The Employment and Distributional Impacts of Nationwide Minimum Wage Changes*

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

May 23, 2022

Abstract

We assess the impact of nation-wide minimum wages on employment throughout the whole wage distribution by exploiting geographical variation in the level of wages. We 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. At the same time, we find small negative effects on employment which are not statistically significant. Combining the estimated change in the wage distribution with a tax and benefit microsimulation model we also estimate the impact on household incomes. The largest gains of the minimum wage go to the middle of the overall working- age income distribution, while the gains to poorer working households are limited by the withdrawal of means tested benefits as earnings increase.

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

*The authors are grateful to the Low Pay Commission (grant number CR20017) for funding and to the Economic and Social Research Council for co-funding through the Centre for the Microeconomic Analysis of Public Policy at IFS (ES/T014334/1). Lindner acknowledges financial support from the Economic and Social Research Council (ES/T008474/1) and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant number 949995). We thank Kirill Borusyak, Alex Bryson, Tim Butcher, Arin Dube, Thomas Le Barbanchon, Rebecca Riley and participants in seminars at the Low Pay Commission, the IAB-LASER workshop, the Milan Labour Lunch Seminar, the KOF Swiss Economic Institute at ETH Zurich and the European Labor Symposium for Early Career Economists for helpful comments. We also thank Jonathan Cribb and Anna Becker for help with data analysis. This work was produced using data from ONS. Data from the Households Below Average Income dataset and Family Resources Survey were made available by the Department for Work and Pensions. We use research datasets which may not exactly reproduce National Statistics aggregates. Use of the ONS and DWP data does not imply endorsement. Any errors and all views expressed are those of the authors.

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

1. Introduction

There has been a revival of interest in minimum wage policy in recent years. Minimum wages are seen as a central policy lever to boost wage growth at the bottom of the distribution, and increasingly viewed as a tool to reduce in-work poverty and support low-income households. Substantial increases in minimum wages have been implemented in several countries, including the UK, Germany, Hungary, Poland and Spain. This makes it imperative to understand the effect of minimum wages on employment, wages and household incomes.

A large number of studies have been conducted on the employment effects of minimum wages (for reviews, see Neumark and Wascher, 2008; Belman and Wolfson, 2014; Dube, 2019*a*).However, there are relatively fewer studies on the effects of minimum wages across the wage distribution, and how this translates into effects on household incomes. Notable recent exceptions include Cengiz et al. (2019) on wages and Dube (2019*b*) on household incomes, but these exploit state-level variation in minimum wages in the US, whereas in most countries there is a single minimum wage that applies across all geographic areas.

This paper proposes a new empirical methodology to estimate the impacts of the minimum wage on employment and wages 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).[1] This allows us to identify how the labour market effects of the minimum wage feed through to households' net incomes, after accounting for taxes and benefits. We apply our method to a large increase to the minimum wage in the UK in 2016, which brought the level of the UK minimum wage close to the international frontier.

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 reflect differences in general price levels, at least at the lower end of the wage distribution. Our method compares trends in employment between groups who would expect to earn the same wage if they lived in the same area, but are differentially exposed to the minimum wage because they live in different areas, which have different regional wage premia. 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,

[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; and Portugal and Cardoso (2006) for Portugal. Appendix B provides a classification of studies of the employment effects of the minimum wage in European countries, by method used for identification.

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 apply this new methodology to study the impacts of the UK's NLW introduction, and its subsequent upratings, on the entire frequency distribution of wages. Introduced in April 2016 with the goal of reaching two thirds of median wages by 2024, the NLW increased the minimum wage by 7.5% in real terms, bringing the bite of the minimum wage close to the the international frontier (see Appendix Figure A1). 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 and corroborates findings of previous work in the UK using other methods (Dube, 2019*a*). 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.

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 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 frequency-distribution method allows us to assess the change in employment across the entire distribution of hourly wages, which has several key advantages. Firstly, disaggregating 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. If the results suggest otherwise then this is a hint that the identification assumptions are not satisfied.[2]

The way in which we identify our frequency-distribution estimates, using regional wage variation, refines the 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 frequency-distribution 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 analyse 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 frequency-distribution 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 Section 4 illustrates the related empirical results. The methodology used to simulate impacts on household incomes is described in Section 5 and the simulation results are reported in Section 6. Section 7 concludes.

[2]See Appendix B in Cengiz et al. (2019).

2. Institutional context

We analyse 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 plots 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 (panel A) and as a percentage of the median wage (panel B). 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.[3]

3. Employment and wage effects of the minimum wage: methodology and data

3.1 Identification strategy

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, in a setting in which a single minimum wage policy applies across the entire country. We identify the effects using a 'regional variation' approach that exploits geographic variation in wage levels, in the tradition of Card (1992). We then nest the 'regional variation' approach into the 'frequency distribution' approach – pioneered by Cengiz et al. (2019) – to trace out the effect of the minimum wage throughout the wage distribution. We start by summarising those approaches and illustrating how we combine features of both. We then describe the empirical implementation.

[3]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).

Regional-variation approach. A common approach for estimating the impacts of minimum wages on employment is to exploit geographic variation in its bite. This approach 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:

$$\Delta E_{rt} = \alpha MIN_{rt-1} + \gamma_t + \mu_{rt} \tag{1}$$

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 looks 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 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, for example 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.

Frequency-distribution approach. The frequency-distribution 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 maintain pay differentials between different occupations, or between supervisory and non-supervisory roles. 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. As we explain in more detail in the next paragraph, for identification, 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) and including 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).

Nesting the regional and frequency-distribution 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. Retrieving an *absolute* one requires additional assumptions that we describe below.

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 calculated. We then use employment changes within each wage bin in lower-wage areas from before to after the minimum wage introduction, net of the counterfactual employment changes identified from what happens in high-wage areas within the same 'real' (but different nominal) wage bin. And we relate those net employment changes to the proximity of the wage bin to the minimum wage. 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 frequency-distribution approach.

Our methodology retains the advantages of the frequency-distribution approach, while adapting it to be applied in a context with uniform national minimum wage policy. Viewed the other way around, we 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, 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 homogeneous treatment effect assumptions (Cengiz et al., 2021).

3.2 Data and sample construction

Our primary data source for the analysis of the impacts of the National Living Wage (NLW) on employment, wages and hours is the Annual Survey of Hours and Earnings (ASHE) for the years from 2010 to 2019. A large-scale, employer-completed survey of earnings and hours of employees in the UK, ASHE provides high-quality data on wages, hours, occupation, industry and basic demographic characteristics at yearly frequency. The survey is collected in April of each year. 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 locally representative Annual Population Survey (APS). 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.

3.3 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 wages for aggregate wage growth. As a first step, we net out overall wage growth from wage levels. What this means in practice is that when studying the impact of the NLW in year t, we 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 using the following regression model, estimated for each year pair from 2015 to 2019:

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

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 Appendix Table A1.

Estimating regional wage effects. As a second step, we define groups of workers who would earn the same wage if they lived in the 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 from 2012 to 2015. Our regression specification is:

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

where X includes individual 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 proxy for being a graduate (based on 4-digit occupation codes). 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 2. Estimates in blue refer to regions in the bottom nine deciles of the distribution of wage premia (later our treatment group), grey ones to the top decile (later our control group). The difference between the smallest and largest premia is of approximately 25 percent. Note that specification 3 assumes that (proportional) regional wage premia are fixed across worker characteristics such as gender, age and so on.

One limitation of our Mincerian estimates is that they do not account for unobservable individual heterogeneity. An alternative is to estimate an 'AKM' specification with two-way fixed effects (Abowd, Kramarz and Margolis, 1999) of the following form:

$$\ln w_{irt} = \gamma_i + \gamma_{r(p)} + Z'_{it}\beta + \eta_{irt} \tag{4}$$

where γ_i are individual fixed effects, γ_r are TTWA fixed effects, Z is a vector of one-digit industry and occupation controls, and η is an error term. Regional wage premia estimated using the AKM specification and the Mincerian specification in equation 3 are highly correlated as shown in panel A of Appendix Figure A2, though estimates using the AKM specification are much less precise due to the small number of movers in our sample. To minimize the risk of limited-mobility bias, we group TTWAs into 30 groups p based on the Mincerian estimates of δ_r from model 3, and estimate model 4 using grouped region fixed effects ($\gamma_{r(p)}$) in lieu of single regional fixed effects (γ_r). Panel B of Appendix Figure A2 shows the correlation between Mincerian regional wage premia δ_r from model 3 (vertical axis) and AKM estimates of $\gamma_{r(p)}$ from model 4 (horizontal axis). The two are strongly positively correlated, though the former are on average larger than the latter. In Section 4.2, we will show robustness of our main estimates of the employment effects of the minimum wage to using AKM and grouped-AKM rather than Mincerian regional wage premia.

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 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 25 pence (£0.25), and later assess the sensitivity of our results to different bin widths.

Sorting bias in the estimation of wage premia. One potential concern for the identification of regional wage premia is sorting of workers across regions based on an idiosyncratic match component of wages, δ_{ir}. In the presence of sorting on δ_{ir}, model 3 would be misspecified and our estimate of δ_r biased. The correct specification would be:

$$\ln w_{irt} = X'_{it} + \delta_{r(i,t)} + \delta_{ir(i,t)} + \theta_t + \nu_{irt} \tag{5}$$

where, for the sake of clarity, we indicate with $r(i,t)$ the region in which individual i is located at time t. Card, Heining and Kline (2013) propose a test to assess the empirical relevance of this type of sorting, looking at the wage changes of individuals moving across regions. The key intuition is that, if workers tend to locate in a given region based on the match component, then we would expect that the *wage gains* of workers who move from one region to another to be different from the *wage losses* of those who make the opposite transition. Ignoring any wage growth arising from changes in job characteristics, experience or year effects, the expected wage change for a worker who moves from region 1 to region 2 between period $t-1$ and t is

$$E[\ln w_{irt} - \ln w_{irt-1} | r(i,t) = 2, r(i,t-1) = 1] = \underbrace{\delta_2 - \delta_1}_{\text{Wage gain}} + \underbrace{E[\delta_{i2} - \delta_{i1} | r(i,t) = 2, r(i,t-1) = 1]}_{\text{Sorting bias}} \tag{6}$$

where we assume that the regional wage premium in region 2 is greater than in region 1. Similarly, the expected wage change for a worker who moves from region 2 to region 1 between period $t-1$ and t is

$$E[\ln w_{irt} - \ln w_{irt-1} | r(i,t) = 1, r(i,t-1) = 2] = \underbrace{\delta_1 - \delta_2}_{\text{Wage loss}} + \underbrace{E[\delta_{i1} - \delta_{i2} | r(i,t) = 1, r(i,t-1) = 2]}_{\text{Sorting bias}} \tag{7}$$

If sorting based on the idiosyncratic match component is relevant, both sorting-bias terms are positive, leading to wage gains for movers in both directions. If, instead, sorting bias is close to zero, then the wage gain from moving from 1 to 2 ($\delta_2 - \delta_1$) is similar in absolute value to the wage loss of moving from 2 to 1 ($\delta_1 - \delta_2$).

We take this test to the data and compare the wage gains and losses of individuals who move across regions with different wage premia. We start by ranking TTWAs into five groups corresponding to quintiles of the distribution of regional wage premia estimated with Mincerian regression 3. We then compute the percentage wage change experienced by individuals moving from a lower (higher) quintile to a higher (lower) one, after controlling for basic demographics, between 2012 and 2015 (the period used to estimate wage premia in equation 3).[4] Results are reported in panel A of Appendix Figure A3. Each dot (o,d) represents the average wage gain associated with moving from origin quintile o to destination quintile d (vertical axis), and the wage loss associated with moving from origin quintile d to destination quintile o (horizontal axis). Having all points lying in a narrow range around the off-diagonal indicates that wage gains and losses are approximately symmetric, suggesting that sorting bias is likely to be of minor relevance in this context. Panel B of Appendix Figure A3 conducts the same symmetry test using AKM-estimated regional wage premia γ_r from model 4, and panel C using grouped-AKM wage premia $\gamma_{r(p)}$. Again, the location of the points in a tight interval around the off-diagonals suggests that idiosyncratic match effects are unlikely to be of major concern.

[4] As in our Mincerian regression of regional wage premia, we control for gender interacted with full-time/part-time status and age, 1-digit occupation, 1-digit industry and a proxy for being a graduate based on 4-digit occupation codes. This controls for the fact that the types of workers who are more likely to move (men, young workers) also tend to experience higher wage growth. Conditional on basic demographics, movers do not experience higher wage growth than stayers.

