The impact of the National Living Wage on productivity

Eduin Latimer; Low Pay Commission

March 2022

Summary

We can find no evidence that the introduction of the National Living Wage (NLW) had a significant effect on productivity at the industry-region level. We estimate the effect of the NLW on average labour productivity by comparing productivity growth in industry-region cells where a high proportion of workers are paid the minimum wage, to productivity growth in industry-region cells with lower minimum wage coverage. We find no statistically significant effect of the NLW on productivity. As a validation check on our method, we use the same method to test whether the NLW increased wages and find a statistically significant effect on wages.

However, the available data does not allow us to estimate the effect of the NLW on productivity with a high level of precision. The range over which we have 95 percent confidence that the true effect falls within is wide. It ranges from a 1 percentage point increase in the share of minimum wage workers decreasing productivity by 0.2 percent to a 1 percentage point increase in the share of minimum wage workers increasing productivity by 0.1 percent. The range is wide as we only have small number of observations at the industry-region level and the minimum wage affects only a small percentage of jobs in each industry-region cell. Future research using firm-level data could produce more precise estimates of the effect of the NLW on productivity.

We also have not carried out an investigation into potential reallocation effects across industries and regions from the NLW. The NLW could have reallocated workers from less productive firms to more productive firms, thus increasing aggregate productivity. There is evidence for the introduction of the German minimum wage caused the reallocation of workers to more productive firms. Our study would capture reallocation within industry-region cells from less productive firms to more productive firms but cannot capture allocation effects which move workers across industries (as categorised in this analysis) or across regions.

We would like to thank Patricia Rice, Jonathan Wadsworth, Rebecca Riley, John Forth and Alex Bryson for providing comments. Any errors are the responsibility of the author.

Introduction:

1 When the UK Government introduced the National Living Wage (NLW) in 2016, one of the key aims was to improve productivity. The Low Pay Commission's remit in 2016 said "A remaining, key economic challenge the Government wants to address is to move away from a low wage, high tax, high welfare society and encourage a model of higher pay and higher productivity – supporting people who work hard and want to get on in life to fulfil their aspirations. As such, the Government wishes to see a higher wage for more experienced workers and so is introducing a premium for workers aged 25 and over."

2 The NLW raised the minimum wage for those aged 25 and over from £6.50 in April 2015 to £7.20 in April 2016, a 10.8 percent increase. The Government asked the Low Pay Commission to recommend a path for the NLW, which would take it to 60 percent of median hourly earnings by 2020. This led to further fast increases in the minimum wage up until April 2020, when it reached £8.70 and hit its target. Between 2015 and 2020, the minimum wage increased by 23.4 percent in real terms for workers aged over 25.

3 A wide body of research has found that the NLW, had large impacts on wages but that so far the impacts on employment have been limited (Aitken, et al., 2018; Low Pay Commission, 2020; Cribb, et al., 2021; Avram & Harkness, 2019; Manning, 2016; Dube, 2019). Studies have found that the NLW increased wages not only for the 7 percent of payrolled jobs which are paid the minimum wage, but that it also had spillover effects on workers paid more. In our review of the first stage of the NLW we estimated that up to around 35 percent of workers in 2019 may have been paid more due to the NLW. Research has found little evidence of any aggregate employment impacts from the NLW, although studies have found negative employment effects in some specifications for female workers or female part-time workers. This raises the question: how else have firms dealt with the increased wage costs?

4 One potential response to a higher minimum wage is for firms and workers to increase productivity. If workers produce more per hour, firms can more easily afford to pay them the increased rate demanded by the minimum wage. This could reduce the need for firms to reduce employment. Alternatively, firms could also absorb the costs into profits (Bell & Machin, 2018) or raise prices to respond to the rises (Frontier Economics, 2020).

5 A rising minimum wage could incentivise a range of different productivity enhancing responses from firms. Bernini & Riley (2016) suggest the following six potential mechanisms:

a. firms adopting more capital-intensive ways of working

b. firms increasing training for staff to boost their skills

c. firms substituting less-skilled labour for a more skilled/experienced workforce

d. firms implementing tougher recruitment criteria

e. greater supervision to encourage more effort

f. outsourcing low-skilled tasks. (This enhances the productivity of the individual firm but not necessarily for the whole economy)

In addition to these mechanisms, firms may ask existing staff to take on additional responsibilities.

6 The National Living Wage could also incentivise workers to increase their productivity. A pay increase could lead workers to feel more valued and/or give workers greater incentive to remain at the current place of work. This could improve productivity through:

a. Increased on-the-job effort

b. Reduced absenteeism

c. Reduced voluntary exit (amongst minimum wage workers) and therefore more experienced, productive workers. There is some evidence of this in the UK and elsewhere (Low Pay Commission, 2022; Brochu & Green, 2013; Dube, et al., 2010). Theoretically, a rising minimum wage could also increase voluntary exit, as it increases the outside option for employed workers.

7 The NLW could also increase productivity through reallocation of workers to more productive firms. Dustmann et al. (2021) found that the introduction of a federal minimum wage in Germany in 2015 led workers to move to more productive firms. The higher costs could force less productive firms out of business and allow more productive firms to win their market share. Engbom & Moser (2021) find evidence of a similar reallocation in Brazil, caused by the uprating of the minimum wage. In the UK, employment has fallen in low-paying industries relative to other industries since the NLW was introduced, but it is hard to assess whether this is due to the reallocation effects of the NLW or other trends. Frontier Economics (2020) also found evidence that employment growth in low-paying firms was weakened by the introduction of the NLW, although they did not find that the NLW induced low-paying firms to exit the market.

8 On the other hand, firms may respond to increased wage costs by cutting back on other costs. In response to a 2019 survey by the Chartered Institute of Personnel and Development, 11 percent of businesses said they would respond to the rise in the National Living Wage by reducing investment and 8 percent said they would respond by reducing training spend. Reduced investment in training or capital spending is likely to reduce productivity in the longer term. The introduction of the National Living Wage may have reduced productivity for some firms.

9 Frontier Economics (2020) is the only existing research we found on the impact of the NLW on productivity. Their results suggested that turnover per employee (a proxy for productivity) actually fell in firms more exposed to the minimum relative to other firms following the introduction of the NLW. However, they suggest caution in interpreting their findings due to limitations in the data and potential selection bias (firms experiencing persistent negative shocks to output may be more likely to be classified as low wage.)

10 There is more evidence on the National Minimum Wage and productivity prior to the NLW's introduction. This evidence tends to show a positive, but not always statistically significant, association between productivity and the NMW. For instance, Forth & O'Mahoney (2003) used industry level data (SIC 92 three digit groupings) to test whether the introduction of the NMW increased productivity in industries with a larger share of minimum wage workers. They define labour productivity of value added (a measure of economic output) per worker. They found a positive but not statistically significant effect. Forth et al. (2009) investigated whether the NMW increased labour productivity both by comparing industries and comparing firms within the most affected industries. This improved on the previous analysis as they use a measure of gross value added per hour rather than per worker. This is a better measure of labour productivity. They did not find a statistically significant effect in either analysis.

11 Galindo-Rueda & Pereira (2004) used linked employee and employer data and also found a positive but not statistically significant positive relationship between exposure to the NMW and productivity, although data limitations meant they had to use an industry-region level treatment measure, which limited the precision of their estimates.

12 Riley & Rosazza-Bondibene (2017) carry out a difference-in-difference analysis on firm level data from the Financial Analysis Made Easy dataset (FAME.) They follow Draca, et al. (2005); Draca, et al. (2011) and use average labour costs as a treatment variable. They use average labour costs as a proxy for minimum wage coverage, which is unavailable at firm level in their data. Under most specifications, they found that the initial introduction of the National Minimum Wage did have a statistically significant positive impact on Gross Value Added (GVA) per job, they also found that the National Minimum Wage had a positive effect on productivity during the 2008-9 recession. The effects are substantive. Depending on the specification, Riley and Rosazza-Bondibene find that firms affected by the minimum wage experienced 3-4 percent greater growth in GVA per job between 1999 and 2002 than those in the control group. They find a similar magnitude effect between 2009 and 2012.

