IFS

Jonathan Cribb
Giulia Giupponi
Robert Joyce
Attila Lindner
Tom Waters
Thomas Wernham
Xiaowei Xu

Working paper

21/48

The distributional and employment impacts of nationwide Minimum Wage changes

The Distributional and Employment Impacts of Nationwide Minimum Wage Changes[*]

Jonathan Cribb[†]   Giulia Giupponi[‡]   Robert Joyce[†]   Attila Lindner[§]
Tom Waters[¶]   Thomas Wernham[¶]   Xiaowei Xu[†]

December 7, 2021

Abstract

We estimate the effect of the introduction of the UK's National Living Wage in 2016, and increases in it up to 2019, using a new empirical method. We apply a bunching approach to a setting with no geographical variation in minimum wage rates. We effectively compare employment changes in each part of the wage distribution in low-wage areas to employment changes among similar workers living in higher-wage areas who are less exposed to increases in the national minimum wage because their nominal wages are further above it. We find substantial positive wage effects, including statistically significant spillovers up to around the 20th percentile of wages. Overall we find small negative effects on employment which are not statistically significant. We combine these estimates with a tax and benefit microsimulation model to estimate the impact on household incomes. The largest gains go to the middle of the overall working-age income distribution, though they are more concentrated within the bottom third if we consider only households with someone in paid work. The gains to poorer working households are limited by the withdrawal of means tested benefits as earnings increase. Effects of minimum wages on household incomes are very sensitive to the size of employment effects.

Keywords: minimum wage, labor demand, income inequality, poverty.
JEL Codes: J23, J38, D31.

[†]Institute for Fiscal Studies
[‡]Bocconi University, Institute for Fiscal Studies and CEPR
[§]University College London, Institute for Fiscal Studies and CEPR
[¶]Institute for Fiscal Studies and University College London

1. Introduction

The National Living Wage (NLW) is currently the flagship policy aimed at helping the low paid in the UK. Applying to employees aged 25+, the NLW brings the minimum wage far higher than it was just a few years ago and close to the international frontier. The UK government has set a target for it to reach two thirds of median wages by 2024, while extending it to workers aged 21-24. The crucial role that minimum wage policy has taken in tackling wage inequality at the bottom of the distribution makes it imperative to understand the impacts that the NLW is having on the labour market and on household incomes. This paper aims at providing a comprehensive assessment of those effects. We first study impacts on hourly wages, hours of work and employment, and then analyse how these labour market effects feed through to families' and households' net incomes, after accounting for taxes and benefits.

A large number of analyses of minimum wage policies have been undertaken in the UK and in other countries (for reviews, see Neumark and Wascher (2008); Belman and Wolfson (2014); Dube (2019*a*)). Taken together, most UK (and other countries') studies have shown that rises in the minimum wage had substantial impacts on wages towards the bottom of the wage distribution, and that associated wage costs were largely absorbed along margins other than employment or hours of work (Manning, 2021). However, the NLW has moved us well beyond previous minimum wage levels, and close to the international frontier when measured as a fraction of average earnings. This limits the applicability of previous evidence derived in the UK or elsewhere.[1]

This paper proposes a new empirical methodology to estimate the impacts of the minimum wage on employment, hours, wages and earnings in a context in which a single minimum wage policy applies to the entire country and no geographical variation in minimum wage rates is available. Our method refines the regional-variation approach pioneered by Card (1992) by tracing out employment changes throughout the whole frequency distribution of wages as in Harasztosi and Lindner (2019) and Cengiz et al. (2019).

The main idea of the paper can be summarised as follows. Similar to Card (1992), we exploit the fact that a significant component of the differences in wages across locations are related to the location itself and not to individual characteristics, at least at the lower end of the wage distribution. For example, a cleaner in London has similar skills to a cleaner in Hull, but in fact, the cleaners in London are paid more – and are therefore less exposed to a national minimum wage policy. We document the contribution of local areas to wages using either a Mincerian regression with demographic controls or a two-way fixed effects regression with person and

---

[1]Most of the US literature exploits methods based on state-level variation (Card and Krueger, 1994; 2000; Dube, Lester and Reich, 2010; Neumark, Salas and Wascher, 2014; Cengiz et al., 2019), and more recently city-level variation (Dube and Lindner, 2021), with only a relatively small number of studies exploiting geographic variation in bite (Card, 1992; Clemens and Wither, 2019). Conversely, methods based on variation in bite across regions or demographic groups are largely applied in Europe, where most countries have no sub-national variation in minimum wage rates (see, for instance, Stewart (2002), Dolton, Bondibene and Wadsworth (2012), Dolton, Bondibene and Stops (2015) and Dube (2019*a*) for the UK; Dustmann et al. (2021) for Germany; Portugal and Cardoso (2006) for Portugal).

region effects which controls for both observed and unobserved differences in the workforce across the country and identifies location effects from movers between different regions.

Our method then compares trends in employment between groups who would earn the same wage if they were located in the same area, but are in fact differentially exposed to the minimum wage because they are living in different areas. In particular, we partition the national wage distribution, net of the regional wage effect, into wage bins – which we label as 'job types'. Within a job type, we use individuals living in higher-wage regions as counterfactuals for individuals living in lower-wage regions. The latter will be closer to the minimum wage, while the former will be further above it. Applying this logic, we estimate the impacts of the minimum wage on the number of jobs in each nominal wage bin.

We implement this new methodology to study the impacts of the UK's NLW introduction, and its subsequent upratings, on the entire frequency distribution of wages. Using data from the Annual Survey of Hours and Earnings (ASHE), a high-quality employer survey on earnings and hours of employees in the UK, we find that, over the 2016-2019 period, the NLW generated strong wage compression at the bottom of the wage distribution, with spillover effects on wages stretching up to at least around the 20th percentile and with little dis-employment effects. We estimate an own-wage elasticity of employment of -0.17, which is in line with many estimates in the literature (see Dube (2019*a*)) and corroborates findings of previous work in the UK using other methods. The vast majority of the estimated 'action' is at or a little above the NLW, giving us confidence that we are picking up the impacts of the NLW. We do not find evidence of significant effects on hours of work.

Besides documenting the change in employment for low wage workers, the estimated change in the number of jobs throughout the entire wage distribution also allows us to assess the impact of the policy on the distribution of household incomes. The relationship between minimum wages and household income is complicated. First, minimum wage workers, and other workers who can be affected by minimum wages through spillovers, vary substantially in terms of the share of their household's income that their earnings make up (primarily because of the presence of incomes from state benefits or from a partner's earnings). This means that they can vary widely both in terms of their position in the household income distribution, and in the proportional effect of a minimum wage on their household income. Second, any simultaneous changes in hours of work or employment would also impact incomes. Third, an increase in a worker's earnings is often met with a rise in tax liability or a fall in benefit entitlements. That means that — for some workers — the increase in net household income might be considerably smaller than the increase in their gross earnings.

We integrate our estimates of the impacts on labour market outcomes on simulations of the impact of the NLW on the distribution of household incomes. Using high-quality household survey data from the Family Resources Survey (FRS), including information on labour market outcomes, household structures and other income sources, we simulate the effects of the NLW through TAXBEN, the IFS tax and benefit microsimulation model and the most detailed model

there is of the UK tax and benefit system (Waters, 2017). In other words, we study the impacts on labour market outcomes and on the distribution of net household incomes, within the same integrated, internally consistent framework. We show that the NLW led to increases in household net incomes up to the eighth decile of the household income distribution, with larger effects in the middle. We show that accounting for the impacts of minimum wages beyond the 'mechanical ones' – i.e. beyond the fact that those with wages below the new minimum have their wages increased – is important for this analysis.

There are multiple advantages of the methods we employ here and various ways in which we refine the approaches used by previous research. The bunching method allows us to assess the change in employment across the entire distribution of hourly wages, which has several key advantages. Firstly, dissagregating the effect of the minimum wage by wage bin increases statistical precision as it allows for an explicit focus on the part of the wage distribution where the minimum wage is plausibly responsible for the changes observed. This is especially important in the context of the UK where empirical strategies commonly employed often have limited statistical power (Brewer, Crossley and Zilio, 2019). Secondly, the method provides an in-built robustness check by revealing what is happening in the upper tail, where the minimum wage would not be expected to have substantial effects (see Appendix B in Cengiz et al. (2019)). If the results suggest otherwise then this is a hint that the identification assumptions are not satisfied.

The way that we identify our bunching estimates, using regional wage variation, refines the more traditional regional variation approach to estimating minimum wage effects. Rather than assuming that all workers in one area offer a good counterfactual for all workers in another area, our approach narrowly defines groups of similar workers – who would earn the same in the absence of geographical variation in wages – and thus enables a more careful comparison across similar workers.

Finally, we demonstrate that the granularity of the bunching approach – where effects on the whole frequency distribution of wages are estimated - brings with it an additional attraction. We can use the estimated changes across the hourly wage distribution to analyses the distributional effects of the policy on household incomes. In applying this approach to the study of the impacts of minimum wages on household incomes, we combine the advantages of two hitherto distinct literatures – simulations and reduced-form econometric approaches. In particular, we retain a key benefit of simulations, which is the ability to easily run counterfactuals (for example, showing how results would be different under alternative tax-transfer systems, which can be very important for generalising across countries), which is not possible with reduced-form approaches. At the same time, we incorporate key advantages of reduced-form methods, since – by integrating the rich information on labour market impacts from the bunching approach – we can account for non-mechanical effects of minimum wages, and in particular employment effects and wage spillovers.

The remainder of the paper is structured as follows. Section 2 describes the institutional context. Section 3 details the methodology and data used for the estimation of the effects on labour market outcomes and the distribution of household incomes. Section 4 illustrates the results and Section 5 concludes.

2.   Institutional context

We analyze the distributional consequences of minimum wages in the context of the UK, which has increased the minimum wage substantially for most adults since 2016. The UK has had a nationwide minimum wage in place since the National Minimum Wage (NMW) introduction on April 1, 1999. As of March 2016, the NMW for adults aged 21+ was £6.70, with separate rates for younger workers and apprentices. From April 2016, a new, higher minimum wage rate was introduced for workers aged 25 and over, branded in the UK as the 'National Living Wage' (NLW) - though it is simply a legal minimum wage in the same sense as previous minimum wages. The minimum rates for younger workers and apprentices were unchanged. Its introduction was announced on July 8, 2015 and it came into force on April 1, 2016. A target for the NLW to achieve 60 percent of median wages by 2020 was also set at the time of announcement.

Figure 1 provides a visual representation of the evolution of the minimum wage applying to workers aged 25 and over – the NMW until March 2016 and the NLW there after – in real terms and as a percentage of the median wage. At the time of its introduction, the NLW was set at £7.20 an hour, an increase of 7.5% from its previous level in both nominal and CPI-adjusted real terms. Overall, the 17% real-terms increase in the minimum wage applying to those aged 25+ betweeen April 2015 and April 2019 led to an increase in its 'bite' relative to median wages of 7.3 percentage points. That is larger than the 6.9 percentage point increase in the bite over the whole prior 16-year period since the UK's minimum wage was introduced in 1999.[2]

3.   Methodology and data

In this section, we describe the methodology and data for our analysis. We start by describing our proposed methodology to identify the effects of the minimum wage on labour market outcomes, and then turn to the impacts on household income and living standards.

---

[2]By April 2019, the NLW was £8.21 per hour. In comparison, the minimum wage for 21-24 year olds was £7.70 (6% lower than the NLW), for 18-20 year olds was £6.15 (25% lower) and £4.35 for 16-17 year olds (47% lower).

Figure 1. Real Minimum Wage Rates and Minimum Wage Bite, 1999-2019
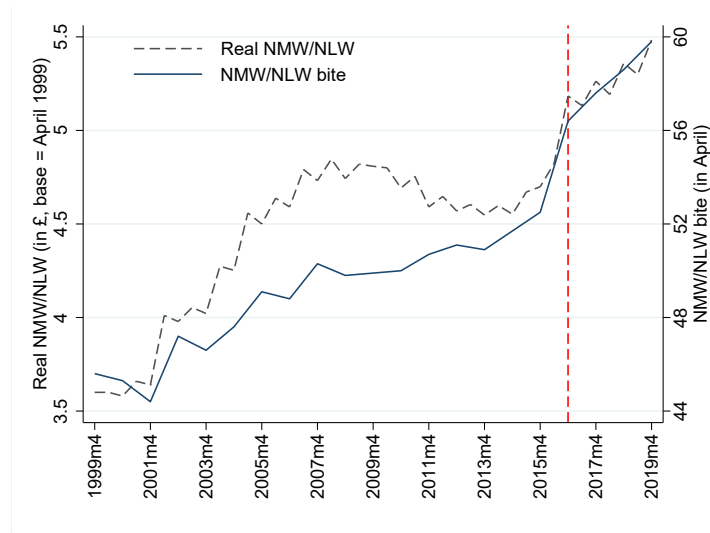
Notes: The graph reports the CPI-adjusted level of the UK minimum wage applying to workers aged 25 and over from April 1999 to April 2019 (dashed gray line, left axis). It also report the minimum wage as a percent of the median wage (solid blue line, right axis). The dashed red line corresponds to the NLW introduction on April 1, 2016.

3.1 Effects on the frequency distribution of wages

3.1.1 Conceptual framework

In this and the following subsections, we illustrate our proposed empirical methodology to estimate the impacts of the minimum wage on the frequency distribution of wages, as pioneered by Cengiz et al. (2019), in a setting in which a single minimum wage policy applies across the entire country. We identify the effects using an approach that exploits geographic variation in wage levels, in the tradition of Card (1992). We start by summarizing those approaches and illustrating how we combine features of both. We then describe the empirical implementation.

Bunching approach. The bunching approach proceeds on the basis that the effects of the minimum wage on wages and employment can be inferred from changes in the frequency distribution of wages at the lower end of the wage distribution. A higher minimum wage will directly affect jobs previously paid below the minimum: some may be destroyed, some pushed at or above the minimum wage. And jobs previously paid at or above the minimum may shift up the wage distribution via spillover effects, for example because of firms' desire to main pay differentials between different occupations, or between supervisory and non-supervisory employees. Thus, a comparison between the frequency distribution of wages observed under a minimum wage policy and a suitably-constructed counterfactual in the absence of the policy will reveal 'missing' mass below and 'excess' mass at or above the new minimum. This implicitly defines the total employment effect, which is the difference between the missing and excess

mass. Using this framework the impacts of the minimum wage on the wage distribution and employment are captured jointly in a fully integrated way.