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 for the years from 2015 to 2019, 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:

$$\Delta \frac{E_{br}}{N_r} = \sum_{j=1}^{J} \mu_j s_{br}^j + \sum_{k=\underline{K}}^{K} \alpha_k \mathbb{I}[b = k | r \in L] + \epsilon_{br} \tag{8}$$

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 model 8 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 (we indicate this top bin by K).

In practice, the estimation of model 8 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 $\times$ region level excluding the bottom 90% of regions, and simply estimate:

$$\Delta \frac{E_{jr}}{N_r} = \sum_{j=1}^{J} \mu_j + \varepsilon_{jr} \quad \text{for } r \in H \tag{9}$$

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 $\times$ region level, containing only the bottom 90% of regions. Given the job type corresponding to any wage bin $\times$ region combination, we can subtract counterfactual employment rate changes based on the applicable job-type trend estimated based on model 9: $\Delta\frac{E_{jr}}{N_r}' = \Delta\frac{E_{jr}}{N_r} - \sum_{j=1}^{J}\hat{\mu}_j s_{br}^j$, 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 $\times$ region combination, relative to the estimated counterfactual in the absence of a minimum wage change. We regress this on nominal wage dummies:

$$\Delta\frac{E_{jr}}{N_r}' = \sum_{k=\underline{K}}^{K} \alpha_k \mathbb{I}[b=k] + \eta_{br} \quad \text{for } r \in L \tag{10}$$

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 $\Delta b = \frac{\sum_{k=\underline{K}}^{0}\hat{\alpha}_k}{(E/N)}$ and the excess mass above it as $\Delta a = \frac{\sum_{k=0}^{\overline{K}}\hat{\alpha}_k}{(E/N)}$, 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 employment due to the NLW. For our baseline estimates we set $\overline{K}$ equal to NLW + 5. We show the sensitivity of our results to alternative choices, and we also routinely report an estimate of $\Delta\text{total} = \Delta a' + \Delta b$, where $\Delta a'$ aggregates the estimated $\hat{\alpha}_k$ over the entire support of the wage distribution (from $\underline{K}$ to K). This is never far from our central estimate of the employment effect (from $\underline{K}$ to $\overline{K}$), which is reassuring evidence in favour of our identifying assumptions: 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, as we explain in the next paragraph, we calculate an elasticity of area employment with respect to area wages. 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 model 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 of the minimum wage on affected employment as the relative change in employment as a share of baseline (given by $\sum_{k=\underline{K}}^{\overline{K}} \hat{\alpha}_k$), divided by the share of the workforce earning below the new minimum wage in the year before treatment ($\overline{b}_{-1}^{P}$). The superscript P indicates that the term is calculated across the whole population (i.e. both high and low wage regions).

$$\%\Delta e_L - \%\Delta e_H \approx \frac{\Delta e_L - \Delta e_H}{\overline{b}_{-1}^{P}} = \frac{\sum_{k=\underline{K}}^{\overline{K}} \hat{\alpha}_k}{\overline{b}_{-1}^{P}} \tag{11}$$

We use the estimated $\hat{\alpha}_k$ coefficients also to compute the proportional impact of the minimum wage on the average wage of affected workers. We first calculate the proportional relative effect of the minimum wage on the 'real' average wage (i.e. wages purged of regional wage premia) of affected workers. We then divide that by pre-policy 'real' average wages among affected workers, as illustrated in the following formula:

$$\%\Delta w_L - \%\Delta w_H \approx \frac{\frac{\overline{rwb}_{-1}^{P} + \sum_{k=\underline{K}}^{\overline{K}} \left(k + \overline{MW}^{P}\right) \hat{\alpha}_k}{\overline{b}_{-1}^{P} + \sum_{k=\underline{K}}^{\overline{K}} \hat{\alpha}_k}}{\frac{\overline{rwb}_{-1}^{P}}{\overline{b}_{-1}^{P}}} - 1 \tag{12}$$

Average wages are computed by taking the ratio of the total wage bill collected by affected workers to the number of such workers. In the pre-policy period, the average wage is computed as the ratio of the real pre-period wage bill among those paid less than the new minimum

$\overline{rwb}^P_{-1}$ divided by the share of the workforce earning below the new minimum $\overline{b}^P_{-1}$.[5] This is the denominator in formula 12.

To understand how the proportional relative effect of the minimum wage on the 'real' average wage of affected workers is computed, it is useful to note that the minimum wage causes both the wage bill and employment to change. The total wage bill collected by affected workers is computed by summing the pre-policy wage bill $\overline{rwb}^P_{-1}$ and the wage bill increase generated by the minimum wage in low-wage regions relative to high-wage ones $\sum_{k=\underline{K}}^{\overline{K}} \left(k + \overline{MW}^P\right) \hat{\alpha}_k$, where $\overline{MW}$ is the average wage in the bin where the minimum wage falls in the post-period. This is then divided by the sum of the pre-policy number of workers paid below the new minimum plus the relative increase in the number of affected workers in low- vs high-wage regions $\overline{b}^P_{-1} + \sum_{k=\underline{K}}^{\overline{K}} \hat{\alpha}_k$. The ratio of these two quantities gives us the numerator in formula 12.

The own-wage elasticity of employment is obtained by dividing the formula in 11 by that in 12.

4. Employment and wage effects of the minimum wage: results

4.1 Main results

Figure 3 reports our main estimates of the effect of the NLW introduction and subsequent uplifts on the frequency distribution of hourly wages from equation 8. Each dot represents our estimate of employment changes – averaged over the four minimum wage increases from 2015 to 2019 – in each £0.25 wage bin relative to the level of the new NLW in each of the years we consider. We normalise all employment changes by baseline employment in the TTWA, so that the sum of the effects across all 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, the grey line represents 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 vertical bars underlying the dots and the shaded area around the grey line show the bootstrapped 95% confidence intervals associated to the relevant estimate. The figure also reports estimates of the terms Δb, $\Delta e = \Delta a + \Delta b$ and $\Delta \text{total} = \Delta a' + \Delta b$, as described in Section 3.3, with bootstrapped standard errors in parenthesis. Estimates of the own-wage employment elasticity and its sub-components – the percentage change in affected employment and affected wages – are also reported.

In interpreting these results, it is important to recall that our methodology effectively allows us to retrieve the estimated effect of the NLW on the distribution of wages in lower-wage regions

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

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.

Our central estimate indicates that, on average, each increase in the minimum wage for those aged 25+ between 2015 and 2019 led to a fall in employment below the NLW of 5.43% (std. error 0.28%) of total employment in the previous year (Δb). Employment at, or within £0.25 of, the new minimum wage, increased by 4.04% (std. error 0.17%) of pre-treatment employment. We find statistically significant increases in the number of jobs at wages slightly higher than the NLW, with spillovers stretching up to around £1.50 above it. 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 and statistically insignificant spillovers up to around £4 above the NLW.

To compute the total employment effect, we add up all employment changes up to £5 above the NLW ($\Delta a + \Delta b$). The missing (Δb) and excess (Δa) masses are of almost identical size in absolute value, so the total employment effect is -0.09% (std. error 0.17%) of pre-treatment employment – a very small decline which is not statistically significant. The estimated employment effect is not sensitive to our choice of cut-off: for example, if instead we looked up to £3.50 above the NLW, the change in employment would be -0.32% (95% CI: -0.61% to -0.02%); if we looked 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 an employment effect that is close to zero, as is our estimate of the effect over the entire wage distribution (Δtotal $= 0.27\%$, std.error 0.33%). The fact that the frequency-distribution approach forces transparency over those changes offers 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. We provide further evidence corroborating the validity of our identification assumption in Section 4.2 below.

We calculate the own-wage employment elasticity using the formula described in Section 3.3. Our central estimate is -0.17 (std. error 0.38) – a small to moderate effect, which is in line with several other estimates in the literature (Dube, 2019*a*), including previous studies in the UK (Stewart, 2004*a*; Dube, 2019*a*; Manning, 2021). Since our estimate of the elasticity depends on estimates of both wage effects and employment effects, the confidence interval around our point estimate is relatively wide, ranging from -0.91 to 0.57.However, it allows us to rule out large negative dis-employment effects.

Appendix Figure A4 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 for each of the four uplifts, we estimate the 'long difference' from 2015 to 2019. Unlike simply pooling the four analyses of year-to-year changes, this specification allows for some lagged adjustments to be captured – for example, delayed effects on firm exit and hence employment from the 2016 NLW which would only 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 (std. error 0.69%), because it uses much less data than our central estimates which effectively pool the results from four different minimum wage increases.

4.2 Identification tests and robustness checks

Identification tests. Our identification strategy rests on a 'common trends' assumption that, absent changes in the minimum wage, employment rates in each 'job type' would evolve in the same way across lower and higher wage regions. As we already noted in Section 4.1, the fact that we do not see differential employment changes between treated and control regions in the upper tail of the wage distribution is reassuring in this respect. Of course, the absence of differential trends at the top of the wage distribution does not entirely rule out differences at the bottom of the wage distribution. Since the UK has had a system of minimum wage rates in place since 1999, with upratings taking place each year albeit with smaller increases than the NLW, there is no perfect placebo. Appendix Figure A5 reproduces estimates from model 8 using pre-NLW years 2011 to 2015, and taking the 2016 to 2019 NLW rates as placebo minima in each year (from 2012) respectively. The analysis yields a close-to-zero employment effect throughout the wage distribution, corroborating the assumption of no differential trends

One potential threat to our identification assumption is a systematic change in outward mobility in low-wage relative to high-wage regions after the minimum wage increase. For example, if low-paid workers in areas strongly affected by the NLW moved to less affected areas, this would bias our results. We investigate the relevance of this issue by comparing the mobility of low-paid and higher-paid workers across out treatment and control regions, before and after the NLW introduction. Using ASHE data on employees aged 25 and over, we estimate a linear probability model of the following form:

$$move_{irt} = \alpha + \beta L_i \cdot P_t + \mu L_i + \gamma_r + \gamma_t + \varepsilon_{irt} \tag{13}$$

where $move_{irt}$ is an indicator taking value 1 if individual i in treatment-group TTWA r at time t is observed moving to a control-group TTWA in $t+1$. The variable L is an indicator for whether the individual is low-paid, which is defined as having an hourly wage in the bottom 10% of the national wage distribution. P is an indicator taking value 1 after the NLW introduction. The parameters γ_r and γ_t indicate region and year fixed effects respectively. We estimate two versions of model 13: one in which the outcome measures movements from a treated (low-wage) to a control (high-wage) region, and one in which it measures movements from a control to a treated region. The parameter β measures the change in the probability of moving for low-paid workers relative to not-low-paid ones from before to after the NLW introduction. Coefficient estimates reported in Appendix Table A2 indicate that low-paid workers are not more likely to move out of, or into, low-wage regions after the introduction of the NLW, relative to their higher-paid counterparts who are unlikely to be affected by the NLW.

Robustness to parametrisation. We now turn to assessing the robustness of our main estimates to different specification choices in the implementation of our frequency distribution approach. Results are shown in Table 1, which reports estimates of the missing mass, the total employment effect, and the own-wage elasticity of employment, for a battery of different parametrisations. For reference, our headline estimates are reported in panel A of Table 1.

Panel B shows robustness to the choice of wage bin width, where we rerun the analysis using real and nominal wage bins – i.e. job types and k bins respectively – of £0.10 and £0.50 instead of £0.25. In panel C, we vary the level of geographical aggregation of TTWAs, changing the sample size threshold below which we group neighbouring travel to work areas to 100 and 400 observations instead of 200. In panel D, we show robustness to changes in the specification used for the estimation of wage premia. In one variant, we estimate wage premia from model 3 using only the bottom half of the wage distribution in each region.[6] In a second variation, we estimate wage premia using an AKM regression on movers across TTWAs – as described in model 4 – rather than a Mincerian regression (Abowd, Kramarz and Margolis, 1999). In a third variation, we estimate wage premia using an AKM regression on grouped TTWAs (see again model 4). In a fourth one, we drop industry and occupation controls from our Mincerian specification in model 3. Panel E shows estimates for different definitions of control regions. We compare regions in the bottom 8 deciles of regional wage premia to regions in the top 2 deciles, instead of comparing the bottom 9 deciles to the top decile. We also run the main specification excluding London from the set of control regions.[7]

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 except for the one using AKM-estimated wage premia on grouped TTWAs, and in all cases the estimated employment effect is small and not statistically significant. 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 premia estimated using only the bottom half of the wage distribution, or using an AKM regression, but these approaches also give less precise estimates than our baseline specification. This is likely due to smaller sample sizes in the estimation of wage premia.