13 Croucher & Rizov (2011) use a similar method to Riley & Rosazza-Bondibene (2017) on unconsolidated firm account data. They also find that labour productivity and labour costs increased more in low-paying firms than in other firms following the introduction of the National Minimum Wage. On the other hand, Crawford et al. (2013) find limited evidence to suggest that the 2008-9 recession affected firms' productivity differently based on their exposure to the minimum wage.

14 Riley & Rosazza-Bondibene (2017) explore some of the potential mechanisms for their measured effect. They find limited evidence that their estimated productivity effects come through increased capital investment or through reduced employment. They suggest the increased productivity could instead be due to increased training, increased worker effort or reduced worker turnover. However, as they measure productivity as output per job rather than per hour, their results could also be caused by an increase in average hours amongst low-paying firms. They also used national level deflators, so they could be picking up the effects of the minimum wage on local prices rather than on productivity effects.

15 Bernini & Riley (2016), building on the findings of Riley & Rosazza-Bondibene (2015), explored further which mechanisms could explain the documented effect of the National Minimum Wage on productivity. As in that previous research they found a positive association between productivity and NMW increases but they could not identify a single channel that explained the productivity effect. Instead, any effects were likely to have arisen through a combination of factors within the firm or for different reasons in different firms. There was evidence in some models that firms may have implemented organisational changes (e.g., replacing unskilled workers with more skilled ones, and increasing the incidence of shift work).

16 Recent productivity growth has been slow in the UK. Aggregate Productivity –as measured by output per hour - grew by on average 0.4 per cent each year between 2010 and 2015 in real terms. This has been coined the 'productivity puzzle' in the UK and has been used to help explain the stagnation in real hourly wages. Productivity grew slightly faster after the NLW was introduced. On average productivity grew by 0.8 percent in real terms each year between 2015 and 2019. However, minimum wage workers make up only 7 percent of the national workforce, so macroeconomic factors rather than increases in the NLW are most likely to have driven the slightly stronger increase in productivity in the second half of the last decade.

Figure 1: Productivity, Median Hourly Earnings and the National Minimum Wage, UK, 2000-2019

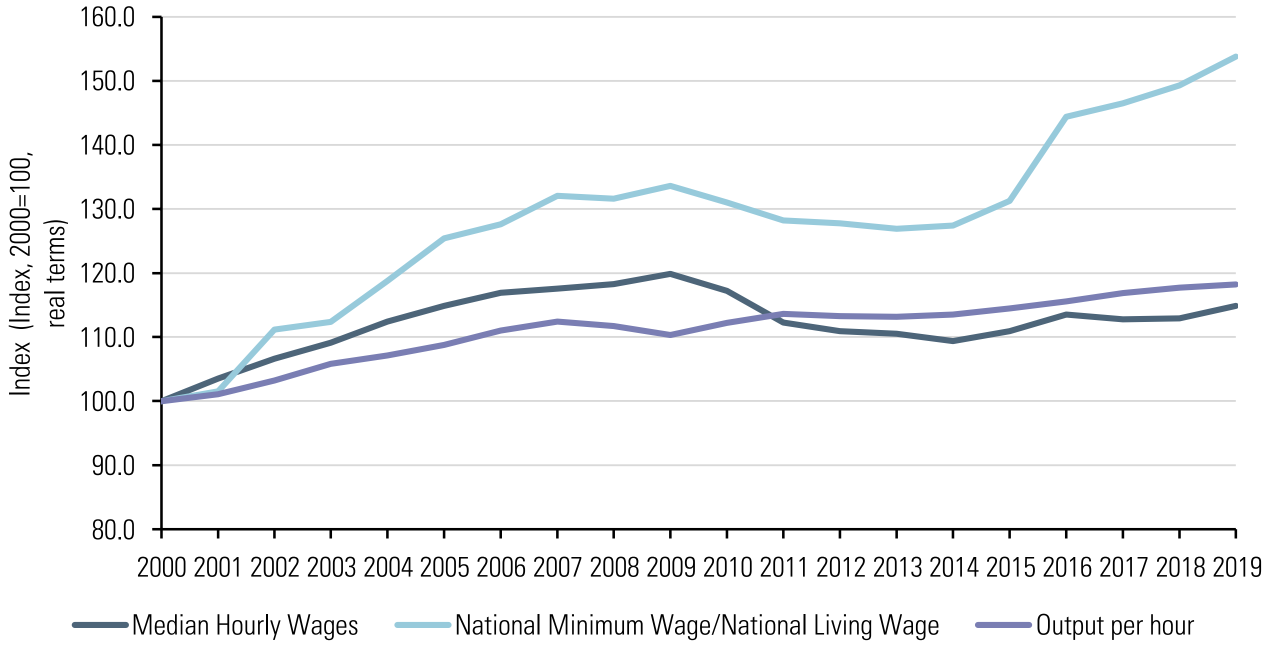

Source: LPC analysis of ONS Labour Productivity Time Series and ASHE, UK, standard weights. Median hourly wages and NMW/NLW are in 2000 prices, CPI adjusted. Output per hour is chained volume measure.

17 In this study, we estimate the effect of the National Living Wage on productivity growth using a difference in difference approach. We test whether productivity growth in industry-regions with a large share of workers are paid the minimum wage is higher than in other industry-regions. If the National Living Wage did drive faster productivity growth, we would expect the industry-region cells with higher minimum wage coverage to experience faster productivity growth. We however find no evidence for this and thus conclude that we find no significant impact of the National Living Wage on productivity.

18 We structure the rest of this paper as follows. We first explain what data we use for our analysis and why we use it, we then explain our method before presenting the results, robustness checks and finally the conclusions we draw from our analysis.

Data:

19 We use data on Gross Value Added per hour by industry and region to measure average labour productivity. Gross Value Added is a measure of economic output, closely related to Gross Domestic Product. It reflects the value of the goods or services sold by a firm net of all costs that can be directly attributed to the production of those goods or services. This measure of Gross Value Added is then divided by a measure of total hours worked (including by self-employed workers) in each industry-region. Ideally, we would also use total factor productivity as an outcome variable, but the data on total factor productivity is not available at a granular enough level to include in this study.

20 The data is based on the UK's 12 ITL1 region/nations and 15 different industry groupings. The industry groups are based on SIC07 and are in most cases equivalent to commonly used industry "sections", although in some cases multiple sections have been grouped into one category. This gives us 180 cells of data each year. We use data from 2012 to 2019 so we have a total sample size of 1440 observations.

21 There are challenges to measuring productivity at a local level. First, the Independent Review of Economic Statistics in the UK noted the challenges in measuring productivity in service sectors, where the nature of the economic product can be hard to define (Bean, 2016). Productivity in public sector services is particularly challenging as there is often not a market price for the output. This is a particular challenge for our study as minimum wage workers disproportionately work in service sectors such as hospitality and social care. We exclude real-estate activities from our dataset, as productivity is high and volatile for this industry, due to imputed rents. We include all other industries in our data, to maintain as large a sample size as possible, but note the additional uncertainty from potential measurement error of productivity. If this measurement error is uncorrelated with our treatment variable (share of workers paid the minimum wage), then it will increase the size of our confidence intervals but will not bias our estimates. If the measurement error is correlated with the treatment variable it could bias our estimates. We run a robustness check excluding non-market sectors, which does not alter our main conclusions.

22 Second, we use productivity data which the ONS generate with national level industry price deflators. This means that improvements in productivity over time could reflect relative price increases in each region. The NLW could causes firms to raise prices of non-tradeables in high coverage areas by more than in low coverage areas. This would be measured as an increase in productivity by our approach, but actually would be a case of price-pass through. This suggests caution in interpreting any positive effect measured by our approach. It could reflect genuine productivity effects, but it also could reflect price effects.

23 Third, productivity data at a regional level requires apportionment of output within large firms. ONS currently uses employment to apportion GVA across different local units of larger firms. If a firm's operation in one area is more productive than its operations in another area, this apportionment will lead to underestimated productivity in the more productive local unit and overestimated productivity in the less productive local unit. This may mean that the data underestimates productivity differences between different regions.

