Technical Annex B: Estimating the Impact of the 2016 Introduction of the National Living Wage

In this Technical Annex, we provide details behind the two different approaches taken in the report to estimating the impact of the 2016 introduction of the National Living Wage (NLW) in the UK.

Grouping estimator

Our first approach uses a grouping estimator, closely following the method in Manning (2016) who evaluated the impact of the 1999 introduction of the National Minimum Wage. A similar approach was also used by Harasztosi and Lindner (2019) in their evaluation of the Hungarian minimum wage. The variation in the bite of the NLW comes from the fact that the affected share of workers who were earning below the 2018 level of the NLW (s_{agr}) varied greatly across regional and demographic groups. We use 12 (NUTS 1) regions, 4 age categories (25-29, 30-39, 40-49, 50-64), and 2 gender categories as in Manning, except we restrict attention to those 25 or older as that is the group covered by the NLW. Across these 96 groups, the affected share varies between 6 and 34 percent.

To construct the cell-specific employment, we used data from the Annual Population Survey for the years 2014-2018. Our outcomes include both estimates using headcount employment, as well as full-time-equivalent (FTE) employment which incorporates information on average weekly hours. To obtain the cell-specific means of weekly hours and the real wages, we used the Annual Survey of Hours and Earnings (AHSE) data between 2014-2018. We used the UK Consumer Prices Index (CPI) to convert all wages into 2018 prices (using the April 2018 CPI value as a benchmark). Note that the ASHE does not cover the self-employed and those who were unpaid during the survey reference period.

The key outcomes we consider are changes in group level log average wages, log headcount employment, and log FTE headcount between 2014 (the year prior to the announcement of the NLW) and 2018 (latest year available).

The regression equation can be written as follows:

$$\Delta Y_{agr} = \beta \times s_{agr} + \alpha_a + \gamma_g + \rho_r + e_{agr}$$

Here the outcomes, ΔY_{agr}, are the change in log average wage, the change in log employment, or the change in log FTE employment between 2014 (the year prior to the announcement of the NLW) and 2018 (latest year available). Just as in Manning (2016), the regressions control for age group, region and gender specific fixed effects ($\alpha_a, \gamma_g, \rho_r$).

Table B1 Impact of NLW on Wages and Employment Changes Between 2014 to 2018

	Change in Log Wage	Change in Log Headcount Employment	Change in Log FTE Employment
Share below 2018 NLW (s_{agr})	0.356	-0.047	-0.005
	(0.150)	(0.112)	(0.126)
Implied OWE		-0.131	-0.014
		(0.318)	(0.325)

Notes. Outcomes are changes between 2014 and 2018. Estimates control for age, gender and region fixed effects. Robust standard errors in parentheses. Implied own-wage employment elasticity (OWE) calculated by dividing the employment coefficient by the wage coefficient; standard errors are calculated using the delta method.

Table B1 reports the estimates. Figure B1 shows these visually by partialing out the fixed effects from both the outcomes, ΔY_{agr}, and the key independent variable, s_{agr}, and then plotting the residuals and fitted lines.

Figure B1 Impact of NLW on Residual Wages and Employment Changes Between 2014 to 2018

