), removing irrelevant structures which are misleading (e.g. high degree of, master *a topic* etc.).

3. identifying jobs that have a mandatory education level using the job title of ads.

4. Classifying ads - the example above is classified as both level 3 (BTEC) and level 4 (HNC).

Using the sample with non-missing information for occupation, wage and location, Table 2 presents summary statistics about the numbers of vacancies and the percentages of non-missing values for the BGT education variable and the education variable we constructed in each year. The last column shows that our classification program allows us to identify education levels for approximately 30% of job vacancies data, while the BGT variable covers about 16%. Even though we double the coverage, the identified ads are not necessarily the same, especially because of the English targeted education system we implement and we are talking of millions of ads, therefore resulting in *a lot more* ads with the education requirement. Not only are we able to cover a larger share of ads, we are also able to target English specific education requirements which are relatively different than the US education system. Table A3 in the Appendix shows that among our sample of ads for which we are able to identify an education requirement, only half of them are covered by the BGT education variable. Moreover, we account for a large set of occupations that legally requires a

[11] The term "degree of" would typically not refer to qualifications. For example it can used to describe characteristics of the future worker such as "the person should display a high degree of integrity".

specific education level which often employers consider as obvious. It is reassuring to see that the coverage of the education variables are very similar than the coverage on the full sample in Table 1, suggesting that we are not introducing a bias by selecting ads with the non-missing wage information. Further descriptives on the education variable are provided in the next section.

Table 2: Descriptive statistics, Jan. 2014 - Dec. 2019 BGT sample

Year	Number of ads	BGT edu (%)	Our Edu Var (%)
2014	2426155	16.12	28.30
2015	3333374	15.50	26.17
2016	3798891	14.59	25.01
2017	3911550	14.98	25.98
2018	3443367	15.84	26.54
2019	2876829	17.45	27.31

Notes: The table reports the number of observations with non-missing information for: Occupation (2-digit UK SOC), Hourly salary (minimum in the ad); TTWA. In the 2 last columns it reports the percentage of ads that the BGT education variable and our education variable cover.

Source: BGT, 2014-2019.

3.2 Construction of the tech level of ads

Irrespective of the occupations to which they belong, we want to identify tech jobs as those whose vacancy descriptions contain "tech keywords" - knowledge and skills that are typically required for a tech job. Intuitively, recruiters using tech keywords when describing the job are certainly looking for a candidate using tech tasks even if they do not explicitly post a tech occupation job. The widely used paper by Harrigan et al. (2018) use "techies" as a measure of Information and Communication Technology (ICT) adoption at the firm level. They define "techies" as occupations that are "closely related to the installation, management, maintenance, and support of ICT, as well as product and process design and longer-term R&D activities". They show that these techies are a good measure for technology adoption within French firms. Those techie occupations are strongly dominated by digital and programming skills which what we are trying to capture in the job vacancies data. We provide in Table A5 in the Appendix the equivalent occupations that are defined as techies at the UK 4-digit SOC code that we use in our data.

We summarises the different steps of the tech classification program in Figure 2. By using their classification of ICT occupations and the occupation information in BGT data, we label ads that are *for sure* looking for a techie. Then, we create a tech index at the skill level. One advantage of the method we describe below is that it allows us to have a tech measure at the ad level conditional on the skills used by the employer to describe the job and not from the occupation classification of the ad. For example, some ads with the following job titles: "Data Architects", "Big Data Engineers" and "Digital Marketing Data Scientist" are classified as high level of technical skills based on their skills and using our algorithm but they are not classified as a techie using the occupation classification.

Our approach consists in two main steps that we describe in more detail below. First, we perform a cluster analysis identifying the skills that are predictors of techies occupation in 2014. Second, we use that pool of skills to identify all other techie ads and compute the probability an employer will hire a tech worker.

Figure 2: Different steps of the tech classification program

Steps:
1. direct matching of occupations information in job vacancies with the occupation classification from Harrigan et al. (2018) (equivalence Table A5).
2. compute a level of technical skills at the keyword level.
3. compute various metrics at the ad level using the information from keywords.
4. compute the probability an employer will hire a tech worker based on a logistic regression using the metrics from **3.**.

3.2.1 How to identify tech skills

The BGT taxonomy of skills contains about 10,000 distinct keywords.[12] The keywords are a mix of knowledge, personal characteristics and skills required to perform various tasks on the job. Some of the keywords correspond to skills and knowledge required to use software and technological devices. We label as "tech" keywords those that refer to knowledge and skills typically required for a techie job. Following Grinis (2019), we use "context mapping" where the key idea is to classify keywords based on their "techiness", i.e. the percentage of techie ads in which a keyword appears. Note, techie ads are defined using the classification by Harrigan et al. (2018) (HRT techies). After computing the "techiness" of all keywords, we use a clustering technique to separate them into techie, neutral, and non-techie. In the results that we show here, we keep in our sample all the set of skills that appear at least twice in 2014. This filter reduces the sample of skills from 10,120 to 9,311 keywords. However, we have performed robustness checks where we keep the whole set of keywords. Also, our results are robust when removing skills that appear less than 2000 times per year.[13] Moreover, the more ads we have the better the classification will be, therefore, we work on a larger set of ads than the analysis, i.e. contain the occupation, education information We decide to perform the skill techie classification using data for 2014 so in the period preceding the increase in the adult minimum wage in 2016. We do so because one might expect a change in the skill description, and so different skills keywords after the policy enters.

We describe below the different steps in more details:

Step 1: We merge *all*[14] online job vacancies from BGT data for 2014 with the classification of techie

[12] See Carnevale et al. (2014) and Grinis (2019) for an explanation on how the original advert text is transformed into sets of keywords, which are stored in the skill variable in the BGT data.

[13] All robustness checks are available in the Appendix

[14] The sub-sample described above.

occupations based on the classification by Harrigan et al. (2018). This allows us to obtain the true value of techie occupations. We compute the average distribution of techie occupations among all the occupations in which skill k appears. Inspired by Grinis (2019), we define this as the techiness of a keyword (x_k) which is simply the proportion of techie occupations in the whole set of ads where k appears.

Step 2: Using the techiness of each keyword (step 1), we implement a k-means clustering where we specify both the number of centers and their initial locations. Researchers tend to not pre-determine the number of clusters and their location. However, in our analysis, the number of clusters and their locations are driven by the type of information we wish to extract. Therefore, we choose them ahead of launching the k-means clustering. In our context, we want to classify keywords into techie, neutral and non-techie. The pre-determined centroids are therefore 0 (100;0), 0.5 (50;50), and 1 (0;100) corresponding to 0% techie (non-techie cluster), 50% techie (neutral cluster) and 100% techie (techie cluster).

Figure 3 shows examples of the top 100 most frequent keywords from the resulting clusters. The method we use is not perfect, nevertheless, it does have the advantage of systematically classifying all the set of keywords in the BGT taxonomy, including many technical terms. More importantly, as Figure 3 shows, the resulting classification does seem fairly plausible. For instance, "Transact-SQL" is a tool used to interact with relational databases, and is correctly classified in the techie cluster. Moreover, "C++" and "Javascript" are programming languages used for web development (and many other applications) and are also allocated to the techie cluster. It is highly probable that firms looking for workers knowing those skills are going to deal with ICT related tasks. Note that keywords like "problem solving", "writing" and "data analysis" all appear in the neutral cluster. This is precisely because employers might not only ask for hard-tech skills even within techie ads but also many other soft skills that complement tech skills and they can also be required within non-techies jobs. Those skills are not tech or non-tech specific, they can be found in various occupations. The non-techie keywords, classified in the non-techie cluster, are mainly soft skills, oriented towards communication, social and people skills, such as "product sales", "people management" and "business planning". Interestingly, "computer literacy", "Microsoft excel" and "spreadsheets" are also included in this category. This could arise because tech ads do not mention those, as they are obviously already learned/operated with more complex tech skills and because nowadays most jobs require these skills.

Figure 3: Examples of neutral (top-left), techie (top-right) and non-techie (bottom-left) keywords

Notes: Graphs show examples of the top 100 most frequent keywords from the the three resulting clusters: neutral (top-left), techie (top-right) and non-techie (bottom-left) collected from English online vacancies and classified using context mapping and clustering. A couple of very frequent keywords are dropped from the plot for scaling issues. Size and color are by frequency of being posted. The bigger the size, the more frequent the keyword appears in ads.
Source: BGT 2014-2019.

Table 3 provides further details on the distribution of tech skills within each cluster identified, based on the whole set of keywords. For instance, the mean techiness of non-techie, techie and neutral keywords are 5.10, 85.30, and 41.50 respectively.[15]

[15] Among the 1914 words identified as in the techie cluster, the average keywords appears 85% of the time in a techie add as defined by the occupation classification, i.e. techiness.

Table 3: Techie, non-techie and neutral clusters (whole set of keywords)

Cluster	Mean	Max	Min	Median	Number of Keywords
Non-techie	5.10	23.20	0.00	2.00	6421
Techie	85.30	100.00	63.40	85.70	1914
Neutral	41.50	63.40	23.30	40.50	1785

Notes: The table reports a set of summary statistics from the classification of 10,120 keywords into the techie, the neutral and the non-techie clusters. The column reports the mean, the maximum, the minimum, the median and the number of keywords in those clusters. We repeat the analysis restricting the set of keywords to words appearing at least twice in the data, and Table A9 shows that the distribution is very similar.
Source: BGT, 2014-2019.

We test the classification carried out above using a test and train sample. Note that the tests conducted here are in-sample, i.e. the sample with explicit techie requirement and keywords used to evaluate our job classification strategy is the same sample as the one used to compute the techiness and classify keywords The verification is used to double check that our keywords are good predictors of techie ads. We do that at this stage because if our set of tech keywords are good predictors of tech occupations, we can then in the next step confidently use the set of tech keywords to estimate a tech level for each ad. To perform the test, we compute at the ad level different metrics, such as the mean of the techiness of all keywords within an ad, the share of keywords in the techie cluster etc. The maximum of correct classification is achieved using the mean techiness of keywords in an ad, which allows us to correctly classify above 89% of ads as non-tech and tech ads (see Figure A1 in the Appendix). Overall, the output of this step is to have for each keyword the share of appearance in a techie ad in 2014 (i.e. "techiness" of the word) and its cluster number. We have 3 clusters, non-tech skills, neutral skills and tech skills.[16]

3.2.2 From tech skills to tech ads

An ad in the BGT data is a set of keywords. There are many ways one can think of going from the tech level of a keyword to the tech level of an ad. Grinis (2019), even though focusing on STEM jobs, describes many different methods and their pros and cons. For example, she mentions that one of the simplest ways of identifying an ad looking for a STEM background is to label those sets that contain at least one STEM keyword as STEM and the rest as non-STEM. Intuitively, since we identified tech keywords in the section above, which are a mix of software use, skills and tasks, the presence of a tech keyword in an ad description could well indicate that the employer is likely to employ someone for a tech position. However, interpretation could be challenging, as all ads with one or more tech keywords would have the same tech level (strong assumption) and second, keywords have not the same level of tech even within the tech cluster, as for example "Microsoft Back Office" and "Data Warehousing" have a lower level than "node.JS", "JavaScript

[16] We then check the classification on the other years and find high accuracy.

Object Notation (JSON)" and "Laravel".[17]

The method we use to identify the level of techniness of an ad is to merge each keyword's techiness and cluster number as defined in the previous section to the whole set of ads for the period 2014-2019 that we use in the analysis. We then produce multiple metrics at the ad level: mean, median, min, max of techniness and share of techie keywords. The mean (median, min and max) techiness of an ad is equal to the mean (median, min and max) techiness of the keywords used in the description of the ad. The share of techie keywords is computed so that it stands between 0 and 1 where 1 means that all keywords in an ad description is a techie skill. Table 4 shows descriptives concerning the distribution of the techie variable at the ad level using again Harrigan et al. (2018)'s classification (HRT techies). Reassuringly, the figures show that the average of the mean techiness of ads in the techie occupations is significantly larger than the non-techie occupations, 48.5 and 12.3 respectively, and this is true also for the other metrics we use. We also provide in the Appendix the distribution of that variable, and we show that it has a long right tail. Yet, it is difficult to interpret and to define the threshold for a low, medium and high tech ad. One reason making interpretation difficult is that the tech level of different keywords is computed using samples of different sizes, therefore, taking a simple average gives equal weight to all keywords in the job description. To avoid this issue, we go further and we compute the probability that the employer is looking for a tech worker using a logistic link function. We then use the estimated relationship to predict probabilities of being a tech ad on the whole sample. Finally, we classify jobs into different tech categories if the probability is above a certain threshold.

Table 4: Distribution of the different techie metrics for techie and non-techie occupations

Techies (HRT)	Median Techiness	Mean Techiness	Percentage Number of Keywords	Max Techiness
Non-Techie	10.64	12.32	1.75	22.71
Techie	47.10	48.51	40.43	73.17

Notes: The table reports descriptives concerning the distribution of the techie variable for two groups of occupations based on the classification by Harrigan et al. (2018): non-techie and techie occupations. The numbers are the average over the ads for the following metrics: median, mean, percentage number of keywords, and maximum techiness in the techie and non-techie occupation sub-samples.

Source: BGT, 2014-2019.

3.2.3 From tech ads to hiring a tech worker

There are many ways we can compute the probability of the employer looking for a tech worker using the skill content of the ad.[18] The method we choose for our study uses the techiness at the ad level. The model utilizes the metrics (mean, median, max and min) computed above as predictors in a logistic regression model.

[17]"Laravel is a free, open-source PHP web framework, created by Taylor Otwell and intended for the development of web applications following the model–view–controller architectural pattern and based on Symfony" - Wikipedia.

[18]We use only one method here, a second method uses the information at the keyword level to estimate the probability of looking for a tech worker.Grinis (2019) compares the different methods in the context of identifying STEM ads. . The method at the keyword level is much more intensive computationally than the method we use.

Recall, the metrics above employ the techiness of all keywords in a vacancy description to compute them. Therefore one of the advantages of this method is that we do not need to have the "document-term" matrix, which in our case would be high dimensional and sparse. The model is simple and not computationally intensive because the predictive relationship is built from just a couple of continuous predictors (i.e. different metrics) instead of several thousands of dummy variables (all the different keywords) if we were to use other methods. Interestingly, Grinis (2019) compares both methods and finds that the logistic regression with all the keywords as predictors has a lower performance than the model with only 2 predictors, i.e. two metrics at the ad level. Yet, Adams-Prassl et al. (2020) use the second method with all the keywords as predictors to retrieve the level of flexibility of jobs.