24 Finally, productivity data is only available for all workers in the economy, although the NLW only applied to workers aged over 25. 85 percent of all workers were over 25 in 2015. Stakeholder evidence suggests that few firms responded to the NLW by switching from 25 and overs to under 25s. For example, around 5 per cent of responders to surveys by the Federation of Small Businesses and the Chartered Institute of Personnel and Development (CIPD) said they had recruited more young workers in response to the introduction of the NLW. Put together this suggests that including productivity of workers under 25 is unlikely to significantly bias our results.

25 The key limitation of our productivity data is that at the data are only available at the industry-region level. This has advantages. It allows us to include reallocation between firms (within-industry-regions). It also means we can use an output per hour measure which more accurately measures productivity. The sample is also large enough within each cell to estimate the share of minimum wage workers, which is often problematic for firm-level analysis. However, we only have 180 observations each year and within each industry-region only a small share of workers are paid the minimum wage. As an alternative to our main data, we use GVA per hour at the local authority district level. This provides a greater number of observations each year (362), although it is only available at current prices, which means it not possible to extract local price effects of the NLW from productivity effects. We adjust the current price local authority productivity data for average inflation across the economy using the GDP deflator. Previous studies have utilised firm level data, which may be able to provide more precise estimates.

26 We use data from the Annual Survey of Hours and Earnings (ASHE) for information on wages. The ASHE is a survey of employees completed by businesses. It collects information on a random sample of one percent of all payrolled employees on HMRC’s Pay-As-You-Earn system in the UK. We use the survey to estimate average hourly wages and the percentage of workers paid the minimum wage at industry-region level. The National Living Wage only applies to workers aged 25 and over, so we estimate minimum wage coverage rates for workers aged over 25 in our analysis. The Office for National Statistics produces two sets of weights for the ASHE: low-pay weights and standard weights. We use the low-pay weights to estimate how many people are paid the minimum wage, and we use the standard weights to calculate average hourly pay. This is in line with ONS best practice. We convert mean basic hourly wages into real terms using the GDP deflator.

27 We measure exposure to the minimum wage in two ways. One is the coverage rate, which is the percent of workers paid below the NMW in 2015 (£6.70). The second is the gap treatment measure. At the individual level, this measure is 0 if a worker is paid more than the incoming NLW in 2015. If a worker is paid less than the incoming NLW in 2015 it is equal to the difference between their pay and the NLW. At the industry-region level, it is the average (mean) for the values of individuals inside the industry-region. This measure has two advantages over the coverage rate measure. First it includes workers paid above the NMW but below the incoming NLW. Second, it measures the extent of underpayment in an industry-region. All things equal we might expect sectors where workers are paid considerably below the incoming NLW, to be more exposed than sectors where workers paid at the incoming rate. However, the gap measure is more sensitive to outliers (e.g., workers with £0 pay), than the coverage rate. We therefore report both measures.

28 Descriptive characteristics of the data (aggregated to industry-region-year cells) are shown in Table 1.

Table 1: Descriptive statistics and correlations for key variables at industry-region-year level

Descriptive statistics:				
	Number of observations	Mean	Median	Standard Deviation
GVA per hour	1440	30.5	27.1	11.7
Mean Wages	1440	15.1	14.8	3.8
Coverage Rate	180	5.6	2.8	7.3
Gap Treatment Measure	180	6.9	4.0	7.5
Correlation Matrix (in 2015):				
	GVA per hour	Mean Wages	Coverage Rate	Gap Treatment Measure
GVA per hour	1.00	0.72	-0.51	-0.55
Mean Wages	0.72	1.00	-0.64	-0.68
Coverage Rate	-0.51	-0.64	1.00	0.98
Gap Treatment Measure	-0.55	-0.68	0.98	1.00
Share of Variation (Sum of Squares as a percent of Total Sum of Squares in 2015)				
	Industry	Region	Residuals	Total
GVA per hour	74	11	15	100
Mean Wages	60	21	19	100
Coverage Rate	86	4	10	100
Gap Treatment Measure	87	4	9	100

Source: LPC analysis of ASHE and ONS Industry-Region Productivity data. Data for the Coverage Rate and Gap Treatment Measure are from 2015, Data for GVA per hour and Mean Wages are for the period 2012-2019.

29 Productivity is lower in the high coverage industry-region cells. This is shown in Figure 2. For instance, in 2015, 36 percent of hospitality workers in Northern Ireland were paid the National Minimum Wage and on average productivity was only £16 per hour. In contrast, in 2015 a manufacturing worker in the South East of England on average produced £42 of output per hour and only 1 percent of them were paid the minimum wage. This follows from the relationship between productivity and pay. Pay tends to be higher for higher productivity workers. In a perfectly competitive market, firms can afford to pay workers more if they are higher productivity. In the real world, factors such as the firms' market power in the product and labour market and worker union membership mediates the relationship between productivity and pay. Nevertheless, pay does tend to be higher for higher productivity workers. The correlation matrix in Table 1 shows that GVA per hour is positively related to mean wages but negatively related to the coverage rate. The wide variation in coverage rate and productivity by industry-region improves the precision of our estimates.

30 Difference between industries account for most of the variance in each measure. Table 1 decomposes the variance in each measure into the variance between industries, the variance between regions and residual variance. For each measure most of the variance comes from differences between industries and much less comes from difference between regions. Differences between industries account for 86 percent of the variation in coverage rate, although only 60 percent of the variation in mean wages. This means that our results are more sensitive to omitted variables which affect different industries differently, than to omitted variables which may affect different regions differently.

Figure 2: Gross Value Added (GVA) per hour and percent of employees paid the minimum wage, by industry-region cell, UK, 2015

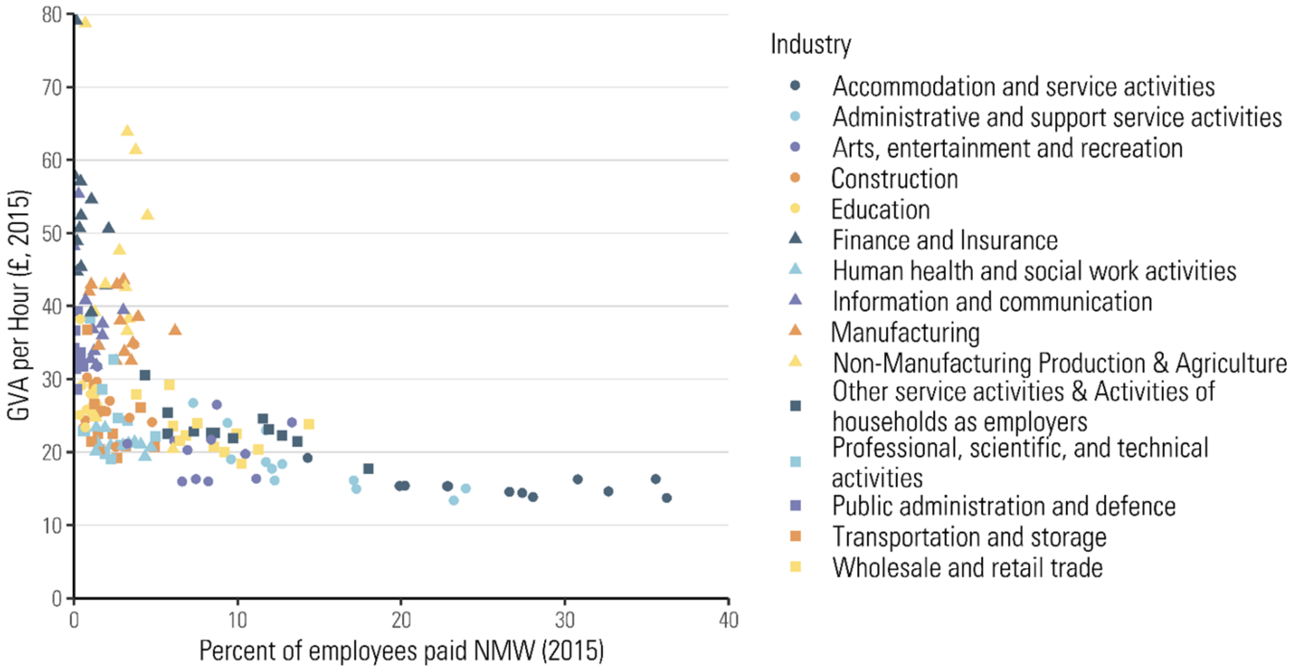

Source: LPC analysis of ONS Productivity by Industry-Region data (current price measure) and ASHE, 2015, UK, using low pay weights, payrolled workers aged 25 and over

Method:

31 We estimate the effect of the introduction of the NLW on productivity using a difference-in-difference approach. If the NLW increased productivity, we would expect industry-region cells with high coverage in April 2015 (before the NLW was introduced) to experience faster growth in productivity between 2015 and 2019 than industry-region cells with lower coverage.