The fact that employment changes can be estimated wage bin by wage bin brings advantages with respect to statistical precision, the verification of identifying assumptions, and the richness of the effects that one is able to estimate. By enabling the researcher to filter out shocks to employment in the upper tail of the distribution, on the basis that they are more likely to be noise than signal with respect to the impacts of the minimum wage, statistical precision is improved. By estimating the impacts on employment in every wage bin, a kind of placebo test is automatically produced: significant estimated effects on the number of jobs far up the wage distribution would act as a red flag that the identification strategy may be conflating impacts of the minimum wage with other differences between treatment and control groups – a check that is not possible with approaches which simply estimate impacts on total employment. Finally, estimating effects wage bin by wage bin paints a richer picture of the minimum wage's effects – in particular by revealing the extent of wage spillover effects on low-wage workers somewhat above the minimum wage. This can be exploited in order to estimate comprehensively the distributional effects of minimum wages, as we show in the latter part of this paper.

For identification, Cengiz et al. (2019) exploit variation in US state-level minimum wage legislation using 138 relatively large minimum wage changes occurring in the US over the 1979-2016 period. They implement a difference-in-differences design comparing changes in the frequency distribution of wages before and after a minimum wage increase between states affected by the policy change and unaffected states. In the UK, like in many other countries (e.g. France, Germany, Greece, Hungary, Ireland, Israel, the Netherlands, New Zealand, Poland and Spain), no geographic variation in minimum wages applies. We exploit regional variation in price levels – and hence in the bite of the minimum wage – in the spirit of a long line of empirical literature stretching back to Card (1992) (Stewart, 2002; Dolton, Bondibene and Wadsworth, 2012; Dolton, Bondibene and Stops, 2015; Ahlfeldt, Roth and Seidel, 2018; Caliendo et al., 2018; Clemens and Wither, 2019; Dube, 2019*a*; Schmitz, 2019; Dustmann et al., 2021).

Regional-variation approach. A common approach for estimating the impacts of minimum wages on employment is to exploit geographic variation in its bite. It can be formalised with a statistical model, where, for any two time periods, employment changes in region *r* are modeled as a function of the bite of the minimum wage in that region:

$$\Delta E_{rt} = \alpha MIN_{rt-1} + \gamma_t + \mu_{rt} \qquad (1)$$

where ΔE_{rt} is the change in employment in region *r* between time $t-1$ and t, MIN_{rt} a measure of the 'bite' of the minimum wage (e.g. the minimum wage as a fraction of the median wage in the region) in *r* at time *t*, γ_t a time fixed effect and μ_{rt} an error term. The key identifying assumption is a 'common trends' assumption that underlying employment trends across regions are unrelated to bite, i.e. they are similar in higher- and lower-bite regions.

A limitation of this approach is that, because it is looking for effects on aggregate employment while the minimum wage typically affects only a small portion of the labour market, statistical power can be low. One can think of the problem as being one of a very weak 'first stage': MIN typically is associated only with very small changes in average wages, and so we should not expect a clear signal when it comes to its impact on aggregate employment. This issue has been addressed by focusing on subpopulations where the minimum wage is known to bite more, e.g. among teenagers (Card, 1992) – though naturally this limits external validity. Another alternative is to further segment the population by demographics such as sex, age and skill level in order to create additional variation in bite (Stewart, 2002; Manning, 2016; Dube, 2019*a*). This introduces additional, potentially strong assumptions: namely that the employment effects of the minimum wage at a given bite are homogeneous across those groups, as well as that underlying employment trends are similar.

Nesting the regional and bunching approaches. Similar to the regional variation approach, our method exploits the fact that a significant component of differences in wage levels between areas, at least at the lower end of the wage distribution, is explained by geographic differences in the general price level, i.e. living costs. This implies that we can define narrow groups of similar workers who would be expected to be paid the same if they lived in the same place, but whose actual wages -– and hence proximity to the minimum wage — vary across areas due to regional differentials. This allows us to use trends in the number of jobs within higher wage bins in high-wage areas as counterfactuals for trends in the number of jobs within lower wage bins in lower-wage areas, effectively matching wage bins across areas that are equivalent in real terms but – due to cost-of-living differences – differentially exposed to the national minimum wage. Since, in practice, all regions are affected – albeit to a different extent – by the minimum wage, our approach shares the characteristic of other "regional-variation" approaches of identifying only a *relative* effect of the minimum wage, on employment in lower-wage areas relative to higher-wage ones.

In our baseline specification, we define as high-wage areas those that are in the top decile of the distribution of regional wage premia. We explain in more detail below how those premia are quantified. We then use employment trends within each wage bin in lower-wage areas, net of the counterfactual employment trends identified from what happens in high-wage areas within the same 'real' (but different nominal) wage bin. Under our assumptions this yields the impact of the minimum wage on the frequency distribution of wages in lower-wage regions (in the relative sense described above), in the spirit of the bunching approach.

In doing this, we retain the advantages of the bunching approach, while adapting it to be applied in a context with uniform national minimum wage policy. Viewed the other way around, we also refine the traditional regional variation approach and its regional-demographic extensions. Traditional approaches implicitly assume that workers living in areas relatively less affected by the minimum wage are a good control group for workers living in more affected areas. When combined with demographic variation, they rest on the assumption that relative differences across demographic groups (e.g. men and women) in one area are a good counterfactual for

relative differences in another one. Our approach relaxes those assumptions because we match narrowly defined subsets of workers living in different areas, who we estimate would genuinely earn the same wage if they lived in the same area or – equivalently – if the location-specific component of wages were removed.

As such the method proposed here is complementary to a recently developed approach that combines the regional-variation approach with machine learning tools citepCengiz2021, with the aim - similar to ours - of ensuring that workers likely to be affected by minimum wage changes in each area are the ones driving results while avoiding homogenous treatment effect assumptions.

3.1.2 Data and sample construction

Our primary data source for the analysis of the impacts of the National Living Wage (NLW) on wages, hours, employment and earnings is the Annual Survey of Hours and Earnings (ASHE). A large-scale, employer-completed survey of earnings and hours of employees in UK, ASHE provides high-quality data on wages, hours, occupation, industry and basic demographic characteristics at yearly frequency. ASHE is (weighted to be) representative at the national level, but not the local level, so to get total employment counts at the Travel-to-Work-Area (TTWA) level we rescale employment counts in ASHE to match employment counts in the Annual Population Survey (APS) – a boosted version of the Labour Force Survey (LFS). We also use APS to get the working age population in each TTWA. TTWAs are statistically-defined geographic units that are constructed by the UK's Office for National Statistics, based on commuting flows, to approximate local labour markets. They identify self-contained areas in which most people both live and work. We group TTWAs with fewer than 200 observations in ASHE with their nearest neighbouring TTWA based on observed commuting flows, so that each TTWA has at least 200 observations in any year in our data. This grouping gives us a total of 139 geographic areas. We check the sensitivity of our results to different degrees of aggregation.

3.1.3 Empirical implementation

The implementation of our approach can be thought as comprising three steps: (i) adjusting wages to account for the aggregate wage growth that would have occurred in the absence of the increase in the minimum wage, (ii) estimating regional wage effects and (iii) estimating the effect of the minimum wage on the wage distribution. The logic of the three steps is summarised below.

Adjusting for overall time trend in wages. As a first step, we net out overall wage growth from wage levels. We do so by putting the base-year's data in end-year's wage terms. What this means in practice is that when studying the impact of the NLW in year t, we will use data on the actual distribution of wages in t and an uprated version of the distribution in $t-1$, corrected

for wage growth. The uprating factor τ_t is estimated based on the following regression model:

$$\ln w_{irt} = \gamma_r + \beta GAP_{rt} + \tau_t + \varepsilon_t \tag{2}$$

where $\ln w_{irt}$ is the log hourly wage of individual i in region r (TTWA) and year t, GAP_{rt} is the mechanical increase in average wages that the higher minimum wage would induce for workers in TTWA r in year t relative to $t-1$, τ_t is a time trend and ε_t an error term. Controlling for GAP_{rt} means that we strip out any association between regional average wage growth and the increase in the minimum wage itself. We uprate $t-1$ wages using the estimated τ_t. Estimates for τ_t and GAP_{rt} in our central specification are shown in Table A1.

Estimating regional wage effects. As a second step, we define groups of workers who would earn the same if they lived in same place. We start by running a Mincer-style regression of log wages ($\ln w_{irt}$) on region (TTWA) fixed effects and individual controls, using pre-NLW data (2012-2014). Our regression specification is:

$$\ln w_{irt} = X'_{it} + \delta_r + \theta_t + \nu_{irt} \tag{3}$$

where X includes individual and firm characteristics, δ_r is an indicator for working in TTWA r, θ_t is a year fixed effect and ν_{irt} an error term. Covariates include gender interacted with full-time/part-time status and age, 1-digit occupation, 1-digit industry and a dummy for being in a graduate job (based on the 4-digit SOC code). We use a Tobit specification to account for the left censoring at the minimum wage. Estimates of regional effects $\hat{\delta}_r$ – which we call 'wage premia' – are shown in Figure A1 in Appendix A. Note that this specification assumes that (proportional) regional wage premia are fixed across worker characteristics such as gender, age and so on.

Having estimated the regional effects $\hat{\delta}_r$, we can purge all wages of those effects. We group those adjusted wages into bins and we refer to them as 'job types'. The interpretation is that people who share a job type would, according to our estimates, earn the same amount if they lived in the same place. But, because of regional wage premia, people can share a job type and yet be differently exposed to the minimum wage if they live in different parts of the country. In our baseline specification, we define 'job types' as adjusted wage bins of size £0.25, and later assess the sensitivity of our results to different widths.

Estimating the effect of the minimum wage on the frequency distribution of wages. Finally we implement a difference-in-differences style specification which compares employment rate trends across nominal wage bins, controlling for trends at the 'job type' level. The regression is run on a dataset of employment rate changes at the region-wage bin level, which we construct from the ASHE micro-data. Employment rates are computed by dividing employment counts by the working age population in the TTWA. For any two periods, we estimate the following

model:

$$\Delta \frac{E_{br}}{N_r} = \sum_{j=1}^{J} \mu_j s_{br}^{j} + \sum_{k=\underline{K}}^{K} \alpha_k \mathbb{I}[b = k | r \in L] + \epsilon_{br} \tag{4}$$

where $\Delta \frac{E_{br}}{N_r}$ is the change in employment in nominal wage bin b in region r (E_{br}), expressed as a share of the region's working-age population (N_r) between year t and $t - 1$; s_{br}^{j} is the share of the nominal b, r cell which in $t - 1$ was in job type j; $\mathbb{I}[b = k]$ is an indicator taking value one if nominal wage bin b in region r falls within k and $k + £x$ of the new NLW in t and zero otherwise, except for regions with high wage premia ($r \in H$), our control group, for which the indicator is always zero; ϵ_{br} is an error term. In our headline estimates, we partition the wage distribution into nominal wage bins b of £0.25 width, and set $x = £0.25$. Also, we define higher wage regions (control group) as those with wage premia in the top decile of the distribution of $\hat{\delta}_r$, and lower-wage regions (treated group) as those with premia in the bottom nine deciles. We assess the sensitivity of our results to different bin widths and definitions of control regions.

The coefficients α_k are the key parameters of interest: they identify the change in the employment rate in each nominal wage bin in lower wage regions (i.e. TTWAs in the bottom nine deciles of the regional wage premia distribution, $r \in L$), relative to the change observed for the same job-type in higher wage regions ($r \in H$). The key identifying assumption is that, absent changes in the NLW, employment rates in each 'job type' would evolve in the same way across lower and higher wage regions. When we present our results, we normalise the estimated α_k from equation 4 by the national pre-treatment employment rate. Hence the estimates of α_k that we report represent (relative) changes in the employment rate in each nominal wage bin as a percentage of the pre-treatment national employment rate. Moreover, we centre our nominal wage bins around the post-reform NLW so that the changes in the distribution of wages relative to the new minimum are easy to visualise. We also aggregate estimates of α_k from $\underline{K} = 2.75$ below the new NLW to just below the new NLW, and all estimates in wage bins more than £15 above the NLW.

In practice, the estimation of model 4 is implemented in two steps. We first estimate the job-type specific trends, based only on high-wage 'control' regions. We take a dataset of employment rate changes at the job type - region level excluding the bottom 90% of regions, and simply estimate:

$$\Delta \frac{E_{jr}}{N_r} = \sum_{j=1}^{J} \mu_j + \varepsilon_{jr} \quad \text{for } r \in H \tag{5}$$

where $\Delta \frac{E_{jr}}{N_r}$ is the change in employment in job type j in region r (E_{jr}), expressed as a share of the region's working-age population (N_r) between year t and $t - 1$, μ_j is a job-type fixed effect and ε_{jr} an error term.

We then take a dataset of employment rate changes at the wage bin - region level, containing only the bottom 90% of regions. Given the job type corresponding to any wage bin - region combination, we can subtract counterfactual employment rate changes based on the applicable

job type trend estimated based on equation 5: $\Delta \frac{E_{jr}}{N_r}' = \Delta \frac{E_{jr}}{N_r} - \sum_{j=1}^{J} \hat{\mu}_j s_{br}^j$, where $\Delta \frac{E_{jr}}{N_r}$ and s_{br}^j are as defined above.

Hence we now have a transformed dependent variable which is simply the change in the employment rate in each wage bin - region combination, relative to the estimated counterfactual in the absence of a minimum wage change. We regress this on nominal wage dummies:

$$\Delta \frac{E_{jr}}{N_r}' = \sum_{k=\underline{K}}^{K} \alpha_k \mathbb{I}[b = k] + \eta_{br} \quad \text{for } r \in L \tag{6}$$

where all variables are defined as above and η_{br} is an error term. We use a bootstrap in order to conduct statistical inference (using 100 bootstrap replications). We allow for clustering at the TTWA-level.