The method uses the techiness metrics as predictors in a regression that models the probability of hiring a tech worker. To ensure that the predicted probabilities lie between 0 and 1, we use a logistic link function and estimate the following regression on our sample of online job postings:

$$logit(\theta_i) = log\left(\frac{\theta_i}{1-\theta_i}\right) = \beta_0 + \beta_{x_i} X_i \tag{1}$$

where θ_i is the probability of looking for a techie for online job posting i and X_i are the different metrics for online job posting i. After estimating the logistic model, we use the estimated coefficients β_{x_i} to estimate the probability $Y = 1$ of looking for a tech worker for a given ad i. We explore different models where we include different sets of regressors. Our preferred specification is:

$$logit(\theta_i) = log\left(\frac{\theta_i}{1-\theta_i}\right) = \beta_0 + \beta_1 mean\_tech_i + \beta_2 max\_tech_i \tag{2}$$

Results from estimating Equation 2 show that the coefficients for "mean techie" and "max techie", β_1 and β_2, are positive and statistically significant (the full table of results can be found in Table A10 in the Appendix). The sign of the estimates implies that an increase in those variables can be translated into an increase in the probability of an employer looking for a tech worker. To test our model, we split the data into a train and test sample, and the classification prediction accuracy is just above 92%.[19] Importantly, the histogram in Figure A3 in the Appendix shows that the predicted probabilities match the distribution of the actual techie occupation much better than raw techie mean (Figure A2).[20]

Overall, the output of this step is that for each job we have its tech level, i.e. the probability that the employer is looking for a techie worker. In the analysis, we will also define thresholds to classify ads: very low, low, medium and high tech level. The underlying assumption is that the higher the probability of looking for a tech worker, the more intense the job is in technical skills/tasks. The thresholds are defined using the distribution of the probability variable. As an illustration, we extract 30 random ads from the dataset with

[19] Using median instead of mean in the logistic regression results in an accuracy not significantly different than the one presented in the main text

[20] All figures showing different distributions (e.g. the "true value", etc.) can be requested to the authors.

their job title and associated probability of hiring a tech worker and the mean techiness in Table A8. Job titles with the lowest levels of technical skills are Social Workers, Shop Managers and others, while job titles with the highest levels are Junior Data Analyst, Electronics Technician Medical Device and Senior Developer. Some occupations that we would find in the middle range of technical skills are Electronics Repair Engineer, Support Technician and Field Service Engineer.

3.2.4 Final Sample for the analysis

We use the sample of ads with non-missing information for the wage, and occupation. We remove the 0.001^{th} percentile and 99.99^{th} percentile of the wage distribution to remove a few outliers. In addition to the above filters, in the analysis we also restrict the sample to firms that are always postings ads throughout the period of analysis, to make sure our results are not driven by firm composition effects for example due to firm exit after the unexpected labour cost increase. As minimum wage changes increases the labour cost for firms employing those workers, it is possible that some firms exit the market. Jha and Rodriguez-Lopez (2021) are the first ones to document a negative relationship between minimum wage and the mass of firms in the US for the restaurant and retail-trade industries. However, focusing on the UK care homes industry, Giupponi and Machin (2018) find no effect of an increase in the NMW on firm closures. By restricting the sample to always-posting firms, our results will speak to the set of firms that can cover the minimum wage increased costs and potentially adapt by changing their hiring behaviour. This further reduces the sample to 2,573 firms each year posting a total of 10,548,821 ads over 2014 to 2019.

We focus on two groups of education, non-graduates and graduates. We define non-graduates as all low-levels of education such as level one, two and three in academic levels (GCSE and A levels) or National Vocational Qualifications as the highest qualification. The graduate category is defined as ads where the description mentions only level six in academic levels (graduate) or above (masters and PhD). In this analysis we do not consider ads of level four and five (HNC, NVQ), neither ads mentioning both graduates and another lower level such as levels four and five (HNC, NVQ). Overall, those ads represent less than 10% of the sample of ads with education requirements.[21]

4 Stylised facts

For the descriptives, we focus on vacancies from always-posting firms that have the wage, occupation, location and education information. We show below that descriptives remain the same when changing the sample.[22] There is considerable variation in the proportion of different education requirements in vacancies over occupations, wages, and space. Turning first to occupation variation, we report in Figure 4 Panel (a) the share of non-graduate vacancies over the whole sample by one-digit SOC codes. The figure shows that non-graduate

[21] Our method allows us to identify more disaggregated education categories, yet we leave this for future research.
[22] Descriptives using only the full set of firms are very similar.

job postings are concentrated in low-middle skilled occupations and graduate ads are located in high skilled occupations.[23] The share of vacancies reporting a non-graduate education requirement ranges from about 75% for "Skilled Trades occupations" to less than 5% for "Professionals occupations". Figure 4 Panel (b) show that online job postings with low levels of education requirements are significantly more likely to be advertised in low-wage jobs. Around 35% of jobs advertising a wage less than £6 per hour were requiring non-graduate levels and the share is even higher for jobs offering a wage between £6 and £9. Those shares are high compared to the top of the wage distribution where less than 10% of vacancies with a wage greater than £14 per hour also require a low education level.

Figure 4: Share of ads with a non-graduate requirement

(a) By 1-digit occupation

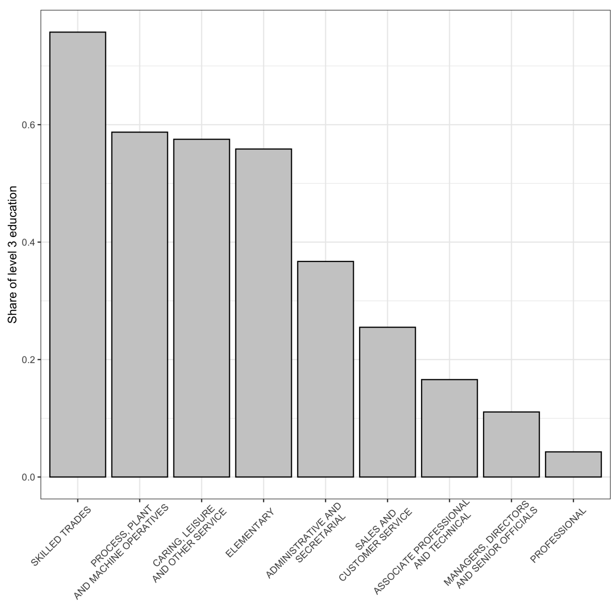

(b) By wage group

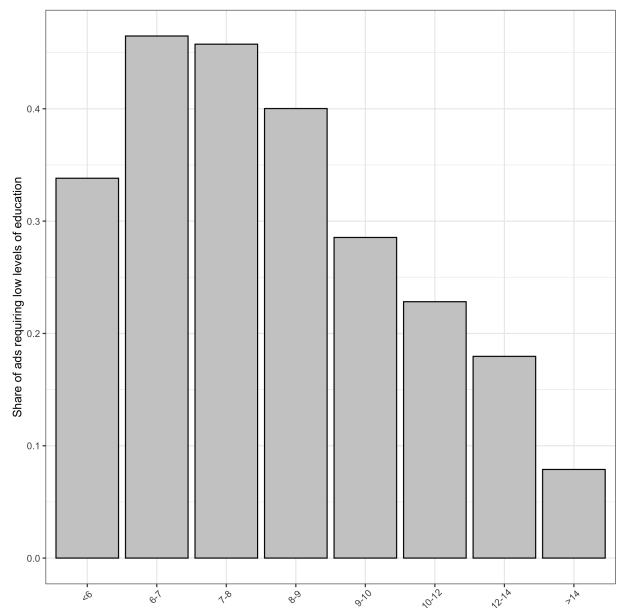

Notes: The figure documents the distribution of ads with a non-graduate requirement across occupations and across wage groups. Panel (a): the bars show the proportion of non-graduate vacancies by one-digit occupations over the 2014 to 2019 BGT sample (see Table A4 for a classification of one-digit occupations). Panel (b): the bars show the proportion of non-graduate vacancies by wage group over the 2014 to 2019 BGT sample.
Source: BGT, 2014-2019.

We explore the variation in the tech level of ads in Figure 5. In Panel (a) we plot the average probability of hiring a worker with technical skills by one-digit SOC occupations. Professional activities have the highest average probability of hiring a tech worker, followed by Associate Professional and Technical activities. While those two occupations belong to the *high skilled* occupation category, the third occupation with the highest average tech level belongs to the *low skilled* category. The occupations with the lowest levels are, as expected the service workers. In Panel (b) we see a clear positive correlation between the probability of hiring a tech workers and the offered wage in the online job postings. While the average probability of hiring a tech workers

[23]Definition of the categories can be found in the Table A4 in the Appendix.

in jobs offering more than 14£is 15%, it is three times lower for jobs offering less than 9£.

Figure 5: Tech level of ads (probability of searching for a tech worker)

(a) By 1-digit occupation

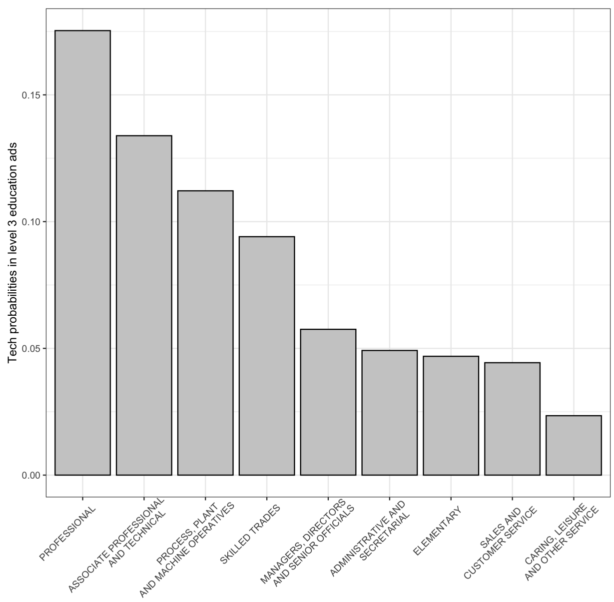

(b) By wage group for non-graduates

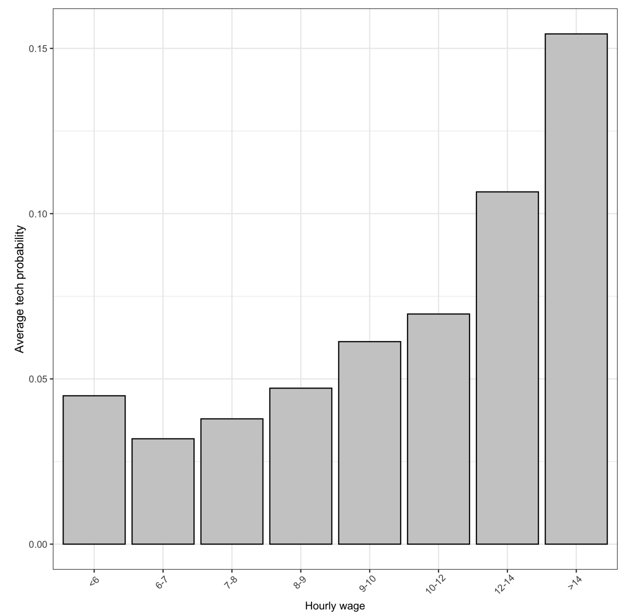

Notes: The figure documents the tech level of ads distribution across occupations and across wage groups. Panel (a): shows the average technology probability of job postings at the one-digit occupation level (see Table A4 for a classification of one-digit occupations) over the 2014 to 2019 BGT sample. Panel (b): shows the average technology probability of job postings for non-graduates by hourly wage bins over the 2014 to 2019 BGT sample. *Source:* BGT, 2014-2019.

We turn to documenting the correlation between exposure of local labour markets to the change in the minimum wage and the areas requirements in terms of education and technology.

Fact 1: *In local labour markets with a high share of low-wage jobs in 2014 (before the change in policy in 2016), education requirements in online jobs postings are relatively lower.*
Fact 2: *In local labour markets with a high share of low-wage jobs in 2014 (before the change in policy in 2016), technology levels in online jobs postings are relatively lower.*

Figure 6 plots the share of ads requiring a non-graduate level, the share of ads offering a wage below the minimum wage and the average tech level before the announcement of the change in the policy in July 2015. Our first fact is that in local labour markets with a high share of low-wage jobs in 2014 (before the change in policy in 2016), are typically areas relying on low educated workers. The correlation is 0.59, meaning that areas with high shares of ads below the minimum wage are also areas with a high share of low education requirements. Second, we show that there is a negative correlation (-0.53) between the share of ads below the minimum wage and the level of technical skill requirements. Local labour markets with a high share of

Figure 6: Share of ads below the minimum wage, with low education requirements and technology level

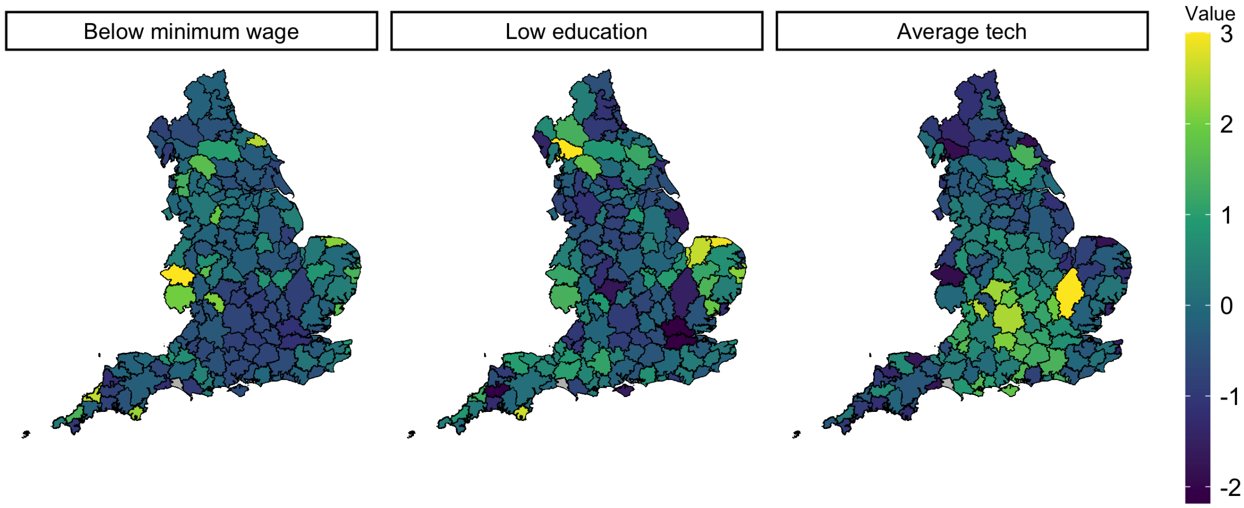

Notes: The maps display, at the TTWA level, the share of ads below the minimum wage (left), the share of ads with non-graduate education requirements (center) and the average probability of hiring a tech worker (right) ahead of the policy change. Values are standardised with an average of 0 and standard deviation of 1 for comparison purposes and are computed on the sample of ads from 2014 to July 2015 as this is when the change in policy was announced.
Source: BGT 2014-2019.

ads exposed to the minimum wage change are also areas that have a low level of technology requirements on average. In the UK, labour skill composition is heterogeneous across local labour markets. It is possible that some areas produce the same goods but with a different combination of non-graduates and graduate workers. The spatial composition of those indicators in Figure 6 is not surprising. The London areas is the area where the share of ads with low education requirements is the lowest. Kollydas and Green (2022) document the regional movements of graduates and find that London is the only area above the average in terms of keeping and attracting young graduates, highlighting the fact that graduate jobs tend to be concentrated in London and the locations around. In terms of the technical skills requirements in ads, we see again that higher levels are concentrated in the Southern locations.