Our main estimates of the treatment effect come from the following regression equation.

Equation 1: $\log{(gva\ per\ hour)_{srt}} = \delta_{sr} + \sigma_t + \beta * NLW\ Exposure_{sr2015} * PostNLW_t + \varepsilon_{srt}$

32 Our dependent variable is log productivity for a given industry *s*, region *r* in time *t.* We take the natural logarithm of productivity, as we want to test whether the introduction of the NLW affects the growth rate of productivity. δ_{sr} is an industry-region fixed effect, it captures any time-invariant differences between different industry-region cells. σ_t is a time fixed effect,

which captures how productivity on average changes over time. β is the coefficient of interest. It measures whether after 2016 when the NLW was introduced, productivity grew faster for industry-region cells more exposed to the minimum wage. $PostNLW_t$ is a dummy variable which equals 1 if it is 2016 (when the NLW was introduced) or later and 0 otherwise. We estimate the model with two measures of exposure to the introduction of the NLW. Firstly, the share of workers paid the current NMW in 2015 and secondly a more complicated measure we refer to as the gap based measure. We provide more details on both measures in paragraph 27.

33 The effects of the NLW on productivity may differ over time so we also estimate the following regression equation, where there is a dynamic treatment effect. The model is the same as equation 1 except we replace the binary time variable $PostNLW_t$ with a full set of time dummies γ_t. This allows us to explore how the effect of the NLW on productivity differs over time. We exclude 2015 from the time dummies to avoid multicollinearity.

Equation 2: $$\log(gva\ per\ hour)_{srt} = \delta_{sr} + \sigma_t + \beta * NLW\ Exposure_{sr2015} * \gamma_t + \varepsilon_{srt}$$

34 There are two key assumptions required for the differences in differences specification above to identify a causal effect. Below I outline them and provide some contextual information on them.

a. Common trends: This assumption requires that if the NLW had not been introduced, growth in log productivity for industry-region cells is uncorrelated with the coverage rate in that industry. If there is another time-varying variable which causes productivity to change at different rates in high productivity industry-region and low productivity industry-regions and is correlated with coverage rates, this assumption would fail. We cannot formally test this assumption, but we can explore the data before the National Living Wage was introduced to provide supporting evidence for the assumption.

   Figure 3 shows how productivity has developed for different industry-region cells since 1998. We group them into quartiles based on the coverage of the National Minimum Wage in the cell in 2015. It shows that high coverage industry-regions have not always followed the same trend as low coverage industry-regions.

   One potential reason for these differing trends could be the National Minimum Wage prior to the National Living Wage's introduction. The National Minimum Wage was introduced in 1999 and was uprated annually until 2015 when the National Living Wage was introduced. While the NMW was at a lower level relative to median wages, it could have driven different trends in the minimum wage prior to the introduction of the NLW. Previous papers have found positive effects of the NMW on productivity on prior to 2012 (Croucher & Rizov, 2011; Riley & Rosazza-Bondibene, 2017). The NMW did not rise relative to median wages between 2012 and 2015, which we use as our pre-treatment period.

   Another potential explanation are macroeconomic events, which appear to have had different effects on different industry-regions. Productivity grew faster for low coverage areas in the period before the 2008 recession, however productivity then fell more steeply following the recession. This probably reflects the greater exposure of high paying industries such as finance to the 2008 recession. Together this meant that between 2002 and 2012 productivity grew by 10 percent in the highest coverage industry-regions and only 7 percent in the lowest coverage industry-regions.

However, our study period runs from 2012 to 2019 and productivity is more stable in this period. There was slow growth on average and there is less evidence of differences in the speed of growth between high coverage industry-regions and low coverage industry-regions. This is the pre-treatment period included in our formal analysis and provides some reassurance that common trends is a reasonable assumption for the period that we study. The NMW also barely grew relative to median wages in this period. The NMW was 51.2 percent of median hourly wages in 2012 and rose slightly to 52.5 of median hourly wages in 2015, before rising much more steeply when the NLW was introduced to 60.6 percent in 2020. It is therefore unlikely that the NMW had significant impact on productivity in our pre-treatment period.

Figure 3: Productivity Index (2015=100) by Industry-Region coverage quartile, UK, 2002-2019

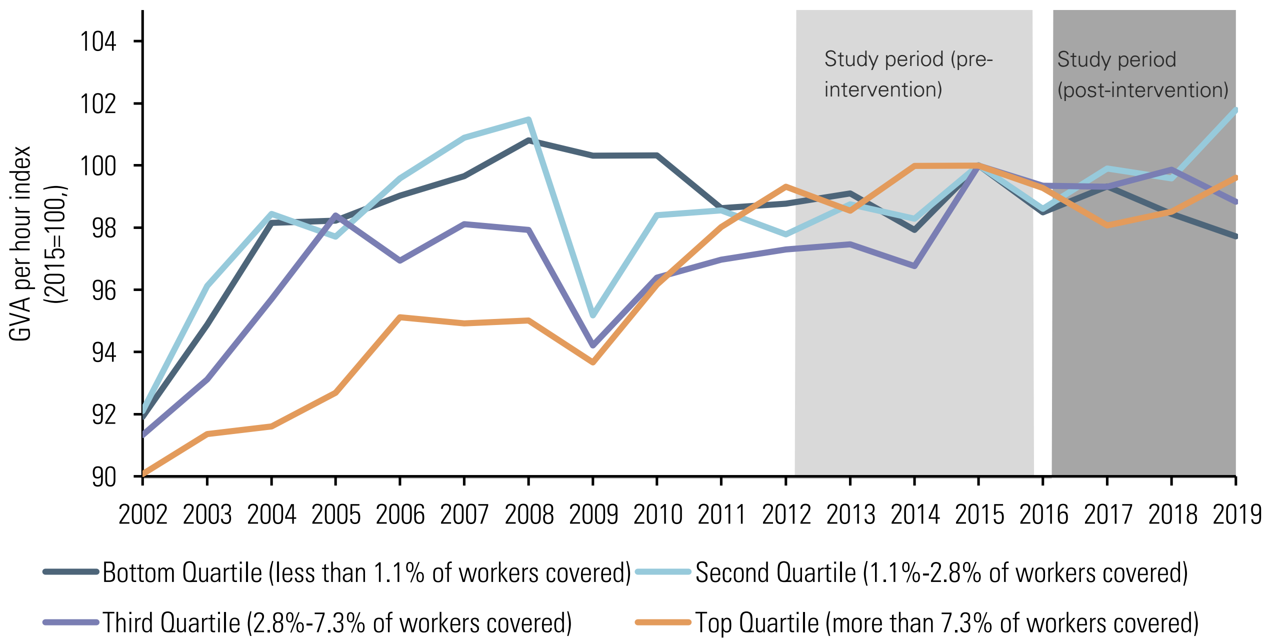

Source: LPC analysis of ONS Productivity by Industry-Region data (chained volume measure) and ASHE, using low pay weights. Industry-region cells are grouped into quartiles based on their 2015 minimum wage coverage for those 25 and over.

We investigate the common trends assumption further in the results section. The event-study plot shows there is some evidence of pre-treatment effects, although the effects are volatile and do not appear to show a clear trend. As a robustness check, we plot the pre-treatment time trend for the relationship between coverage and productivity and test whether post treatment the estimated effect differs from this trend. We discuss the event-study plot in more detail in Paragraphs 45-48.

We also add in industry-region specific time trends to our regression model, as another robustness check. These should capture any existing difference in the trend of productivity growth between high and low productivity areas, although they may also absorb part of the treatment effect. Our main findings are qualitatively unchanged by adding in these time trends.