[6] To define the subgroup of workers with wages in the bottom half of the wage distribution, we estimate model 3 using observations over the entire wage distribution and use the estimated coefficients – with the exception of the estimated region fixed effects δ_r – to predict individual wages. We then estimate our new δ_r from model 3 on the bottom half of the distribution of predicted wages.

[7] Note that 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.

Effects on 16-64 year olds. All the estimates we illustrated so far are based on the sample of individuals aged 25-64, that is the age group which the NLW legally applies to. Yet, there are reasons to believe that people under the age of 25 could be impacted too. 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 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 employment in that age group 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 and, in turn, their wages. The choice between education and work is also important for young people and may be impacted by minimum wage policy.

To jointly capture this wide range of factors, we apply our approach to examine effects on individuals aged 16 to 64. Appendix Figure A6 reports estimates of employment changes around the NLW as a share of the pre-treatment employment rate among 16-64s, using regional wage premia from model 3 and time trends from model 2 also estimated on the 16-64 population. Our estimates of the fall in employment below the NLW (5.34%, std. error 0.22%) and of the total employment effect up to 5 above the NLW (-0.06%, std. error 0.20%) are marginally smaller in absolute value than those found for 25-64 year olds. The own wage employment elasticity is -0.10 (std. error 0.34). These results suggest that the wages of under-25s were positively impacted by the introduction of the NLW and subsequent uplifts. 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, 2022). Our estimates also suggest that the overall employment effect of the NLW for the under-25s was either broadly neutral or positive.

5. Effects on the household income distribution: methodology and data

5.1 From hourly wages to household income

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, yet a difficult one to answer (Dube, 2019*b*). 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 minimum wage increases may cause changes in hours worked and in the likelihood of remaining employed. The impact of the minimum wage on the household income distribution is sensitive to whom individuals affected by the minimum live with, for two reasons. Firstly, households with more affected earners will be more impacted by changes to wages than households with only one. Secondly, the net incomes of all household members, including earnings after tax, benefits and investment income, will partly determine where affected earners rank in the household income distribution.

The UK has an individually assessed system of income and earnings taxation, and a system of cash transfers which is – for those of working age – overwhelmingly means-tested against family-level income and financial assets. This includes an extensive system of in-work, but means-tested transfers, mostly through tax credits which, in the UK, are really just cash transfers by another name. While the UK has by no means the most generous set of transfer entitlements in the developed world, the safety net is considerably more comprehensive than in the US, where the distributional impacts of minimum wages on net incomes have been studied previously (Dube, 2019*b*). This is important context for the analysis that will follow: many of those minimum wage workers who have low household incomes are in receipt of income-related transfers which get reduced when earnings increase, muting the gains to them of rises in the minimum wage.

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 Appendix Figure A7. 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 25-64 year old). Whilst the highest wage earners are very likely to have high household incomes, with more than 90% being in the top three household income deciles, the lowest decile of wage earners are spread across most of the household income distribution, with approximately 35% lying in the bottom third and over 50% 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 (see Appendix Figure A8).[8]

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 pass-through. Similarly, a decrease in earnings caused by any dis-employment effects will be partially mitigated by tax decreases and benefit increases. Appendix Figure A9 shows the median marginal tax rate for low-wage workers in each household income decile (defined among households with at least one 25-64 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.

[8]Working households are defined as households with at least one member with positive earnings from employment.

Previous studies have often used simulation approaches in order to estimate the impacts of minimum wage increases on household incomes (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 – i.e. post taxes and transfers – income. This is particularly important in institutional settings where income-related transfers, particularly those for working households, 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 existing tax-transfer system, or simulate the effect under an alternative one. That could be particularly useful in addressing external validity concerns, for example 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 have three main limitations (Dube, 2019*b*). First, they must make an assumption about the impact of the minimum wage on employment and hours worked. Typically the assumption is that there is no effect. A notable exception is Sabia and Burkhauser (2010), who refine this aspect, by importing an out-of-sample employment elasticity from previous literature. Whether an out-of-sample elasticity is appropriate in the setting where one is simulating a minimum wage increase for is, of course, an open question. Second, simulation approaches 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. Third, measurement error in hourly wages (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. We can address these limitations by taking advantage of one key, yet previously unexploited feature of the frequency-distribution approach. Specifically, we can use estimates of the impact of the minimum wage on the whole frequency distribution of wages to simulate non-mechanical effects on employment and wages. In combination with a careful strategy for addressing measurement error in hourly wages (described further in Section 5.3 below), this means we can address all the traditional limitations of simulation-based approaches while retaining their advantages.

The basic steps we take are the following: (i) we take detailed survey data on households' income from before the introduction of the NLW; (ii) we impute hourly wages in the data to account for measurement error; (iii) we change some workers' status to unemployed, reflecting the dis-employment effects of the NLW that we estimate with our frequency-distribution approach; (iv) using the same estimates, we change hourly wages to account for estimated wage effects of the NLW; and, finally, (v) we use a tax-benefit microsimulator to calculate net household incomes. We describe these steps in more detail below.

5.2 Data and sample construction

Our main data source is 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.[9] The national minimum wage was constant over that period, at the same level observed in the 2015 ASHE data used as baseline year in our frequency-distribution estimates. We use only households with at least one person aged 25-64, leaving us with 13,463 households.

5.3 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 frequency-distribution at precisely the minimum wage, just as one would expect if there is measurement error in the derived hourly wage (Skinner et al., 2002). Appendix Figure A10 shows the hourly wage distribution in ASHE (April 2015) – an employer-reported survey of earnings and hours – and Appendix Figure A11 reports the hourly wage distribution in the FRS (October 2014 to September 2015). A comparison of the two distributions highlights the presence of measurement error in the FRS.

This poses a challenge for simulations of the effects of the minimum wage by household income. Measurement error in wages will mean that the apportionment of the minimum wage 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 downward biased. In

[9] To be consistent with the frequency-distribution analysis, we use τ_t from model 2 to uprate earnings. For other financial variables we use official price indices, such as average rents.

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

We address this challenge using the 'donor method' approach developed in Skinner et al. (2002) and Harkness and Avram (2019*b*). The donor approach exploits the fact that, in the data, workers paid by the hour *directly* report their hourly wage *as well as* their earnings and hours, allowing us to compare their directly reported (or 'direct') and derived hourly wages.[10] 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, whether the family has young or old children, region of residence, and household income decile). Including household income decile in this model is an augmentation of previous implementations of the donor method, 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 coefficients to predict an hourly wage for all workers in the data. Third, for each 'donee' we find their ten nearest donor neighbours based on predicted hourly wages, randomly pick one of them, and impute their direct hourly wage to the donee.[11] Such a process introduces simulation error, so we repeat it 40 times and average across repetitions when calculating the final results. See Appendix Figure A11 for a comparison of derived and measurement-error corrected hourly wages in the FRS.

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 the Working Tax Credit. In those cases we keep hours unchanged and adjust earnings in order to resolve the inconsistency.

5.4 Simulating impacts on household income

Imputing employment effects. Our central estimates from the frequency-distribution analysis implied a small, though not statistically significant, dis-employment 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 set their earnings to zero. This assumes that a worker who would

[10] In the data, 14% of workers report an hourly wage.

[11] Only neighbours with a predicted wage within £0.50 of the donee can be selected as nearest neighbours.

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.

Imputing wage effects. Having simulated employment effects, we then simulate wage changes to account for the mechanical effect of the NLW – bringing those who earn below the NLW to the new minimum – and spillover effects – causing some wages to increase beyond the NLW. The first step is to calculate the post-policy cumulative distribution function of wages which is induced by the NLW. This distribution follows mathematically from the baseline FRS distribution of wages (after adjusting for measurement error) and the frequency-distribution estimates from the estimation of employment and wage effects.[12,13] We call this the 'target' distribution. We then modify the wages of workers in our sample to conform to this target distribution. To do this we make a 'no re-ranking' assumption, meaning we assume that the NLW does not cause a worker who would otherwise have had a wage strictly lower than another worker to end up with a wage strictly higher 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.

Calculating 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, which is the most detailed micro-simulation model of the UK tax-transfer system (Waters, 2017). 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 households 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.[14] In a relatively small number of cases,

[12]For this exercise we require absolute, rather than relative, estimates of the effect of the NLW on each wage bin. To get this, we multiply our estimates of α_k from model 8 by the ratio of the overall absolute employment effect to the overall relative employment effect. See Section 3.3 under 'Estimating the effect of the minimum wage on the frequency distribution of wages' for details on how to retrieve the absolute 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.

[13]We assume that *within* each £0.25 wage bin the distribution of wages stays the same, except for the bin that spans the range between the NLW and the NLW +£0.25. We base the distribution of the latter on the observed hourly wages in the bin around the NLW in the 2018-19 FRS.

[14]The exception is that we assume full take-up of child benefits, since child benefit take-up rates are over 95%.

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. Instead, we obtain a predicted probability of take-up based on parameter estimates from a logistic regression of take-up status on entitlement amount, work status, family type and age.[15] We then randomly assign take-up using that household-specific probability.[16]

Our estimates should be interpreted as partial equilibrium effects: given that we find that firms' wage bills have gone up (even after dis-employment 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.

A *household* consists of all people who occupy a housing unit regardless of relationship. We show how results differ if we take the income sharing unit to be narrower than the whole household. Specifically, we replicate our analysis using 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. We define this alternative income sharing unit as *family*. 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) show that the distributional impact of minimum wages can differ somewhat depending on what the income sharing unit is assumed to be.

6. Effects on the household income distribution: results

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 estimated from our baseline frequency-distribution specification for those aged 25-64 (shown in Figure 3). 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 that occurred between 2016 and 2019 inclusive.[17]

Figure 4 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.[18] The bars separately show the effect on net income (income

[15] We classify families by four categories: couples and singles, with and without children.

[16] 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.

[17] 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.

[18] 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.

after taxes and benefits) and net tax payments (taxes minus benefits). The two together sum to the effect on gross household earnings. The line plotted on the right hand axis shows the proportional impact on net income. The graph also shows the average level and proportional impacts across all households in the sample (most rightward estimates).

On average each NLW increase raises net household incomes across 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 of the total increase in pre-tax earnings, in the second and third income deciles, where many workers' households are receiving means-tested benefits which are quite rapidly withdrawn as earnings increase.[19] 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 fourth and fifth deciles. 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 4 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 dis-employment effects. After accounting for taxes and reductions in means-tested benefits, this leads to an increase in net household income of £5.75 per week. Average household income among minimum wage workers is £588 per week, meaning that the increase in income resulting from the minimum wage is just under 1% on average. But only 5% of working-age households contain a minimum wage worker. Even in the third and fourth deciles, 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, illustrated in Figure 4.

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 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. Among working households, 7% have a minimum wage worker. The highest concentration of minimum wage workers is in the lowest decile, where the number rises to 15%. As can be seen in Figure 5, 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* or *benefit units* are the unit of income sharing, and analyse effects at the

[19] 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 gross earnings gain clawed back via lower benefits when their earnings increase.

family level, as discussed in Section 5.4. 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 wage family accounts for 52% of the household's income. Appendix Figure A12 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 those seen at the household level in Figure 4. 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.

One advantage of our simulation approach is that we are able to decompose the total effect on incomes into different components. Figure 6 builds up to the overall effect seen in Figure 4 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). To do this, we essentially apply the procedure described in Section 5.4 to impute a post-policy frequency distribution of wages based on the pre-policy distribution and the parameter estimates from Section 4. However, we make a simple adjustment to those parameter estimates in order to purge them of the implied dis-employment effects and isolate only the marginal effect of the spillovers. Namely, we add the estimated number of displaced workers back in to the wage bin that starts at the level of the post-policy NLW. This number is found by summing over all employment changes up to £5 above the NLW. We then add in the dis-employment effect to recover the total estimated effect shown in Figure 4.

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 dis-employment 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. dis-employment effects could be larger or smaller, and it turns out that this matters quite significantly. Figure 6 shows what the distributional results would look like if the dis-employment effects were as large as the lower 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 dis-employment 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).

Comparison with evidence on impacts on household income in the United States. A comparison of the distributional impacts of minimum wages in the US and the UK reveals that minimum wage policies have a much more progressive impact in the US than in the UK. Our evidence for the UK indicates that the largest gains from the NLW introduction and uplifts go to the middle of the working-age household income distribution, both in cash and percentage terms (Figure 4). This is different from what has been found for the US. Using US Current Population Survey (CPS) data from 1984 to 2013, Dube (2019*b*) documents substantial and statistically significant positive effects of minimum wage increases on family income post taxes and transfers for quantiles between the seventh and twentieth, declining sharply to a statistical zero by the thirty-fifth quantile.[20] In Appendix C, we provide some context on minimum wage workers and their location in the household income distribution in the two countries, which can help explain these differences. We summarise the evidence here and refer the reader to Appendix C for more details.