As in Cengiz et al. (2019), the set of α_k coefficients can be used to compute total employment effects of the minimum wage. The missing mass below the new minimum wage can be computed as $\Delta a = \frac{\sum_{k=\underline{K}}^{0} \hat{\alpha}_k}{(E/N)}$ and the excess mass above it as $\Delta b = \frac{\sum_{k=0}^{\bar{K}} \hat{\alpha}_k}{(E/N)}$, where E/N is pre-treatment national employment divided by the working age population. By dividing employment rate changes by the pre-treatment national employment rate, we calculate the missing and excess mass as a share of the national employment rate. Their sum, which we define as $\Delta e = \Delta a + \Delta b$, represents the total employment effect, or more precisely the percentage change in the employment rate due to the NLW. For an approximately constant working age population over the pre-treatment and treatment periods, Δe estimates the change in employment as a percent of pre-treatment workforce due to the NLW. For our baseline estimates we set $\bar{K}$ equal to NLW + 5. We show the sensitivity of our results to alternative choices, and we also routinely report an estimate of $\Delta \text{empl} = \Delta a + \Delta b'$, where $\Delta b'$ aggregates the estimated $\hat{\alpha}_k$ over the entire support of the wage distribution. This is never far from our central estimate of the employment effect, which is reassuring evidence in favour of our identifying assumptions because it implies that employment rate changes within given job types were very similar between treatment and control regions whenever we look beyond the lower portion of the wage distribution, where the minimum wage should not be having meaningful impacts.

A conceptual difference between our α_k coefficients and those of Cengiz et al. (2019) is that we estimate the effect on the employment rate in lower-wage regions *relative* to higher-wage regions - not the *absolute* effect. This follows directly from the fact that the UK does not provide geographic variation in minimum wage policy, so there are no geographic areas that are completely 'untreated' which can be used as controls in order to identify absolute effects. One can however recover absolute effects across the whole economy with some extrapolation. First, we calculate an elasticity of area employment with respect to area wages (see below). Second, we estimate the absolute wage effect of the NLW by comparing the wage distribution before and after an NLW increase (uprating the earlier year using the τ_t from equation 2). Third, we multiply these figures to get an estimate of the absolute change in employment.

Calculating the employment elasticity. We compute the own-wage elasticity of employment as the proportional change in employment for affected workers divided by the proportional change in wages for affected workers. Our estimated $\hat{\alpha}_k$ coefficients are key inputs for this calculation.

We approximate the proportional impact on affected employment as the relative change in employment as a share of baseline (given directly by summing our $\hat{\alpha}_k$ coefficients), divided by the share of the workforce earning below the new minimum wage in the year before treatment. Terms superscripted P are calculated across the whole population (i.e. both high and low wage regions).

$$\%\Delta e_L - \%\Delta e_H \approx \frac{\Delta e_L - \Delta e_H}{\overline{b}_{-1}^{P}} = \frac{\sum_{k=\underline{K}}^{K} \hat{\alpha}_k}{\overline{b}_{-1}^{P}} \tag{7}$$

For wages, we first calculate from our $\hat{\alpha}_k$ coefficients the proportional relative effect of the minimum wage on 'real' average wages (i.e. wages purged of regional wage premia). We then divide that by pre-policy average wages among affected workers:

$$\%\Delta w_L - \%\Delta w_H \approx \frac{\frac{\overline{rwb}_{-1}^{P} + \sum_{k=\underline{K}}^{\bar{K}}\left(k + \overline{MW}^{P}\right)\hat{\alpha}_k}{\overline{b}_{-1}^{P} + \sum_{k=\underline{K}}^{\bar{K}} \hat{\alpha}_k}}{\frac{\overline{rwb}_{-1}^{P}}{\overline{b}_{-1}^{P}}} - 1 \tag{8}$$

where $\overline{rwb}_{-1}^{P}$ is the real pre-period wage bill among those paid less than the new minimum.[3] $\overline{MW}$ is the average wage in the bin where the minimum wage falls in the post-period. Elasticities are obtained by dividing equation 7 by equation 8.

3.1.4 Effects on younger workers and other subgroups

As discussed, an advantage of our approach is that it allows for a unified examination of effects on both employment and the wage distribution. In the context of workers aged under 25, who are not directly affected by the NLW, this means that we can jointly capture the wide range of factors which might nevertheless result in a relationship between the NLW and their own labour market outcomes. These include impacts of labour substitution between older and younger workers, 'downward' wage spillovers of the NLW onto younger workers (e.g. if firms avoid age-related pay differentials created by different minimum wages), and any impacts of those downward wage spillovers on the employment of younger workers.

The approach described above can be extended straightforwardly to examine the impacts of the NLW on younger workers. We repeat all the main steps of that methodology specifically for the sample of workers aged under 25, including estimating underlying wage trends and the

---

[3]Specifically, we deflate wages in each region with the reference being the average wage premium in the low wage regions. In other words, $\overline{rwb}_{-1}^{P}$ is in the price (wage) of the low wage regions. This makes it consistent with $\hat{\alpha}_k$.

regional wage premia specifically for that group. The coefficients of interest, $\sum_{k=\underline{K}}^{K} \alpha_k \mathbb{I}[b = k]$, do however still indicate the distance to the prevailing NLW rate for those aged 25+.

Among those aged 25+, we perform two kinds of analysis to provide greater granularity to our results. The first of these is a simple decomposition of the aggregate effect on employment from our main specification. Here we estimate equation 4 separately for each subgroup, replacing $\Delta(E_{br}/N_r)$ with the change in employment *in a given subgroup* in nominal wage bin b in region r, still as a share of the region's (total) working-age population. We use this to study how the overall effect is driven by impacts on the number of part-time and full-time workers, different genders and different age-groups. Note that this does not allow the different subgroups to have different wage trends or wage premia - it merely decomposes the estimated aggregate effect. In a second set of subgroup analysis we examine heterogeneous employment effects by gender and age, allowing for subgroup-specific wage trends and regional wage premia..

3.2   Effects on the household and family income distribution

There is another key, but previously unexploited, advantage of the bunching approach for estimating minimum wage effects on employment and wages: its applicability to analyses of the distributional effects of minimum wages on household incomes. The role of minimum wage policy in tackling poverty or inequality in living standards, as opposed to just individual labour market outcomes, is a central policy question and one that has attracted significant amounts of research.

The relationship between changes in wages and changes in the household income distribution is complicated by a range of factors, including hours of work, incomes of other household members, and interactions with the tax and benefits system. Hours of work determine how a change in wages will translate into a change in earnings, though the relationship is complicated by the fact that the introduction of the NLW may cause changes in hours worked. The impact of the NLW on the household income distribution is also sensitive to whom individuals affected by the introduction of the NLW live with, for two reasons. Firstly, households with more earners will be more affected by changes to wages than households with only one. Secondly, the net incomes of all household members, including earnings after tax, investment income and benefits, will partly determine where affected earners rank in the household income distribution. The mapping from the individual wage distribution to the household income distribution, influenced by hours of work, the incomes of other household members, and the tax and benefit system, is illustrated in figure 2. This figure shows, for each individual wage decile, the proportion of workers living in each household income decile (defined among households with at least one 26-64 year old). Whilst the highest wage earners are very likely to have high household incomes, the lowest decile of wage earners are spread across most of the household income distribution, with over 50% lying in the middle 40% of the distribution. However, if we restrict the sample to working households, a majority of the lowest decile of wage earners lie in the bottom 30% of the household income distribution.

Figure 2. Proportion of individuals in each household income decile, by wage decile (FRS October 2014 to September 2015, using donor method
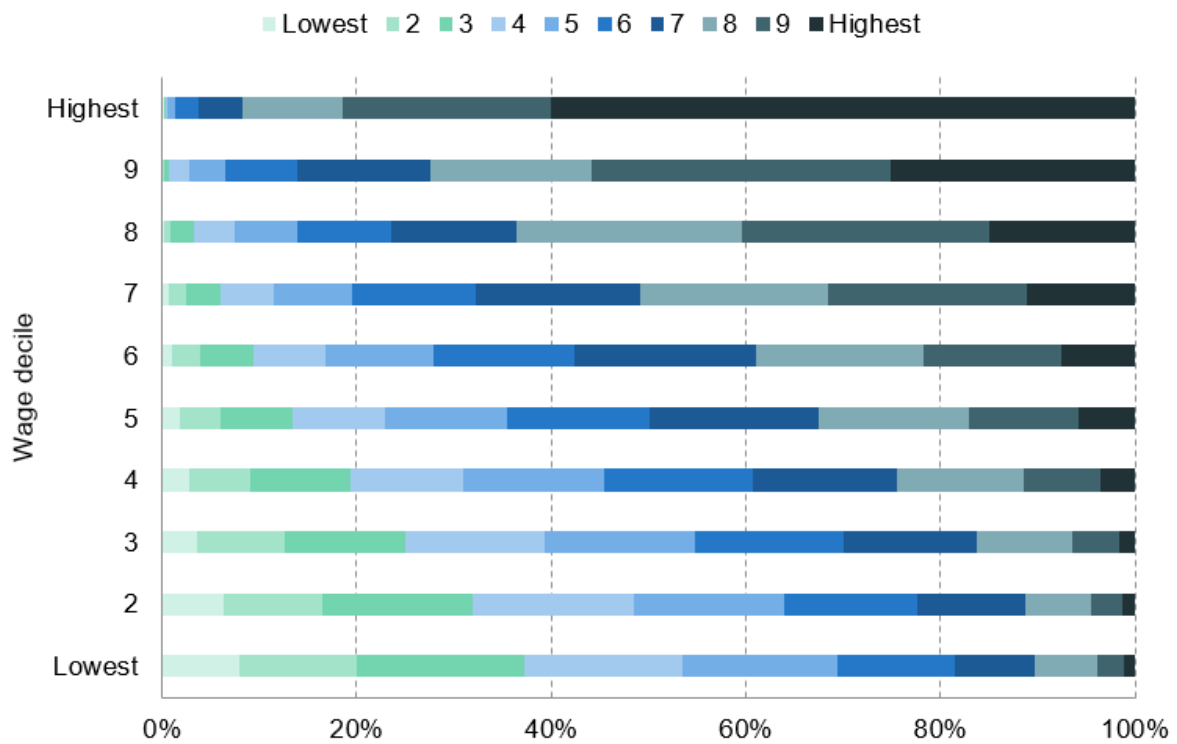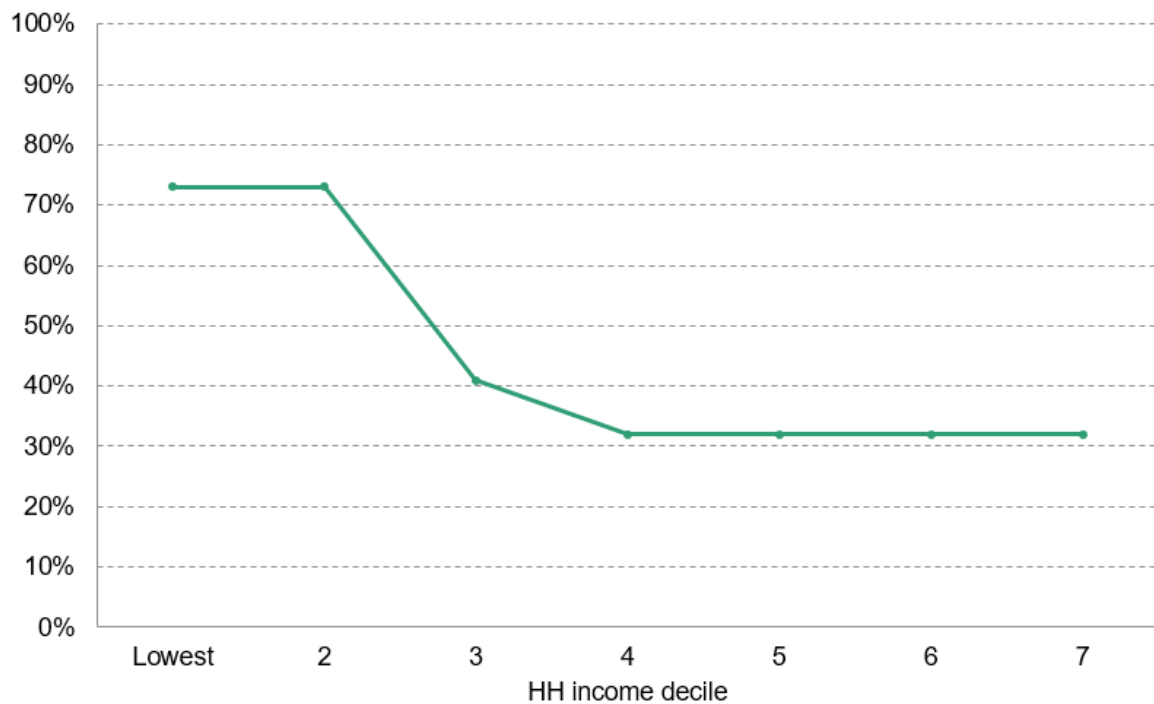

Notes: Sample is employees aged 25-64. Household income deciles are defined among all households with at least one 25-64 year old. Income is equivalised and net of taxes and benefits.

The tax and benefit system plays an additional role in determining the impact of the NLW on household incomes. An increase in earnings caused by an NLW increase will not all feed into household income, as taxes and the withdrawal of benefits reduce the impact. Similarly, a decrease in earnings caused by any disemployment effects will be partially mitigated by tax decreases and benefit increases. Figure 3 shows the median marginal tax rate for low-wage workers in each household income decile (defined among households with at least one 26-64 year old). Those in the lower net household income deciles, which contain high proportions of low wage earners, have higher marginal tax rates, due to withdrawal of means-tested benefits. Therefore any given wage increase for workers in those deciles will, on average, result in lower income rises than for workers in higher net household income deciles.

Figure 3. MARGINAL TAX RATES FOR MINIMUM WAGE EARNERS, BY HOUSEHOLD INCOME DECILE

Notes: Sample is employees aged 25-64 earning the national minimum wage. Household income deciles are defined among all households with at least one 25-64 year old. Income is equivalised and net of taxes and benefits. Deciles above the 7th are omitted due to small sample sizes.

Previous studies have often used simulation approaches in order to estimate the impacts of minimum wage increases on household incomes (e.g. Brewer and Agostini (2017); Sabia and Burkhauser (2010)). The typical approach is to take household survey data collected shortly prior to a minimum wage hike, and to simulate an increase in some workers' earnings based on the assumption that those with a wage below the new minimum will see their wage rise to that level. Because these workers are observed together with the rest of their household,

and their income sources, this allows for a simulation of the effects by household income. Often a tax-benefit micro-simulation tool is used in order to account for interactions between earnings and the tax and transfer system, arriving at a more accurate estimate of impacts on net (post tax/transfer) income. This is particularly important in institutional settings where income-related transfers (particularly those for working families) are widespread, as in the UK.