We cannot test whether unobserved variables alter the trends in productivity post-treatment. The EU referendum is one key event, which could have influenced productivity growth differently across different industry-region cells. The EU referendum caused a sudden and sustained drop in the value of the pound relative to other currencies. Costa, Dhingra and Machin (2019) find that industries with greater exposure to that devaluation (measured by a greater share of intermediate imports) saw greater reductions in wages and a reduction in training. This could lead to a reduction in productivity in these more exposed industries. We do not control for the EU referendum in this study. A potential extension of the study is to add in a time-varying control which captures exposure to intermediate imports. We also cannot rule out other time-varying unobserved variables biasing the results.

b. No anticipation: The introduction of the National Living Wage did not impact productivity before it was introduced. This would be broken if firms anticipated that the minimum wage would rise and adjusted beforehand, which led to increased productivity in low-paying industries and regions.

   This assumption is likely to hold. The National Living Wage was a surprise announcement in the 2015 Budget in July 2015. This is nine months before it came into effect in April 2016. There was a window where firms could have adjusted in advance of the National Living Wage came into effect. However, this is unlikely to bias our results for two reasons. Firstly, we use the annual average of productivity for each year. If productivity increased in the first quarter of 2016 before the National Living Wage was formally introduced, this would be captured as a treatment effect in our regression. Second, changes in firm behaviour are likely to affect productivity with a lag. For instance, if firms increased training for workers, this will only gradually feed through into productivity.

35 We use standard errors clustered at the industry-region level. This allows for serial correlation in the error term from a given industry-region. (Bertrand, et al., 2004) show that serial correlation of errors can lead researchers to underestimate the standard deviation of their estimates in difference-in-difference studies. To avoid issues with serial correlation we cluster our standard errors at the industry-region level (across time periods.) We use CR2 errors, which include a small-sample adjustment, due to our relatively small number of clusters (180). Our standard errors are also robust to heteroskedasticity in the error term. Our standard errors do not allow for spatial auto-correlation across regions, if this does occur our standard errors could underestimate the uncertainty around our estimates.

Results:

36 The results provide no evidence that the introduction of the NLW affected productivity. Table 2 shows the results from three regressions with log GVA per hour as the explanatory variable. Column 1 shows the estimated coefficient using the 2015 coverage rate. We find no statistically significant effect of higher minimum wage coverage on productivity in the treatment period (2016-2019.) Our point estimate is slightly negative (-0.06), which would suggest that a 1 percentage point increase in coverage in an industry-region reduces productivity by approximately 0.06 percent. This would be a small negative effect. As coverage (for those aged 25 and over) increased from 4.3 percent in 2015 to 6.6 percent in 2016, our estimate would imply that the introduction of the NLW reduced productivity by 0.08 percent. However, the estimated effect is not statistically significant at the 5 percent level (or the 10 percent level.)

37 The 95 percent confidence interval from our main estimate covers a wide range of effects. A 1 percentage point increase in coverage in an industry-region could increase productivity by 0.1 percent in an area or reduce it by 0.22 percent. While this range may appear narrow, it is worth remembering that only 7 percent of workers are paid the minimum wage. If aggregate productivity increases by 0.1 percent, then productivity for affected workers needs to increase by a much larger amount. Our confidence intervals are wide, due to the limitations of our study design imposed by the available data. We only have data on a small number of industry-region cells (180) and within most cells a small percentage of workers are covered by the minimum wage. This means that if firms responded to the rise in the NLW by making productivity enhancements which just affected minimum wage workers (e.g., increased individual training), the affected workers would need large productivity increases for them to produce statistically significant estimates in our regressions.

38 However, the NLW could theoretically also have led firms in the affected industries to increase productivity for all workers. If the affected firms responded to the increase in the NLW by increasing capital investment or improved organisational structures, this would increase productivity for all workers in the affected industries (not just minimum wage workers). If this were the case, we would expect to find faster productivity growth in the industries more exposed to the minimum wage, but we do not find this. Our findings suggest the introduction of the NLW did not lead to large scale productivity enhancing investments in the affected industries and regions.

39 Column 2 shows the results using an alternative treatment measure, based on the gap between workers' pay and the forthcoming NLW rate. Our results are similar when using this alternative treatment measure. There is still a slightly negative point estimate, which is not statistically significant. This means that including workers paid between the 2015 National Minimum Wage and 2016 National Living Wage, does not affect our results.

40 Column 3 adds linear time trends to the model. This is to control for other factors which may mean that productivity grows faster in certain industry-region cells than in others. The results again look like our main model, a small negative effect which is not statistically significant. This provides some assurance that our results are not being biased by other time-varying variables, although it only controls for linear time trends. Our estimates could still be biased by omitted variables, which have stronger effects in some years than others, for instance the Brexit referendum in 2016.

Table 2: Regression results with Log GVA per Hour as outcome variable

	(1) Main model	(2) Main model with gap treatment measure	(3) Main model with industry-region time trends
Coverage Rate*Post-NLW	-0.06 (0.08), [-0.21,0.08]		-0.04 (0.1), [-0.21,0.13]
Gap Measure*Post-NLW		-0.07 (0.08), [-0.21,0.08]	
Year fixed effects	Yes	Yes	Yes
Industry-region fixed effects	Yes	Yes	Yes
Industry-region specific time trends	No	No	Yes
Number of observations	1440	1440	1440
R squared	0.96	0.96	0.98

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.
Standard errors are clustered at industry-region level and use CR2 method.
Standard errors are reported in circular brackets, 95 percent confidence intervals are reported in square brackets.

41 In Table 3 we present the results from regressions with mean wages rather than productivity as the treatment variable. This is a useful validation check of our approach. Many studies have shown that the introduction of the National Living Wage in 2016 did have substantively and statistically significant effects on the wages of affected workers (Aitken, et al., 2018; Cribb, et al., 2021; Avram & Harkness, 2019). Reassuringly, column 1 shows that if we apply our main model to wages rather than productivity we find a positive effect, statistically significant at the 0.01 percent level. This suggests that the introduction of the NLW in 2016 increased wages in high coverage industry-regions relative to low coverage areas. For each 1 percentage point in higher coverage, wages grew 0.25 percent faster.

42 However, in the model with industry-region time trends, we do not find a statistically significant effect of the introduction of the NLW on wages. This suggests caution in interpreting the results with industry-region time trends. As the industry-region time trends are estimated using post-treatment as well as pre-treatment data, part of the treatment effect could be absorbed into them.

Table 3: Regression results with Log Mean Wages as Outcome Variable

	(1) Main model	(2) Alternative treatment measure	(3) Main model with industry-region time trends
Coverage Rate*Post-NLW	0.25***		0.05
	(0.04), [0.17,0.33]		(0.06), [-0.08,0.18]
Gap Measure*Post-NLW		0.26***	
		(0.04), [0.17,0.34]	
Year fixed effects	Yes	Yes	Yes
Industry-region fixed effects	Yes	Yes	Yes
Industry-region specific time trends	No	No	Yes
Number of observations	1440	1440	1440
R squared	0.96	0.96	0.97

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.
Standard errors are clustered at industry-region level and use CR2 method.
Standard errors are reported in circular brackets, 95 percent confidence intervals are reported in square brackets.

43 We explore the effect of the NLW on productivity further by looking at how the estimated effect varies across the treatment period. Given firms may only be able to increase productivity gradually, we could imagine that productivity starts to grow faster a few years after the NLW was introduced. We can test this by estimating a model with dynamic treatment effects. We estimate the same regression as our main model, but we replace the treatment variable (2015 coverage rate interacted with a post-treatment indicator) with interactions of seven year-specific dummies with 2015 coverage rate. We put things in event time. This means 2016 is equivalent to year 0 as this was when the change was implemented. We use 2015 (the year before treatment) as the reference year for the time dummies, in line with standard practice. A positive coefficient for year t means that coverage is more strongly correlated with the outcome (wages or productivity) in year t than it was in 2015.

44 Figure 4 provides no evidence of a lagged effect of the NLW's introduction on productivity. The estimated treatment effect remains centred around zero throughout the treatment period (2016-2019). This provides more evidence to support our conclusion that the NLW did not affect productivity performance.