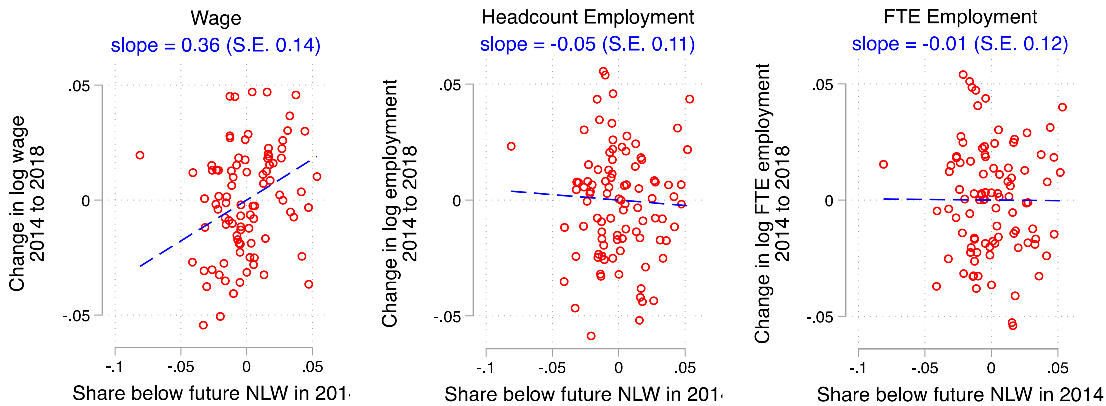

The left panel shows that average wage growth was much faster following the introduction of the NLW in groups with larger share of affected workers: the regression coefficient is 0.36 (S.E. 0.14). In contrast, there is no statistically significant or sizable relationship between the affected share and employment. The regression coefficient is -0.05 (S.E. 0.11). Together, the implied OWE (own-wage employment elasticity) is -0.131 (S.E. 0.318), which suggests little impact on employment following the introduction of the NLW, and is similar to the estimate from Aitken, Dolton and Riley (2019). To account for possible hours effects, we can use full-time-equivalent (FTE) employment instead of headcount employment. The OWE using FTE employment is -0.014 (S.E. 0.325), which does not suggest a negative impact on headcount or hours.

A key assumption in this exercise was that conditional on region, age and gender fixed effects, the outcomes would have moving in a parallel trend across the 96 cells were it not for the NLW. While this assumption is not directly testable in the post 2014 period, we can assess whether this was true in the period immediately prior to 2014. We do this by considering a placebo outcome $\widetilde{\Delta Y_{agr}}$ which are the differences in outcomes between 2012 and 2014. We regress these prior changes on the same actual shares below the 2014 NLW, s_{agr}. The parallel trends assumption suggests the coefficients should be zero. As Table B2 shows, that is indeed the case. The coefficient on the wage and employment outcomes are all negative in sign, and statistically indistinguishable from zero, consistent with a lack of any pre-existing trends, thereby supporting the validity of the research design.

Table B2 Impact of NLW on Placebo Wages and Employment Change between 2012 to 2014

	Change in Log Wage	Change in Log Headcount Employment	Change in Log FTE Employment
Share below 2018 NLW (s_{agr})	-0.192	-0.146	-0.073
	(0.144)	(0.104)	(0.125)

Bunching Estimator

An alternative approach for estimating the impact of the NLW introduction considers the number of low-wage jobs, following the approach developed in Cengiz et al. (2019). The key idea is to look at the number of jobs below versus at-or-just-above the minimum wage. A fall in the jobs paying below the minimum suggests the policy was binding. Comparing the number of these "missing jobs" to the increase in jobs paying at-or-just-above the minimum wage ("excess jobs") provides an estimate for the employment effect of minimum wage increases.

This method requires us to form a "counterfactual" frequency distribution of wages. Cengiz et al. (2019), and the results using high minimum wage US states in this report, use employment changes by wage bins in states that do not raise the minimum wage to form this counterfactual. An obvious limitation in the UK context is that there is no "control jurisdiction" that did not raise the statutory minimum wage; as a result, we have to use a before-after comparison in the frequency distribution of wages to infer the impact on low-wage jobs. Namely, we use the 2014 frequency distribution of jobs as the counterfactual, after adjusting for wage inflation and population growth. This lack of a control group makes the evaluation using this bunching approach less convincing as we consider longer run horizons, as the 2014 frequency distribution of the deflated wages is less likely to remain a valid counterfactual over the longer horizon. We return to this point below.

We use data from ASHE to construct the number of jobs held by those 25 or over by £1 wage bins relative to the 2018 NLW. Since nominal wages are rising each year, this would spuriously suggest a fall in the number of jobs below a nominal wage threshold due to wage inflation. To correct for that, we deflate the wages using a measure of wage inflation, namely the annual rate of growth in the mean wage in the ASHE. Additionally, we account for population growth by dividing the job counts by the national population of those 25 or over (using data from the Annual Population Survey (APS), from the Office of National Statistics).