Simulation has advantages, such as the ability to explicitly decompose the impacts on net household incomes, or to explore alternative scenarios, by changing the inputs to the simulation. For example, one can isolate the impact of the tax-transfer system, and indeed one can simulate the effect under an alternative tax-transfer system. That could be particularly useful in addressing external validity concerns, e.g. when trying to understand the implications of results in one country for another, or how a potential reform to taxes or transfers would interact with the impacts of minimum wages. A reduced form empirical approach that tried to directly estimate the impacts of minimum wages on household incomes could not do this.

However, the simulation approaches used thus far also have limitations, summarised by Dube (2019*b*):

- They must make an assumption about the impact of the minimum wage on employment or hours worked. Typically the assumption is that there is no effect. Sabia and Burkhauser (2010) refine this, by importing an out-of-sample employment elasticity from previous literature. Whether that elasticity is appropriate in the setting they are simulating a minimum wage increase for is, of course, an open question.

- They must also make an assumption about wage spillovers above the new minimum, and non-compliance below it. Again usually the assumption is that there are none of either.

- Measurement error in hourly wages of workers (or in other sources of household income), which is common in the household survey data on which these studies typically rely, can weaken the measured relationship between a worker's hourly wage and household income. This will tend to attenuate any distributional impact of minimum wages by household income.

The first two of these limitations are similar: essentially, simulation approaches have only captured the mechanical effects of minimum wage increases - or have had to introduce further assumptions in order to try to capture non-mechanical effects. But by integrating this endeavour with the bunching approach, we can use estimates of the impact of the minimum wage on the whole frequency distribution of wages to simulate these non-mechanical effects. In combination with a careful strategy for addressing measurement error in hourly wages (described further below), this means we can address the traditional limitations of simulation-based approaches while retaining their advantages.

The basic steps we follow are:

1. Take detailed survey data on households' income from before the introduction of the NLW.

2. Impute hourly wages in the data to account for measurement error (details below).

3. Change some workers' status to unemployed, reflecting disemployment effects of the NLW.

4. Change hourly wages to account for estimated wage effects of the NLW.

5. Use a tax-benefit microsimulator to calculate net household incomes.

We discuss these points in more detail below.

3.2.1 Data and sample construction

We use the Family Resources Survey (FRS), an annual cross-sectional survey of around 20,000 households which forms the basis of the UK's official household income statistics and contains detailed information on household characteristics and incomes. We use FRS data from October 2014 to September 2015, and uprate financial variables (principally earnings and rent) to 2019-20 prices.[4] The national minimum wage was constant over that period, at the same April 2015 level observed in the pre-NLW ASHE data used to produce our bunching estimates. We use only households with at least one person aged 25-64, leaving us with 13,463 households.

### 3.2.2 Addressing measurement error in hourly wages

For most employees in the FRS, we observe weekly or monthly earnings and weekly hours of work. One can compute a 'derived' hourly wage by simply dividing one by the other. As is well known, the distribution of derived hourly wages in survey data often contains an implausibly large number of low values, and a limited amount of bunching at precisely the minimum wage (see for example Skinner et al. (2002))- just as one would expect if there is measurement error in the derived hourly wage. Figure 4, which compares the wage distribution in the FRS (Oct 2014 - September 2015) and in ASHE (April 2015) (which uses employer-reported earnings and hours), shows this phenomenon is evident in the FRS.

This presents a challenge for simulations of the minimum wage's effects by household income. Measurement error in wages will mean that apportionment of the minimum wage's effects to different parts of the household income distribution will be noisy. It is likely to systematically bias the distributional patterns: for example, a classical measurement error process would weaken the observed relationship between wages and any variable, including household income, meaning that the measured distributional impact would be attenuated. In addition,

---

[4]To be consistent with the bunching analysis, we use τ_t from equation 2 to uprate earnings. For other financial variables we use official price indices, such as average rents.

Figure 4. DISTRIBUTION OF WAGES IN FRS (OCT 2014 - SEPTEMBER 2015) AND ASHE (APRIL 2015)
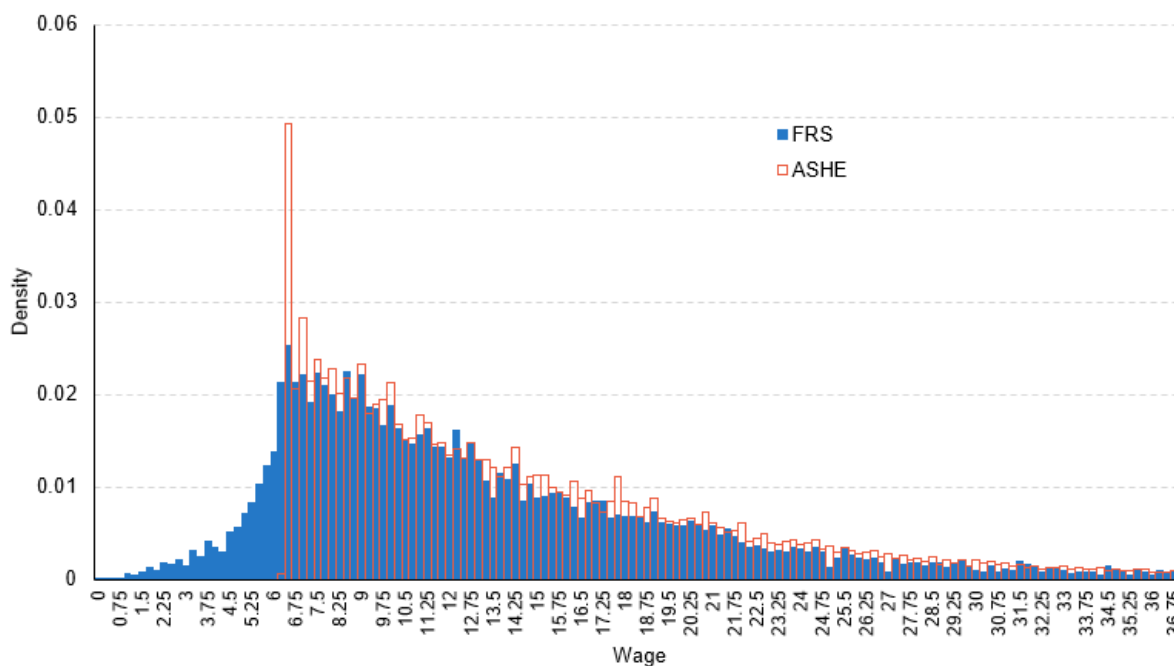
Notes: Wage measured in £ per hour. Wages bins are 25p wide.

since we will be simulating wage changes based on bunching estimates obtained from ASHE, internal consistency demands that the underlying wage distributions of ASHE and the FRS are similar.

We address this using the 'donor method' approach used in Skinner et al. (2002) and Harkness and Avram (2019*b*). This exploits the fact that a subset of workers in the data *directly* report their hourly wage if they are paid by the hour (i.e. rather than being salaried), [5] *as well as* reporting their earnings and hours - allowing us to compare their derived and direct hourly wages. The direct wage distribution of such workers is much more plausible, with bunching at the minimum wage, and few workers paid beneath the minimum wage.

We use the joint distribution of derived and direct hourly wages for these workers ('donors') to impute hourly wages for those not paid by the hour ('donees'), as follows. First, for workers paid by the hour, we regress the directly measured hourly wage on the derived wage and a set of controls (education, industry, age, children, region, and household income decile). Including household income decile in this equation is an augmentation of previous implementations of the donor method (such as in Harkness and Avram (2019*b*)), reflecting the fact that we must preserve the relationship between wages and household income in our context where the aim is to study the distributional impact of the minimum wage. Second, we use the estimated

---

[5]14% of workers in the data report an hourly wage, and 20% of workers with whose highest qualification level is GCSE or lower

coefficients to predict an hourly wage for all workers in the data. Third, for each 'donee' we find their 10 nearest donor neighbours [6] (based on predict hourly wages), randomly pick one of them, and impute their direct hourly wage to the donee. Such a process introduces simulation error, so we repeat it 40 times and average across repetitions when calculating the final results.

Adjusting hourly wages in this way creates an inconsistency between a worker's updated wage and their stated hours and earnings. Since weekly/monthly earnings reports in the FRS are often given with reference to payslips, we believe that there is less chance of measurement error in earnings than in hours. We therefore adjust hours (rather than earnings) to resolve the inconsistency. The one exception is in cases where doing so would cause us to make a working lone parent's hours fall below 16 hours a week, as there are strong financial incentives for lone parents to work at least 16 hours per week due to the operation of Working Tax Credit. In those cases we keep hours unchanged and adjust earnings in order to resolve the inconsistency.

3.2.3   Apply employment effects

Our central estimates from the bunching analysis implied a small, though not statisically significant, disemployment effect. To simulate the effect of this, we randomly select the applicable fraction of workers who earn at or below the new minimum wage, and remove their earnings. This assumes that a worker who would have earned just £0.01 less than the new minimum wage is as likely to lose their job as a worker who would have been on the pre-reform minimum wage. We test the sensitivity of our results to instead randomly selecting only from the workers who would have earned no more than the previous minimum wage. The results are essentially unchanged when we do this.

3.2.4   Apply wage effects

The baseline distribution of wages in combination with our bunching estimates[7] imply a post-policy distribution of wages, which one might call the 'target' distribution for our simulation exercise.[8] We make a no re-ranking assumption - that is, we assume that the NLW does not cause a worker who would otherwise have had a wage strictly less than another worker to end up with a wage strictly more than her. Hence, given their baseline wage rank, we simply change each worker's wage to be equal to the wage level at that same rank in the target distribution.

---

[6]Only neighbours with a predicted wage within 50p of the donee can be selected as nearest neighbours.

[7]For this exercise we require absolute, rather than relative, estimates of the NLW on each wage bin. To get this we multiply our estimates of α_k from equation 4 by the ratio of the overall absolute employment effect (see Section 3.1.3) to the overall relative employment effect. This is equivalent to assuming that the shape of the effect of the NLW (but not the magnitude) is the same across high and low wage regions.

[8]We assume that *within* the 25p wage bins the distribution of wages stays the same, except for the bin that spans the NLW to NLW+25p, the distribution of which we base on the observed hourly wages in the bin around the NLW in the 2018-19 FRS.

3.2.5 Calculate net household incomes

The above steps simulate the impact of the minimum wage on individuals' employment status and wages, in a household survey dataset. One can then use tax-transfer micro-simulation to account for the knock-on effects of earnings changes on taxes paid and transfers received, accounting for all the relevant demographic and economic characteristics of the household. We do this using TAXBEN, the IFS tax-benefit microsimulator (Waters (2017)), which is the most detailed micro-simulation model of the UK tax-transfer system. We use the parameters of the 2019-20 system, as we are simulating the impacts of the NLW reforms between 2015 and 2019.

Not all families claim the means-tested transfers that they are entitled to. A simulation that took no account of that would overstate the interactions between minimum wages and the transfer system. Therefore, if a household did not report receiving a benefit in the survey even though they appear to have been entitled based on their characteristics, we assume that they continue not to take up that benefit in our simulation.[9] In a relatively small number of cases, households gain entitlement to a transfer as a result of the simulated impacts of the minimum wage on labour market outcomes. In those cases we cannot use reported take-up as a guide, so we randomly decide whether they take up their entitlement based on the predicted probability obtained from a logistic regression of take-up status on entitlement amount, work status, family type and age.

A caveat is that self-reported take-up in survey data tends to imply lower overall benefits spending than administrative records. In recent years about 18% of all benefit spending is estimated to be 'missing' in the FRS data (Corlett (2021)). As a robustness check, we also present results under the assumption of full take-up. The key conclusions are unchanged.

Our estimates should be interpreted as partial equilibrium effects: given that we find that firms' wage bills have gone up (even after disemployment effects), unless there is an offsetting increase in productivity then product prices must be raised or profits must be reduced, or both. Either of these would reduce real incomes for some households. It is therefore likely that the true general equilibrium effect on household incomes will be lower than we show, with ambiguous distributional implications .

We show how results differ if we take the income sharing unit to be narrower than the whole household - specifically what is sometimes known in the UK as a 'benefit unit', or more commonly in the US as a 'tax unit', which is an individual, any cohabiting or married partner, and any children. Under this definition, for example, students living together would not be assumed to share income, and neither would an adult living with their parents. Brewer and Agostini (2017) showed that the distributional impact of minimum wages can differ somewhat depending on what the income sharing unit is assumed to be.
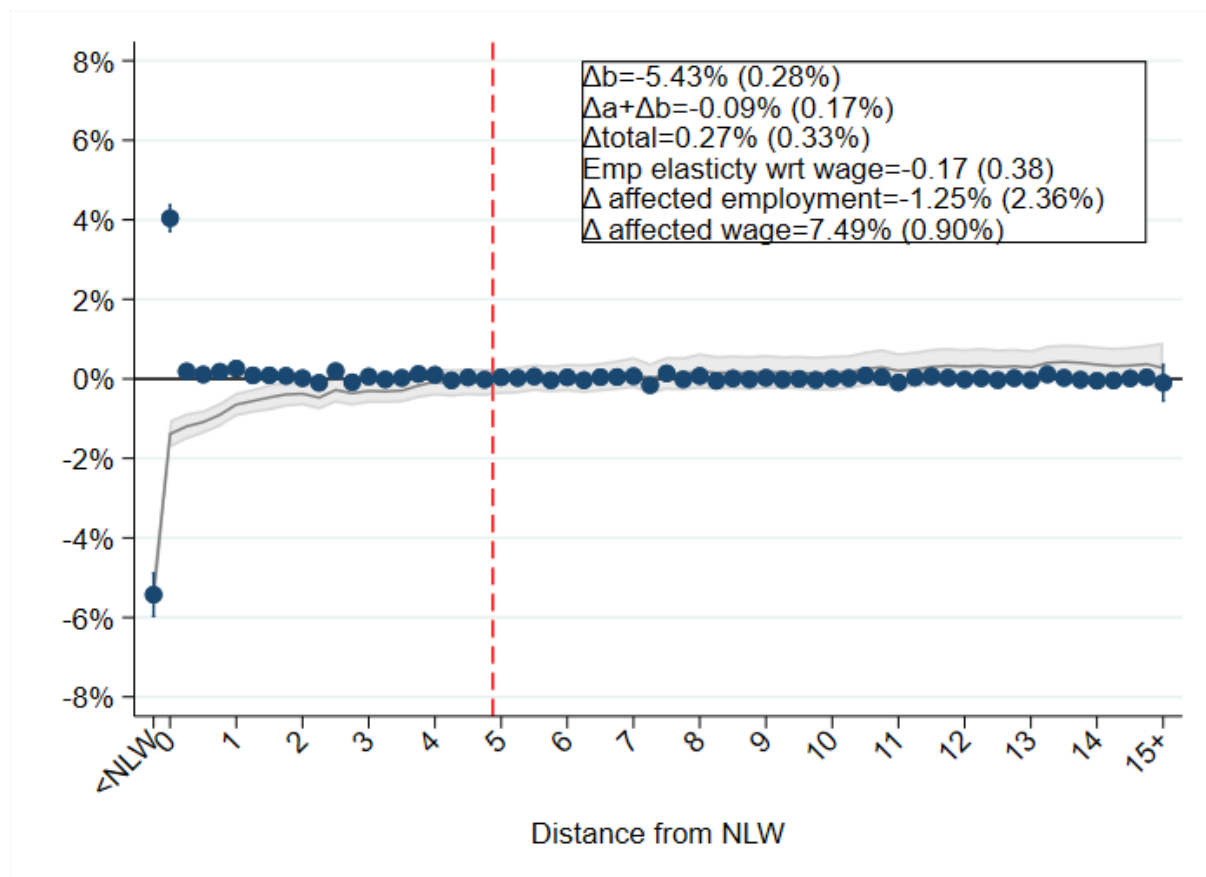
---

[9]We assume full take-up of child benefit. Child benefit take-up rates are over 95%.