45 Event study plots are also useful ways to investigate the assumption of common trends. Figure 4 does provide some evidence that the common trends assumption may not hold for productivity. It appears productivity may have been growing slower for high coverage areas before the NLW was introduced. In two of the periods before the NLW was introduced, NMW coverage appears more positively related to productivity than it was immediately before the NLW was introduced. This suggests productivity growth could have been on different trends for high coverage areas than low coverage areas. If this is the case, then our estimates of the NLW's effect on productivity would be biased downwards.

Figure 4: Estimated dynamic impacts of introduction of NLW on productivity

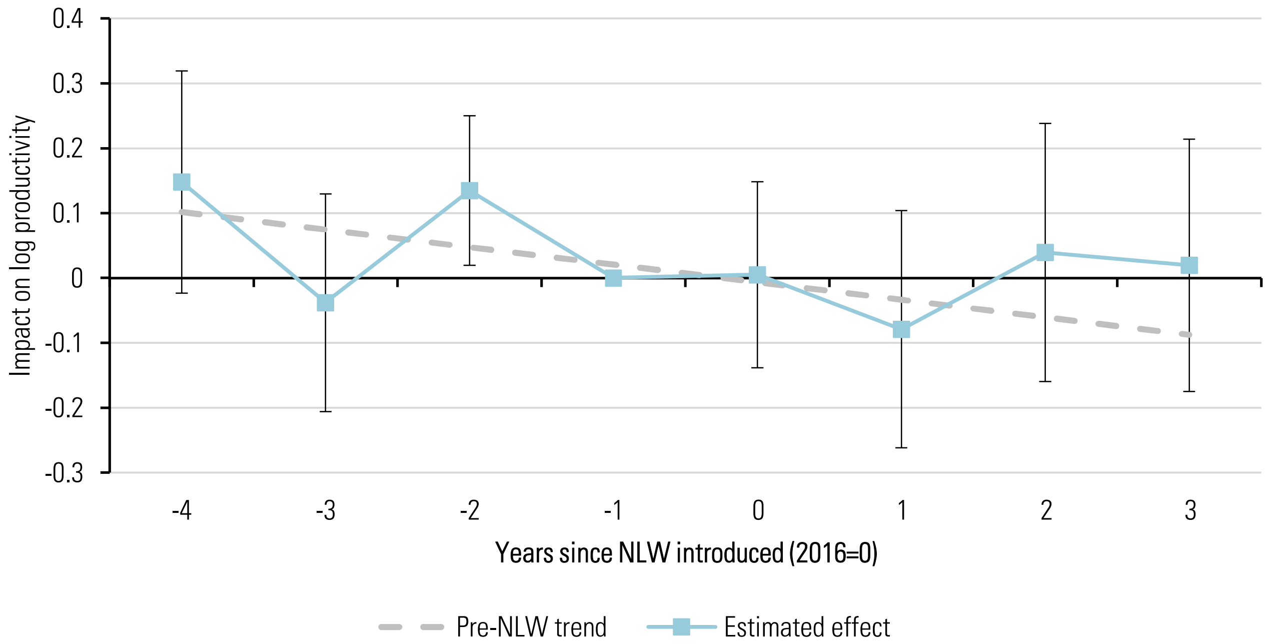

Source: LPC analysis using ONS productivity by Industry-Region data (chained volume measures) and ASHE, 2015, UK, low pay weights. Minimum wage coverage only includes workers aged 25 and over. Excludes real estate activities.

46 However, we can adjust for a pre-existing trend. If we assume there is a pre-existing trend in the relationship between minimum wage coverage and productivity, we still do not find a statistically significant effect of the NLW on productivity. The grey line shows the linear trend based on the four years before the NLW was introduced. We test whether the estimated effect is statistically significantly different from the pre-existing trend and find that in no year is the difference statistically significant at the 5 percent level. This is shown by the 95 percent confidence intervals crossing the trend line in each of the four post-treatment periods. The estimated effect does move above the trend line in periods 2 and 3 (2018 and 2019), which would make the estimated effect positive (relative to the trend) rather than negative. However, even if the estimated effect is positive, there is no statistically significant difference from the pre-existing trend. We still have found no clear evidence of the NLW improving productivity.

47 We only adjust for a linear trend and the relationship between productivity and minimum wage coverage could also be affected by one-off factors which only affect certain years e.g., the EU referendum in 2016. We can also see in the pre-treatment period, the estimated effect is volatile and does not follow closely to the linear trend, which suggests additional caution in interpreting the results. We cannot rule out time-varying variables not included in our model biasing our results.

48 Figure 5 shows the dynamic of the introduction on wages. It provides a good comparison to the evidence on productivity. There are no statistically significant pre-treatment effects in the case of wages, which provides reassurance that common trends assumption holds. Then following treatment, the estimated coefficient gradually increases. The coefficients are statistically significant from one year after the NLW was introduced. This reflects the fact that the NLW continued to increase at a faster rate than average wages between 2017 and 2019, as rates were recommended so that the government hit its target of the National Living Wage equalling 60 percent of National Living Wage in 2020.

49 The fact that the 2016 effect of NMW coverage on average wages is not statistically significant shows that our estimates have fairly low power. The introduction of the NLW led to the largest increase in the minimum wage in that period. Other studies have shown that the NLW did have an immediate impact on wages. We do not find this effect statistically significant as we use a fairly small sample of industry-regions and many of the industry-regions have a small share of minimum wage workers. If we had firm-level data or high quality data at smaller geography we may find evidence of productivity effects, which we cannot find with our current approach.

50 Figure 6 summarises our results. It shows the total growth in productivity and wages between 2015 and 2019 by industry-region. While wages grew faster between 2015 and 2019 in industry-regions more exposed to the NLW, the same trend does not hold for productivity. The introduction of the NLW appears to have increased wages, but not productivity. This suggests that firms absorbed the increased labour costs through other channels than productivity, such as by raising prices or reducing profits. These alternative channels are discussed in the LPC's review of the first stage of the NLW. (Low Pay Commission, 2022)

Figure 5: Estimated dynamic impact of introduction of NLW on wages

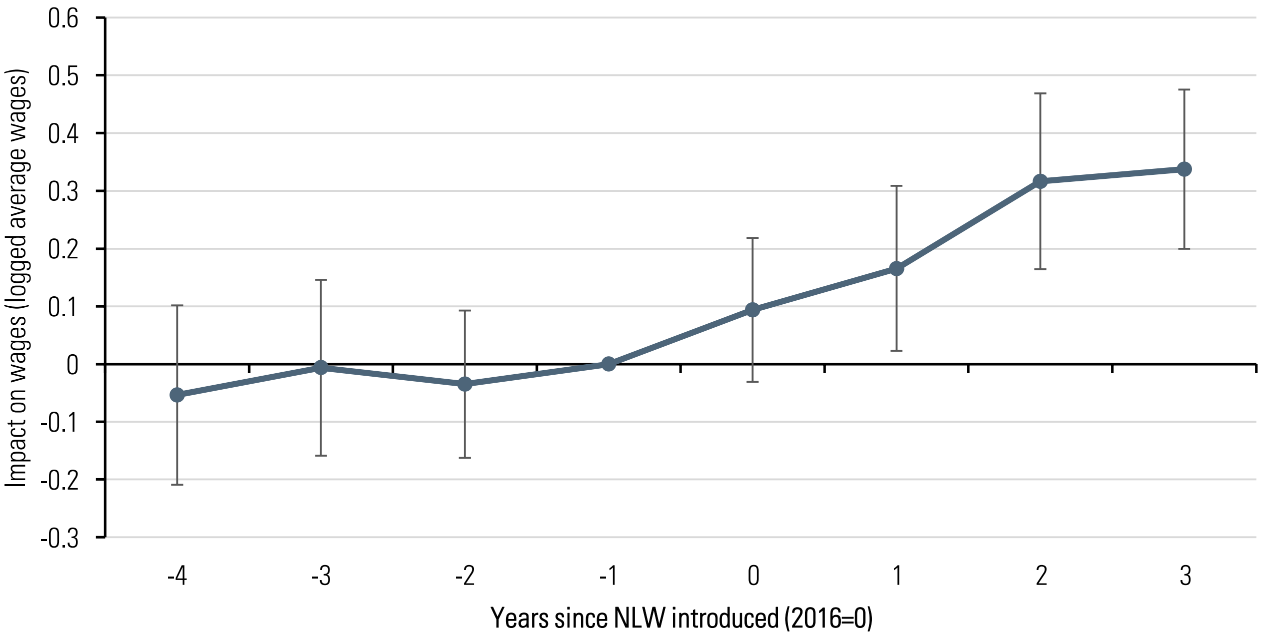

Source: LPC analysis using ASHE, UK, low pay weights used for coverage, standard weights for mean hourly wages. Minimum wage coverage only includes workers aged 25 and over. Excludes real estate activities.