A first reason for the observed differences is that, in the UK, minimum wage workers tend to be concentrated in the middle of the household income distribution, while in the US they are predominantly located towards the bottom of the distribution (Appendix Figure C1). This seems to be explained by the fact that – because of transfers and other sources of household incomes – there is less of a correspondence between household earnings and household income in the UK compared to the US (Appendix Figure C2).[21] A second reason for the discrepancy is that, in the UK (but not the US), minimum wage workers at the bottom of the income distribution are less likely to gain from minimum wage increases, because they work fewer hours (Appendix Figure C3) and face higher marginal tax rates (Appendix Figure C4) than those higher up the income distribution. Finally, individuals at the bottom of the income distribution in the UK get less of their income from work (Appendix Figure C5), and a much higher share of their total earnings comes from self-employment, which is not covered by the minimum wage (Appendix Figure C6).

7. 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 and employment. We have built on the 'frequency-distribution' 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 wage policy, by exploiting wage differences between different parts of the country. We estimate the impacts of the NLW on the number of jobs within

[20] See Figures 5 and 6 in Dube (2019*b*).

[21] Conversely, it does not seem to be the case that minimum wage workers are more likely to live with someone with higher hourly wages in the UK than in the US.

each wage bin, 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 income. 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 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 of 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, 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.

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 dis-employment.

References

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333.

Abowd, John M, Francis Kramarz, David N Margolis, and Thomas Philippon. 2000. "The tail of two countries: minimum wages and employment in France and the United States." IZA Discussion Papers.

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.

Bazen, Stephen, and Nicolas Skourias. 1997. "Is there a negative effect of minimum wages on youth employment in France?" *European Economic Review*, 41(3-5): 723–732.

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

Böckerman, Petri, and Roope Uusitalo. 2009. "Minimum wages and youth employment: Evidence from the Finnish retail trade sector." *British Journal of Industrial Relations*, 47(2): 388–405.

Bossler, Mario, and Hans-Dieter Gerner. 2020. "Employment effects of the new German minimum wage: Evidence from establishment-level microdata." *ILR Review*, 73(5): 1070–1094.

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.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality*." *The Quarterly Journal of Economics*, 128(3): 967–1015.

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.

Cengiz, Doruk, Arindrajit Dube, Attila S Lindner, and David Zentler-Munro. 2021. "Seeing Beyond the Trees: Using Machine Learning to Estimate the Impact of Minimum Wages on Labor Market Outcomes." National Bureau of Economic Research Working Paper 28399.

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.

Connolly, Sara, and Mary Gregory. 2002. "The national minimum wage and hours of work: implications for low paid women." *Oxford Bulletin of Economics and Statistics*, 64: 607–631.

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

Resolution Foundation Spotlight.
Dickens, Richard, Rebecca Riley, and David Wilkinson. 2014. "The UK minimum wage at 22 years of age: a regression discontinuity approach." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 177(1): 95–114.
Dickens, Richard, Rebecca Riley, and David Wilkinson. 2015. "A Re-examination of the Impact of the UK National Minimum Wage on Employment." *Economica*, 82(328): 841–864.
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.
Feenberg, Daniel, and Elisabeth Coutts. 1993. "An introduction to the TAXSIM model." *Journal of Policy Analysis and Management*, 12(1): 189–194.
Frings, Hanna. 2013. "The employment effect of industry-specific, collectively bargained minimum wages." *German Economic Review*, 14(3): 258–281.
Georgiadis, Andreas, and Alan Manning. 2020. "The impact of the uk minimum wage: Evidence from high-frequency firm-level data." *Available at SSRN 3560897*.
Giupponi, G., and S. Machin. 2022. "Minimum Wages and the Wage Policy of Firms." Working Paper.
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.*
Kabátek, Jan. 2015. "Happy birthday, you're fired! the effects of an age-dependent minimum wage on youth employment flows in the netherlands."
Kertesi, Gábor, and Janos Kollo. 2004. "Fighting 'low equilibria' by doubling the minimum wage? Hungary's experiment." William Davidson Institute at the University of Michigan William Davidson Institute Working Papers Series 2004-644.
König, Marion, and Joachim Möller. 2009. "Impacts of minimum wages: a microdata analysis for the German construction sector." *International Journal of Manpower*, 30(7): 716–741.
Kreiner, Claus Thustrup, Daniel Reck, and Peer Ebbesen Skov. 2020. "Do lower minimum wages for young workers raise their employment? Evidence from a Danish discontinuity." *Review of Economics and Statistics*, 102(2): 339–354.
Machin, Stephen, Alan Manning, and Lupin Rahman. 2003. "Where the Minimum Wage

Bites Hard: Introduction of Minimum Wages to a Low Wage Sector." *Journal of the European Economic Association*, 1(1): 154–180.

Machin, Stephen, and Joan Wilson. 2004. "Minimum wages in a low-wage labour market: care homes in the UK."

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.

Pereira, Sonia C. 2003. "The impact of minimum wages on youth employment in Portugal." *European Economic Review*, 47(2): 229–244.

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

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

Vadean, Florin, and Stephen Allan. 2021. "The effects of minimum wage policy on the long-term care sector in England." *British Journal of Industrial Relations*, 59(2): 307–334.

Vom Berge, Philipp, and Hanna Frings. 2020. "High-impact minimum wages and heterogeneous regions." *Empirical Economics*, 59(2): 701–729.

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

Yannelis, Constantine. 2014. "The minimum wage and employment dynamics: evidence from an age based reform in Greece."

Figures and tables

Figure 1. Real minimum wage rate and minimum wage bite, 1999-2019

A. Real minimum wage rate

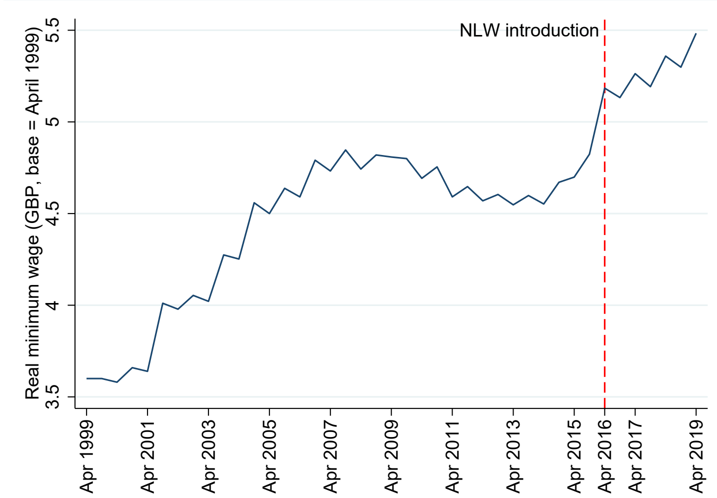

B. Minimum wage bite (minimum wage as a percent of median wage)

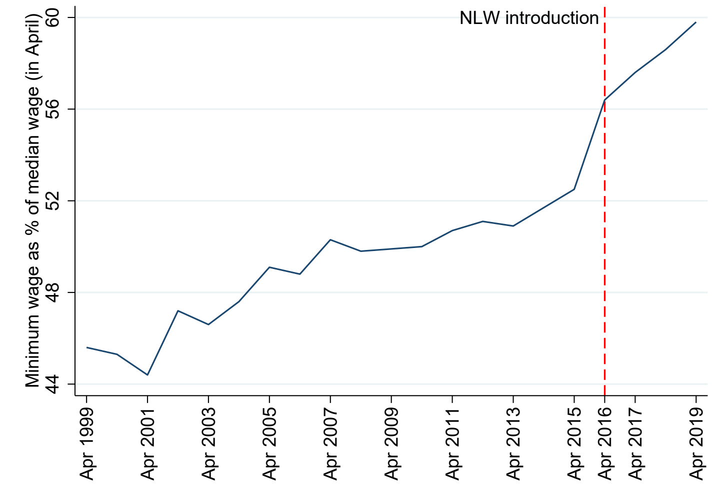

Notes: Panel A reports the CPI-adjusted level of the UK adult minimum wage applying to workers aged 25 and over from April 1999 to April 2019. Panel B reports the adult minimum wage bite, i.e. the adult minimum wage rate as a percent of the median wage. The dashed red line corresponds to the National Living Wage (NLW) introduction on April 1, 2016.

Figure 2. ESTIMATED WAGE PREMIA

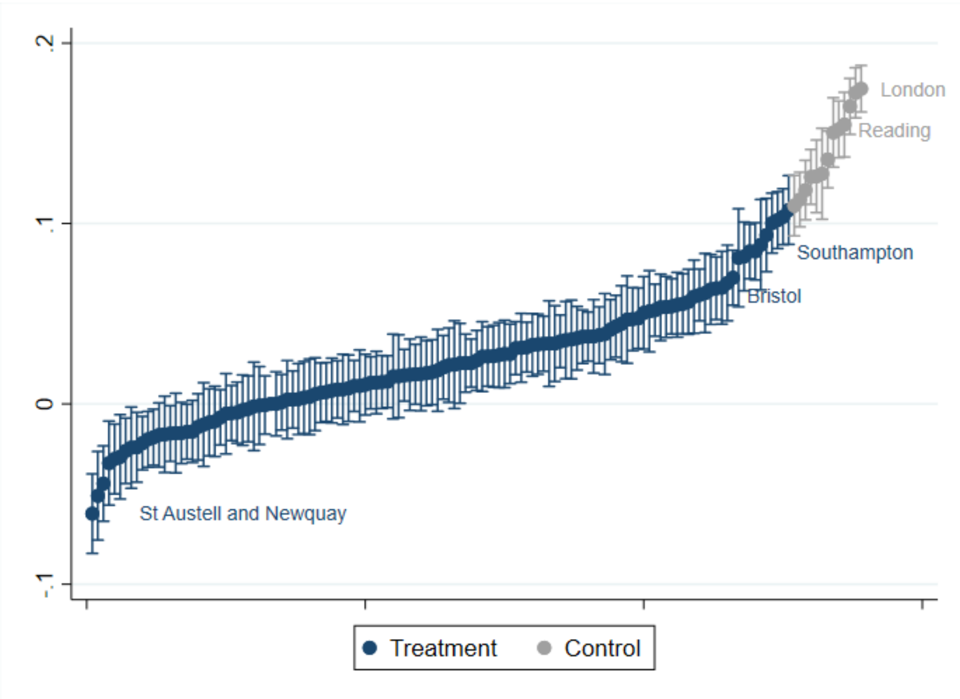

Notes: The graph reports estimates of δ_r from model 3 (solid circles) using data from 2012 to 2014. The capped vertical bars show the 95% confidence interval based on robust standard errors clustered at the TTWA level. Estimates in blue refer to regions in the bottom nine deciles of the distribution of wage premia (treatment group), grey ones to regions in the the top decile of the distribution (control group).