4. Results

4.1 Employment and wages

4.1.1 Main results

Figure 5. NLW introduction and uplifts, 2015-2019

Notes: Δb is the effect of the NLW on change in employment below the NLW as a fraction of employment in the previous year. $\Delta a + \Delta b$ is the effect of the NLW on change in employment up to £5 above the NLW. $\Delta total$ total is the estimated effect on total employment over the whole wage distribution. We also report the percentage change in employment for affected workers, and the percentage change in wages for them. The ratio of these two is the own-wage elasticity of employment, also shown.

Figure 5 shows the estimated effect of the NLW introduction and subsequent uplifts on the frequency distribution of hourly wages. In interpreting this and the subsequent discussion based on it, it is important to recall the discussion in Section 3: we are effectively showing thee estimated effect of the NLW on the distribution of real wages in lower-wage regions *relative* to high-wage regions, and the implied impact on aggregate employment is, similarly, an estimate of the effect on employment in lower-wage areas relative to higher-wage ones.

We report employment changes averaged over the four minimum wage increases (2015-2019) for each 25p wage bin, defined each year relative to the level of the new NLW. We normalise all employment changes by baseline employment in the TTWA, so that the sum of all the effects across wage bins can be interpreted as the total percentage (not percentage point) change in employment arising from the change in the minimum wage. We aggregate estimated effects in wage bins below the new NLW in each year, as well as in wage bins more than £15 above the NLW. The grey line shows the running total of employment changes up to that point in the distribution. For example, at £5 above the NLW it shows the implied estimate of the impact of an increase in the NLW on the number of jobs paid at or below £5 above the new NLW.

The central estimate is that, on average, each increase in the minimum wage for those aged 25+ between 2015 and 2019 led to substantial fall in employment below the NLW (Δb), of 5.43% (std. error 0.28%) of total employment in the previous year. It led to a rise in employment at, or within 25p of, the new minimum wage, of 4.04% (std. error 0.17%) of pre-treatment employment. We also find statistically significant wage increases in the number of jobs on slightly higher wages, with spillovers stretching up to around £1.50 above the NLW. This is around the 20th percentile of hourly wages, which is broadly consistent with evidence of wage spillovers from minimum wages found previously (Cengiz et al. (2019), Harkness and Avram (2019*a*), Autor, Manning and Smith (2016)). Our point estimates also indicate very small spillovers up to around £4 above the NLW, though these are not statistically significant.

To compute the total employment effect, we add up all employment changes up to £5 above the NLW ($\Delta a + \Delta b$). Doing this, the missing and excess masses are of almost identical size and so we obtain an estimate of -0.09% of pre-treatment employment, a very small decline which is not statistically significant. The 95% confidence interval is between -0.43% and 0.25%. Of course, £5 above the NLW is an arbitrary cut-off, and we could choose others. For example, if we instead calculated employment changes up to £3.50 above the NLW, the change in employment would be -0.32% (95% CI: -0.61% to -0.02%). If we calculated employment changes up to £6.50 above the NLW, the change in employment would be 0.00% (95% CI: -0.037% to 0.037%). These figures are all consistent with no more than a small disemployment effect.

Estimated employment changes over £5 above the NLW are very close to zero. As discussed, the fact that the bunching approach forces transparency over those changes acts rather like a placebo test, given that the minimum wage would not be expected to have material effects far up the wage distribution. In short, this is reassuring with respect to our identifying assumption of parallel trends between lower-wage and high-wage regions, increasing confidence that the effects we obtain at the bottom of the distribution are just driven by the NLW. The placebo checks shown in Appendix A using pre-NLW data provide further reassurance of our identification strategy.

We calculate the own-wage employment elasticity using the formula set out in Section 3. We estimate an elasticity of -0.17 (the ratio of the percentage change in employment for affected workers, divided by the percentage in the average wage for affected workers, both shown in

Figure 5 which is in line with many estimates in the literature (see Dube (2019*a*)), including some previous studies in the UK (Manning (2021), Dube (2019*a*), Stewart (2004)). As this depends on estimates of both wage effects and employment effects, the confidence interval is however relatively wide, at -1.25 to 0.47.[10]

As the NLW potentially continues on its trajectory towards two thirds of median wages, it is important to consider whether the impacts of the minimum wage are changing materially as the bite increases. As ever, it is inevitably harder to answer that confidently, since it requires no longer pooling minimum wage changes together but studying them separately and then comparing them, meaning statistical precision is reduced considerably. But separate bunching estimates for each of the four minimum wage changes in question (the introduction of the NLW in 2016, and its uprating in 2017, 2018 and 2019) show that in none of those years do we have a statistically significant employment effect, as shown in Appendix A. Appendix A also shows the estimated effect of the NLW introduction and all subsequent uplifts using data from just 2015 and 2019. That is, instead of pooling data from each of the four uplifts, we estimate the 'long difference' from 2015 to 2019. Unlike simply pooling the 4 analyses of year-to-year changes, this would allow some lagged adjustments to be captured - for example, delayed effects on firm exit and hence employment from the 2016 NLW which show up in 2018 or 2019. The employment change up to £5 of the NLW is estimated at -0.43%, which is in fact very close to four times the estimated average effect from pooling the four consecutive-year periods. This estimate is more imprecise however (standard error 0.69%), because it uses much less data than our central estimates which effectively pool the results from 4 different minimum wage increases.

4.1.2 Robustness checks

Our conclusions are robust to an array of alternative specification choices, as shown in Table 1, which reports estimates of the missing mass, the total employment effect, and the own-wage elasticity of employment.

We test robustness to the choice of:

- **Wage bin width.** We run the analysis using (real and nominal) wage bins of 10p and 50p instead of 25p.

- **Geographic aggregation.** We vary the sample size threshold below which we group neighbouring travel to work areas, using thresholds of 100 and 400 observations instead of 200.

- **Wage premium estimation.** In one variant, we estimate wage premiums using only the bottom half of the wage distribution in each region. In a second, we estimate wage

---

[10]Here and throughout, we report the confidence interval of elasticities by calculating the 0.025 and 0.975 quantiles of the bootstrap replications, rather than using the estimated standard error. We choose this approach because the elasticity statistic is likely to be non-normally distributed - as it is the ratio of the employment and wage effects, even if the statistics estimating those two effects are normally distributed, their ratio will not be.

premiums using an AKM regression on movers rather than a Mincerian regression (Abowd et al. 1999). In a third variation, we drop industry and occupation controls from our Mincerian specification.

- **Choice of control regions.** We compare regions in the bottom 8 deciles of regional wage premiums to regions in the top 2 deciles, instead of comparing the bottom 9 deciles to the top decile.

Overall, the estimates from the alternative specifications are similar to our baseline estimates. Point estimates for missing jobs below the new NLW are within 0.5 percentage points of our main estimate across all specifications, and in all cases the estimated employment effect is small and not statistically significant.[11] Estimates of the own-wage elasticity almost always allow us to rule out very large elasticities (for example, Neumark and Wascher (2008) argue that the own-wage elasticity can easily be -1 or -2). The biggest difference to the point estimate of the elasticity comes when we use regional wage premiums estimated using only the bottom half of the wage distribution, or using an AKM regression (identifying premia only from workers who move areas), but these approaches also give considerably less precise estimates than our baseline specification. The AKM specification in particular gives extremely imprecise results due to the small number of movers between regions (confidence interval -3.01% to 0.65%).

Table 1. ROBUSTNESS CHECKS

| | Δb | | $\Delta a + \Delta b$ | | Elasticity | |
	Mean	S.E.	Mean	S.E.	Mean	C.I.
Main specification	-5.43%	0.28%	-0.09%	0.17%	-0.17	[-1.25,0.47]
Bin width						
10p	-5.47%	0.28%	-0.10%	0.16%	-0.22	[-1.20,0.55]
50p	-5.51%	0.24%	-0.18%	0.20%	-0.26	[-1.18,0.02]
Threshold for grouping TTWAs						
100 observations	-5.48%	0.21%	-0.07%	0.16%	-0.13	[-1.03,0.23]
400 observations	-5.32%	0.29%	-0.05%	0.19%	-0.10	[-1.09,0.51]
Wage premium estimation						
Bottom half of distribution	-4.91%	0.36%	-0.16%	0.16%	-0.39	[-1.57,0.39]
AKM	-5.15%	0.42%	-0.20%	0.26%	-0.48	[-3.01,0.65]
No industry/occupation controls	-5.61%	0.28%	-0.20%	0.18%	-0.37	[-1.09,0.31]
Control regions						
Top 2 deciles	-5.60%	0.26%	-0.15%	0.14%	-0.28	[-1.04,0.19]

Notes: Δb is the effect of the NLW on change in employment below the NLW as a fraction of employment in the previous year . $\Delta a + \Delta b$ is the effect of the NLW on change in employment up to £5 above the NLW. Elasticity is the estimated employment elasticity with respect to the wage. Standard errors and confidence intervals are estimated by bootstrapping. The specifications with different control regions would not be expected to have the same Δa and Δb as the main specification since the difference between treatment and control regions is smaller in the alternative specification than in the main. Therefore this robustness check is mainly informative for the elasticity it delivers.

---

[11]Note that the specification with different control regions would not be expected to have the same Δa and $\Delta a + b$ as the main specification.

4.1.3 Decomposition of aggregate results

Table 2 decomposes our main employment effects into those accounted for by part-time and full-time workers, and by graduate and non-graduate jobs. As set out in Section 3, this is done by using subgroup-specific employment in the numerator of the dependent variable, expressed as a share of the total working-age population and with the rest of the specification kept the same. The purpose is simply to highlight where any effect on employment delivered by our main specification is coming from. In the next subsection we will study heterogeneity in the treatment effects by subgroup, allowing key parameters to vary by subgroup, which we do not do here.

Table 2 shows that the majority (3.30%) of the 5.43% fall in employment below the NLW was seen among part-time workers, in particular women working part-time who accounted for nearly half (2.47%) of the fall. Virtually all of the fall in jobs paid below the new NLW is among workers in non-graduate occupations, as defined using the classification in Aghion et al. (2019). This largely mirrors the baseline characteristics of low-wage workers, as one would expect.[12] The point estimate of the overall employment effect is not statistically significant for any the subgroups.

Table 2. DECOMPOSITION BY SUBGROUP

| | Δb | | $\Delta a + b$ | |
	Mean	S.E.	Mean	S.E.
Main specification	-5.43%	0.28%	-0.09%	0.17%
Hours				
Full-time	-2.12%	0.13%	-0.04%	0.12%
Part-time	-3.30%	0.16%	-0.04%	0.11%
Hours by gender				
Full-time men	-0.98%	0.08%	0.06%	0.10%
Full-time women	-1.15%	0.07%	-0.10%	0.08%
Part-time men	-0.83%	0.05%	0.03%	0.04%
Part-time women	-2.47%	0.14%	-0.08%	0.10%
Graduate job				
Graduate job	-0.10%	0.02%	-0.03%	0.06%
Non-graduate job	-5.33%	0.26%	-0.06%	0.18%

Notes: Δb is the effect of the NLW on change in employment below the NLW as a fraction of employment in the previous year . $\Delta a + \Delta b$ is the effect of the NLW on change in employment up to £5 above the NLW. Full-time defined as working 30 or more hours a week. Graduate jobs defined at the 4-digit SOC code level using the classification in Aghion et al. (2019).

---

[12]Between 2016 and 2019, 65% of all minimum-wage workers aged 25-64 were in non-graduate occupations. 24% worked part-time, and 19% were women working part-time.

4.1.4 Heterogeneous effects by gender and age

In this section we consider heterogeneity in the effects that the minimum wage has on wages and employment by different gender and age group. We adapt equation 4 in Section 3.1.3 to estimate subgroup-specific regional wage premia δ_{rg} and time trends τ_{tg} for each demographic group g. We then use these to estimate the effect of the NLW on the frequency distribution of wages for each subgroup, normalised to the pre-treatment employment rate of that subgroup.

Figure 6 shows results estimated separately by gender. The fall in employment below the NLW, and the corresponding rise at (or just above) the NLW, is more pronounced for women than for men, as expected given that women are more likely to be on low wages. The point estimate of the total employment change up to £5 of the NLW is slightly positive for men (0.20%, confidence interval -0.14% to 0.44%) and slightly negative for women (-0.44%, confidence interval -0.85% to -0.03%). The negative effect for women is (just) statistically significant at the 95% level. Like our aggregate results, this pattern - of a small and not statistically significant positive effect for men, and a small negative effect for women with a p-value close to 0.05 - is robust across a number of alternative specifications (though the p-value for women is sometimes a little above, and sometimes a little below, 0.05 - highlighting how one should not fixate on binary standards of statistical significance).