Figure B2: Impact of the Introduction of the NLW on the Counts of Jobs by Wage Bins, 2014 to 2018

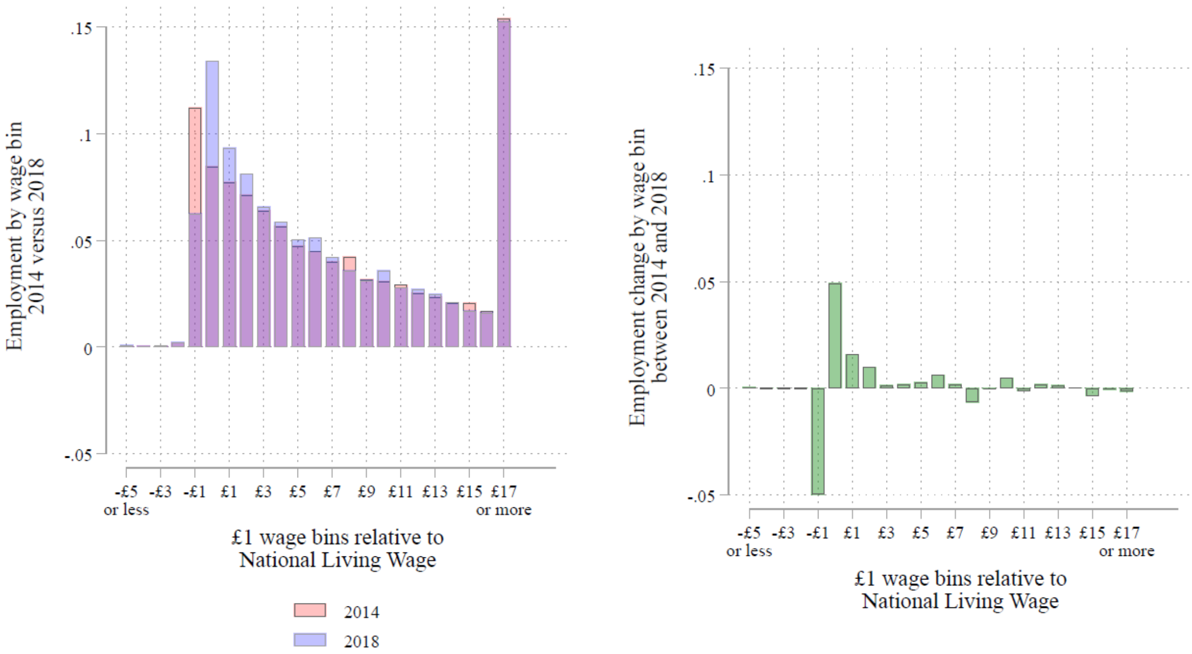

Source: Author's calculations based on ONS data (ASHE and APS data, 2014-2018).

Figure B2 (reproduced from chart 4.J in the main report) plots the per-capita employment rate by wage bins in 2014 (the year prior to the announcement of the NLW) and 2018. The wage bins are relative to the 2018 National Living Wage. We also divide the per-capita change in employment by wage bin by the 2014 employment-to-population-ratio (EPOP). This normalization allows us to interpret the changes by wage bin as changes in the number of jobs relative to the 2014 total, but adjusted for population growth. (This is analogous to the procedure used in Cengiz et al. 2019.)

Comparing the job counts at wages in the upper tail—say a few pounds above the NLW—the figure suggests the 2014 frequency distribution of wages was likely a reliable counterfactual for the 2018 distribution: the distribution of employment by wage bins were quite similar except just around the new NLW. There was a clear reduction in jobs paying below the 2018 NLW, consistent with a clear bite of the policy. At the same time, we see additional jobs paying at or above the NLW which sum to slightly more than the missing jobs, indicating no net job loss.