Figure 3. IMPACT OF THE MINIMUM WAGE ON THE WAGE DISTRIBUTION: BASELINE ESTIMATES ON WORKERS AGED 25-64

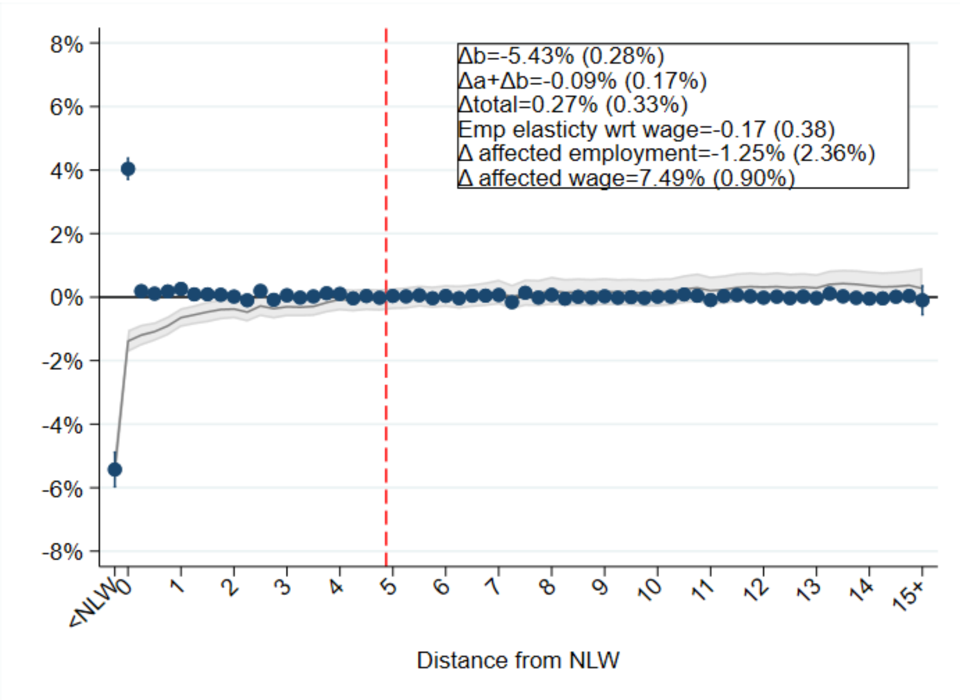

Notes: The graph reports estimates of the coefficients α_k from model 8 of the effect of the NLW introduction and subsequent uplifts on the frequency distribution of hourly wages. The sample includes individuals aged 25-64. Each dot (corresponding to α_k) represents our estimate of employment changes – averaged over the four minimum wage increases from 2015 to 2019 – in each £0.25 wage bin relative to the level of the new NLW in each of those. Employment changes are normalised by baseline employment in the TTWA, so that the sum of the effects across all wage bins can be interpreted as the total percentage change in employment arising from the change in the minimum wage. Estimated effects in wage bins below the new NLW, as well as in wage bins more than £15 above the NLW, are aggregated in one single point estimate. The grey line shows the running total of employment changes up to that point in the distribution. The vertical bars underlying the dots and the shaded area around the grey line show the bootstrapped 95% confidence intervals associated to the relevant estimate. The graph also reports estimates of the terms Δb (the percent change in employment below the new NLW), $\Delta e = \Delta a + \Delta b$ (the percent change in employment up to £5 above the new NLW) and $\Delta \text{total} = \Delta a' + \Delta b$ (the percent change in employment over the entire wage distribution), with bootstrapped standard errors in parenthesis. Estimates of the own-wage employment elasticity and its sub-components – the percentage change in affected employment and affected wages – are also reported. See Section 3.3 for further details on these statistics.

Figure 4. IMPACT OF THE MINIMUM WAGE ON HOUSEHOLD INCOMES: DECOMPOSITION BY INCOME SOURCE

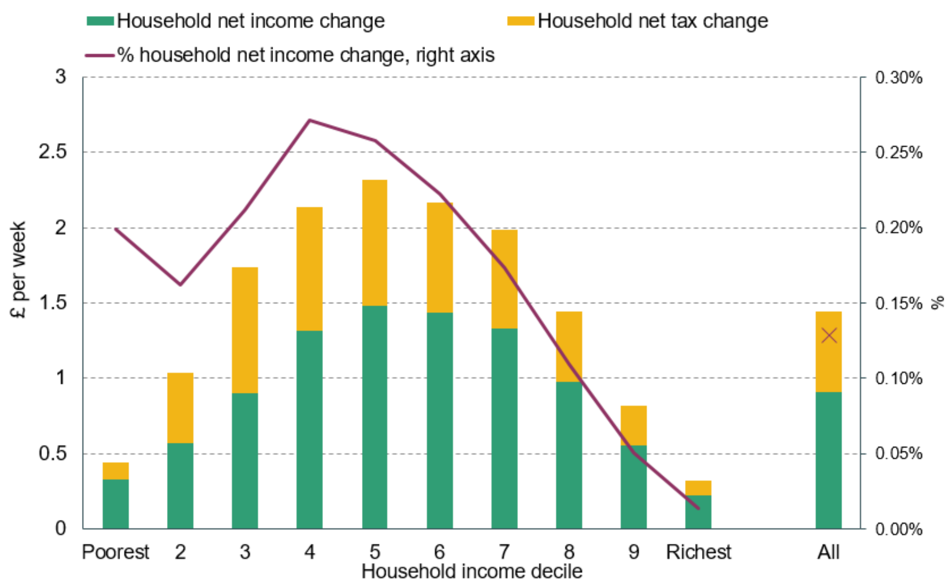

Notes: The graph reports the simulated distributional effect of the average increase of the NLW from 2015 to 2019 on household income. The vertical bars separately show the cash effect on net income (income after taxes and benefits in green) and net tax payments (taxes minus benefits in yellow). The two together sum to the cash effect on gross household earnings (left axis). The line plotted on the right axis shows the proportional impact on net income. The graph also shows the average level and proportional impacts across all households in the sample (most rightward bar and cross respectively). The sample includes households with at least one person aged 25-64. Households are ranked based on pre-NLW income in this sample. Income is equivalised and net of taxes and benefits.

Figure 5. Impact of the minimum wage on household incomes of working households: Decomposition by income source

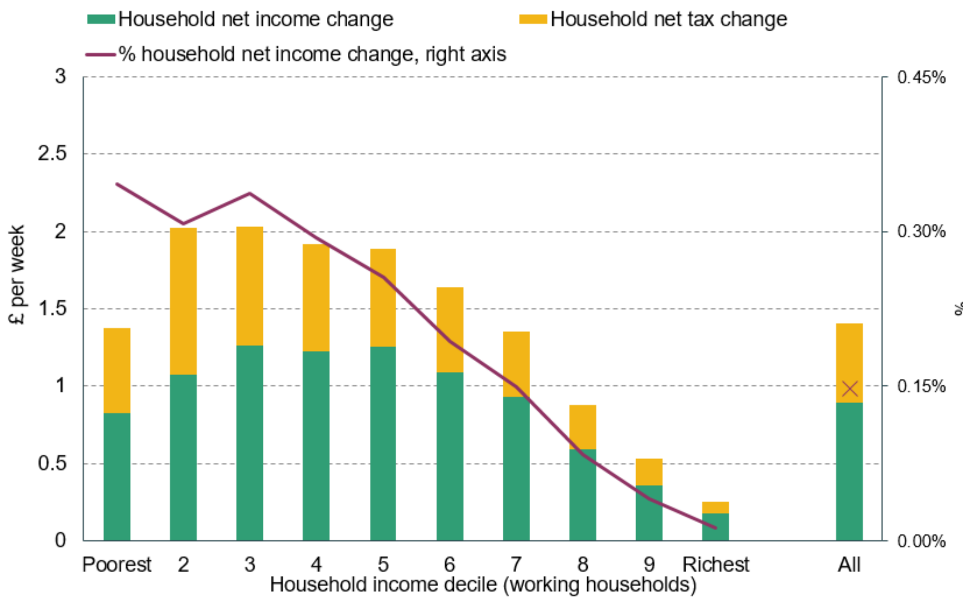

Notes: The graph reports the simulated distributional effect of the average increase of the NLW from 2015 to 2019 on household income. The vertical bars separately show the cash effect on net income (income after taxes and benefits in green) and net tax payments (taxes minus benefits in yellow). The two together sum to the cash effect on gross household earnings (left axis). The line plotted on the right axis shows the proportional impact on net income. The graph also shows the average level and proportional impacts across all households in the sample (most rightward bar and cross respectively). The sample includes 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 based on pre-NLW income in this sample. Income is equivalised and net of taxes and benefits.

Figure 6. IMPACT OF THE MINIMUM WAGE ON HOUSEHOLD INCOMES: DECOMPOSITION BY SOURCE OF RESPONSE

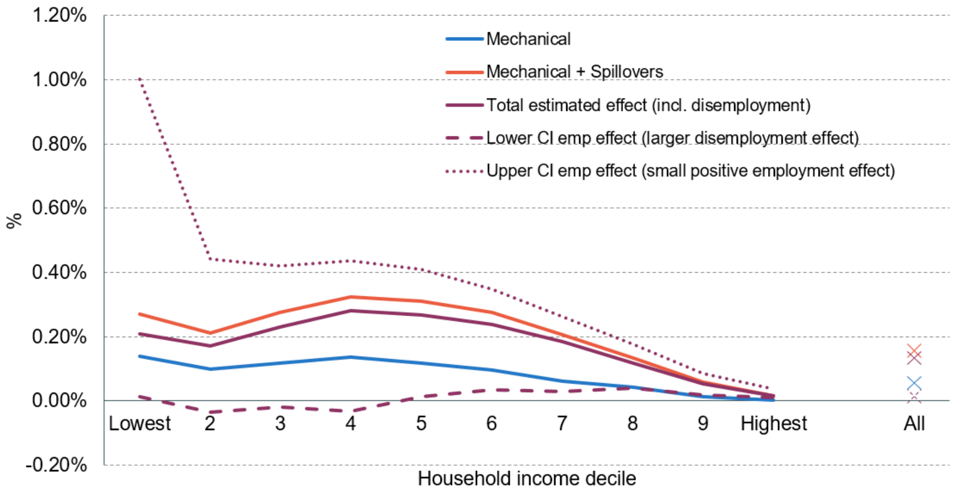

Notes: The graph reports the simulated distributional effect of the average increase of the NLW from 2015 to 2019 on household income by source of response. The **mechanical** change is the result of increasing wages of those previously earning below the NLW to the NLW. The **mechanical + spillovers** effect accounts for changes in the wage distribution as a result of the NLW, stripping out dis-employment effects. The **total** change incorporates the full set of effects as estimated in 3, that is the mechanical effect, spillovers and dis-employment effects. The **larger dis-employment effect** change incorporates dis-employment effects as negative as the lower bound of the 95 % confidence interval. The **small positive employment effect** change incorporates employment effects as positive as the upper bound of the 95 % confidence interval. The graph also shows proportional impacts across all households in the sample (most rightward crossed). The sample includes households with at least one person aged 25-64. Households are ranked based on pre-NLW income in this sample. Income is equivalised and net of taxes and benefits.

Table 1. IMPACT OF THE MINIMUM WAGE ON THE WAGE DISTRIBUTION: ROBUSTNESS CHECKS

	Δb		$\Delta a + \Delta b$		Elasticity	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Main specification						
Main specification	-5.43%	0.28%	-0.09%	0.17%	-0.17	0.38
Panel B. Bin width						
£0.10	-5.47%	0.28%	-0.10%	0.16%	-0.22	0.41
£0.50	-5.51%	0.24%	-0.18%	0.20%	-0.26	0.34
Panel C. Threshold for grouping TTWAs						
100 observations	-5.48%	0.21%	-0.07%	0.16%	-0.13	0.35
400 observations	-5.32%	0.29%	-0.05%	0.19%	-0.10	0.40
Panel D. Estimation of wage premia						
Bottom half of distribution	-4.91%	0.36%	-0.16%	0.16%	-0.39	0.60
AKM	-5.15%	0.42%	-0.20%	0.26%	-0.48	4.08
Grouped-AKM	-6.87%	1.09%	-0.23%	0.46%	-0.36	0.75
No industry/occupation controls	-5.61%	0.28%	-0.20%	0.18%	-0.37	0.37
Panel E. Control regions						
Top 2 deciles	-5.60%	0.26%	-0.15%	0.14%	-0.28	0.31
Excluding London	-5.36%	0.30%	-0.10%	0.23%	-0.19	0.50

Notes: The table reports estimates of Δb (the percent change in employment below the new NLW), $\Delta e = \Delta a + \Delta b$ (the percent change in employment up to £5 above the new NLW) and the own-wage employment elasticity for a set of different parametrisations of model 8. Estimates are averaged over the four minimum wage increases from 2015 to 2019. Columns (1), (3) and (5) report our central estimates; columns (2), (4) and (6) the bootstrapped standard errors. See Section 3.3 for further details on these statistics. Panel A reports baseline estimates from Figure 3. Panel B shows robustness to the choice of wage bin width, where we rerun the analysis using real and nominal wage bins – i.e. job types and k bins respectively – of £0.10 and £0.50 instead of £0.25. Panel C varies the level of geographical aggregation of TTWAs, changing the sample size threshold below which we group neighbouring travel to work areas to 100 and 400 observations instead of 200. Panel D shows robustness to changes in the specification used for the estimation of wage premia. In one variant, we estimate wage premia from model 3 using only the bottom half of the wage distribution in each region. In a second, we estimate wage premia using an AKM regression on movers across TTWAs – illustrated in model 4 – rather than a Mincerian regression (Abowd, Kramarz and Margolis, 1999). In a third variation, we estimate wage premia using an AKM regression on grouped TTWAs (see again model 4). In a fourth one, we drop industry and occupation controls from our Mincerian specification in model 3. Panel E shows estimates for different definitions of control regions. We compare regions in the bottom 8 deciles of regional wage premia to regions in the top 2 deciles, instead of comparing the bottom 9 deciles to the top decile. We also run the main specification excluding London from the set of control regions.