Figure 6. GENDER

(a) Men　　　　　　　　　　　　　　　(b) Women
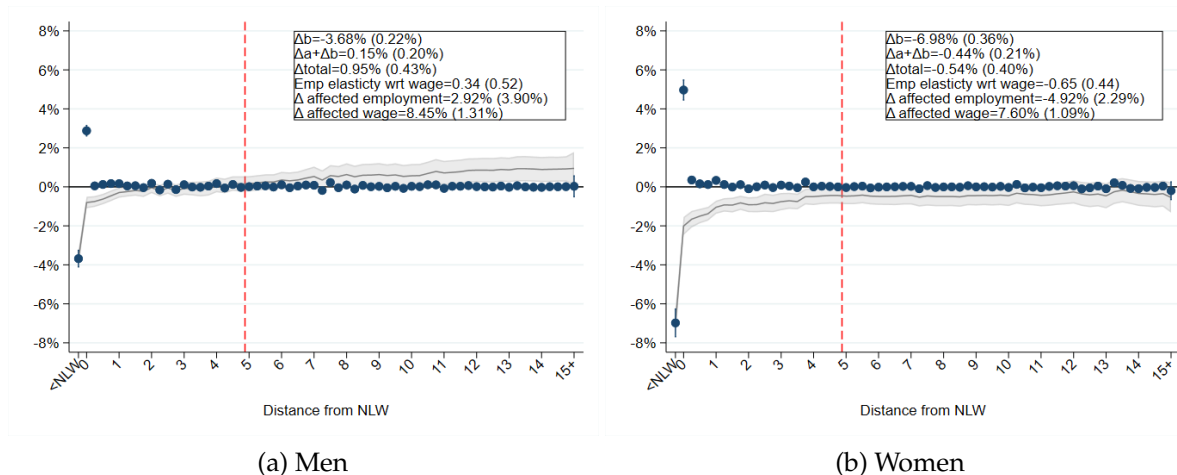
Notes: See Notes to 5.

Results by different age groups among the 25+ population are shown in Figure 7. Estimated effects on employment are small and not statistically significant for any age group shown.

We can also apply our approach to examine effects on young people under the age of 25. They were not legally affected by the NLW over the period studied here, but there are various ways in which they could be affected in practice. These could include 'downward wage spillovers' if firms avoid implementing the age-related pay differentials that the legal minima would

Figure 7. AGE
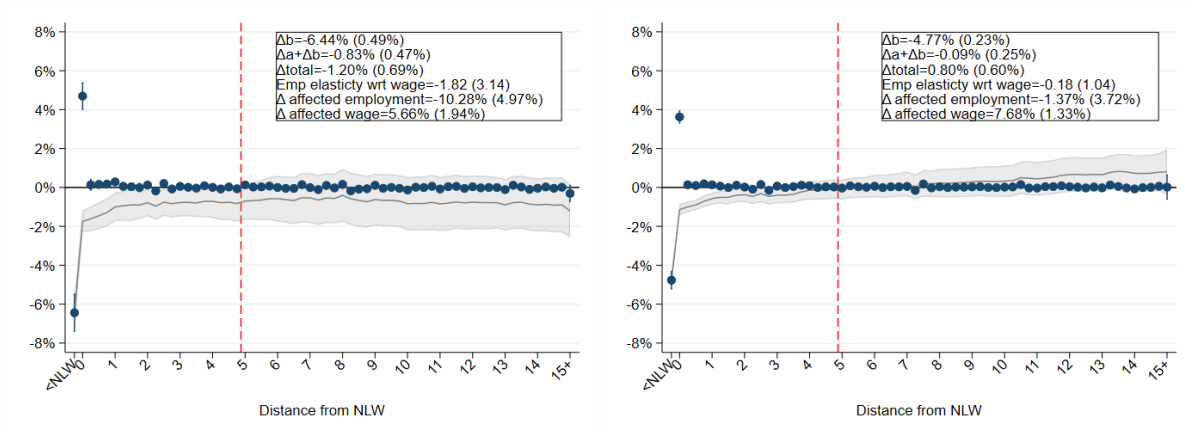
(a) 25-34

(b) 35-54
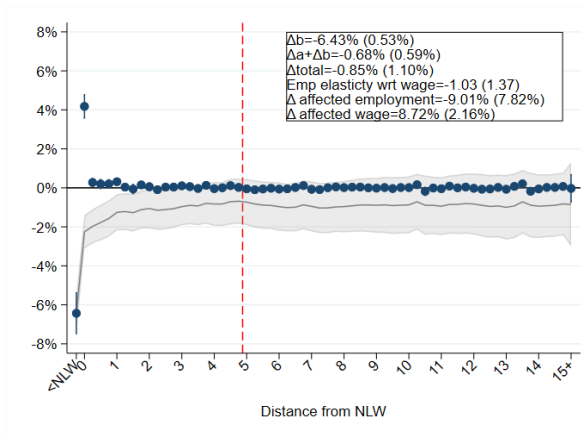
(c) 55-64

Notes: See Notes to 5.

allow, due for example to administrative costs or constraints or fairness concerns. As this would effectively represent an increase in labour cost for the under-25s, one might see impacts on the numbers of that age-group employed as a result. Alternatively, to the extent that the NLW makes under-25s cheaper to employ than older workers, labour substitution might act to increase their employment rates. The choice between education and work is also important for young people and may be impacted by minimum wage policy.

Following the same methodology as above, we estimate changes in employment around the NLW as a share of the pre-treatment employment rate among under-25s, using geographic wage premia and time trends estimated on the under-25 population. Our results suggest that downward wage spillovers onto 16-24 year olds are important: the wages of under-25s were substantially positively impacted by the introduction of the NLW and subsequent uplifts. As Figure 8 shows, the number of jobs for this age-group paid below the new NLW fell by 4.9% (std. error 1.0%) . This is consistent with previous studies that show positive wage spillovers of the NLW for younger workers, potentially reflecting employer preferences for fairness (Giupponi and Machin, 2018).[13]

Our estimates suggest it is more likely than not that the overall employment effect of the NLW for the under-25s was either broadly neutral or positive. Given that the policy clearly increased their wages, this paints a fairly positive picture for the group. It should, however, be noted that the estimates for this group are relatively imprecise due to the smaller sample size: we cannot confidently rule out substantial employment effects in either direction, as the confidence interval runs from -1.80% up to 3.10%.

4.2 Household and family incomes

We now turn to the effects of the National Living Wage on household incomes. We simulate the impact on net household incomes of the employment and wage effects estimating from our baseline bunching specification. As such, our results should be interpreted as an estimate of the average effect on household incomes of each of the four increases in the minimum wage for those aged 25-64 that occurred between 2016 and 2019 inclusive.[14]

Figure 9 shows the simulated distributional effect of the NLW across the household income distribution. We focus on households containing someone aged 25 to 64, and partition those households into deciles of income. [15] The bars separately show the effect on net income (income after taxes and benefits) and net tax payments (taxes minus benefits). The two together sum

---

[13]Simply comparing how the distribution of wages for under-25s changed before and after a NLW hike (uprating the distribution before the rise and pooling together all four years) suggests that around 250,000 under-25s saw an increase in wages to above the NLW as a result of the policy.

[14]The FRS data we use cover the period when the National Minimum Wage was £6.50. We simulate an increase to £6.93 - the 43p increase being the average year-to-year increase seen between 2016 and 2019.

[15]We assign households to income deciles based on their household equivalised income, using the OECD-modified equivalence scale, before the introduction of the NLW. We compute changes in household incomes, with the household as the unit of analysis.

Figure 8. NLW INTRODUCTION AND UPLIFTS, 2015-2019, UNDER-25S
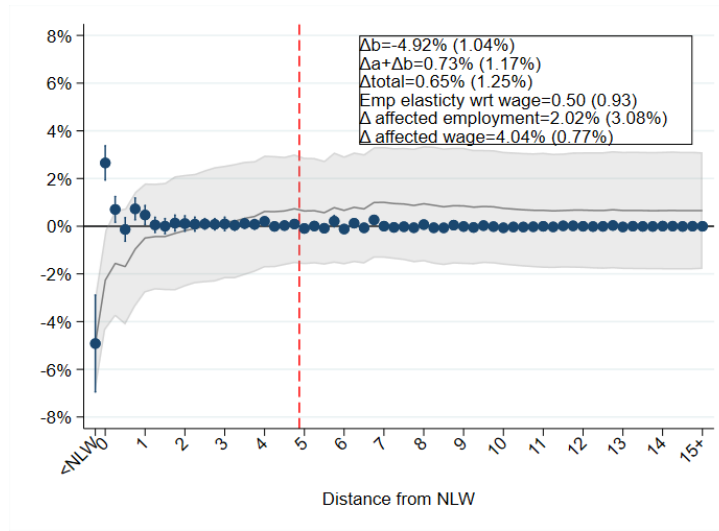
Notes:Δb is the effect of the NLW on change in employment below the NLW as a fraction of employment in the previous year . $\Delta a + \Delta b$ is the effect of the NLW on change in employment up to £5 above the NLW. $\Delta total$ is the estimated effect on total employment over the whole wage distribution. Apprentices are excluded from this analysis.

to the effect on gross household earnings. The line, plotted on the right hand axis, shows the proportional impact on net income.

On average each NLW increase raises net household incomes across the population (of households with someone aged between 25 and 64) by 0.13%. Around a third of the increase in pre-tax earnings is offset by reduced income-related benefits or higher taxes. Those clawbacks are even higher, reaching almost half, in the second and third income deciles, where many workers' households are receiving means-tested benefits which are quite rapidly withdrawn as earnings increase.[16] This limits the effect of the NLW on poorer households' income. The proportional effect of the NLW is broadly increasing across the bottom half or so of the distribution, with the strongest effects around the middle. Effects taper off fairly quickly as we move above the middle of the distribution, but it is worth noting that even in the eighth decile the proportional effect is still about half that seen in the poorest decile.

A few basic numbers from the simulation underlying Figure 9 help to illustrate the key mechanisms at work and to explain the scale of effects. The average increase in earnings among existing minimum wage workers is £8.15 per week, not accounting for spillovers or disemployment effects. After accounting for taxes and reductions in means-tested benefits, this leads to an increase in net household income of £5.75. Average household income among minimum wage workers is £588 per week, meaning that the increase resulting from the minimum wage

---

[16]The bottom decile includes a significant number of households who are not in receipt of benefits - perhaps because they are not entitled in virtue of having significant assets, or because they are not claiming benefits they are entitled to. This means that they see relatively less of the NLW gain clawed back via lower benefits when their earnings increase.

increases income by just under 1% on average. But only 5% of working-age households contain a minimum wage worker. Even in deciles 3 and 4, where minimum wage workers are most common, only 9% of households have a minimum wage worker. This explains the much more modest effects on household incomes when averaged across the population. [17]

Figure 9. NLW IMPACT ON HOUSEHOLD INCOMES, DECOMPOSED BY INCOME SOURCE
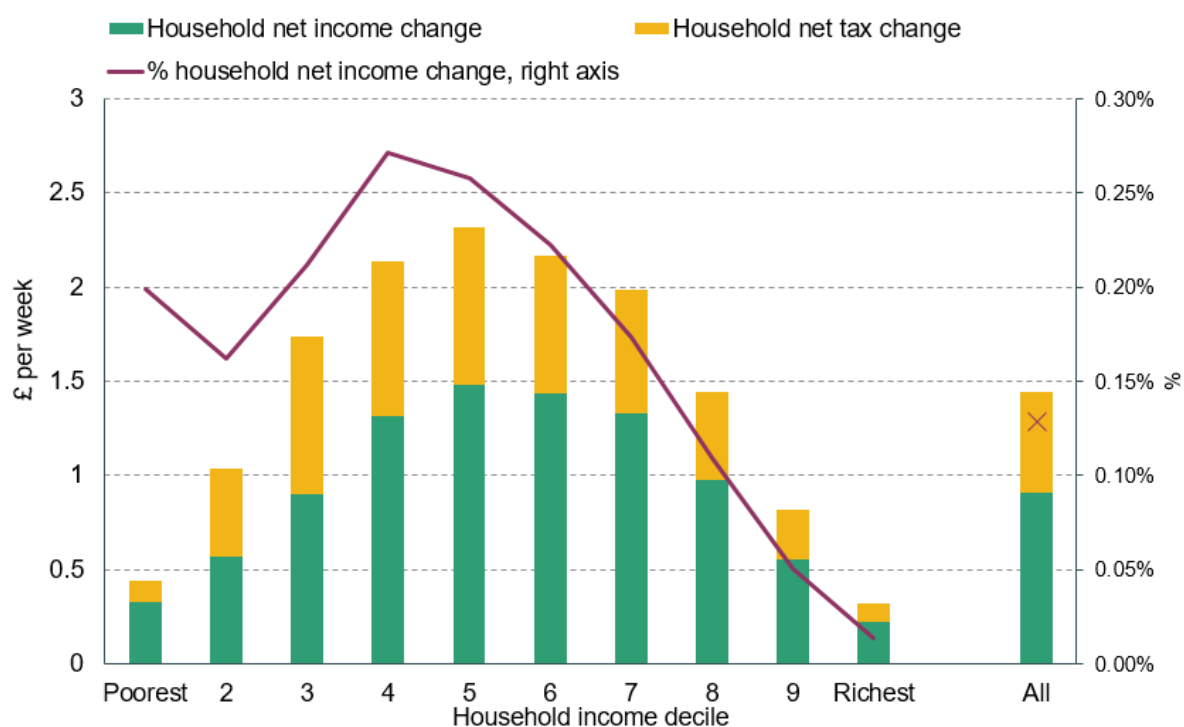
Notes: Effect is of average increase in NLW from 2015 to 2019. Cash effects on LH axis, % effects on RH axis. Sample is households with at least one person aged 25-64. Households are ranked on pre-NLW income, among this sample. Income is equivalised and net of taxes and benefits.

Part of the reason that the impact of the NLW is somewhat muted among poorer households is that many do not have anyone in work and so cannot gain from the NLW increase. If one looks only at working households - which, depending on the context, may be the more relevant population for policy-makers thinking specifically about minimum wage policy, especially if employment effects of the minimum wage are small - then a more progressive picture emerges, as seen in Figure 10. The poorest 30% of working households each see proportional gains of around 0.35%. Effects then steadily decline as one moves further up the distribution.