Figure B3 shows the changes in jobs across the distribution for all years between 2012 and 2018. Reassuringly, there is virtually no change in the distribution in the years prior to the announcement of the NLW (i.e., 2012, 2013). There is a very slight indication of upgrading wages in 2015 around the future NLW. In contrast, 2016 (and to certain extent 2017) saw a major reduction in jobs paying below the NLW, and a concomitant increase in jobs paying at or slightly above. Finally, we see that the spike at the NLW was somewhat larger in 2018 than before. We also see some spillovers going up to around £3 above the new NLW.

Figure B3: Changes from 2014 in Counts of Jobs by Wage Bins, 2012 to 2018

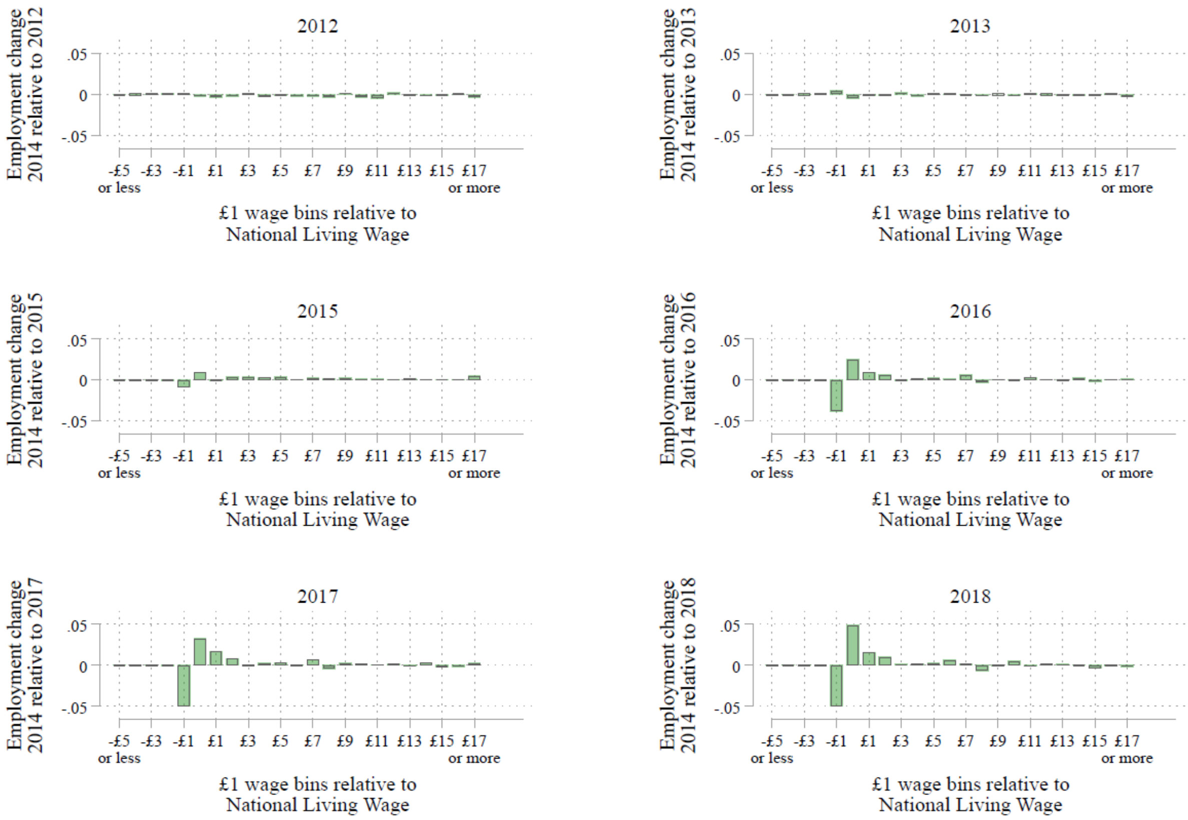

Source: Author's calculations based on ONS data (ASHE and APS data, 2012-2018).