Appendix A. Additional figures and tables

Figure A1. Minimum wage as percent of full-time median wage, 2019

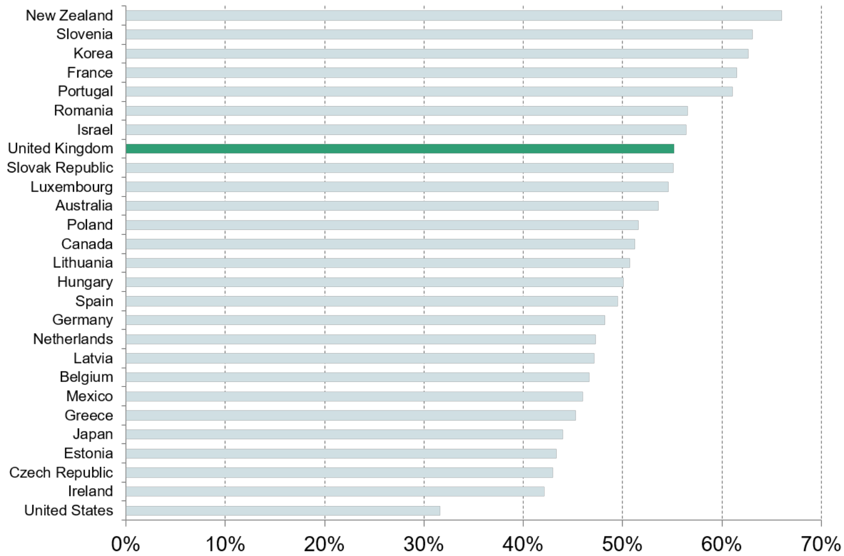

Notes: The graph reports an international comparison of the ratio of minimum wages to median earnings of full-time employees across a set of OECD countries for 2019. The data is sourced from OECD.

Figure A2. CORRELATION BETWEEN MINCER AND AKM WAGE PREMIA

A. Correlation between Mincer and AKM wage premia

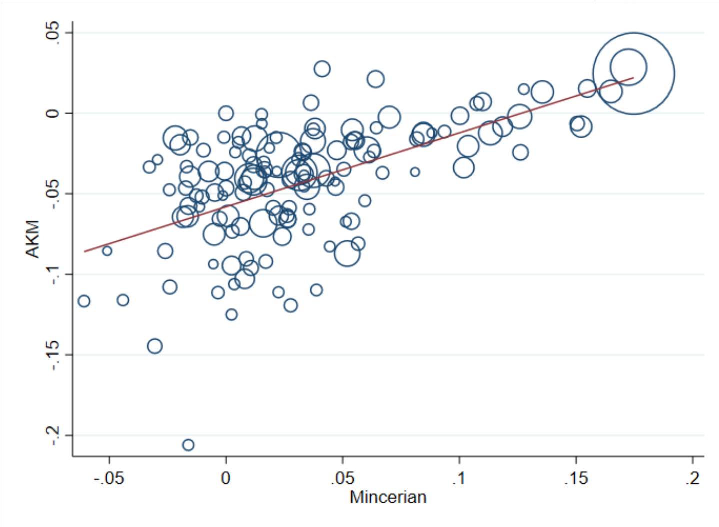

B. Correlation between Mincer and grouped-AKM wage premia

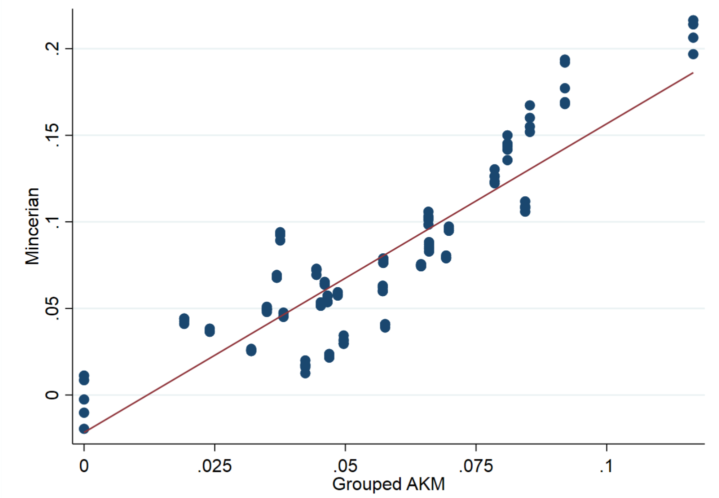

Notes: Panel A reports the correlation between Mincerian regional wage premia δ_r from model 3 (horizontal axis) and AKM estimates of γ_r from model 4 (vertical axis). Dots are weighted by the number of workers in each TTWA in 2012-2015. Panel B reports the correlation between Mincerian regional wage premia δ_r from model 3 (vertical axis) and grouped-AKM estimates of $\gamma_{r(p)}$ from model 4 (horizontal axis).

Figure A3. SORTING BIAS: SYMMETRY OF WAGE GAINS AND LOSSES

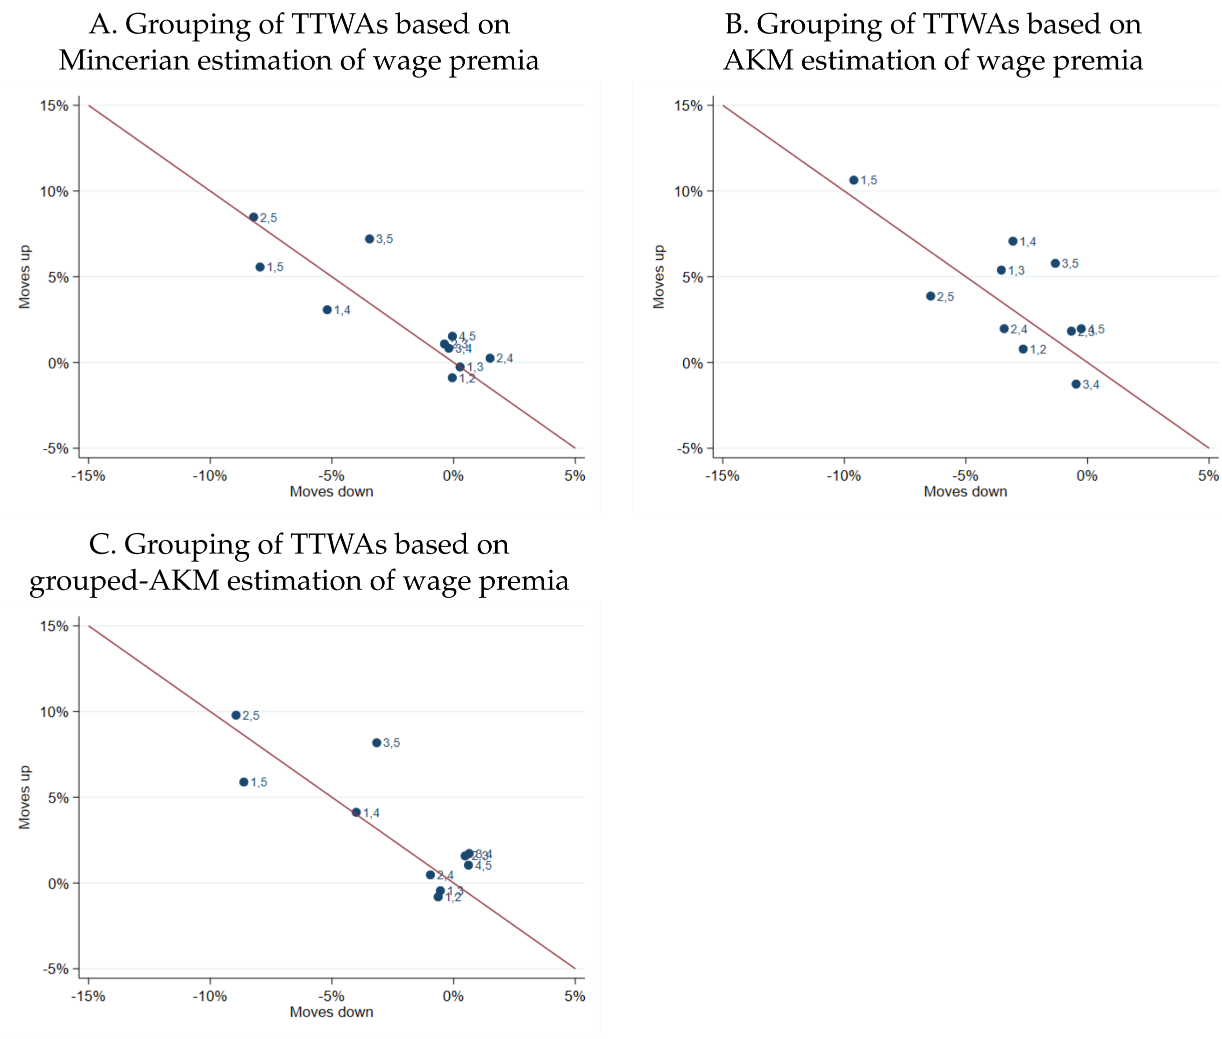

Notes: Following Card, Heining and Kline (2013), the graphs report tests of the symmetry of wage gains and losses of individuals who move across regions with different wage premia. TTWAs are ranked into five groups corresponding to the quintiles of the distribution of regional wage premia. In panel A, TTWAs are ranked based on the estimated δ_r from Mincerian regression 3; in panel B, the ranking is based on estimates of γ_r from AKM model 4 and, in panel C, on estimates of $\gamma_{r(p)}$ from a version of AKM model 4 estimated on groups of TTWAs. Each dot (o, d) represents the average wage gain associated with moving from origin quintile o to destination quintile d (vertical axis), and the wage loss associated with moving from origin quintile d to destination quintile o (horizontal axis). Wage changes are residualised on basic demographic characteristics. The red line represents the off-diagonal.

Figure A4. IMPACT OF THE MINIMUM WAGE ON THE WAGE DISTRIBUTION: NLW INTRODUCTION AND UPLIFTS, 2015 TO 2019

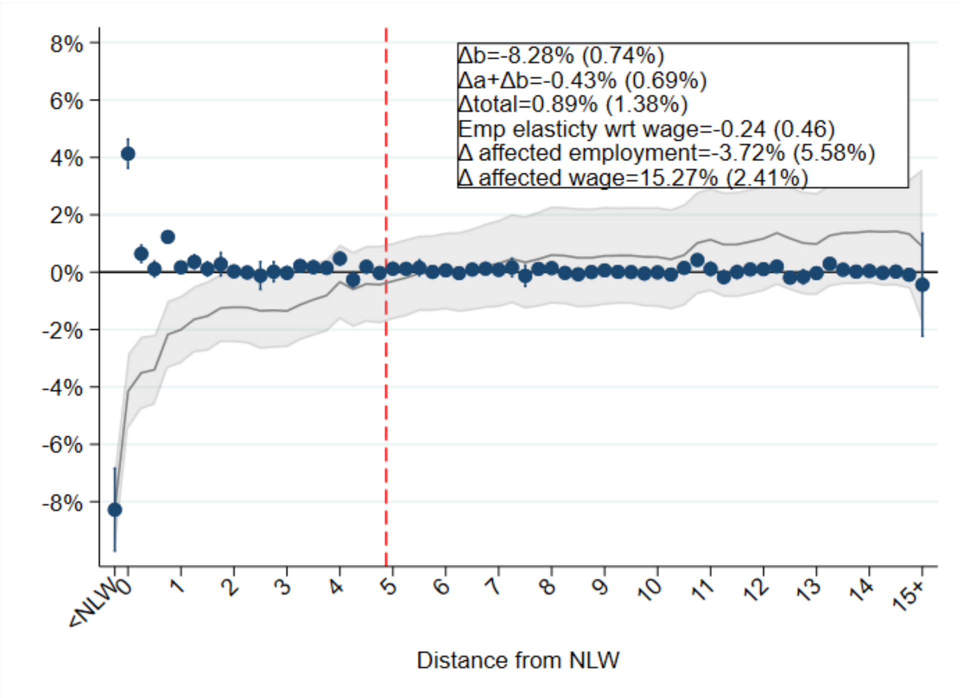

Notes: The graph reports estimates of the coefficients α_k from model 8 of the effect of the NLW introduction and subsequent uplifts on the frequency distribution of hourly wages using data from 2015 and 2019 only. Each dot (corresponding to α_k) represents our estimate of employment changes in each £0.25 wage bin relative to the level of the new NLW in each of those. Employment changes are normalised by baseline employment in the TTWA, so that the sum of the effects across all wage bins can be interpreted as the total percentage change in employment arising from the change in the minimum wage. Estimated effects in wage bins below the new NLW, as well as in wage bins more than £15 above the NLW, are aggregated in one single point estimate. The grey line shows the running total of employment changes up to that point in the distribution. The vertical bars underlying the dots and the shaded area around the grey line show the bootstrapped 95% confidence intervals associated to the relevant estimate. The graph also reports estimates of the terms Δb (the percent change in employment below the new NLW), $\Delta e = \Delta a + \Delta b$ (the percent change in employment up to £5 above the new NLW) and $\Delta \text{total} = \Delta a' + \Delta b$ (the percent change in employment over the entire wage distribution), with bootstrapped standard errors in parenthesis. Estimates of the own-wage employment elasticity and its sub-components – the percentage change in affected employment and affected wages – are also reported. See Section 3.3 for further details on these statistics.