Figure 6: Growth in Productivity and Wages by minimum wage coverage at industry-region level (real terms, 2015-2019)

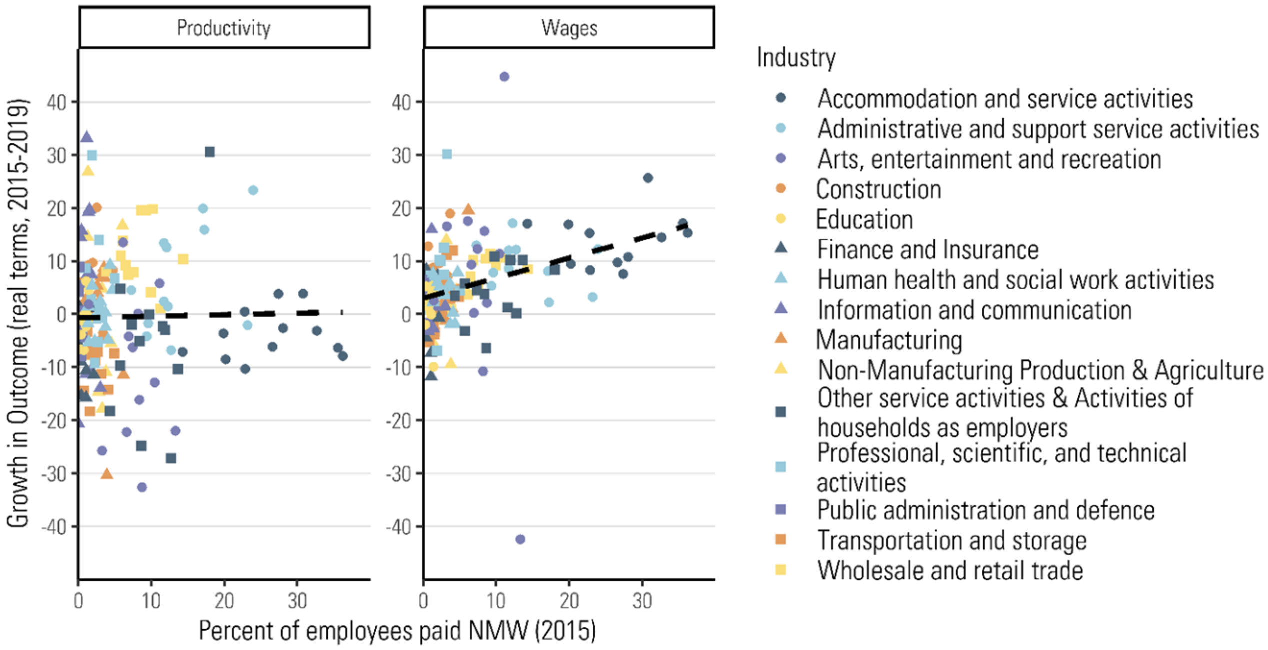

Source: LPC analysis using ONS Productivity by Industry-Region data and ASHE, UK, low pay weights used for coverage, standard weights for mean hourly wages. Minimum wage coverage only includes workers aged 25 and over. Excludes real estate activities.

Robustness Checks:

51 We also carry out robustness checks for a non-linear effect of minimum wage coverage on productivity. Our main results use a continuous treatment variable and test whether there is a linear relationship between the coverage rate in 2015 and the growth of productivity from 2015 to 2019. This may not capture the effects of the NLW on productivity if there are non-linear effects. For instance, if the NLW increased productivity in all industry-regions beyond a given threshold but had no impact below that threshold. We test for these non-linear treatment effects using three binary treatment variables:

1) Binary Treatment Variable 1: There are six industry-regions in 2015, which had a coverage rate of 0. We treat these six industry-regions as the control group and all other industry-regions as the treatment group

2) Binary Treatment Variable 2: We include all industry-regions in the bottom quartile for coverage (2.8 percent) in the control group and all other authorities as the treatment group.

3) Binary Treatment Variable 3: We test whether the NLW increased productivity in the highest exposure quartile of authorities. We use the fourth quartile of coverage as a treatment group and the third quartile as a control group. We use the third quartile as a control group they are likely to be more similar in their other characteristics to the fourth quartile industry-regions than other than the bottom two quartiles.

52 We still find no evidence of the NLW impacting productivity, when using binary treatment variables. Table 4 summarises the results from a regression using the same model as in equation 1 but substituting the continuous treatment variable for binary treatment variables. The binary treatment variables are equal to 1 if an observation is in the treatment group and 0 otherwise. We find that for each of the specifications, the effect remains centred around 0 and is not statistically significant. This provides more evidence to support our conclusion that the introduction of the NLW did not substantively alter productivity for affected workers.

Table 4: Impact of NLW on Productivity using Binary Treatment Variables

	(1)	(2)	(3)
Binary Treatment Variable 1	0.02		
	(0.02), [-0.02,0.06]		
Binary Treatment Variable 2		0.00	
		(0.01), [-0.03,0.02]	
Binary Treatment Variable 3			-0.02
			(0.02), [-0.07,0.02]
Year fixed effects	Yes	Yes	Yes
Industry-region fixed effects	Yes	Yes	Yes
Industry-region specific time trends	No	No	Yes
Number of observations	1440	1440	720
R squared	0.96	0.96	0.93

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.
Standard errors are clustered at industry-region level and use CR2 method.
Standard errors are reported in circular brackets, 95 percent confidence intervals are reported in square brackets.

53 We also tested the impact of the NLW on productivity using data on GVA per hour at the local authority district level. The estimated effects are small and positive, but not statistically significant effect when using this approach. The results are shown in Table 5. The approach is the same as our main approach, except rather than using industry-region-year cells as observations, we use local authority-year cells as observations. We cluster at the local authority level to adjust for serial correlation in the residuals within groups. Productivity is only available in current prices for local authorities, so we use this data and put it into real prices using the GDP deflator. This means that these results (like our main results) might also capture local price effects rather than productivity effects.

54 The confidence intervals are wider using this approach, despite there being more observations. This is due to a lower level of variation in coverage between local authority areas than between different industry-region cells. The maximum level of coverage of any authority in 2015 was 11 percent, whereas some industry-regions cells had over 30 percent coverage.

Table 5: Regression results using local authority level data

	(1) Main model	(2) Alternative treatment measure	(3) Main model with industry-region time trends
Coverage Rate*Post-NLW	0.14 (0.15), [-0.15,0.43]		0.04 (0.07), [-0.11,0.19]
Gap Measure*Post-NLW		0.19 (0.13), [-0.07,0.46]	
Year fixed effects	Yes	Yes	Yes
Industry-region fixed effects	Yes	Yes	Yes
Industry-region specific time trends	No	No	Yes
Number of observations	2896	2896	2896
R squared	0.97	0.97	0.99

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.
Standard errors are clustered at local authority level and use CR2 method.
Standard errors are reported in circular brackets, 95 percent confidence intervals are reported in square brackets.

55 Table 6 presents some further robustness checks. In column 1 we add industry-year fixed effects into the main model. This means that we are only exploiting variation in coverage and productivity growth between different regions within the same industry. Similar to the results from Table 5, the estimated effect of the NLW on productivity is positive but not statistically significant. Column 2 excludes non-market sectors (Public Administration, Health, Education and Household Activities) from the analysis. Measuring productivity in these sectors is difficult, and systematic measurement error could bias our main results. Reassuringly, the estimated effect remains negative and not statistically significant after excluding these sectors. Column 3 replicates the main analysis on only the 6 industry groups with highest coverage. This addresses the concern that we may have different effects in high coverage sectors than low coverage sectors (e.g., finance). The estimated effect remains negative and not statistically significant at the five percent level.