Figure B4: Counts of Jobs paying Below and At-or-just-above the Real 2018 NLW – 2011 to 2018

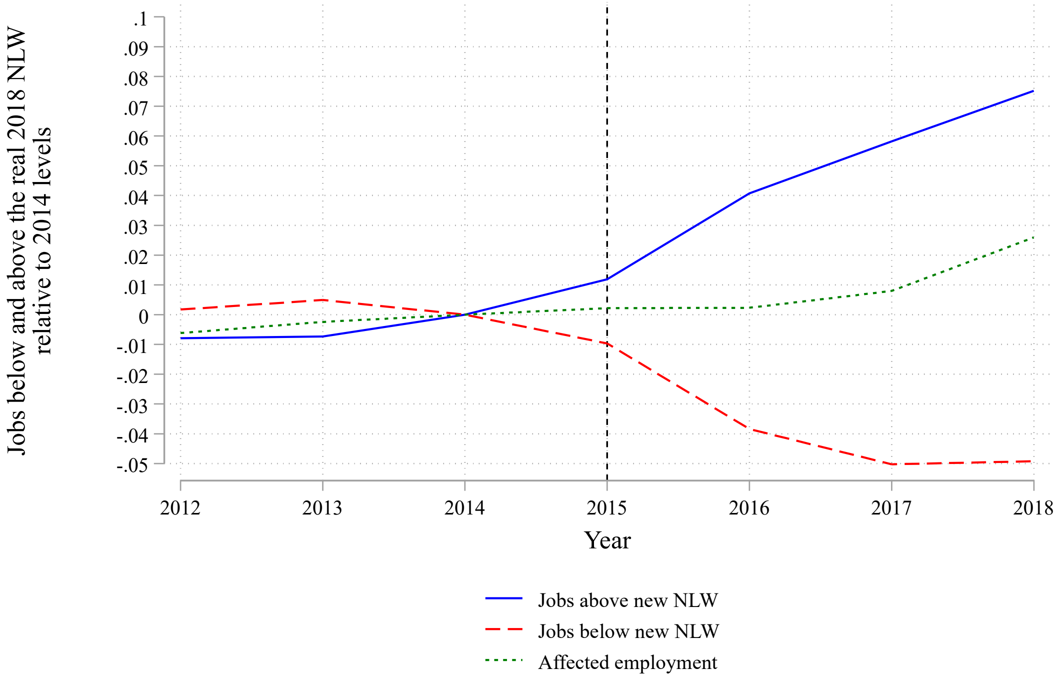

Source: Author's calculations based on ONS data (ASHE and APS data, 2012-2018).

Figure B3 (reproduced from Chart 4.K in the main report) shows the same information in a slightly different way. The figure plots the evolution of jobs paying right below the NLW ("missing jobs below"), and jobs paying at or up to £3 above the NLW ("excess jobs above") over time. The sum of the two produces the overall employment effect. First, we again see relatively stable patterns in the missing and excess jobs (as well as their sum) in the years through 2014, consistent with Figure B3. This is reassuring and suggests there were not sizable pre-existing trends. Second, we see a sharp reduction in the jobs paying below the NLW in 2016 and 2017. At the same time, there was an equally sized increase in jobs paying at or just above the NLW. Summing the two, the overall employment impact is moderately positive if we use the 2018 counts, or right around zero if we use 2016 or 2017 counts. As the real frequency distribution is unlikely to remain stable for an extended period of time, there is an argument to focus on the first few years of implementation. At the same time, we note that there does not seem to be changes anywhere else in the frequency distribution of wages through 2018.

Overall, we can draw several conclusions. First, the NLW was quite binding for those 25 or older, as indicated by the missing jobs below the NLW level. Second, there was a moderate amount of wage spillovers, going up to around £3 above the NLW, as indicated by the excess jobs right above the NLW. Third, the overall number of low-wage jobs (e.g. number of jobs paying below NLW+£3) was virtually unchanged between 2011 and 2017, including after the announcement of the NLW in 2015, while it rose slightly in 2018. Finally, and importantly, the changes in jobs in bins +£4 or higher remained virtually unchanged through 2018. This lack of any movements in the upper tail is crucial for a valid research design, as it suggests the (deflated) wage frequency distribution was fairly stable over the period under study, which is a necessary condition to ensure that the 2014 distribution was a valid counterfactual.