Figure A5. IMPACT OF THE MINIMUM WAGE ON THE WAGE DISTRIBUTION: PLACEBO CHECK

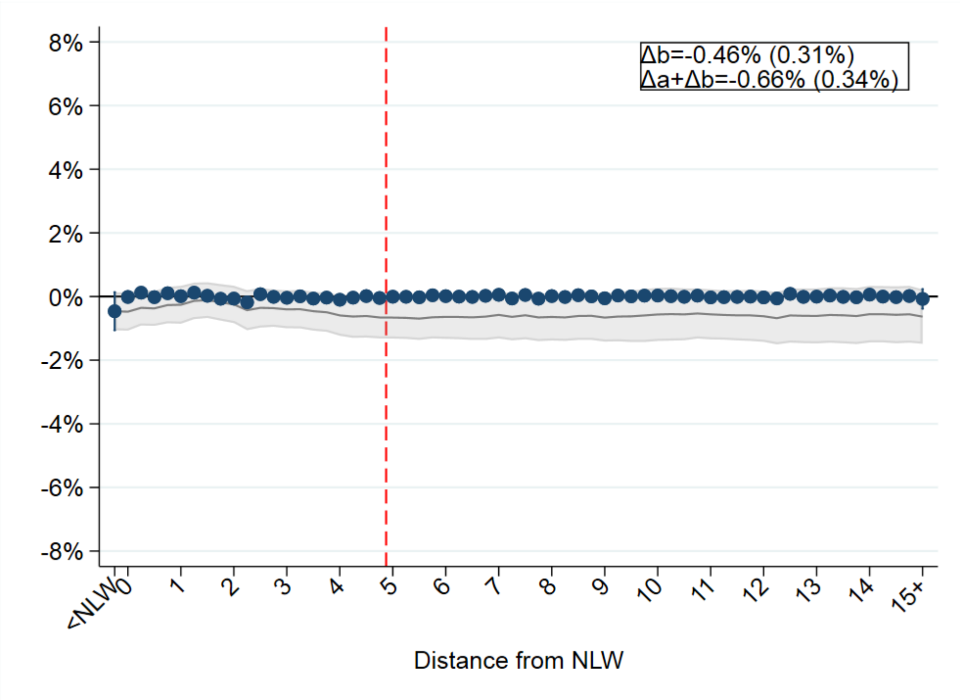

Notes: The graph reports estimates of the coefficients α_k from model 8 using data from 2011 to 2015 – before the introduction of the NLW – and the level of the NLW as a placebo minimum wage rate. We apply the 2016 NLW to 2012, the 2017 NLW to 2013, the 2018 NLW to 2014 and the 2019 NLW to 2015. Each dot (corresponding to α_k) represents our estimate of employment changes in each £0.25 wage bin relative to the level of the new placebo NLW in each of those. Employment changes are normalised by baseline employment in the TTWA, so that the sum of the effects across all wage bins can be interpreted as the total percentage change in employment arising from the change in the minimum wage. Estimated effects in wage bins below the placebo NLW, as well as in wage bins more than £15 above the placebo NLW, are aggregated in one single point estimate. The grey line shows the running total of employment changes up to that point in the distribution. The vertical bars underlying the dots and the shaded area around the grey line show the bootstrapped 95% confidence intervals associated to the relevant estimate. The graph also reports estimates of the terms Δb (the percent change in employment below the new NLW) and $\Delta e = \Delta a + \Delta b$ (the percent change in employment up to £5 above the new NLW), with bootstrapped standard errors in parenthesis. See Section 3.3 for further details on these statistics.

Figure A6. IMPACT OF THE MINIMUM WAGE ON THE WAGE DISTRIBUTION: BASELINE ESTIMATES ON WORKERS AGED 16-64

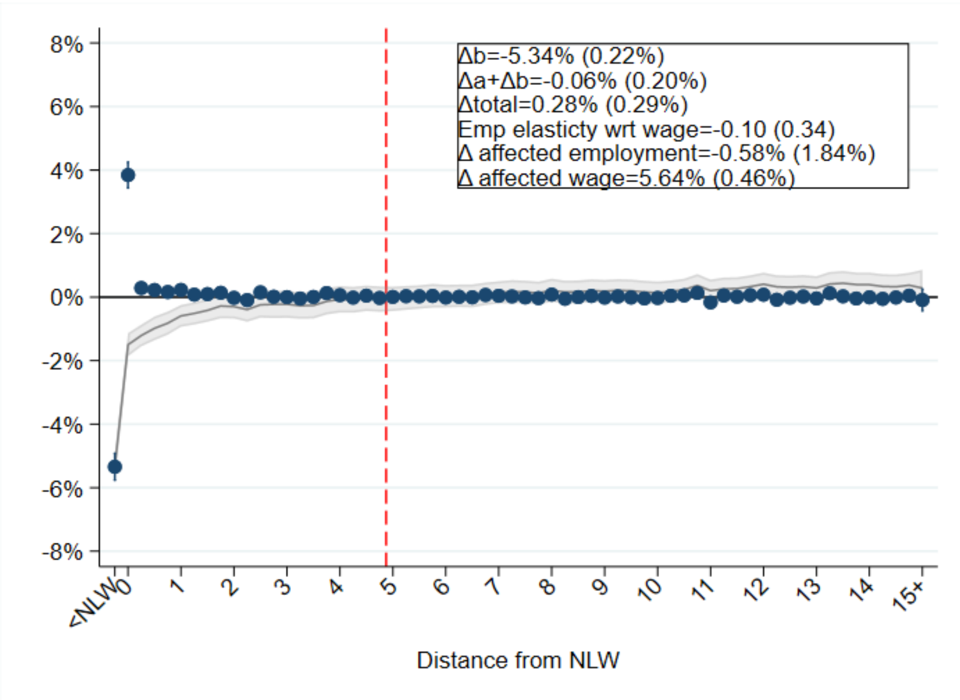

Notes: The graph reports estimates of the coefficients α_k from model 8 of the effect of the NLW introduction and subsequent uplifts on the frequency distribution of hourly wages. The sample includes individuals aged 16-64. Each dot (corresponding to α_k) represents our estimate of employment changes – averaged over the four minimum wage increases from 2015 to 2019 – in each £0.25 wage bin relative to the level of the new NLW in each of those. Employment changes are normalised by baseline employment in the TTWA, so that the sum of the effects across all wage bins can be interpreted as the total percentage change in employment arising from the change in the minimum wage. Estimated effects in wage bins below the new NLW, as well as in wage bins more than £15 above the NLW, are aggregated in one single point estimate. The grey line shows the running total of employment changes up to that point in the distribution. The vertical bars underlying the dots and the shaded area around the grey line show the bootstrapped 95% confidence intervals associated to the relevant estimate. The graph also reports estimates of the terms Δb (the percent change in employment below the new NLW), $\Delta e = \Delta a + \Delta b$ (the percent change in employment up to £5 above the new NLW) and $\Delta \text{total} = \Delta a' + \Delta b$ (the percent change in employment over the entire wage distribution), with bootstrapped standard errors in parenthesis. Estimates of the own-wage employment elasticity and its sub-components – the percentage change in affected employment and affected wages – are also reported. See Section 3.3 for further details on these statistics.

Figure A7. PROPORTION OF INDIVIDUALS IN EACH HOUSEHOLD INCOME DECILE, BY WAGE DECILE

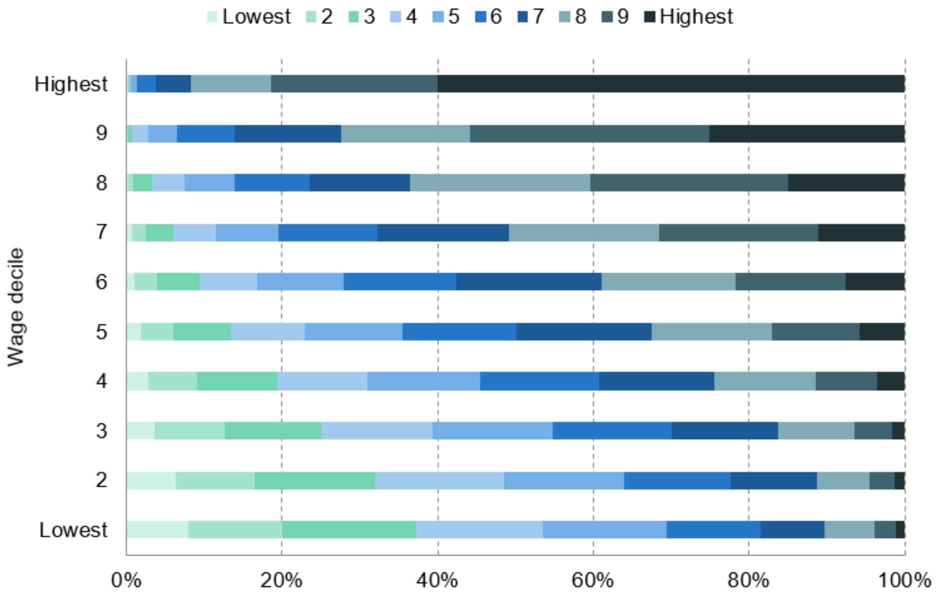

Notes: The graph shows the proportion of individuals in each household decile (horizontal axis) by individual wage decile (vertical axis). The sample includes employees aged 25-64. Household income deciles are defined among all households with at least one 25-64 year old. Income is equivalised and net of taxes and benefits. The data are sourced from the Family Resources Survey for October 2014 to September 2015. Wages are adjusted using the donor method illustrated in Section 5.3.

Figure A8. Proportion of individuals in each household income decile (defined among working households only), by wage decile

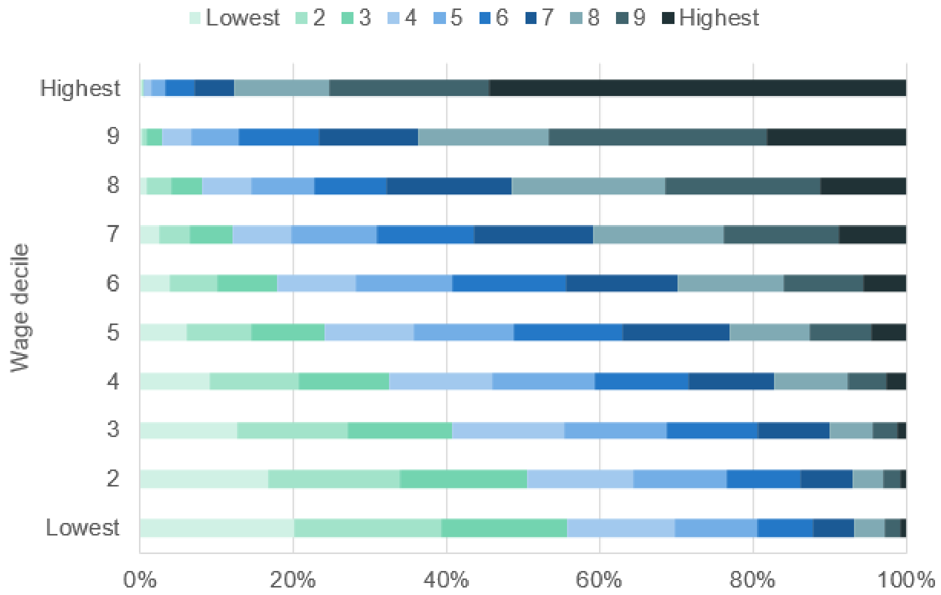

Notes: The graph shows the proportion of individuals in each household decile (horizontal axis) by individual wage decile (vertical axis). The sample includes employees aged 25-64. Household income deciles are defined among all households with at least one 25-64 year old receiving employment income. Income is equivalised and net of taxes and benefits. The data are sourced from the Family Resources Survey for October 2014 to September 2015. Wages are adjusted using the donor method illustrated in Section 5.3.