Figure 7 plots the share of ads requiring graduates and non-graduates over time for low and middle skill occupations. The graph suggests that when the policy started, the share of ads looking for non-graduates decreased. In the period before the policy, it seems like there was a small decrease in the share of ads requiring a graduate level.[24]

[24]The minimum wage policy change occurred almost simultaneously with the results of the 2016 Brexit referendum, which might affect our results. We investigated the extent to which regions exposed to the policy might also have been affected by Brexit, and find a very low correlation between the two measures of exposure. If anything, Brexit induced a decline in the posting of adverts for higher-skilled jobs, that would shift the composition of adverts in favour of lower-skilled jobs. Hence, our results are under-estimated if Brexit acts as a confounder. See Appendix D for a detailed discussion.

Figure 7: Share of ads with graduate and non-graduate requirements for middle and low skilled occupations

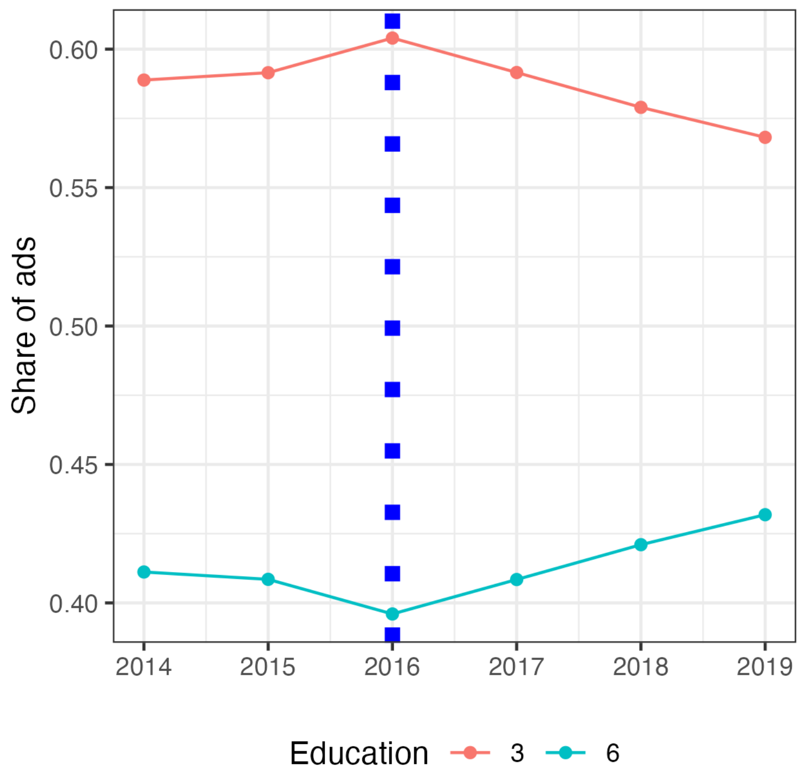

Notes: The figure plots the share of ads with graduate (in blue) and non-graduates (in red) education requirements over middle and low skilled occupations (SOC one-digit 4 to 9). We remove from the sample ads with an hourly rate above £12. The dashed blue line stands for the implementation of the minimum wage policy.
Source: BGT 2014-2019.

5 Education, technology and the minimum wage

The section above highlighted that vacancies requiring a low level of education are concentrated in low skilled occupations and in low wages. Also, we show that there is spatial variation in the share of ads requiring low wage, low education and technical skill levels. To explore the relationship between firms' labour costs and the demand for education and tech skills, we analyse the impact of a large and unexpected increase in the national minimum wage in April 2016 on the prevalence of graduates vacancies. The National Minimum Wage Act 1998 established a binding minimum wage across the UK, with the National Minimum Wage (NMW) taking effect on 1 April 1999. Our analysis uses the large increase in the NMW in April 2016. The minimum wage rose by 7.5% from £6.70 to £7.20.[25] Datta et al. (2019) show that when the policy got announced in July 2015 it was largely unexpected. This unanticipated and significant wage shock provides a quasi-experiment to analyse its consequences. We suspect that the sudden increase in labour costs for firms will have an effect on the employment structure, both in terms of the level of education and the skills required.

[25]See Table A1 for a detailed evolution of the NMW since 2014.

5.1 Impact on offered wages

We start by confirming that the 2016 increase in the minimum wage led to a shift at the right side of the distribution of posted wages. After the policy implementation, firms were not allowed to offer wages below £7.20. Therefore, we look if we see in the data a drop in the share of vacancies advertising a wage below £7, and an increase in the share of vacancies offering more than £7. We regress the share of different wages categories within each 1-digit SOC occupations over time (regressions are not at the ad level) following this specification:

$$S_{o,w,t} = \beta_1 post * I_w + \text{control}_t + \alpha_o + \alpha_s + \alpha_t + \epsilon_{o,w,t} \tag{3}$$

The variable of outcome $S_{t,o,w}$ is the share of ads in each o occupation (1-digit SOC) and wage categories w at time t (monthly). We use the shares of vacancies at different education levels rather than their absolute number given our interest in the distribution rather than the overall labour demand. control_t stands for a dummy variable for the time between the announcement and beginning of the policy. We add occupation α_o and season α_s by year α_t fixed effects as well as a control for the time between the announcement and the start of the policy.[26]

[26] We run specifications with α_o and month α_m by year α_t fixed effects. Results show very similar results.

Figure 8: NMW impact on posted wages, by occupation group

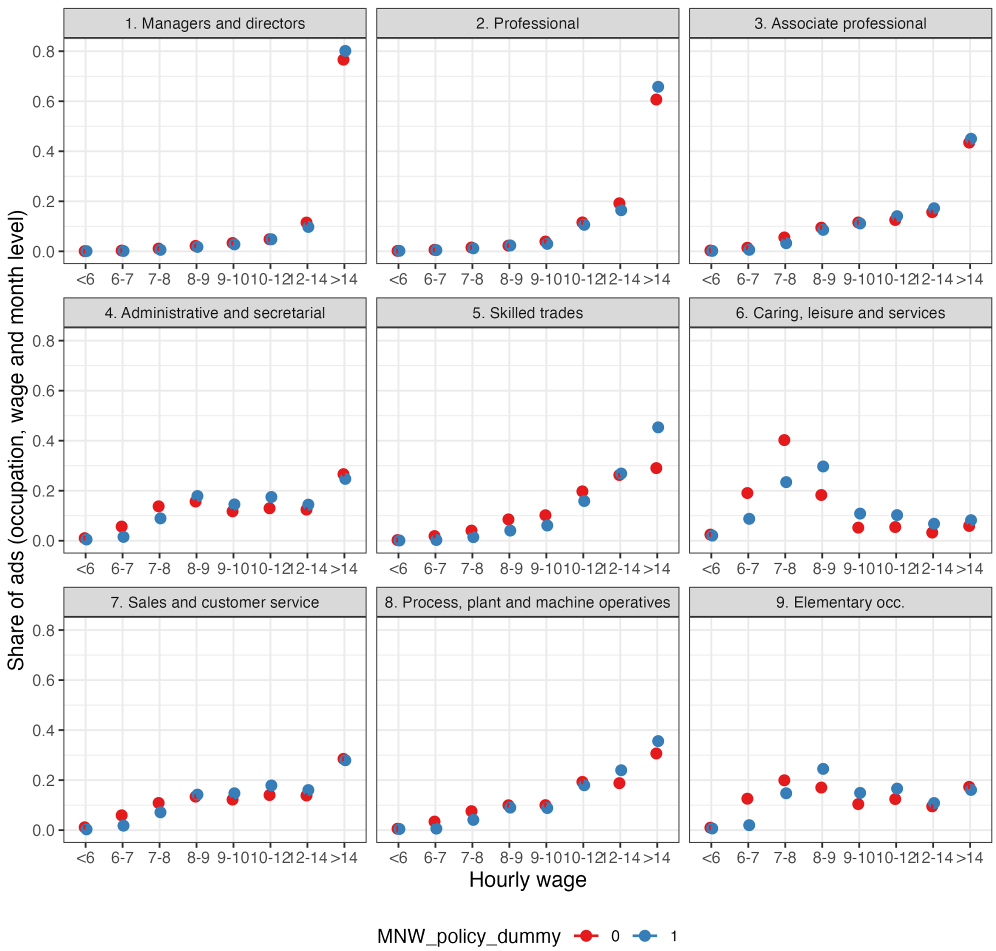

Notes: The figure gives estimates of β_1 from Equation 3: estimates of the treatment effect of NMW increase on wage bins within each occupation group. The blue dot stands for post-policy period, and the different between the blue and the red dot is the estimated coefficient. The error bars for the 95% confidence interval are not visible because estimates are precise. Each panels are a different occupation group at the one-digit level. The first three panels in the top row are high skilled, second row the middle skilled and the last three panels are for low skilled occupations. The estimates are identified by comparing the share of vacancies in a given wage bin before and after the NMW introduction, controlling for occupation and season by time fixed effects for each wage bin, as well as for the time between the announcement and beginning of the policy.
Source: BGT 2014-2019.

Figure 8 shows that there is a clear and significant drop in the number of ads below the new minimum wage, especially for low-skilled jobs. Quite surprising, there is not a clear and significant increase in ads just at the new minimum wage (7-8 category). However, there is also a statistically significant increase in ads in the wage bin £2 above the new minimum and modest increases in the £3 and £4 bins above the new minimum. This pattern of wage changes is consistent with the analysis by Adams-Prassl et al. (2020). We have done the same analysis on the reduced sample with education requirements and on different wage bins and we get similar results. This is reassuring as it confirms that our results are not driven by the wage bins we have chosen and the size of the sample.

Figure 9: NMW impact on posted wages, by education group

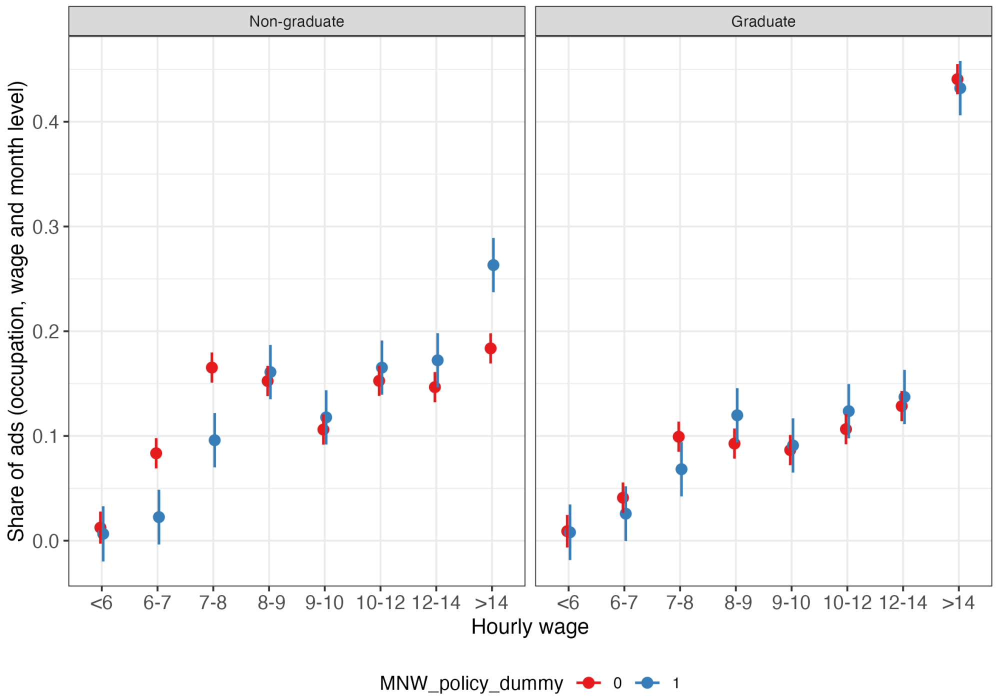

Notes: The figure gives estimates of β_1 from Equation 3 estimates of the treatment effect of NMW increase on wage bins within each educational group. The left panel is for the non-graduate and the right panel the graduates. The blue dot stands for post-policy period, and the different between the blue and the red dot is the estimated coefficient. The error bars show the 95% confidence interval. The estimates are identified by comparing the share of vacancies in a given wage bin before and after the NMW introduction, controlling for occupation, season and time fixed effects for each wage bin, as well as for the time between the announcement and beginning of the policy.
Source: BGT 2014-2019.

We also perform this analysis at the education level, by using the same specification in Equation 3 but we change the variable of outcome to the share of ads ($S_{e,w,t}$) in each education e and wage categories w at time t. Figure 9 plots the results. The figure shows no effect of the minimum wage policy for ads looking for graduates (right panel), i.e. we do not see a significant shift in the distribution rightward. However, the left panel shows an increase in posted wages for non-graduate jobs, which is what the NMW aims to induce. However, those aggregated figures mask high levels of heterogeneity and it is difficult to draw any conclusions. It could be that the observed decrease is not fully compensated by an increase in higher wage ads. Moreover, Figure 8 and Figure 9 are graphical evidence of the compliance with the minimum wage policy. They show that employers actually comply with the change in minimum wage policy as we see that the increase in the minimum wage has an effect on the wage distribution, bringing to 0 the share of ads paying below £7. In the literature, the minimum wage is considered to be binding if it causes a spike around the minimum wage in the wage distribution (Lemos (2009); Dickens and Manning (2004)).

5.2 Impact on education requirements

Strategy. Following many papers' methods, we exploit the considerable variation in the exposure to the minimum wage across TTWA areas in the UK. A similar method was used on German data in Dustmann et al. (2022), Ahlfeldt et al. (2018) and on UK BGT data in Adams-Prassl et al. (2020). The method consists of tracing how local share of demand for education evolve in TTWA differently exposed to the minimum wage, due to differences in pre-policy local wage levels.

To capture the extent to which a local labour market has been affected by the substantial increase in the minimum wage, we rely on an exposure measure where we use the initial share of online jobs posting offering a wage below £7 an hour in local labour markets (intensity), under the assumption that localities that are more intensive in low-wage jobs should be more affected by the minimum wage change. We can then construct a degree of exposure measure of each occupation in each local labour markets in England.

The minimum wage bite is the percent of workers paid below the minimum wage prior to the policy change (Card and Krueger (1993); Machin et al. (2003)):

$$NMW_{o,l,2014} = \frac{adsHwage < £7_{o,l,2014}}{ads_{o,l,2014}} \tag{4}$$

Equation 4 is equal to the share of ads paying below the minimum wage in 2014 in a two-digit occupation-location. It gives us a measure for each occupation by TTWA of its exposure to the change in the minimum wage policy. We repeat the same methodology using LFS data for robustness. LFS data allows us to calculate the share of employment with a wage below £7 an hour in regional labour markets, but, due to data constraints, we can only do so at the regional and 1-digit SOC occupation.

We leverage the variation in exposure across occupations and local areas for our identification strategy. The intuition of the DiD framework is that we are comparing online job postings which are offering wages below the minimum wage with job postings offering above, and therefore the latter is not affected by the change in policy. It is important for our empirical strategy that there are no geographic spillovers across TTWA, yet we believe that this is plausible in our framework as our outcomes are firms' hiring. The BGT data has the information on the location of the job and not the location of the HQ of the firm hiring. Therefore, geographic spillovers should be very small.