56 Column 4 repeats our main model excluding the hospitality (accommodation and food services) sector. Excluding hospitality reverses the sign of the estimated effect from negative to positive, but the effect remains not statistically significant at the 5% level. The coverage rate is the highest in the hospitality sector, and particular high in some regions. Productivity growth was weak across the sector. Due to the high coverage in the hospitality sector, it has a large influence on our results. Productivity improvements are likely to be particularly difficult to make in hospitality, as it is difficult to substitute capital for workers in an industry based on interpersonal service (e.g., having a meal served to you by waiters, a barperson pouring you a drink). However, even when we exclude hospitality, we do not find a statistically significant effect. These robustness checks provide reassurance that our main conclusions are not sensitive to how we specify the regression model.

Table 6: Further Robustness Checks

	(1) Main model with industry-year fixed effects	(2) Main model excluding non-market sectors	(3) Main model on high coverage sectors	(4) Main model excluding hospitality sector
2015 Coverage Rate * Post NLW	0.12	-0.07	-0.12	0.20
	(0.23)	(0.07)	(0.11)	(0.19)
	[-0.33,0.57]	[-0.22,0.07]	[-0.33,0.09]	[-0.16,0.57]
Year fixed effects	Yes	Yes	Yes	Yes
Industry-region fixed effects	Yes	Yes	Yes	Yes
Industry-year fixed effects	Yes	No	No	No
Number of observations	1440	1056	576	1344
R squared	0.97	0.97	0.94	0.96

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.
Standard errors are clustered at local authority level and use CR2 method.
Standard errors are reported in circular brackets, 95 percent confidence intervals are reported in square brackets.

Conclusions:

57 In conclusion, we have found no evidence that the introduction of the National Living Wage changed productivity in affected industries. We tested whether the NLW increased productivity, by comparing the growth in productivity for industry-region cells with a high share of minimum wage workers against the growth in productivity for industry-region cells with a lower share of minimum wage workers. We found that productivity grew at a similar rate between 2016 and 2019 for areas with high or low minimum wage coverage. We have conducted a wide range of robustness checks and in none of our specifications does the NLW have a statistically significant effect on productivity.

58 Our results contrast with some previous studies which found that the National Minimum Wage was associated with increased productivity (Croucher & Rizov, 2011; Riley & Rosazza-Bondibene, 2017). The National Living Wage provided sustained large increases to the minimum wage in the UK, so if the previous increases had raised productivity, we would expect the National Living Wage to further that increase.

59 One potential explanation for this tension is that the productivity effects from the National Minimum Wage came from intensification. Intensification involves working workers harder, for instance shortening breaks. The Chartered Institute of Personnel and Development survey data shows that some employers affected by the NLW have focused on increasing worker effort (23 per cent of private sector firms affected by the NLW and 30 per cent in the public sector). Intensification could lead to improvements in the level of productivity for low-paid workers, but it cannot permanently increase the growth rate of productivity. There is a limit to how hard workers can work. Firms may have already achieved the gains from intensification in the early years of the National Minimum Wage, leaving little room for further benefits.

60 Another potential explanation is that recent macroeconomic events could have biased our results. Our results are robust to the existence of a pre-existing trend, but they could be affected by shocks to productivity following the introduction of the NLW. One potential shock that could bias our results is the 2016 EU referendum. (Costa, et al., 2019) found that the depreciation in the pound following the EU referendum reduced wages and training in industries which were more exposed to the depreciation. This referendum effect could mask the effect of the NLW, although the fact we find wage effects provides some validation for our approach. One extension to this work would be to add in a measure of exposure to exchange rates as a control, to test whether this alters our results.

61 There are also challenges to measuring productivity. Our study improves on previous papers by using a measure of output per hour rather than output per job. This means our study would not construe increased hours as increased productivity. However, there are still potential issues with the data that could affect our results. Our productivity data uses national level deflators so cannot separate changes in local prices from changes in productivity. There also could be measurement error in our measures of productivity.

62 Finally, it is important to remember that there are fairly wide confidence intervals around our estimate. Our 95 percent confidence intervals for our estimate ranges from a 1 percentage point increase in the share of minimum wage workers decreasing productivity by 0.2 percent to a 1 percentage point increase in the share of minimum wage workers increasing productivity by 0.1 percent. Future studies which use firm level data may be able to identify the effect of the NLW on productivity with more precision.

Works Cited

Aitken, A., Dolton, P. & Riley, R., 2018. The Impact of the Introduction of the National Living Wage on Employment, Hours and Wages. *Research Report for the Low Pay Commision.*

Avram, S. & Harkness, S., 2019. The impact of minimum wage upratings on wage growth and on the wage distribution. *Research Report for the Low Pay Commission.*

Bean, C., 2016. *Independent Review of UK Economic Statistics: final report,* s.l.: s.n.

Bell, B. & Machin, S., 2018. Minimum Wages and Firm Value. *Journal of Labour Economics.*

Bernini, M. & Riley, R., 2016. Exploring the relationships between the NMW and productivity. *Report to the Low Pay Commission.*

Bernini, M. & Riley, R., 2016. Exploring the Relationships between the NMW and Productivity. *Report to the Low Pay Commission.*

Bertrand, M., Duflo, E. & Mullainathan, S., 2004. How Much Should We Trust Differences-in-Differences Estimates?. *The Quarterly Journal of Economics.*

Brochu, P. & Green, D., 2013. The Impact of Minimum Wages on Labour Market Transitions. *The Economic Journal.*

Costa, R., Dhingra, S. & Machin, S., 2019. Trade and Worker Deskilling. *NBER Working Paper.*

Crawford, C., Jin, W. & Simpson, H., 2013. Firm's productivity, investment and training: What happened during the recession and how was it affected by the national minimum wage?. *Report to the Low Pay Commission .*

Cribb, J. et al., 2021. *The impact of the NLW on employment, wages and household incomes,* s.l.: s.n.

Croucher, R. & Rizov, M., 2011. The Impact of the National Minimum Wage on Labour Productivity in Britain. *E-Journal of International and Comparative Labour Studies.*

Draca, M., Machin, S., Reenen, v. & John, 2011. Minimum Wages and Firm Profitability. *American Economic Journal: Applied Economics.*

Draca, M., Machin, S. & van Reenen, J., 2005. The Impact of National Minimum Wage on Profits and Prices. *Report to the Low Pay Commission.*

Dube, A., 2019. *Impacts of minimum wages: review of the international evidence,* s.l.: HM Treasury.

Dube, A., Lester, W. T. & Reich, A., 2010. Minimum Wage Effects Across State Borders: Estimates using Contiguous Counties. *Review of Economics and Statistics,* Volume 92, pp. 945-964.

Dustmann, C. et al., 2021. Reallocation Effects of the Minimum Wage. *Quarterly Journal of Economics.*

Engbom, N. & Moser, C., 2021. Earnings Inequality and the Minimum Wage: Evidence from Brazil. *NBER Working Paper.*

Forth, J. & O'Mahoney, M., 2003. The Impact of the National Minimum Wage on Labour Productivity and Unit Labour Costs. *Report to the Low Pay Commission.*

Forth, J., Rincon-Aznar, A. & Robinson, C., 2009. The Impact of Recent Upratings of the National Minimum Wage on Competitiveness, Business Performance and Sector Dynamics. *Report to the Low Pay Commission .*

Frontier Economics, 2020. The Impact of National Living Wage on Businesses. *Report for the Low Pay Commission.*

Galindo-Rueda, F. & Pereira, S., 2004. The Impact of the National Minimum Wage on Business.

Low Pay Commission, 2020. *National Minimum Wage. Low Pay Commission Report 2020.,* s.l.: s.n.

Low Pay Commission, 2022. *Review of the First Stage of the National Living Wage,* s.l.: s.n.

Low Pay Commission, 2022. *Review of the First Stage of the National Living Wage,* s.l.: s.n.

Manning, A., 2016. *The Elusive Employment Effect of the Minimum Wage,* s.l.: s.n.

Riley, R. & Rosazza-Bondibene, 2., 2017. Raising the Standard: Minimum wages and firm productivity. *Labour Economics .*