Thus far we have been analysing effects at the household level. An alternative approach is to assume that families are the unit of income sharing, and analyse effects at the family level, as discussed in Section 3.2.5. This matters because 35% of families with a minimum wage worker live in a household with another family. Among this group, on average the minimum

---

[17] Among working households only - a group we examine in Figure 10 - 7% of households have a minimum wage worker. The highest concentration of minimum wage workers is in the lowest decile, where the number rises to 15%

Figure 10. NLW IMPACT ON HOUSEHOLD INCOMES, AMONG WORKING HOUSEHOLDS, DECOM-POSED BY INCOME SOURCE
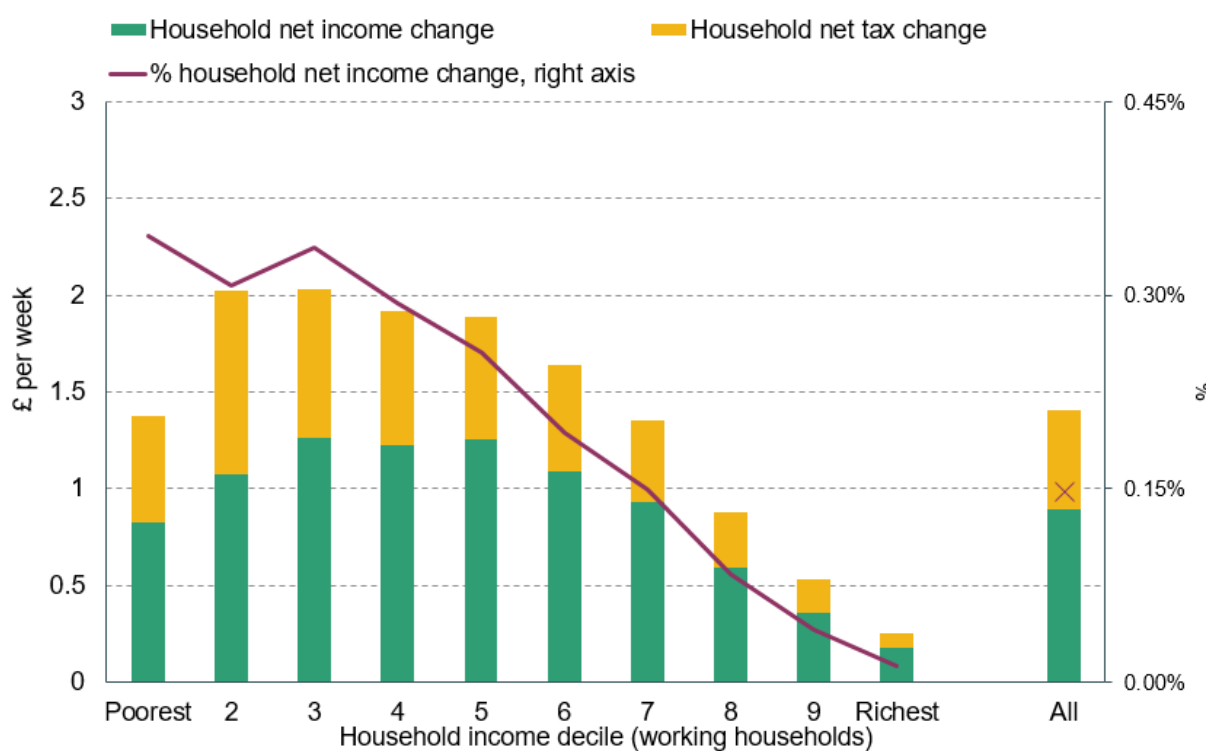
Notes: Effect is of average increase in NLW from 2015 to 2019. Cash effects on LH axis, % effects on RH axis. Sample is households with at least one person aged 25-64, and at least one person who is in work prior to the introduction of the NLW. Households are ranked on pre-NLW income, among this sample. Income is equivalised and net of taxes and benefits.

wage family accounts for 52% of the household's income. Figure 11 shows the average effect of NLW increases on family incomes (among all families, not just those in work). In cash terms the patterns are little different to that seen at the household level in Figure 9. However the proportional effect is substantially more progressive. This reflects the fact that the lowest income families have considerably less income than the lowest income households on average.

Figure 11. NLW IMPACT ON FAMILY INCOMES, DECOMPOSED BY INCOME SOURCE
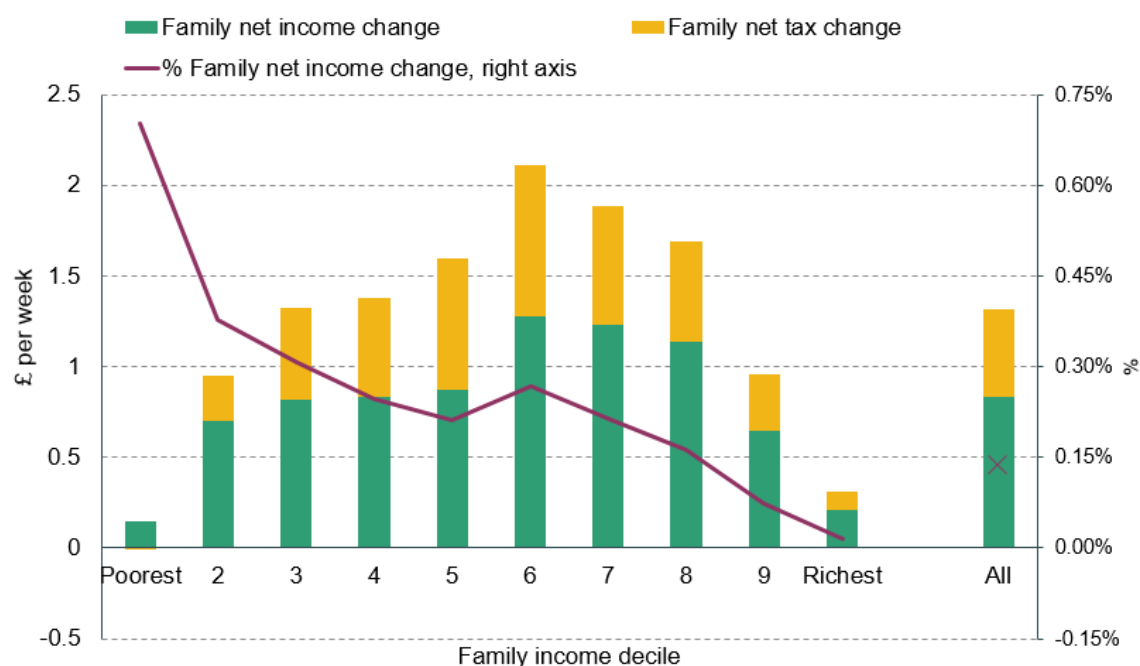
Notes: Effect is of average increase in NLW from 2015 to 2019. Cash effects on LH axis, % effects on RH axis. Sample is families with at least one person aged 25-64. Families are ranked on pre-NLW income, among this sample. Income is equivalised and net of taxes and benefits.

One advantage of our simulation approach is that we are able to decompose the effect on incomes. Figure 12 builds up to the overall effect seen in Figure 9 in several stages. We begin with the 'mechanical' effect: the impact on incomes from simply increasing wages for those paid under the NLW up to the NLW level (this is what is typically done in extant simulation exercises). We then add in the effect of wage spillovers, using the labour market results but adjusting the estimated change for the bin at the NLW to remove disemployment effects. We then add in the disemployment effect to recover the total estimated effect shown in Figure 9. The spillovers are estimated to have a large effect, a similar size or bigger than the direct mechanical change in most deciles. Spillover effects are larger further up the distribution, reflecting the fact that higher wage earners tend to be in higher income households. This turns the flat distributional effect across the bottom half from the mechanical impact to one that is biggest for middle income households. The disemployment effects have reasonably similar effects across the distribution, though slightly bigger towards the bottom where more workers

are directly affected by the NLW (rather than benefiting from spillovers) and so at risk of job loss.

Of course, these results are based on our central estimates of the labour market effects. Disemployment effects could be larger or smaller, and it turns out that this matters quite significantly. Figure 12 shows what the distributional results would look like if the disemployment effects were as large as the upper bound of the 95% confidence interval. The effects on incomes would be much lower, with only small gains across the distribution. It is worth noting that even disemployment effects this large - which, by definition, are very unlikely according to our estimates - do not fully eliminate the income gains on average (notwithstanding the caveat given earlier that we are unable to account for effects on household incomes via non-labour market channels such as prices or profits).

Figure 12. NLW IMPACT ON HOUSEHOLD INCOMES, DECOMPOSED BY SOURCE OF RESPONSE
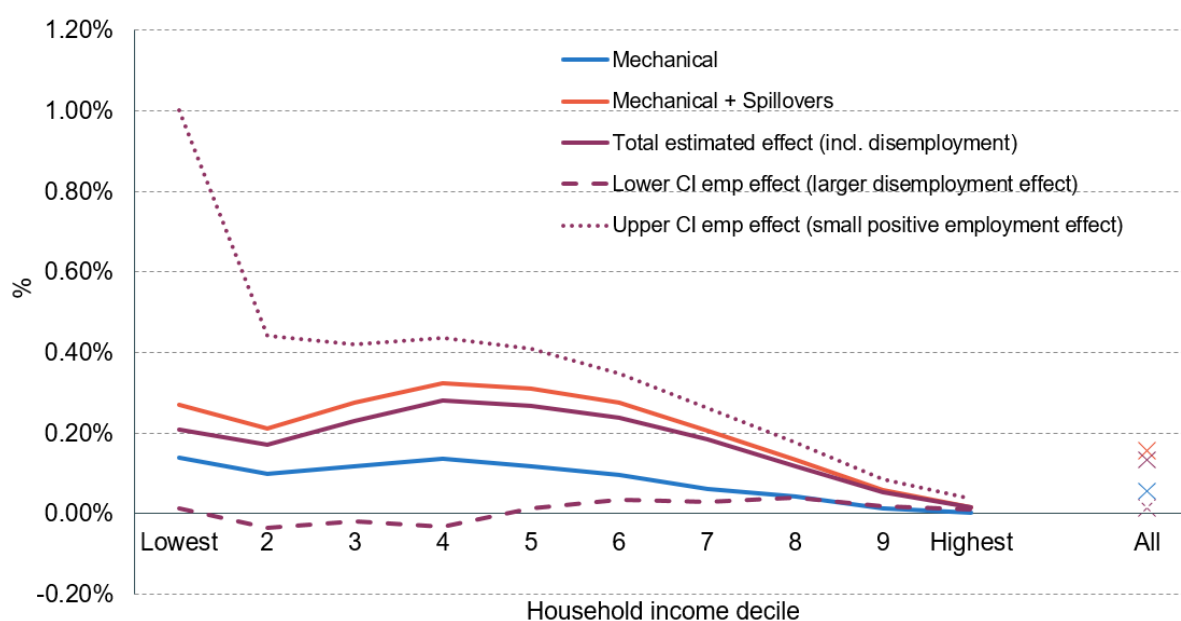
Notes: Mechanical change is that as a result of increasing wages of those previously earning below NLW to the NLW. **Mechanical + Spillovers** accounts for changes in wage distribution as a result of NLW, stripping out dis-employment effects. **Total** incorporates the full set of effects as estimated in 5. **With larger disemployment effect** incorporates employment effectswere as large as the upper bound of the 95 % confidence interval. Sample is households with at least one person aged 25-64. Households are ranked on pre-NLW income, among this sample.

5. Conclusion

We have examined the effects that the introduction of the UK's National Living Wage has had on wages, employment, and households' incomes, covering the period between the introduction of the NLW and the last pre-pandemic uprating – that is, 2015 to 2019. To do this, we have

developed a new approach to estimating the effects of a minimum wage on wages, employment and hours. We have built on the 'bunching' approach pioneered by Harasztosi and Lindner (2019) and Cengiz et al. (2019), and have applied it to a context where there is no within-country variation in minimum wages, by exploiting wage differences between different parts of the country. We estimate the impacts of the NLW on the number of jobs within each wage band, meaning that we jointly capture both employment and wage effects in a single, internally consistent framework.

In addition, the estimates of the effects of a higher minimum wage on employment and wages, combined with a tax and benefit microsimulation model, and household survey data, allow us to study the impacts of the NLW on the distribution of household incomes. Our approach enables us to account not only for employment and spillover effects onto those with higher wages, but the interactions between wages, taxes paid, and benefits and tax credits received. We can identify the relative importance of each of these mechanisms in terms of the effect of the minimum wage on household incomes. Having said that, we cannot incorporate the distributional effects potential impacts of a higher minimum wage on either profits or consumer prices.

We find that the NLW and its increases up to 2019 had substantial effects on wages towards the bottom of the wage distribution. Averaging across the four increases to the minimum wage for those aged 25+ that we consider (i.e. in April of 2016, 2017, 2018 and 2019), we estimate that each increase caused a reduction in the number of people paid below the new NLW of a magnitude equivalent to around 5.4% of employees. We find statistically significant increases in the number of jobs not only at the new NLW, but also up to around £1.50 per hour above it (approximately the 20th percentile of hourly wages) - indicating 'spillover' effects on the wages of some employees above the minimum.

Our central estimate of the impact of these minimum wage rises on employment is negative but small and not statistically significant. Averaging across each of the four increases to the minimum wage for those aged 25+ between April 2016 and 2019, we estimate that each increase reduced employment by 0.1% of the pre-policy workforce in lower-wage regions relative to high-wage regions, with a 95% confidence interval spanning -0.4% to +0.2%. Hence we can rule out large effects with high confidence. The finding of small, negative and statistically insignificant employment effects is consistent across alternative specifications. There is some evidence of more negative impacts on employment of women than men. Under 25 year olds were also affected, with large positive 'spillover effects' onto their wages. Our estimates also suggest that the impact on the employment of the group was most likely broadly neutral or positive, though this estimate is relatively imprecise due to the lower sample size.

Looking at the effects of the NLW and its increases up to 2019 on household incomes, the biggest gains go to the middle of the working-age household income distribution, both in cash and percentage terms. If we look only at households working before the introduction of the NLW, however, the impact is more progressive: each NLW increase on average raised incomes among

the bottom 30% of working households by about 0.35%, with effects steadily declining above that. Our results demonstrate that the distributional effects are very sensitive to the size of any (dis)employment effect, but they also suggest that it is highly unlikely that the net income gains arising from wage increases are, in aggregate, offset by income losses from disemployment.

References

Aghion, P., Bergeaud, R. A., Blundell, and R. Griffith. 2019. "The innovation premium to soft skills in low-skilled occuptions."

Ahlfeldt, Gabriel M., Duncan Roth, and Tobias Seidel. 2018. "The regional effects of Germany's national minimum wage." *Economics Letters*, 172: 127–130.

Autor, David H., Alan Manning, and Christopher L. Smith. 2016. "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment." *American Economic Journal: Applied Economics*, 8(1): 58–99.

Belman, Dale, and Paul J Wolfson. 2014. *What does the minimum wage do?* WE Upjohn Institute.