We acknowledge the existence of the new Difference-in-Differences (DiD) estimators (Callaway and Sant'Anna (2021), de Chaisemartin and D'Haultfœuille (2020))[27], but our framework fits well in the standard DiD strategy where we have treated occupation-TTWA cells (highly exposed) receiving the treatment (change in minimum wage) and our comparison group (not exposed cells) which are not receiving the treatment. We further show in the results that there are no signs of pre-trends, as our outcomes in both groups are not

[27] Non-exhaustive list, refer to Roth et al. (2023) for a full discussion on the new estimators.

followin significant different paths before the policy change. The recent DiD literature aims at relaxing some assumptions, such as multiple treatment periods, and in different time periods. Those are not applicable here where the change in the minimum wage happens once and is national. The new estimators also allows researchers to relax the parallel trend assumption. While this could be very useful, we test this assumption and we cannot find evidence of violation.

Figure 10: The variation in NMW bite

(a) Across TTWA, by occupations

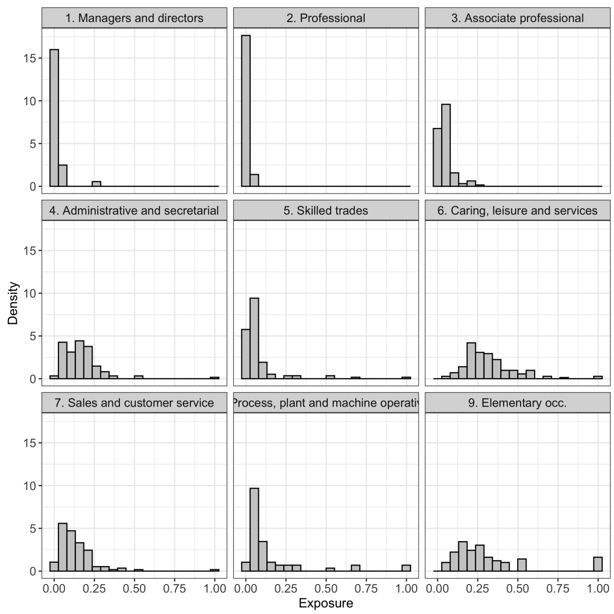

(b) Across TTWA

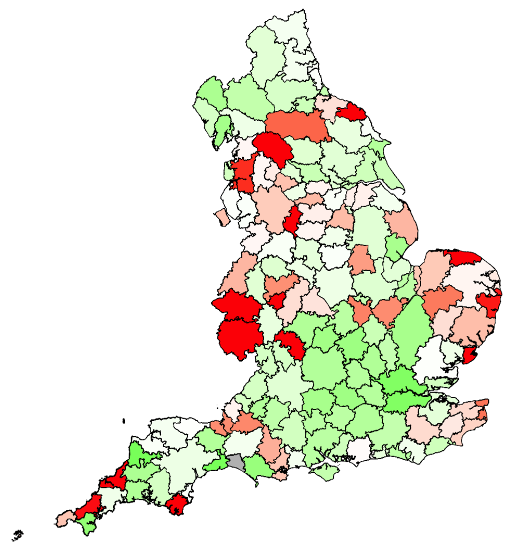

Notes: The figures are descriptives of the minimum wage bite constructed in Equation 4. In Panel (a) we compute the bite at the one-digit level and each panel shows the distribution across TTWA of the share of vacancies that are paid below the NMW (minimum wage bite) for a given occupation. In Panel (b) we compute the measure at the TTWA level and the map displays the spatial distribution, the greener the area, the lower the share of online ads posting wages below the minimum wage in 2014, on the contrary, the redder the larger that share is.
Source: BGT, 2014.

Figure 10 plots the variation in NMW bite across TTWA, by occupation. The NMW bite is close to 0 for many TTWA for managerial and technical occupations. However, it is higher among lower-skilled occupations that are more exposed to the minimum wage increase as they are often lower wage occupations, as we saw in the descriptives. Panel (b) shows that London and surrounding TTWA are areas where the share of job vacancies offering a wage below the minimum wage in 2014 is low. However, regions in the North and the South West are highly exposed to the change in the minimum wage policy as the share of ads posting wages below the minimum wage threshold is high relative to the other areas. This variation is similar to the considerable spatial heterogeneity in education, wage and technical levels we described in the descriptives (Figure 6).

Specifications. Using the identification strategy detailed above, our main specification is:

$$S_{e,t,o,l} = \beta_1 NMW_{o,l,2014} * Dpolicy_t + \beta_x X_t + FE + \epsilon \tag{5}$$

The dependent variable $S_{e,t,o,l}$ captures the proportion of vacancies in an occupation-education-TTWA-year-month cell that mentions an education requirement $e \in \{graduate, non-graduate\}$. The variable NMW stands for the exposure to the change in the minimum wage at the TTWA-2-digit occupation level in 2014. *Dpolicy* is a dummy that equals 1 from April 2016 onward and 0 otherwise. β_1 is the estimated average treatment effect on the treated for a particular education level. It is identified from the interaction between our measure of treatment intensity $NMW_{o,l,2014}$ and our temporal dummy *Dpolicy*. The regression includes a control ($\beta_x X_t$) for the time between announcement of the policy and its implementation in order to account for employers expectations. We add year×season×TTWA and occupation fixed effects and different combinations of them depending on the specification. Our preferred specification includes TTWA×year×season and occupation fixed effects as it absorbs any common variation within a TTWA-time that is due to the shock. The standard errors are two-way clustered at the occupation-TTWA level.

Because our identification strategy is based on the idea that once we control for region fixed effects and occupation-specific time shocks, hiring within a given occupation in high-wage regions is a good control for hiring in the same occupation category in a low-wage region. To improve the comparability of occupation-specific hiring, we use 3-digit SOC occupation categories.

In addition to the aggregated result above, we run an event-study type of regression. By denoting $\delta_{q,t}$ (i.e. year-quarter dummies) the dynamic response of the outcome $S_{t,o,l}$ at time $t+k$ to the policy event at time t, we run the following:

$$S_{t,o,l} = \beta_1 NMW_{o,l,2014} \times \delta_{q,t} + \beta_x X + FE + \epsilon \tag{6}$$

Our event study allows us to confirm the causality of our results if we find no specific trends in the periods pre-policy. Our outcomes should have no significant different trends in high exposed area relative to low exposed areas. To make sure our results are not driven by firm composition effects for example due to firm exit after the unexpected labour cost increase, we perform the analysis on the sample of firms that are always posting ads throughout the period of the analysis. This further reduces the sample to 2,573 firms each year posting a total of 10,548,821 ads over 2014 to 2019.

Results

Table 5: Impact of the change in the minimum wage policy: overall occupations

	Share of non-graduate ads (month, TTWA, SOC2)			
	(1)	(2)	(3)	(4)
Constant	0.3495***			
	(0.0643)			
Exposure (NMW)	1.259***	0.1511***	0.1489***	
	(0.1915)	(0.0445)	(0.0442)	
MNW Policy Dummy	0.0038	0.0532***	0.0504***	0.0326***
	(0.0079)	(0.0150)	(0.0148)	(0.0086)
Expect Policy	-0.0189***	-0.0044	-0.0047	-0.0103*
	(0.0051)	(0.0073)	(0.0071)	(0.0054)
Exposure (NMW)× NMW Policy Dummy	-0.1306***	-0.1040***	-0.0934***	-0.1086***
	(0.0283)	(0.0329)	(0.0326)	(0.0330)
Year×Season FE		Yes		
TTWA FE		Yes		
Occupation FE		Yes	Yes	
TTWA×Year×Season			Yes	
Season×TTWA×Occ				Yes
Year				Yes
Observations	88,212	88,212	88,212	88,212
R^2	0.25294	0.68525	0.70129	0.75563

Notes: The table displays the results from Equation 5 of the monthly share of ads at the occupation-education-TTWA level on the policy dummy interacted with the exposure to the minimum wage with different sets of fixed effects. Column (1) has no fixed effects, column (2) includes time, TTWA and occupation fixed effects, column (3) includes TTWA time trends and occupation fixed effects and column (4) year and TTWA by year by occupation fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.
Sources: Burning Glass Data 2014-2019.

Specification 1. The results in Table 5 show that the more the occupation-TTWA cell is exposed to the minimum wage increase, the larger the adjustment is in required education, i.e. reduction in the share of ads requiring less than a graduate level. In column (1) we do not include any fixed effects and the coefficient shows that the change in the minimum wage led to a 13 percentage point decrease in the proportion of non-graduate vacancies. In the subsequent columns, we gradually add different combinations of fixed effects which reduces slightly the coefficient relative to Column (1). We see that the increase in the minimum wage decreases the share of non-graduate ads by 10 percentage points in Column (2), 9 percentage points in Column (3) and 11 percentage points in Column (4). The stability of the coefficient of interest to the inclusion of different fixed effects signals that we are able to capture a causal estimate of the minimum wage. The coefficients for the other variables included have the expected sign. For example, the persistent negative and significant coefficient for the exposure suggests that the more exposed areas-occupations have a higher share of non-graduate ads. In terms of magnitude, our estimate suggests that for a sum of 133,979 ads in highly exposed areas, of which 18% are non graduate ads, therefore, a reduction of 9 percentage points means that the share of non-grad ads is now 16.4% suggesting a drop of 2144 ads (which is of same amplitude as dropping an average highly exposed TTWA per year).

The descriptives in the previous section show that ads paying around the minimum wage are concentrated in low and medium skilled occupations. Therefore, we perform the same analysis but on the sub-sample of middle and low-skill occupations.[28] Those categories of occupations are most likely to be impacted by the change in the minimum wage. Table 6 show that the effect of the minimum wage policy is more pronounced for non-graduates in middle and low skilled occupations. The coefficients shows that the minimum wage increase reduced the share of job postings for non-graduate workers by 12 percentage points in Columns (2) and (3) and 13 percentage points in Column (5). Those results are not surprising as low-middle skilled occupations are in general populated by workers with lower level of education due to the task content of the job.

Overall, those results show that the increase in minimum wage resulted in a simultaneous increase in wages (above section), and a decrease in the demand for non-graduate workers. There is evidence in the US that suggests minimum wages reduce the jobs available to low-skill workers (Clemens et al. (2021)). The results suggest that employers move away from less-skilled workers, by raising the educational requirements in online job vacancies after the minimum wage change. This finding is relevant from a policy point of view as the minimum wage is intended to help the least-skilled workers who are often receiving low wages. If the demand for non-graduate workers declines substantially then the aim of the policy can be questioned.

[28] Definition of the categories can be found in the Table A4 in the Appendix.

Table 6: Impact of the change in the minimum wage policy: low and middle skilled occupations

	Share of non-graduate ads			
	(1)	(2)	(3)	(4)
Constant	0.6482***			
	(0.0741)			
Exposure (NMW)	0.3516*	0.1157**	0.1222**	
	(0.1641)	(0.0408)	(0.0422)	
MNW Policy Dummy	0.0095	0.0537**	0.0518**	0.0485***
	(0.0122)	(0.0210)	(0.0213)	(0.0109)
Expect Policy	-0.0109**	-0.0036	-0.0041	-0.0060
	(0.0047)	(0.0121)	(0.0117)	(0.0050)
Exposure (NMW) $\times$ NMW Policy Dummy	-0.1072**	-0.1169***	-0.1154**	-0.1286***
	(0.0393)	(0.0361)	(0.0384)	(0.0381)
Year$\times$Season FE		Yes		
TTWA FE		Yes		
Occupation FE		Yes	Yes	
TTWA$\times$Year$\times$Season			Yes	
Season$\times$TTWA$\times$Occ				Yes
Year				Yes
Observations	49,365	49,365	49,365	49,365
R^2	0.02614	0.39683	0.44170	0.51141

Notes: The table displays the results from Equation 5 of the monthly share of ads at the occupation-education-TTWA level on the policy dummy interacted with the exposure to the minimum wage with different sets of fixed effects for the sample of low and middle skilled occupations only. Column (1) has no fixed effects, column (2) includes time, TTWA and occupation fixed effects, column (3) includes TTWA time trends and occupation fixed effects and column (4) year and TTWA by year by occupation fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.
Sources: Burning Glass Data 2014-2019.

Specification 2: event-study. Figure 11 reports the results of the location-occupation-level event studies. We find that the share of non-graduate ads decreases in location-occupations after the announcement of the policy using the variation in exposure to the change in the policy. Panel (a) implements the event study with Season$\times$TTWA$\times$2-digit Occupation and year fixed effects over all occupations while panel (b) is the same specification but on low and middle-skill occupations.[29] Focusing on Panel (b) where the effects are mainly taking place, the response of the share of non-graduate ads is non-significant but shows a declining trend from the quarter when the policy was announced. The response is then negative, significant and amplified when we reach the quarter when the policy was actually put into place 2016 of about 15 percentage points. Considering that low educated workers are less productive, those results are in line with Horton (2017) showing firms shifted from low-productivity workers towards higher-productivity workers. Also, our findings are similar to the ones in Clemens et al. (2021), where using BGT data for the US, they conclude that increasing the minimum wage has shifted firms demand towards older workers and higher levels of education in low-wage occupations.

The non-significant coefficients before the dates of the announcement and the actual policy suggests that there is no sign of pre-trends conditional on the controls and fixed effects included in our regression. Location-occupations that were more exposed to the potential change (i.e. larger share of below-minimum wage ads) at a given time were on a comparable path in years prior to the event and started diverging only afterwards. This finding restricts the potential set of confounders that could explain the decrease in the share of non-graduate ads. Therefore, we can interpret our estimates as causal effects.

Figure 11: Graphical results

(a) All occupations (b) Low and Middle Occupations

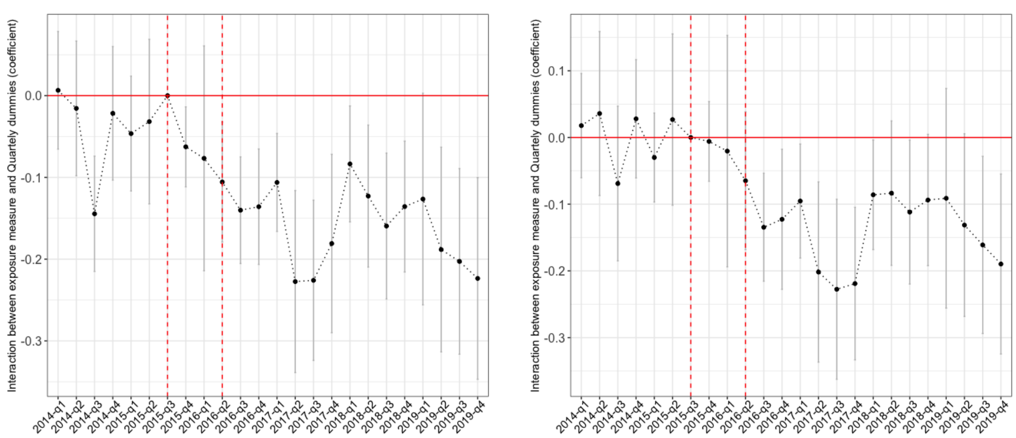

Notes: Results of an event study analysis (see Equation 6), using the set of fixed effects of specification (4) in Table 5 with quarterly dummy interactions. The error bars show the 95% confidence interval using clustered standard errors at the location and two-digit occupation level. The red dash bars show the time of announcement of the change in policy and when it was implemented. Same figure but with a different set of fixed effects is available in Figure A5.
Sources: BGT 2014-2019.

Robustness Checks. We conduct a series of robustness checks in order to examine how our coefficients of interest behave when the regression specification is modified. First, we perform the same analysis using LFS data to construct the minimum wage bite. More precisely, we compute the share of employment for each occupation-region that have a wage below the minimum wage rate before the policy. The intuition

[29]The same figure but with a different set of fixed effects is available in Figure A5.

remains the same, the higher the share, the more exposed an occupation is in a specific region to the change in minimum wage. Due to data constraints, we compute the bites at the regional-SOC 1-digit occupation level. We compare this variable to the same level of aggregation using BGT data, and find that they are highly correlated (0.81).[30]

Moreover, in order to address the concern that our results are driven by time varying unobservables that are not absorbed by time fixed effects, we conduct a placebo test. We introduce a break 2 years after the actual timing of the policy, which took place in April 2016. We then run the same exercise and look at our coefficient of interest: the interaction between a post-April 2018 (placebo policy) dummy with the local exposure measure. Table A11 in Appendix shows the results. The coefficient of interest is not statistically significant for all specifications suggesting that the effect on the share of non-graduate ads was not observed after (prior for the placebo introduced before) to the policy, but driven primarily by it. We also perform the analysis removing London TTWA from the sample to check that our results are not mainly driven by this area. Results are presented in Table A12 in the Appendix and the coefficients are extremely close to the baseline results in Table 6. Therefore, we conclude that results are not driven by London. Moreover, we perform a random allocation of the location-occupation exposure to the minimum wage policy. Our results are in Table A13 where we see that none of the coefficients are statistically significant. Finally, one might think that the sample coverage of BGT increases over time. However, BGT updates the historical data when updating the codes. If anything, the representativeness for non-graduates should be improving and BGT should be capturing more of those ads.

Our results could relate to spatial inequalities across UK TTWAs. Figure 10 shows that the intensity of treatment is heterogeneous across the UK. Our results suggest that the higher the intensity, the more labour-labour substitution. Those areas are generally populated by workers with low levels of education, so the minimum wage increase could exacerbate unemployment of workers from the bottom part of the education distribution. Therefore, inequalities could be accentuated. Our results confirm theories suggesting that employers shift towards more productive workers to compensate for the increase in labour costs. We show the first mechanism that we consider in this section, where the change in productivity is captured by the education level of workers. In the next section, we turn to our second mechanism which is the change in the level of technical skills demanded after an increase in labour costs.

5.3 Impact on technical skill requirements

In the previous sections, we showed that the increase in the minimum wage resulted in a simultaneous increase in wages, and a shift away from lower educational levels. In this section, we document the second mechanism which is the substitution towards technology in order to compensate for the increase in labour costs. Therefore, we want to investigate the change in the demand for technical skills. We apply the same

[30] Descriptives comparing the variation in NMW bite across region, by occupation for LFS and BGT can be found in Figure A6 in the Appendix.

restriction to the data as for the previous section; we restrict the sample to firms that are always posting throughout the period of analysis.

Relationship between technical skill jobs and wages In this section we first examine whether or not tech jobs are associated with higher wages in the labour market. We run simple linear regressions on the overall sample, and including interactions with our education variable:

$$log(w_i) = \beta_0 + \beta_1 Pr(tech|k_i) + \beta_2 Educ + \beta_3 Pr(tech|k_i) \times Educ + \beta_X X_i + \epsilon \tag{7}$$

Where w_i is the hourly wage of job posting i, $Pr(tech|k_i)$ is the probability that the employer for vacancy i seeks a tech worker conditional on the keywords k_i collected from i's online job advert and $Educ$ are dummy variables for the education requirement of the ad (non-graduates and graduates). We include different sets of fixed effects in X_i such as the month and year of the posting, the location and the occupation. In some specifications we include time × location (occupation) fixed effects to control for location (occupation) time trends.

The results are presented in Table 7. The coefficient in Column (1) suggests that a non-graduate ad will offer a lower wage relative to a graduate ad of about 28 percentage points for a same 3-digit occupation job. There is a wage premium of being a graduate as already documented in the literature.[31] This finding is robust to the inclusion of occupation time trends. Column (2) suggests that technical skills are positively correlated with offered wages within occupations. The higher the probability of looking for a techie based on the keywords it contains the higher the wage. We further explore the tech wage premium in column (3) and (4) where we split the sample into graduate job postings and non-graduate job postings, respectively. Finally, in Column (5) we repeat the same exercise on the full sample but we add an interaction term for the probability of looking for a tech worker (based on tech keywords) with a dummy for the graduate requirement in the ad. The coefficient for the interaction term is not statistically significant, suggesting that there are no wage premium in ads requiring both a graduate education requirement and technical skills. A few explanations could account for this finding. For example, only non-graduate workers get a wage premium when performing jobs requesting technical skills because the set of technical skills within an occupation for the different levels of education are not so different. Those results are not driven by the occupation as we run the specifications including occupations by time fixed effects. Our results are comparing ads within the same occupation controlling for time trends.

We create four categories for the ads based on their tech level: very low tech are ads with a tech probability in the first quartile of the distribution (p25: 2%), low level ads are in ads with a level above the p25 and below p75 (p75: 11%), medium tech corresponds to ads between the 75th and 90th percentile (p90: 71%), and high tech is defined as the top 10% of observations. The ads with probability of looking for a techie below

[31] The wage premium increased for some time in the UK in 1950-1970s (Machin and McNally (2007)), but is now flat (Blundell et al. (2021)).

Table 7: Relationship between probability of hiring a tech worker and wages

	Hourly Wages (log)				
	Full Sample (1)	Full Sample (2)	Graduate Sample (3)	Non-Graduate Sample (4)	Full Sample (5)
Graduate Dummy (GD)	0.1778***				0.1770***
	(0.0229)				(0.0271)
$Pr(tech\|k_i)$ (log)		0.3008***	0.2528***	0.3300***	0.3539***
		(0.0810)	(0.0886)	(0.0655)	(0.0679)
$Pr(tech\|k_i)$ (log) $\times$GD					-0.1012
					(0.0956)
Occ.$\times$Month-Yr FE	Yes	Yes	Yes	Yes	Yes
TTWA FE	Yes	Yes	Yes	Yes	Yes
Observations	2,871,345	2,871,345	2,350,414	520,931	2,871,345
R^2	0.36812	0.35997	0.31257	0.43217	0.37389

Notes: The table displays the results from Equation 7 of the log hourly wage at the ad level on a dummy for graduate ads, the log probability of hiring a tech workers and the interaction of both. All specifications include occupation time trends and TTWA fixed effects. Columns (1) and (2) are over the full sample of ads, column (3) the graduate ads, column (4) the non-graduate set of ads and the last column is including the interaction term over the full sample. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation (three-digit SOC) level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.
Sources: Burning Glass Data 2014-2019.

the 25^{th} are treated as the omitted group in the table presenting the results. These results are summarised in Table 8. For all specifications, the estimated coefficient of interest increases in magnitude as the tech level of the ad increases: amongst all vacancies, the relationship between the wage offered and tech level is positive and significant. We split the sample between graduates and non-graduates in columns (2) and (3) respectively. Interestingly, there is not much of a change in wage for graduates in ads with a medium or high tech level. However, in the non-graduate sample, the relationship between the wage offered and tech level is almost twice more as we move from the third quartile of the probability of being a tech distribution to the top 10%. In the last column, we repeat the same but interacting the different tech categories with a dummy for the education requirement in the ad. Similarly to the results in Table 7, there are no significant effects on wages when an ad is both looking for a graduate workers with technical skills.

All the results of the association between technical skill levels in ads and offered wages are robust to different set of fixed effects, different threshold for the tech categories and when using the average techiness of the ad rather than the probability of hiring a tech worker. Note that these results are not causal as there could be an unobserved omitted variable, but they are in line with the literature on the graduate/college wage premium (Lindley and Machin (2016)).

Minimum wage impact Now we turn to our main question where we want to document if the minimum wage increase gives rise to a shift away from jobs with low levels of technical skills. Is the minimum wage increase shaping the type of work people do by accelerating automation and consequently changing their tech

Table 8: Relationship between probability of hiring a tech worker (categories) and wages

	Hourly Wages (log)			
	Full Sample (1)	Graduate Sample (2)	Non-Graduate Sample (3)	Full Sample (4)
Low Tech Prob	-0.0021	0.0005	0.0077	
	(0.0113)	(0.0116)	(0.0071)	
Medium Tech Prob	0.1191***	0.1200***	0.0922***	
	(0.0177)	(0.0183)	(0.0153)	
High Tech Prob	0.1448***	0.1238***	0.1780***	
	(0.0397)	(0.0405)	(0.0540)	
Very Low Tech Prob				-0.0074
				(0.0183)
Medium Tech Prob				0.0927***
				(0.0157)
High Tech Prob				0.1816***
				(0.0486)
Graduate Dummy (GD)				0.1672***
				(0.0207)
Very Low Tech Prob × (GD)				0.0062
				(0.0285)
Medium Tech Prob × (GD)				0.0281
				(0.0198)
High Tech Prob × (GD)				-0.0589
				(0.0626)
3-digit Occ.×Month-Yr FE	Yes	Yes	Yes	Yes
TTWA FE	Yes	Yes	Yes	Yes
Observations	2,871,345	2,350,414	520,931	2,871,345
R^2	0.36091	0.31431	0.43166	0.37536

Notes: The table displays the results from Equation 7 of the log hourly wage at the ad level on a dummy for graduate ads, categorical variables for the probability of hiring a tech worker and the interaction of both. All specifications include occupation time trends and TTWA fixed effects. Columns (1) is for the full sample of ads, column (2) the graduate ads, column (3) the non-graduate set of ads and the last column is including the interaction term over the full sample. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation (three-digit SOC) level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.
Sources: Burning Glass Data 2014-2019.

requirement in ads? If this is the case, it is likely that the effects will be captured more in the graduate set of online job postings as in our sample the techie jobs are more populated by graduate ads. The increase in the labour costs firms are facing can increase their capital expenditures towards more tech-oriented capital, therefore leading them to hire less workers with low levels of technical skills. Therefore, this mechanism suggests a decreasing share of low level of technical skills.

Here we use a similar specification equivalent to Equation 6 where the outcome variable is the share of ads with a low probability of searching for a tech worker among the graduates. We report results in Figure 12 and the corresponding table in the Appendix (Table A14). From the figures, we see that the minimum wage policy decreases the demand for jobs with low levels of technical skills workers among the set of ads for graduate. This finding is more convincing in Panel (b) among the low and middle skilled occupations where there are no signs of pre-trends. The response of the share of low technical skill level ads is non-significant but shows a negative and significant point estimate of 15 percentage point one quarter after the policy was implemented. The response is then amplified one year later in 2017 with a point estimate almost twice as large reaching nearly 30 percentage points. The figures also suggest that the effect is immediate and lasting. The estimate for the share of low-tech level ads within high-skilled occupations is non-significant. This finding is reassuring as it suggests that our results are not capturing a structural shift across the occupational distribution. Those results are in line with the predictions made in previous sections: areas facing a labour cost increase are shifting towards more productive workers or workers with the minimum set of technical skills to works with more capital/ICT intensive firms.

Figure 12: Graphical results

(a) All occupations (b) Low and Middle Occupations

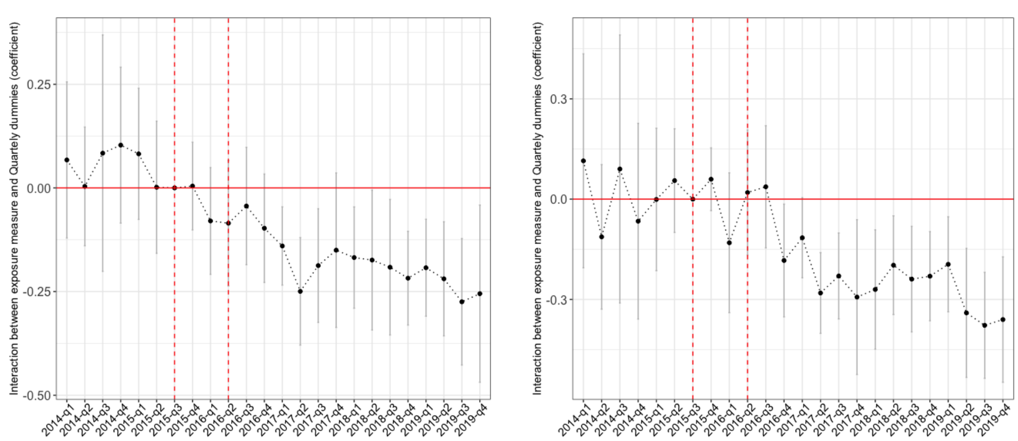

Notes: Results of an event study analysis (see Equation 6 with tech outcomes), using the set of fixed effects of specification (3) for Panel (a) and specification 6 for Panel (b) in Table A14 with quarterly dummy interactions. The error bars show the 95% confidence interval using clustered standard errors at the location and two-digit occupation level. The red dash bars show the time of announcement of the change in policy and when it was implemented.
Source: BGT 2014-2019.

For robustness, we explore the same specification in levels. We do not display the results in the paper but there is a robust increase in high tech ads in levels, while no significant effects from medium to low levels of tech ads. Moreover, in order to emphasise that our results support the idea that firms shift towards technological tools we show that the minimum wage increase impacts jobs at risk of automation. We find that the share of online job postings in occupations at high risk of automation decreases after the minimum wage is introduced, while no effect is found on the other set of occupations. This results supports our labour-capital mechanism, where firms substitute their production process with technology.[32] This result is in line with findings in Lordan (2019) where she uses UK employment data and shows that the change in the NMW decreases the share of low-skilled workers in low-wage jobs in automatable employment. Our result can be seen as complementary to hers as we show that firms' hiring behaviour changes in response to the change in the NMW.

Figure 9 shows that graduate ads are not directly concerned by the change in the minimum wage. However, we find a negative effect of the increase in the minimum wage on low technical skill ads within graduate online job postings. We conclude that there seems to be an indirect effect on graduates' skill requirements within those graduate roles via the technical skill shift in demand.

6 Discussion

It is natural to ask how these analyses connect. Our analysis shows that the minimum wage increase causes a shift in the composition of labour demand. Results show a shift towards more educated workers and, for graduate ads in low-middle occupations, a shift towards more tech jobs. The effect on the share of non-graduate workers seems to take place immediately after the policy announcement even though we see significant effects only from the actual implementation of the policy. There is a stronger effect taking place one year after the policy. Our results suggest that the effect on the education requirements are long lasting. Comparing to the results on the technical skill requirements, we can see that the decrease in low technical skill level takes slightly longer to materialise, however, we see robustly long-lasting effects.

Those findings are in line with recent studies showing that to compensate a labour cost increase, firms are increasing productivity. We focus on two channels: increasing education levels and/or shifting towards technology. On the one hand, the less educated workers are impacted by a reduction in the set of ads posted for their levels of education, and on the other hand, graduate workers are not necessarily impacted in the set of ads offered but in the changing technical skill requirements. In a recent paper, Vogel (2023)'s model shows that their is a direct effect for workers earning below the new minimum wage which we also find and nicely fit the story of the change in education requirements. And an indirect channel for workers higher up the wage distribution as there are spillovers, which we document by showing a change in the technical skill requirements of graduate workers. Overall, firms seem to compensate this increase in labour cost across the skill distribution.