Brewer, Mike, and Paola De Agostini. 2017. "The National Minimum Wage, the National Living Wage and the tax and benefit system." ISER mimeo.

Brewer, Mike, Thomas Crossley, and Federico Zilio. 2019. "What do we really know about the employment effects of the UK's National Minimum Wage?" Institute for Fiscal Studies IFS Working Papers W19/14.

Caliendo, Marco, Alexandra Fedorets, Malte Preuss, Carsten Schröder, and Linda Wittbrodt. 2018. "The short-run employment effects of the German minimum wage reform." *Labour Economics*, 53: 46–62. European Association of Labour Economists 29th annual conference, St.Gallen, Switzerland, 21-23 September 2017.

Card, David. 1992. "Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage." *ILR Review*, 46(1): 22–37.

Card, David, and Alan B Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *American Economic Review*, 84(4): 772–793.

Card, David, and Alan B. Krueger. 2000. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply." *American Economic Review*, 90(5): 1397–1420.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *The Quarterly Journal of Economics*, 134(3): 1405–1454.

Clemens, Jeffrey, and Michael Wither. 2019. "The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers." *Journal of Public Economics*, 170: 53–67.

Corlett, Adam. 2021. "Improving our understanding of UK poverty will require better data." Resolution Foundation Spotlight.

Dolton, Peter, Chiara Rosazza Bondibene, and Jonathan Wadsworth. 2012. "Employment, Inequality and the UK National Minimum Wage over the Medium-Term*." *Oxford Bulletin of Economics and Statistics*, 74(1): 78–106.

Dolton, Peter, Chiara Rosazza Bondibene, and Michael Stops. 2015. "Identifying the employment effect of invoking and changing the minimum wage: A spatial analysis of the UK." *Labour Economics*, 37: 54–76.

Dube, Arindrajit. 2019*a*. "Impacts of minimum wages: review of the international evidence." H.M. Treasury, London.

Dube, Arindrajit. 2019*b*. "Minimum Wages and the Distribution of Family Incomes." *American Economic Journal: Applied Economics*, 11(4): 268–304.

Dube, Arindrajit, and Attila Lindner. 2021. "City Limits: What Do Local-Area Minimum Wages Do?" *Journal of Economic Perspectives*, 35(1): 27–50.

Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. "MINIMUM WAGE EFFECTS ACROSS STATE BORDERS: ESTIMATES USING CONTIGUOUS COUNTIES." *The Review of Economics and Statistics*, 92(4): 945–964.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge. 2021. "Reallocation Effects of the Minimum Wage." *The Quarterly Journal of Economics.* qjab028.

Giupponi, G., and S. Machin. 2018. "Changing the Structure of Minimum Wages: Firm Adjustment and Wage Spillovers."

Harasztosi, Peter, and Attila Lindner. 2019. "Who Pays for the Minimum Wage?" *American Economic Review*, 109(8): 2693–2727.

Harkness, Susan E, and Silvia Avram. 2019*a. The impact of minimum wage upratings on wage growth and the wage distribution. A report prepared for the Low Pay Commission.*

Harkness, Susan E, and Silvia Avram. 2019*b. The NMW/NLW and progression out of minimum wage jobs in the UK: Final Report.*

Manning, Alan. 2016. "The Elusive Employment Effect of the Minimum Wage." Centre for Economic Performance, LSE CEP Discussion Papers dp1428.

Manning, Alan. 2021. "The Elusive Employment Effect of the Minimum Wage." *Journal of Economic Perspectives*, 35(1): 3–26.

Neumark, David, and William L. Wascher. 2008. *Minimum Wages.* Cambridge, MA, USA:MIT Press.

Neumark, David, J. M. Ian Salas, and William Wascher. 2014. "Revisiting the Minimum Wage—Employment Debate: Throwing Out the Baby with the Bathwater?" *ILR Review*, 67(3_suppl): 608–648.

Portugal, Pedro, and Ana Rute Cardoso. 2006. "Disentangling the Minimum Wage Puzzle: An Analysis of Worker Accessions and Separations." *Journal of the European Economic Association*, 4(5): 988–1013.

Sabia, Joseph J., and Richard V. Burkhauser. 2010. "Minimum Wages and Poverty: Will a $9.50 Federal Minimum Wage Really Help the Working Poor?" *Southern Economic Journal*, 76(3): 592–623.

Schmitz, Sebastian. 2019. "The Effects of Germany's Statutory Minimum Wage on Employment and Welfare Dependency." *German Economic Review*, 20(3): 330–355.

Skinner, Chris, Nigel Stuttard, Gabriele Beissel-Durrant, and James Jenkins. 2002. "The Measurement of Low Pay in the UK Labour Force Survey*." *Oxford Bulletin of Economics and Statistics*, 64(supplement): 653–676.

Stewart, Mark B. 2002. "Estimating the Impact of the Minimum Wage Using Geographical Wage Variation*." *Oxford Bulletin of Economics and Statistics*, 64(supplement): 583–605.

Stewart, Mark B. 2004. "The Employment Effects of the National Minimum Wage." *The Economic Journal*, 114(494): C110–C116. Publisher: [Royal Economic Society, Wiley].

Waters, Tom. 2017. "TAXBEN: The IFS tax and benefit microsimulation model." Institute for Fiscal Studies IFS User Guides.

Appendix A. Additional figures and tables

Figure A1. ESTIMATED WAGE PREMIUMS
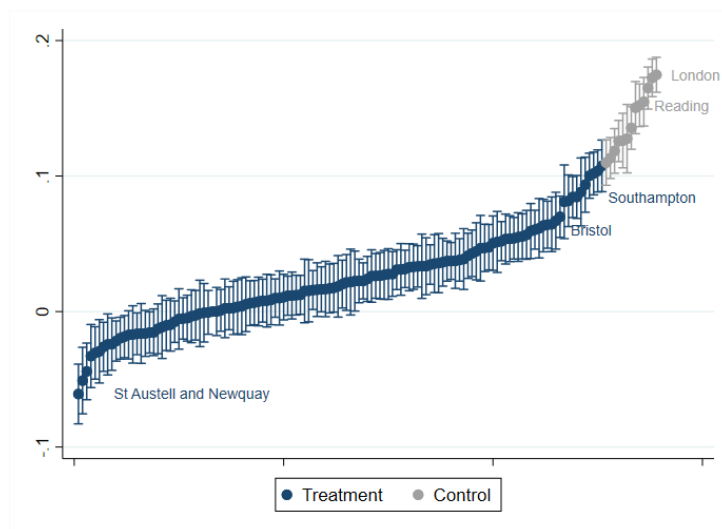
Notes: Wages premiums for each region are estimated using Equation 3. The solid dots show point estimates, and the error bars show the 95% confidence interval.

Table A1. ESTIMATES OF EQUATION 2

	2015 to 2016	2016 to 2017	2017 to 2018	2018 to 2019
τ_t	0.0172***	0.0274***	0.0234***	0.0215***
	(-5.35)	(-8.04)	(-6.87)	(-6.11)
GAP_{rt}	2.930**	0.732	2.32	4.350**
	(-2.87)	(-0.35)	(-1.13)	(-2.66)
Observations	288059	284815	286512	284552

Notes: Parametric standard errors in parentheses. ** denotes statistical significance at the 95% level; *** at the 99% level.

Figure A2. Placebo checks

| (a) 2013-2015	(b) 2012-2015
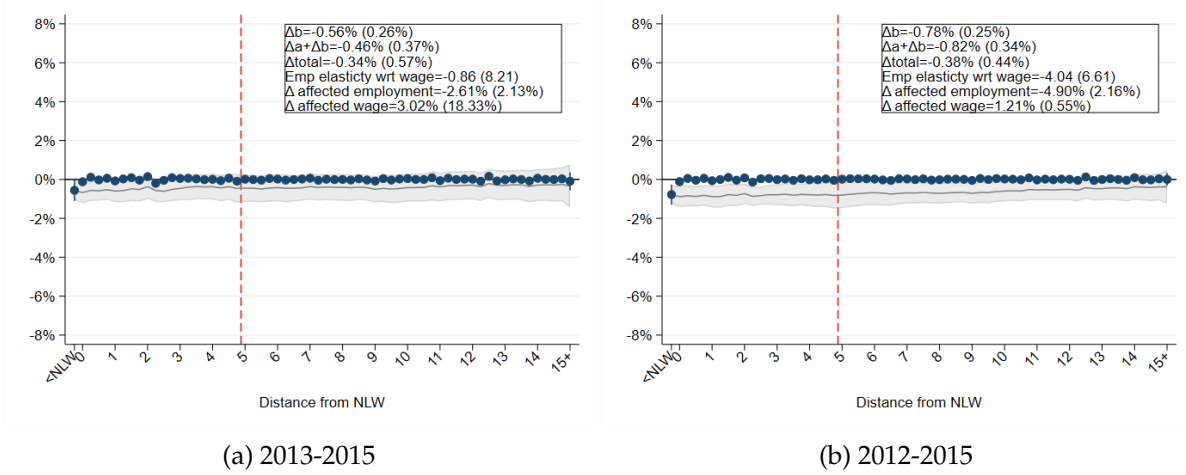
Notes: These figures apply our method to data from before the introduction of the NLW. Panel (a) looks at changes around the 2016 NLW in 2013-14, and the 2017 NLW in 2014-15. Panel (b) applies the 2016 NLW to 2012, the 2017 NLW to 2013, and the 2018 NLW to 2014. The figures show no difference in employment by wage bin in lower-wage regions compared to their equivalent real wage bins in high-wage regions, with the exception of the bottom wage bin, particularly in the 2012-2015 placebo. Given that employment was still recovering from the Great Recession in 2012, this could have led to differential trends between higher- and lower-wage regions at the bottom of the wage distribution. Also see notes to Figure 5.

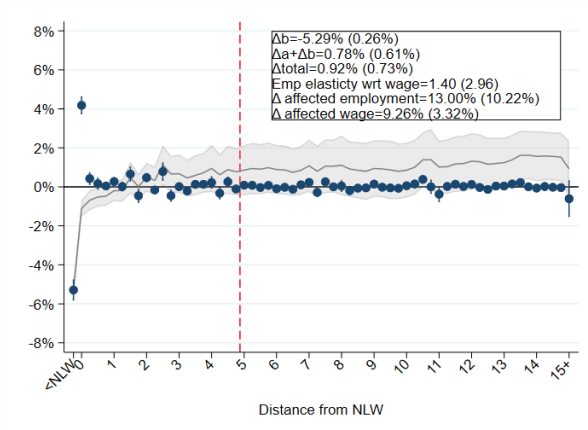## Figure A3. NLW introduction and uplifts, big jump 2015-2019
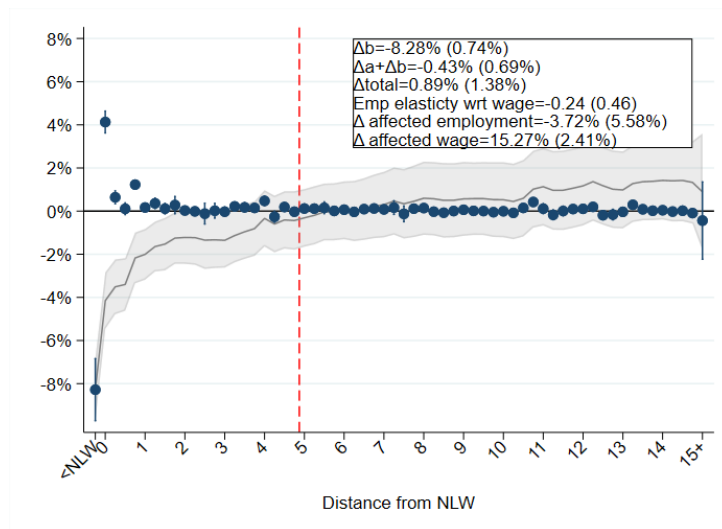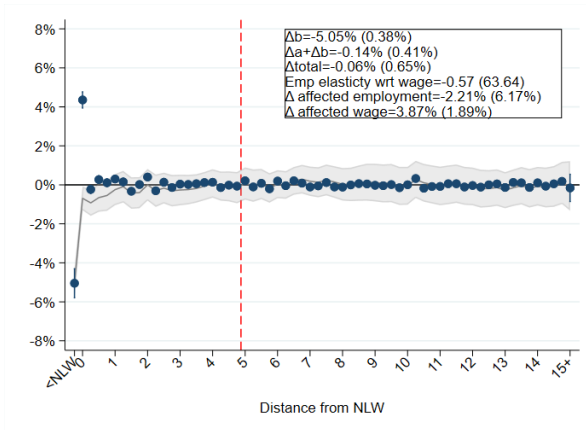

Notes: See notes to Figure 5.
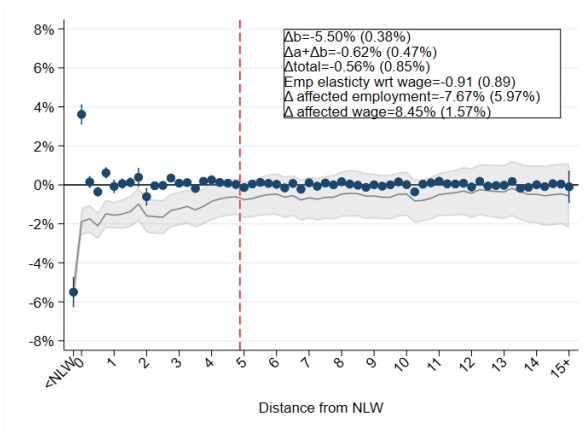
Figure A4. SINGLE-YEAR ESTIMATES
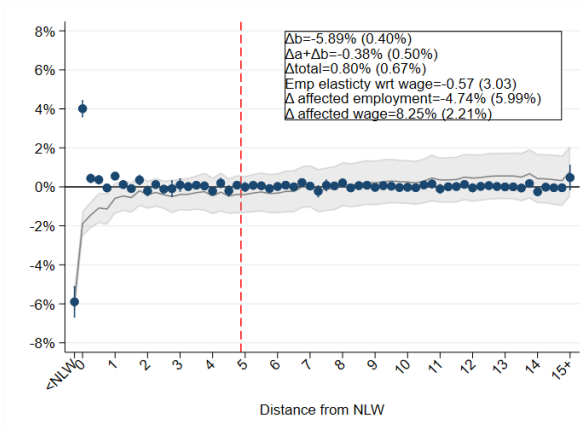
(a) 2015-2016

(b) 2016-2017

(c) 2017-2018

(d) 2018-2019

Notes: See notes to Figure 5.