[32] Detailed methodology and results can be found in Table A15 the Appendix.

Our framework does not allow us to disentangle between the extensive and intensive margins of the effect because we do not use firm-level data[33], yet we try to speak to the intensive margin only by restricting the sample of ads posted by firms always active in hiring over the period of the analysis. Results suggest that effects are concentrated in low-medium skilled occupations and those are occupations where the minimum wage is likely to have more important effects because a larger share of workers/ads are paid below the minimum wage before the policy change relative to high skilled occupations, i.e. they are more exposed.

There are several potential explanations for the changes we observe in the posted education requirements and technical skills. The main mechanism on the demand-side explanation is that firms are looking for more productive workers in response to increased labor costs (Mayneris et al. (2018)). A second channel that we do not consider in our study is that there may also be a change in the composition of firms that hire, due to firms exiting and the ones remaining employing higher skilled workers. Yet, even if we do not disentangle those two effects, the conclusion from the employee's side remains the same: they are facing an increase in education and tech level requirements from employers. Another main concern is the change in graduates supply and share of people with more tech skills during our period of analysis. We argue that firms do not anticipate the availability of graduates vs non-graduate that comes from the increase in supply. Firms post the vacancy with the adequate level of education and technical skills for the position and they might hire workers with higher level of education or tech as job applicants are more skilled than expected. If this holds, using the requirements from the job postings would not be driven by changes in the supply.

Note that due to the increase in the labour cost, the least productive firms may exit the market and incentives for new firms to enter may decrease. Jha and Rodriguez-Lopez (2021) are the first ones to document a negative relationship between minimum wage and the mass of firms in the US for the restaurant and retail-trade industries. Similar findings were apparent for China. Yet, even though the pool of firms shrinks, the education requirements for jobs will still be observed. Again, we try to tackle this issue by restricting the set of ads from always-posting ads throughout 2014 to 2019.

Another limitation of our study is that we do not consider that employers can replace workers without a degree with slightly younger, and cheaper, workers who are still below the relevant age threshold for the minimum wage. Yet, firms need to anticipate that the cost of its young workforce will jump (+20-25%) when employees reach the threshold age.

7 Conclusion

We document the response to a labour cost increase using an unexpected change in the minimum wage policy. We find that a higher national minimum wage led to a decrease in the share of non-graduate ads and in low levels of technical skills requirements in low and middle skilled occupations. Our paper has several

[33] We leave this for future research where we will look at the impact on the minimum wage on the density and variety of firms.

implications for policymakers. Our results suggest that policymakers need to be conscious of changes in hiring patterns within occupations even though most studies do not show any aggregate impacts on employment. Minimum wage policies aim at supporting low-income and therefore low-educated individuals. We do not document effects for workers already on the labour market, but our results show that on the hiring side of the labour market the targeted population might be harmed by the policy change. Future research using the text in the add could look at non-wage compensation: opportunities for training and access to insurance.

References

Acemoglu, D. (2001). Good jobs versus bad jobs. *Journal of labor Economics 19*(1), 1–21.

Adams, A., J. Prassl, et al. (2018). Zero-hours work in the united kingdom. *Conditions of Work and Employment Series 101*.

Adams-Prassl, A., M. Balgova, and M. Qian (2020). Flexible work arrangements in low wage jobs: Evidence from job vacancy data.

Ahlfeldt, G. M., D. Roth, and T. Seidel (2018). The regional effects of germany's national minimum wage. *Economics Letters 172*, 127–130.

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

Azar, J., I. Marinescu, M. Steinbaum, and B. Taska (2020). Concentration in us labor markets: Evidence from online vacancy data. *Labour Economics 66*, 101886.

Black, S., J. Furman, E. Rackstraw, and N. Rao (2016). The long-term decline in us prime-age male labour force participation. *VoxEU. org*.

Blundell, R., D. A. Green, and W. Jin (2021, 07). The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium. *The Review of Economic Studies 89*(1), 142–180.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230. Themed Issue: Treatment Effect 1.

Cammeraat, E. and M. Squicciarini (2021). Burning glass technologies' data use in policy-relevant analysis.

Card, D. (1992). Do minimum wages reduce employment? a case study of california, 1987–89. *ILR Review 46*(1), 38–54.

Card, D., J. Heining, and P. Kline (2013, 05). Workplace Heterogeneity and the Rise of West German Wage Inequality*. *The Quarterly Journal of Economics 128*(3), 967–1015.

Card, D. and A. B. Krueger (1993). Minimum wages and employment: A case study of the fast food industry in new jersey and pennsylvania.

Carnevale, A. P., T. Jayasundera, and D. Repnikov (2014). The online college labor market: Where the jobs are. *Georgetown University Center on Education and the Workforce*.

Clemens, J. (2021). How do firms respond to minimum wage increases? understanding the relevance of non-employment margins. *Journal of Economic Perspectives 35*(1), 51–72.

Clemens, J., L. B. Kahn, and J. Meer (2021). Dropouts need not apply? the minimum wage and skill upgrading. *Journal of Labor Economics 39*(S1), S107–S149.

Datta, N., G. Giupponi, and S. Machin (2019). Zero-hours contracts and labour market policy. *Economic Policy 34*(99), 369–427.

de Chaisemartin, C. and X. D'Haultfœuille (2020, September). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–96.

Delaney, J. and K. L. Papps (2022). Hiring behaviour and the national minimum wage.

Deming, D. and L. B. Kahn (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics 36*(S1), S337–S369.

Dickens, R. and A. Manning (2004). Spikes and spill-overs: the impact of the national minimum wage on the wage distribution in a low-wage sector.

Dickens, R., R. Riley, and D. Wilkinson (2015). A re-examination of the impact of the uk national minimum wage on employment. *Economica 82*(328), 841–864.

Dustmann, C., A. Lindner, U. Schönberg, M. Umkehrer, and P. Vom Berge (2022). Reallocation effects of the minimum wage. *The Quarterly Journal of Economics 137*(1), 267–328.

Elias, P. and K. Purcell (2013). Classifying graduate occupations for the knowledge society. *Institute for employment research. University of Warwick. Futuretrack working paper 5.*

Flinn, C. J. (2006). Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica 74*(4), 1013–1062.

Giupponi, G. and S. J. Machin (2018). Changing the structure of minimum wages: firm adjustment and wage spillovers.

Gopalan, R., B. H. Hamilton, A. Kalda, and D. Sovich (2021). State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data. *Journal of Labor Economics 39*(3), 000–000.

Green, F. and G. Henseke (2016). The changing graduate labour market: analysis using a new indicator of graduate jobs. *IZA Journal of Labor Policy 5*(1), 1–25.

Grinis, I. (2019). The stem requirements of "non-stem" jobs: Evidence from uk online vacancy postings. *Economics of Education Review 70*, 144–158.

Harrigan, J., A. Reshef, and F. Toubal (2018). Techies, trade, and skill-biased productivity. Technical report, National Bureau of Economic Research.

Hershbein, B. and L. B. Kahn (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review 108*(7), 1737–72.

Hershbein, B., C. Macaluso, and C. Yeh (2018). Concentration in us local labor markets: evidence from vacancy and employment data. Technical report, Working paper.

Horton, J. J. (2017). Price floors and employer preferences: Evidence from a minimum wage experiment. *Available at SSRN 2898827*.

Jardim, E., M. C. Long, R. Plotnick, E. Van Inwegen, J. Vigdor, and H. Wething (2018). Minimum wage increases and individual employment trajectories. Technical report, National Bureau of Economic Research.

Javorcik, B., B. Kett, K. Stapleton, and L. O'Kane (2020). Unravelling deep integration.

Jha, P. and A. Rodriguez-Lopez (2021). Minimum wage and firm variety.

Katz, L. F. and A. B. Krueger (1992). The effect of the minimum wage on the fast-food industry. *ILR Review 46*(1), 6–21.

Kollydas, K. and A. Green (2022). Graduate pathways: Identifying patterns of regional retention and attraction.

Ku, H. (2020). Does minimum wage increase labor productivity? evidence from piece rate workers.

Lemos, S. (2009). Minimum wage effects in a developing country. *Labour Economics 16*(2), 224–237.

Lindley, J. and S. Machin (2016). The rising postgraduate wage premium. *Economica 83*(330), 281–306.

Lordan, G. (2019). People versus machines in the uk: Minimum wages, labor reallocation and automatable jobs. *Plos one 14*(12), e0224789.

Luca, D. L. and M. Luca (2019). Survival of the fittest: the impact of the minimum wage on firm exit. Technical report, National Bureau of Economic Research.

Machin, S., A. Manning, and L. Rahman (2003). Where the minimum wage bites hard: Introduction of minimum wages to a low wage sector. *Journal of the European Economic Association 1*(1), 154–180.

Machin, S. and S. McNally (2007). Tertiary education systems and labour markets. *Education and Training Policy Division, OECD 6*.

Manning, A. (2012). Minimum wage: maximum impact. *The Resolution Foundation*.

Mayneris, F., S. Poncet, and T. Zhang (2018). Improving or disappearing: Firm-level adjustments to minimum wages in china. *Journal of Development Economics 135*, 20–42.

McNeil, C. and H. Bonella (2016, September). Brexit and the uk labour market. *IPPR Policy Report*.

Neumark, D. (2001). The employment effects of minimum wages: Evidence from a prespecified research design the employment effects of minimumwages. *Industrial Relations: A Journal of Economy and Society 40*(1), 121–144.

Neumark, D., W. L. Wascher, W. L. Wascher, et al. (2008). *Minimum wages*. MIT press.

Roth, J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2018, 10). Firming Up Inequality*. *The Quarterly Journal of Economics 134*(1), 1–50.

Stewart, M. B. (2004). The impact of the introduction of the uk minimum wage on the employment probabilities of low-wage workers. *Journal of the European Economic Association 2*(1), 67–97.

Vogel, J. (2023, March). The race between education, technology, and the minimum wage. Working Paper 31028, National Bureau of Economic Research.

A Appendix: History NMW rates

Table A1: National Minimum Wage History

	Age				
Year	25 and over	21 to 24	18 to 20	Under 18	Apprentice
April 2019 to March 2020	£8.21	£7.70	£6.15	£4.35	£3.90
April 2018 to March 2019	£7.83	£7.38	£5.90	£4.20	£3.70
April 2017 to March 2018	£7.50	£7.05	£5.60	£4.05	£3.50
October 2016 to March 2017	£7.20	£6.70	£5.30	£3.87	£3.30
April 2016 to September 2016	£7.20	£6.70	£5.30	£3.87	£3.30
Year	21 and over		18 to 20	Under 18	Apprentice
2015	£6.70		£5.30	£3.87	£3.30
2014	£6.50		£5.13	£3.79	£2.73

Notes: The table reports the UK NMW rates.

Source: https://www.gov.uk/national-minimum-wage-rates

B Appendix: Sample Descriptives

Table A2: Descriptive statistics of the coverage before restricting to England, Jan. 2014 - Dec. 2019

	2014	2015	2016	2017	2018	2019
# ads	6090821	7772632	8639924	9372975	8697955	6945605
Job Title	1.00	1.00	1.00	1.00	1.00	1.00
TTWA	0.74	0.76	0.77	0.76	0.73	0.76
Experience	0.13	0.13	0.12	0.13	0.13	0.15
UK SOC Code	0.98	0.98	0.98	0.98	1.00	1.00
SIC Code	0.51	0.49	0.49	0.50	0.54	0.57
Hourly Salary	0.63	0.64	0.65	0.62	0.61	0.62
Qualification	0.18	0.17	0.17	0.18	0.17	0.18

Notes: The table reports descriptives statistics of the coverage of BGT data before restricting to England only. The first row is number of ads, then Job Title of the ad, Occupation at the 4-digit UK SOC, Industry (SIC at division or section levels), experience requirements in years, Hourly salary (minimum in the ad); TTWA stands for Travel To Work Areas and qualification is the BGT measure of education groups.

Source: BGT, 2014-2019.

Table A3: Descriptive statistics of the BGT education variables, Jan. 2014 - Dec. 2019

Year	Number of ads	Share (BGT education)
2014	686576	51.29
2015	872481	52.37
2016	950264	51.91
2017	1016082	51.56
2018	913952	54.45
2019	785547	59.11

Notes: This table reports the share of ads with non-missing information for the BGT education variable on the sample of ads where our education variable is non-missing. The sample also has non-missing information for: Occupation (2-digit UK SOC), Hourly salary (minimum in the ad); TTWA.
Source: BGT, 2014-2019.

Table A4: UK SOC 1-digit

	Name	Skill level
SOC 1	Managers, directors and senior officials	High
SOC 2	Professional occupations	High
SOC 3	Associate professional and technical occupations	High
SOC 4	Administrative and secretarial occupations	Middle
SOC 5	Skilled trades occupations	Middle
SOC 6	Caring, leisure and other service occupations	Middle
SOC 7	Sales and customer service occupations	Low
SOC 8	Process, plant and machine operatives	Low
SOC 9	Elementary occupations	Low

Notes: The table reports the skill level by occupation at the one-digit UK SOC.

C Appendix: Techie Construction

Table A5: Techie occupations

4-digit SOC code	Title
2122	Mechanical engineers
2123	Electrical engineers
2124	Electronics engineers
2126	Design and development engineers
2127	Production and process engineers
2133	IT specialist managers
2134	IT project and programme managers
2135	IT business analysts, architects and systems designers
2136	Programmers and software development professionals
2137	Web design and development professionals
2139	Information technology and telecommunications professionals n.e.c.
2150	Research and development managers
3112	Electrical and electronics technicians
3113	Engineering technicians
3131	IT operations technicians
3132	IT user support technicians
5242	Telecommunications engineers
5245	IT engineers

Notes: The table reports the UK SOC code equivalent to techie occupations as defined by Harrigan et al. (2018). First column is the occupational code of the SOC classification, and the second column is the title of the occupation.

Table A6: Random selection of techie ads

	Job Title	Occupation Code	Techie Mean
1	Infrastructure Architect	2135	64.89
2	Field Maintenance Engineer	3113	17.57
3	Web Developer Html/Css/Javascript/Php	2137	88.62
4	Graduate Technical It Support	3132	69.45
5	Mobile Web Developer - Global Media - Role - K	2137	74.66
6	Principal Software Engineer C++/Java	2136	64.17
7	Graduate Product Development Support Engineer Composites	2126	35.09
8	Process Lead Pensions Quality Monitoring	2139	15.01
9	Electronic Design Engineer	2124	38.48
10	Production Engineer	2127	25.37
11	Senior Software Test Engineer	2139	78.09
12	C++ Developer Python Senior Finance Premium	2136	61.58
13	Mechanical Design Engineer	2122	62.87
14	Development Delivery Lead	2134	27.25
15	Process Development Engineer	2126	23.32
16	Training Contract Pls	3131	19.54
17	Graduate/Junior Java Engineer	2136	76.09
18	Mobile Plant Fitter/Mobile Plant Engineer	2127	19.06
19	Net Developer	2136	79.90
20	Electrical Engineering Fitter	2123	15.68

Notes: The table lists a random selection of techie ads with the Job Title in the first column, the second column reports the occupation code at the four-digit UK SOC Code and the last column the mean techiness of the ad constructed by our program.

Table A7: Random selection of techie ads

	Job Title	Occupation	Techie Mean
1	Consultant In Omf With Specialist Interest In Orthognathic And Def	2139.00	8.37
2	Pharmaceutical Consultant Pharmacist Qp - Quality And Compli	2139.00	11.00
3	Senior Group Auditor - It	2135.00	14.63
4	Research And Data Analyst | Advertising |	2135.00	16.26
5	Medical Engineering Technician	3112.00	22.99
6	Specialist Medical Technical Officer - Electronics	3113.00	28.15
7	Design & Application Engineer	2136.00	29.44
8	Electrical Warranty Engineer 277 5	2123.00	32.02
9	Shift Production Engineer	2127.00	34.06
10	Senior Design Engineer	2126.00	37.53
11	Manufacturing Engineer	3113.00	41.51
12	Mechanical Engineer Surveyor	2122.00	44.16
13	Machine Learning Engineer | Julia | Python	2135.00	46.88
14	Lead Systems Engineer	2135.00	48.88
15	Mobile Electrical Engineer	2123.00	54.82
16	27032 Design & Development Engineer	2126.00	57.18
17	Graduate Developer - 28K - 30K P/A	2136.00	58.90
18	Java Team Lead - - Ajax/Jquery - 80000	2136.00	72.58
19	Software Developer/.Net Developer	2136.00	77.46
20	Java Developer Renowned Software House	2136.00	88.74

Notes: The table lists a random selection of techie ads with the Job Title in the first column, the second column reports the occupation code at the four-digit UK SOC Code and the last column the mean techiness of the ad constructed by our program.

Table A8: Example of 30 randomly selected ads

	Job Title	Occupation	Probability	Techie Mean
1	Senior Staff Nurse	2231	0.04	13.25
2	Social Worker	2442	0.01	5.14
3	Quantity Surveyor	2433	0.05	15.97
4	Financial Accountant	2421	0.03	11.48
5	Graduate Software Developer	2136	0.99	90.53
6	Project Support Coordinator	4215	0.16	16.91
7	Collections Agent	7122	0.02	10.24
8	Marketing Executive	3543	0.05	14.86
9	Graduate/Junior Data Analyst	2135	0.88	59.01
10	Shop Manager	1254	0.02	10.37
11	Family Solicitor	2413	0.01	5.91
12	Electrical Engineer	2123	0.90	65.88
13	Post-Doc Stats Physics/Comp. Mod.	2112	0.57	39.36
14	Lecturer In Hospitality Management	2311	0.13	23.73
15	Graduate Landscape Architect	2431	0.14	26.75
16	Electronics Repair Engineer	2124	0.44	41.73
17	Electronics Technician Medical Device	3112	0.93	70.07
18	Senior Loyalty Analyst	2423	0.35	29.64
19	Java Developer Sql Spring.	2136	0.99	89.24
20	1St/2Nd Line Support Engineer	3132	0.65	44.33
21	Sales Support/Junior Administrator	7129	0.17	21.64
22	Linux Technical Lead Solutions Engineer	2136	0.86	58.64
23	Senior/Principal Structural Engineer	2121	0.31	33.52
24	Senior Developer Sharepoint, .Net	2136	0.96	75.34
25	Senior It Support Analyst	3132	0.82	57.71
26	Senior C# .Net Developer, 80K	2136	0.97	78.05
27	Field Service Engineer	5249	0.54	49.50
28	National Media Manager	3416	0.05	14.98
29	Social Worker	3233	0.01	4.53
30	Support Technician	3132	0.27	29.08

Notes: The table lists a random selection of 30 ads with the Job Title in the first column, the second column reports the occupation code at the four-digit UK SOC Code, Column 3 is the probability from our logistic model of hiring a tech workers, and the last column the mean techiness of the ad constructed by our program.

Table A9: Techie, non-techie and neutral clusters (dropping keywords appearing strictly less than twice)

Cluster	Mean	Max	Min	Median	Number of Keywords
Non-techie	5.70	23.50	0.00	3.20	5787
Techie	83.70	100.00	62.70	84.20	1781
Neutral	41.40	62.50	23.60	40.50	1743

Notes: The table reports a set of summary statistics from the classification of 9,311 keywords into the techie, the neutral and the non-techie clusters. The column reports the mean, the maximum, the minimum, the median and the number of keywords in those clusters. The Table shows the distribution keeping only words appearing at least twice, and we see that they are very close with the distribution keeping all keywords.
Source: BGT, 2014-2019.

Figure A1: Using average techiness above a certain threshold to classify jobs as techie

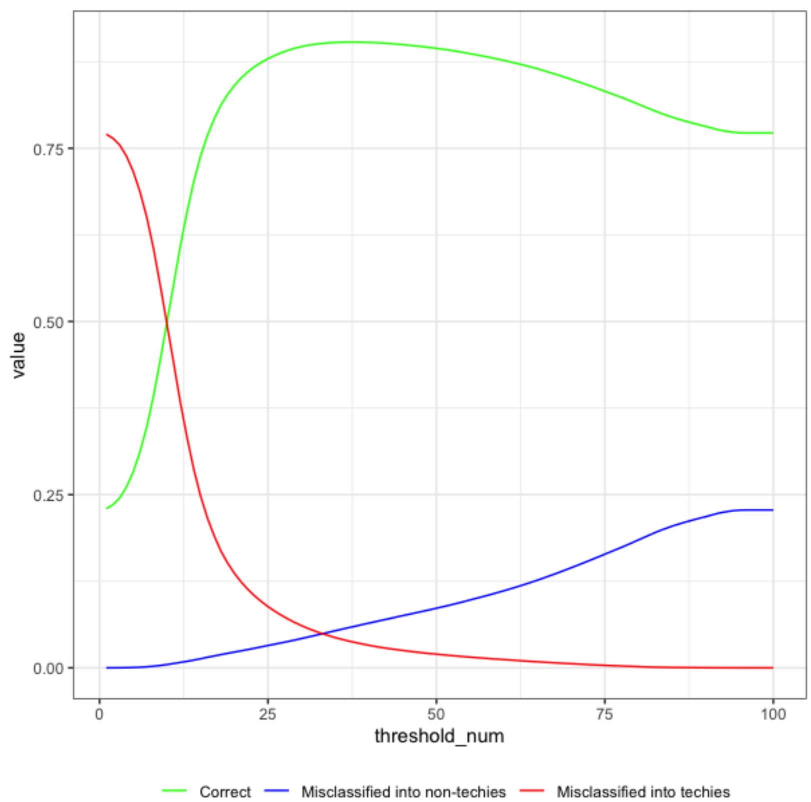

Notes: The correct classification rate peaks at 89% for a threshold of an average techiness equal to 0.30. It means that we correctly classify above 89% of ads as non-tech and tech ads.
Source: BGT 2014.

Figure A2: Distribution of the Techie Mean Variable

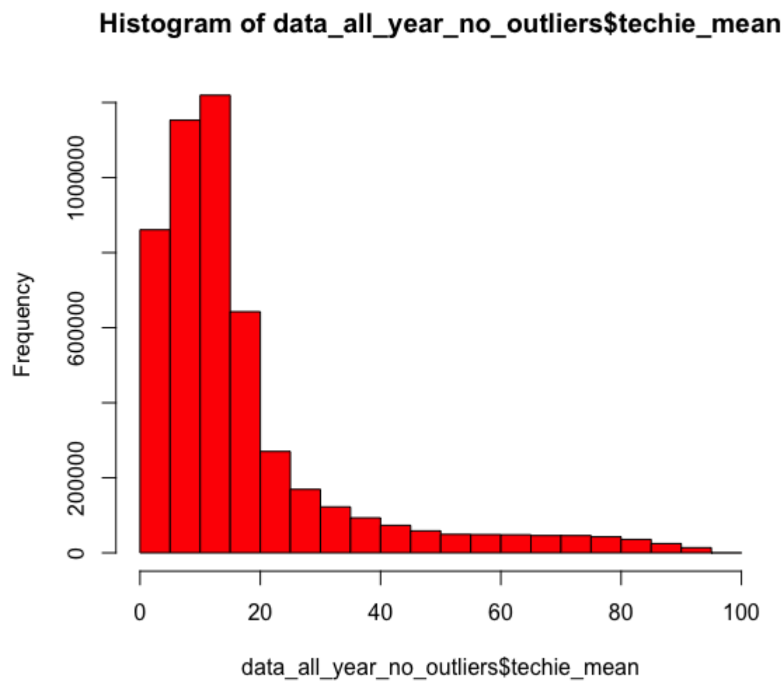

Notes: The figure shows the distribution of the techie mean variable at the ad level, using the log transforms it as a normal distribution.

Source: BGT 2014.

Figure A3: Predicted Probabilities Distribution

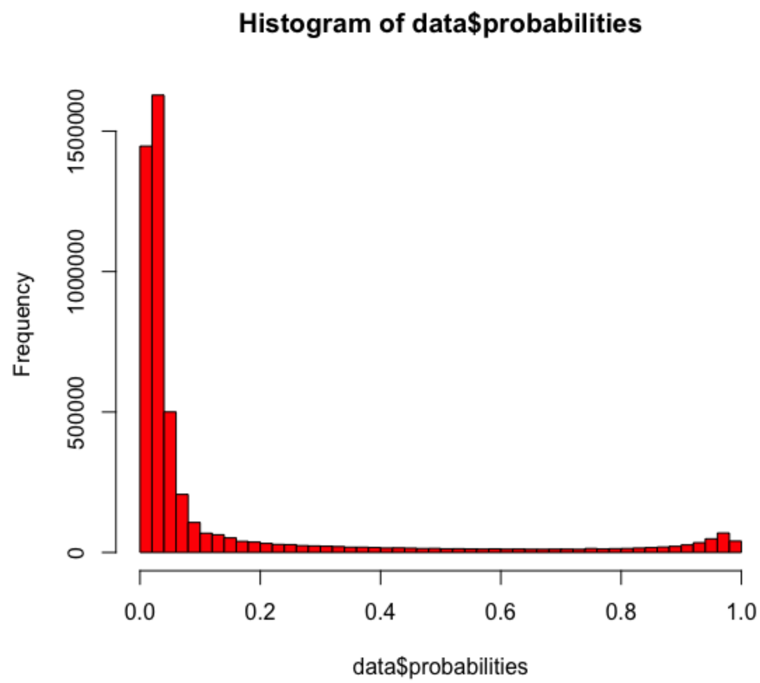

Notes: The figure shows the distribution of the estimated probabilities at the ad level using logit model.
Source: BGT 2014-2019.

Table A10: Logit regression

	Techies Dummy
Mean Techiness	0.077***
	(0.0002)
Max Techiness	0.024***
	(0.0001)
Constant	−4.871***
	(0.003)
Observations	7,748,298
Log Likelihood	−1,607,424.000
Akaike Inf. Crit.	3,214,853.000

Notes: The table displays the results from estimating Equation 2 in a logistic model. The coefficients for "mean techie" and "max techie", β_1 and β_2, are positive and statistically significant. Significance codes: ***: 0.01, **: 0.05, *: 0.1.
Sources: Burning Glass Data 2014-2019.

D Appendix: Brexit Discussion

One might be concerned about Brexit happening at the same time as the change in the minimum wage. The UK officially left the EU on March 29 2019 but with a transition period until the 31 December 2020. It is only from January 2021 that Brexit agreements were enacted. However, the vote in June 2016 might have had an impact on labour markets, given that there was an impact on other markets such as foreign exchange. In a policy report also using BGT data, McNeil and Bonella (2016) study the behaviour of employers in posting online job vacancies before and after the vote in June 2016. They find that hiring trends were similar to the current patterns except for the financial sector pre- and post the EU referendum vote. Javorcik et al. (2020) construct a measure of exposure to barriers on professional services exports and find that the more exposed areas experienced a significant decline in online job postings after the referendum. However, they show that this decline was mainly driven by higher skilled jobs. The effect that they document could potentially explain the small decrease in graduate job postings that we see in Figure 7 from 2015 to 2016. On the contrary, the minimum wage increase has impacts primarily on jobs that are around the minimum wage threshold. As our descriptives show, low wage jobs are concentrated in low skilled occupations which makes the Brexit effect less meaningful in our study.

We further argue that the minimum wage shock does not affect the same jobs as the Brexit shock. We construct a local labour market exposure measure following the approach widely used in the literature on local labour markets (Autor et al. (2013)). We calculate how exposed each local labour market is to potential future barriers to trade induced by both the drop in global demand and the increased uncertainty generated by the Brexit vote. First, we use a sector level employment exposure measure which was calculated before the Brexit vote in 2016. Then, we weight this measure by the pre-sample online job posting composition for each TTWA using BGT data for 2014.

The sectoral exposure to global shocks measure comes from industry level international input-output linkages and are the number of FTE jobs in 2016 supported by exporting activities for each trading country in different sectors, taken from the ONS Department for International Trade (Equation 8). We rely on the IO tables as they capture all links between sectors, inputs and sources of final demand. Moreover, IO tables take into account not only exporting sectors, but also domestic links with other sectors in the economy. Therefore, using IO tables enables us to capture the full chain of domestic activities connected to exporting activities.

$$\text{Industry Exposure to Global Shock}_{i,c} = \text{FTE jobs exports}_{i,c,2016} \qquad (\text{IEGS}_{i,c}) \tag{8}$$

where FTE jobs $\text{exports}_{i,c,2016}$ are UK Full Time Equivalent jobs directly and indirectly linked to exports in sector i.

We then weight this sectoral variable using the pre-sample industry online vacancies composition for each

TTWA (Equation 9) to obtain a local labour market exposure measure:

$$\text{Local Labour Market Exposure to Global Shock}_{i,l,c} = \sum_{i \in l} \text{vacancies share}_{i,l,2014} \times \text{IEGS}_{i,c} \tag{9}$$

where vacancies share$_{i,l,2014}$ is equal to $\frac{\sum \text{vacancies}_{i,l,2014}}{\sum \text{vacancies}_{l,2014}}$

In order to compare the local labour market exposure to global shock to the exposure to minimum wage changes computed in Equation 4 at the TTWA level, we average across industries at the TTWA. We graphically represent the exposure measure in Figure A4, with red colour representing a greater exposure. The London area and the South East are areas highly exposed to potential trade barriers, as are some areas in the North. Comparing those exposure levels to the ones in Figure 10, we clearly see that the more exposed TTWA from the global shock are different than the ones from the minimum wage shock. The correlation between the two exposure measures is 0.21. Therefore, it is unlikely that the Brexit vote is driving our results.

Figure A4: TTWA Exposure to Global Shocks

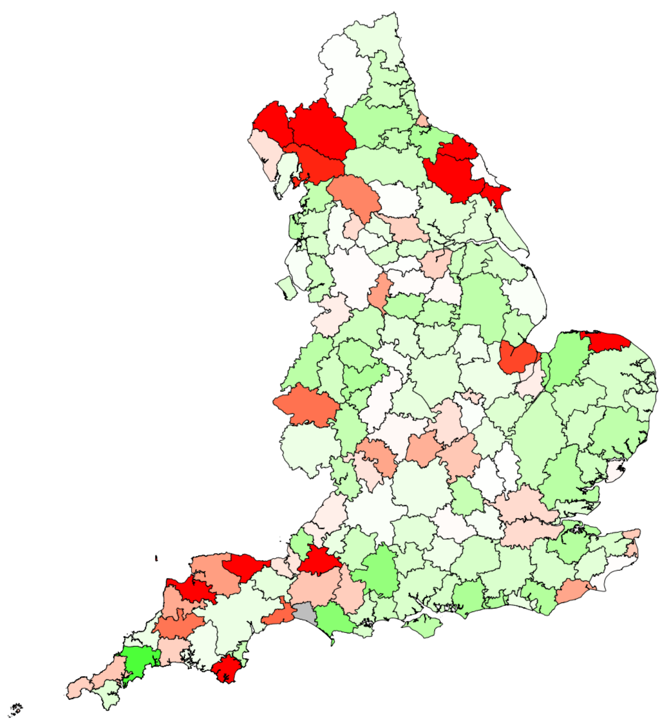

Notes: The map displays the spatial distribution of the variable computed in Equation 9. The greener the area, the lower the exposure of the local area to global shocks in 2016. On the contrary, the redder the higher the exposure. For the plot, the exposure measure is scaled and capped at 3 for visualisation.

Source: BGT 2014 and ONS Department for International Trade (2016).

E Appendix: Results

Figure A5: Graphical results

(a) All occupations (b) Low and Middle Occupations

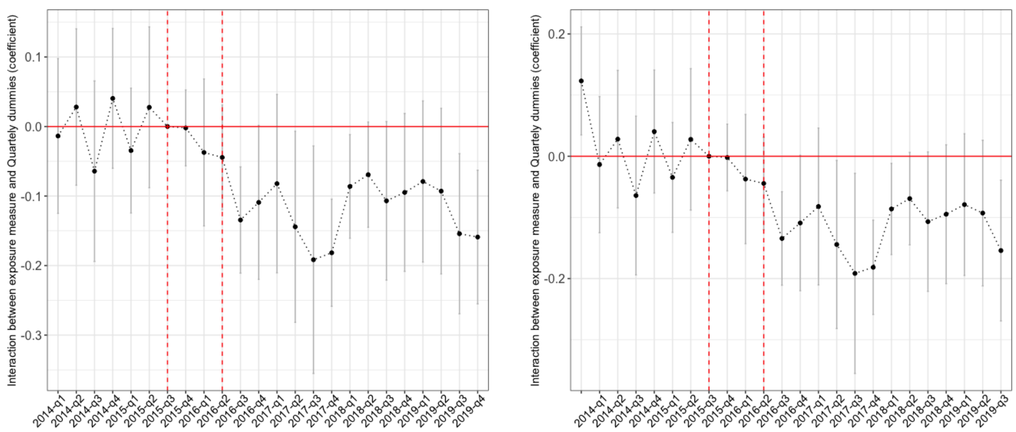

Notes: Results of an event study analysis (see Equation 6), using the set of fixed effects of specification (3) in Table 5 with quarterly dummy interactions. The error bars show the 95% confidence interval using clustered standard errors at the location and two-digit occupation level. The red dash bars show the time of announcement of the change in policy and when it was implemented.

Sources: BGT 2014-2019.

Figure A6: The variation in NMW bite across regions, by occupation

(a) BGT data **(b)** LFS

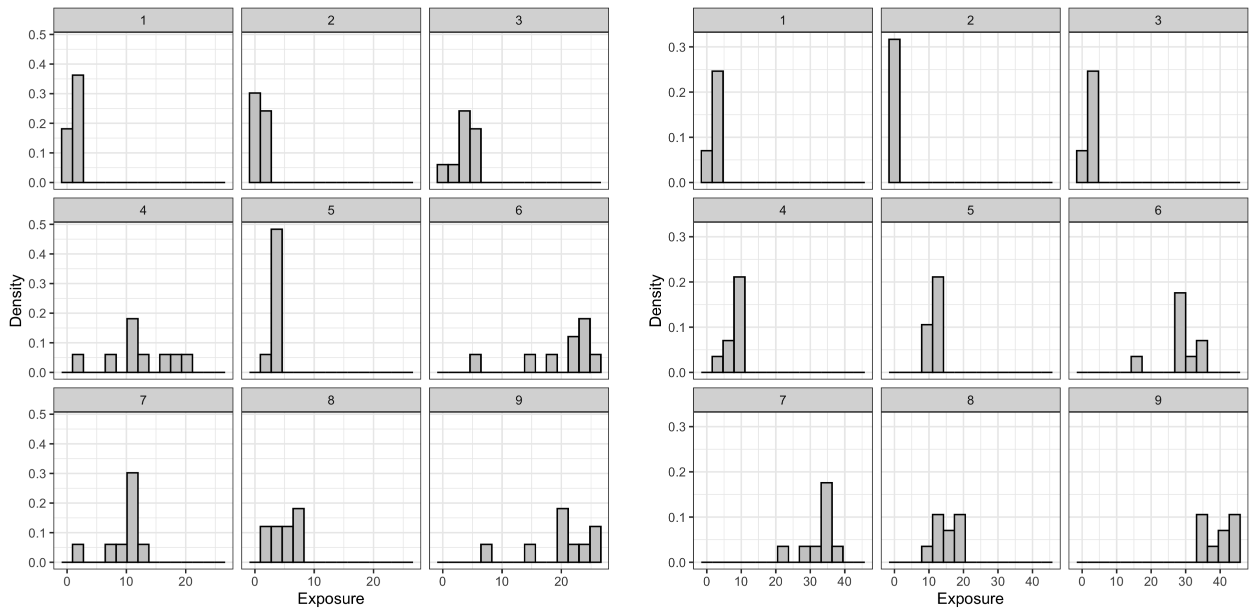

Notes: The figures are descriptives of the minimum wage bite constructed in Equation 4. In Panel (a) we compute the bite at the one-digit level and each panel shows the distribution across TTWA of the share of vacancies that are paid below the NMW (minimum wage bite) for a given occupation (see Table A4 for the corresponding occupation). In Panel (b) we compute the bite at the one-digit level but using the LFS data and each panel shows the distribution across TTWA of the employment paid below the NMW (minimum wage bite) for a given occupation (see Table A4 for the corresponding occupation).
Source: BGT, 2014.

Table A11: Low and Middle skilled occupations - Placebo break

	Share of non-graduate ads (month, TTWA, SOC2)			
	(1)	(2)	(3)	(4)
Constant	0.6519***			
	(0.0735)			
Exposure (NMW)	0.2922*	0.0498	0.0594*	
	(0.1605)	(0.0288)	(0.0293)	
MNW Placebo Dummy	0.0044	0.0118	0.0129	0.0045
	(0.0142)	(0.0118)	(0.0121)	(0.0097)
Exposure (NMW) × MNW Placebo Dummy	-0.0400	-0.0415	-0.0484	-0.0406
	(0.0408)	(0.0343)	(0.0360)	(0.0364)
Year×Season FE		Yes		
TTWA FE		Yes		
Occupation FE		Yes	Yes	
TTWA×Year×Season			Yes	
Season×TTWA×Occ				Yes
Year				Yes
Observations	49,365	49,365	49,365	49,365
R^2	0.02510	0.39583	0.44101	0.51022

Notes: The table displays the results from Equation 5 of the monthly share of ads at the occupation-education-TTWA level on the placebo policy dummy (introduce the placebo break in April 2018) interacted with the exposure to the minimum wage with different sets of fixed effects for the sample of middle and low skilled occupations. Column (1) has no fixed effects, Column (2) includes time, TTWA and occupation fixed effects, Column (3) includes TTWA time trends and occupation fixed effects and Column (4) year and TTWA by year by occupation fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.

Sources: Burning Glass Data 2014-2019.

Table A12: Low and middle skilled occupations - without London

	Share of non-graduate ads (month, TTWA, SOC2)			
	(1)	(2)	(3)	(4)
Constant	0.6617***			
	(0.0742)			
Exposure (NMW)	0.3152*	0.1006**	0.1087***	
	(0.1612)	(0.0347)	(0.0356)	
MNW Policy Dummy	0.0070	0.0539**	0.0529**	0.0486***
	(0.0124)	(0.0214)	(0.0216)	(0.0111)
Expect Policy	-0.0117**	-0.0030	-0.0035	-0.0058
	(0.0046)	(0.0122)	(0.0117)	(0.0051)
Exposure (NMW) × MNW Policy Dummy	-0.0999**	-0.1122***	-0.1143**	-0.1242***
	(0.0403)	(0.0362)	(0.0381)	(0.0385)
Year×Season FE		Yes		
TTWA FE		Yes		
Occupation FE		Yes	Yes	
TTWA×Year×Season			Yes	
Season×TTWA×Occ				Yes
Year				Yes
Observations	48,362	48,362	48,362	48,362
R^2	0.02106	0.38237	0.42861	0.49713

Notes: The table displays the results from Equation 5 of the monthly share of ads at the occupation-education-TTWA level on the policy dummy interacted with the exposure to the minimum wage with different sets of fixed effects. The sample is the middle and low skilled occupations and we remove London in all specifications. Column (1) has no fixed effects, column (2) includes time, TTWA and occupation fixed effects, column (3) includes TTWA time trends and occupation fixed effects and column (4) year and TTWA by year by occupation fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.

Sources: Burning Glass Data 2014-2019.

Table A13: Low and middle skilled occupations - random allocation

	Share of non-graduate ads (month, TTWA, SOC2)			
	(1)	(2)	(3)	(4)
Constant	0.7156***			
	(0.0500)			
Exposure (random)	0.0297	0.0122	0.0199	
	(0.0522)	(0.0143)	(0.0146)	
MNW Policy Dummy	-0.0100	0.0295	0.0309	0.0227
	(0.0122)	(0.0258)	(0.0242)	(0.0155)
Expect Policy	-0.0089*	-0.0036	-0.0039	-0.0059
	(0.0046)	(0.0122)	(0.0117)	(0.0050)
Exposure (random) $\times$ MNW Policy Dummy	-0.0042	-0.0044	-0.0172	-0.0036
	(0.0135)	(0.0127)	(0.0107)	(0.0135)
Year$\times$Season FE		Yes		
TTWA FE		Yes		
Occupation FE		Yes	Yes	
TTWA$\times$Year$\times$Season			Yes	
Season$\times$TTWA$\times$Occ				Yes
Year				Yes
Observations	49,365	49,365	49,365	49,365
R^2	0.00067	0.39574	0.44088	0.51043

Notes: The table displays the results from Equation 5 of the monthly share of ads at the occupation-education-TTWA level on the policy dummy interacted with the exposure to the minimum wage with different sets of fixed effects. In this table we randomly allocated the exposure measures to the different TTWA. Column (1) has no fixed effects, column (2) includes time, TTWA and occupation fixed effects, column (3) includes TTWA time trends and occupation fixed effects and column (4) year and TTWA by year by occupation fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.

Sources: Burning Glass Data 2014-2019.

Table A14: Impact of the minimum wage on probability of hiring a tech worker

	Share of ads with low tech probabilities					
	Overall Occupations			Low and Middle Skilled Occupations		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.5369***			0.7666***		
	(0.0604)			(0.0374)		
Exposure (NMW)	1.032***	0.2211***	0.2357***	0.2286	0.0563	0.0836
	(0.2759)	(0.0606)	(0.0609)	(0.2101)	(0.0447)	(0.0613)
MNW Policy Dummy	0.0266**	0.0430***	0.0432***	0.0076	0.0461*	0.0524**
	(0.0104)	(0.0153)	(0.0145)	(0.0139)	(0.0213)	(0.0227)
Expect Policy	-0.0215***	-0.0039	-0.0047	-0.0352**	-0.0109	-0.0108
	(0.0074)	(0.0084)	(0.0082)	(0.0143)	(0.0169)	(0.0171)
Exposure (NMW) ×	-0.2799***	-0.2058**	-0.2053**	-0.1930**	-0.1950**	-0.2161**
MNW Policy Dummy	(0.0808)	(0.0854)	(0.0875)	(0.0655)	(0.0899)	(0.0984)
Season ×Year		Yes			Yes	
TTWA		Yes			Yes	
Occupation		Yes	Yes		Yes	Yes
Season × TTWA ×Year			Yes			Yes
Observations	77,008	77,008	77,008	26,737	26,737	26,737
R^2	0.12375	0.50462	0.52853	0.00716	0.29756	0.36589

Notes: The table displays the results from Equation 5 of the monthly share of ads with low tech probabilities for graduates only at the occupation-TTWA level on the policy dummy interacted with the exposure to the minimum wage with different sets of fixed effects. The first 3 columns are the overall sample and the last three colums are over the low and middle skilled occupations only. Columns (1) and (4) have no fixed effects, Columns (2) and (5) includes time, TTWA and occupation fixed effects and Columns (3) and (6) year and TTWA by year by occupation fixed effects. Standard errors (in parentheses) are two-way clustered at the TTWA and occupation level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among occupations and locations.

Sources: Burning Glass Data 2014-2019.

Table A15: Impact of the minimum wage on the share of ads at risk of being automated

	Share of ads					
	High Risk of Automation			Low Risk of Automation		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.1789***			0.8201***		
	(0.0033)			(0.0032)		
Exposure (NMW)	0.0453***	0.0125***	0.0129***	0.0006	-0.0036	-0.0041
	(0.0121)	(0.0047)	(0.0049)	(0.0077)	(0.0045)	(0.0046)
MNW Policy Dummy	0.0117***	0.0047	0.0171***	-0.0121***	0.0019	-0.0128***
	(0.0024)	(0.0063)	(0.0046)	(0.0024)	(0.0069)	(0.0047)
Expect Policy	-0.0056**	-0.0115***	-0.0062**	0.0032	0.0133***	0.0073**
	(0.0027)	(0.0035)	(0.0030)	(0.0030)	(0.0042)	(0.0035)
Exposure (NMW) ×	-0.0201**	-0.0198***	-0.0204***	0.0059	0.0057	0.0066
MNW Policy Dummy	(0.0088)	(0.0075)	(0.0077)	(0.0072)	(0.0072)	(0.0073)
Season × Year		Yes			Yes	
TTWA		Yes			Yes	
Year			Yes			Yes
Season × TTWA			Yes			Yes
Observations	163,224	163,224	163,224	167,026	167,026	167,026
R^2	0.01393	0.30835	0.33491	0.00640	0.19736	0.22747

Notes: The table displays the results from Equation 5 of the monthly share of ads for 2-digit occupations at high risk of automation (ROA) (Columns 1 to 3) and at low risk of automation (Columns 4 to 6) on the policy dummy interacted with the exposure to the minimum wage with different sets of fixed effects. Columns (1) and (4) have no fixed effects, Columns (2) and (5) includes time and TTWA fixed effects and Columns (3) and (6) year and TTWA by season fixed effects. Standard errors (in parentheses) are one-way clustered at the TTWA level. Significance codes: ***: 0.01, **: 0.05, *: 0.1. The reported standard errors are robust to correlation in the errors among locations. Occupations at ROA are defined following the ONS (`https://www.ons.gov.uk/visualisations/dvc599/beeswarm/data.xlsx` where high ROA are occupations with a share of automation above 50%.
Sources: Burning Glass Data 2014-2019.