Figure A9. MARGINAL TAX RATES FOR MINIMUM WAGE EARNERS BY HOUSEHOLD INCOME DECILE

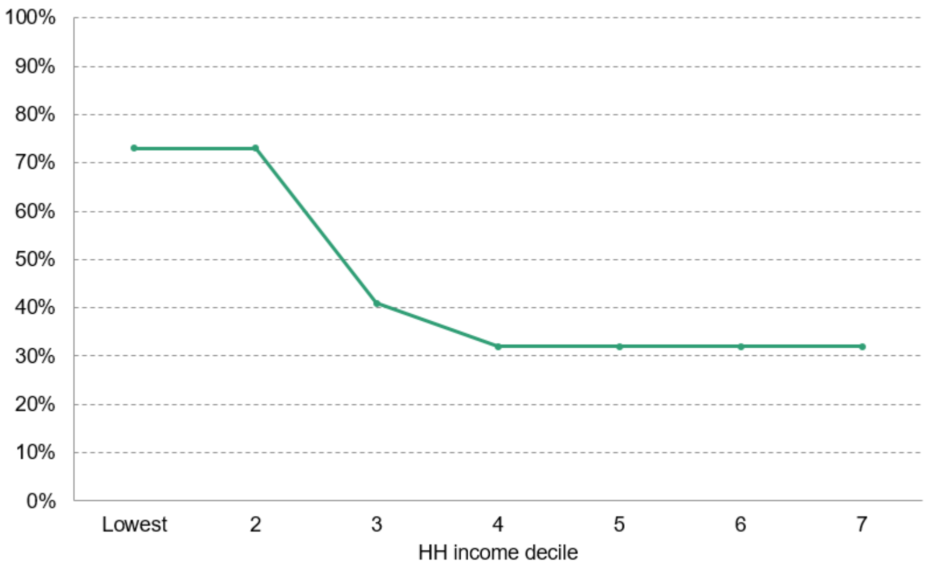

Notes: The graph shows the median marginal tax rate for minimum wage workers by household income decile. The sample includes employees aged 25-64 earning the national minimum wage. Household income deciles are defined among all households with at least one 25-64 year old. Income is equivalised and net of taxes and benefits. Deciles above the 7th are omitted due to small sample sizes.

Figure A10. DISTRIBUTION OF WAGES IN THE ANNUAL SURVEY OF HOURS AND EARNINGS (ASHE, APRIL 2015)

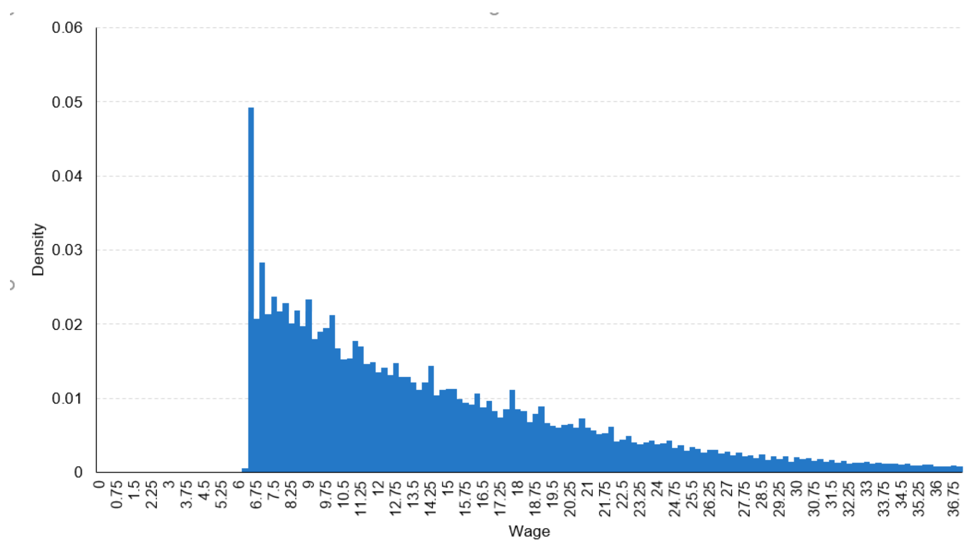

Notes: The graph plots the empirical distribution of hourly wages in the the Annual Survey of Hours and Earnings (ASHE, April 2015). Hourly wages are measured in £ per hour and binned in wage bins £0.25 wide. Hourly wages can be either directly reported or derived using weekly or monthly earnings and weekly hours of work, and dividing one by the other.

Figure A11. DISTRIBUTION OF WAGES IN THE FAMILY RESOURCES SURVEY (FRS, OCTOBER 2014 - SEPTEMBER 2015), BEFORE AND AFTER CORRECTING FOR MEASUREMENT ERROR

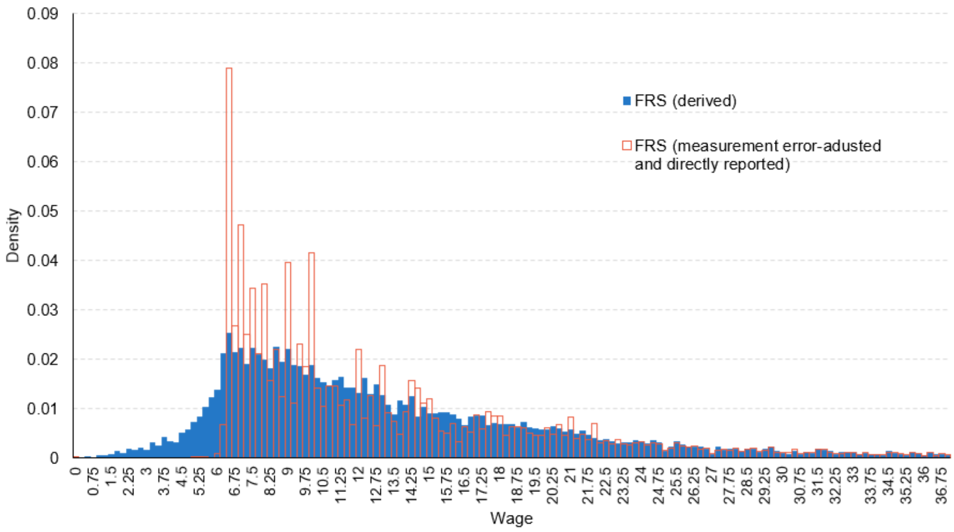

Notes: The graph plots the empirical distribution of hourly wages in the Family Resources Survey (FRS, from October 2014 to September 2015) before and after correcting for measurement error. Hourly wages are measured in £ per hour and binned in wage bins £0.25 wide. Hourly wages before correcting for measurement error are derived using weekly or monthly earnings and weekly hours of work, and dividing one by the other. Hourly wages after correcting for measurement error are adjusted using the donor method illustrated in Section 5.3, or include directly reported wages where these are available.

Figure A12. IMPACT OF THE MINIMUM WAGE ON FAMILY INCOMES: DECOMPOSITION BY INCOME SOURCE

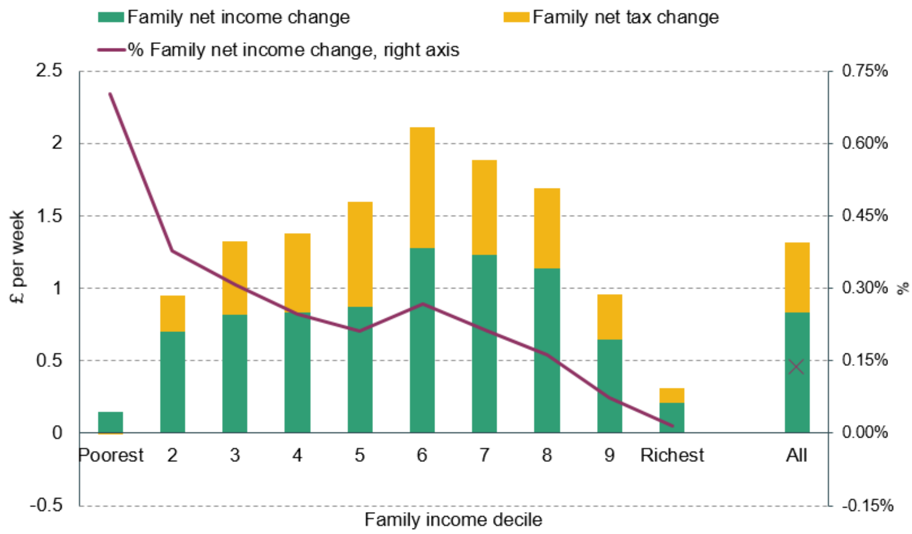

Notes: The graph reports the simulated distributional effect of the average increase of the NLW from 2015 to 2019 on family income. The vertical bars separately show the cash effect on net income (income after taxes and benefits in green) and net tax payments (taxes minus benefits in yellow). The two together sum to the cash effect on gross family earnings (left axis). The line plotted on the right axis shows the proportional impact on net income. The graph also shows the average level and proportional impacts across all family in the sample (most rightward bar and cross respectively). The sample includes family with at least one person aged 25-64. Families are ranked based on pre-NLW income in this sample. Income is equivalised and net of taxes and benefits.

Table A1. ESTIMATES OF WAGE TRENDS FROM MODEL 2

	Log hourly wage ($\ln w_{irt}$)			
	2015-2016	2016-2017	2017-2018	2018-2019
	(1)	(2)	(3)	(4)
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: The table reports the coefficient estimates of τ_t and β from model 2. Robust standard errors clustered at the TTWA level are reported in parenthesis.

Table A2. MOBILITY OF LOW-WAGE WORKERS ACROSS REGIONS

	Probability of moving region	
	Low-wage to high-wage (T to C)	High-wage to low-wage (C to T)
	(1)	(2)
β	-0.000	-0.003
	(0.000)	(0.002)
μ	-0.002***	-0.002*
	(0.000)	(0.001)
Observations	445825	135998
Share of low-paid movers 2012-2015	0.002	0.014
Share of not-low-paid movers 2012-2015	0.004	0.016

Notes: The table reports the coefficient estimates of β and μ from model 13. Column (1) reports estimates of mobility from treatment to control regions (bottom 9 deciles to top decile of wage premia), while column (2) for mobility from control to treatment regions. Robust standard errors clustered at the TTWA level are reported in parenthesis. The table also reports the share of low-paid and not-low-paid workers moving from treatment to control regions (and vice versa) in 2012-2015. Low-paid workers are defined as those with hourly wages in the bottom 10% of the national distribution.

Appendix B. Classification of European minimum wage studies by empirical approach

Table B1. EUROPEAN MINIMUM WAGE STUDIES BY EMPIRICAL APPROACH

Method	Paper	Country
Regional variation approach	Ahlfeldt, Roth and Seidel (2018)	Germany
	Caliendo et al. (2018)	Germany
	Dustmann et al. (2021)	Germany
	Schmitz (2019)	Germany
	Vom Berge and Frings (2020)	Germany
	Kertesi and Kollo (2004)	Hungary
	Pereira (2003)	Portugal
	Portugal and Cardoso (2006)	Portugal
	Dolton, Bondibene and Stops (2015)	UK
	Dolton, Bondibene and Wadsworth (2012)	UK
	Dube (2019*b*)	UK
	Georgiadis and Manning (2020)	UK
	Stewart (2002)	UK
Individual-level bite	Abowd et al. (2000)	France
	Böckerman and Uusitalo (2009)	Finland
	König and Möller (2009)	Germany
	Yannelis (2014)	Greece
	Connolly and Gregory (2002)	UK
	Dickens, Riley and Wilkinson (2015)	UK
	Stewart (2004*a*) and Stewart (2004*b*)	UK
Firm- or sector-level bite	Bazen and Skourias (1997)	France
	Bossler and Gerner (2020)	Germany
	Frings (2013)	Germany
	Harasztosi and Lindner (2019)	Hungary
	Machin, Manning and Rahman (2003)	UK
	Machin and Wilson (2004)	UK
	Vadean and Allan (2021)	UK
Age discontinuity	Kreiner, Reck and Skov (2020)	Denmark
	Kabátek (2015)	Netherlands
	Dickens, Riley and Wilkinson (2014)	UK

Notes: The table provides a classification of studies of the employment effects of the minimum wage in European countries, by method used for identification. The regional variation approach exploits differences in the minimum wage bite across geographical units. It can be combined with differences in bite across demographic or skills groups. Studies classified under the 'individual-level bite' methodology are based on difference-in-differences strategies, comparing individuals directly affected by the minimum, with some group of workers who are higher up the pay distribution thus not directly affected. Strategies based on firm- or sector-level bite exploit variation in the minimum wage bite across firms or sectors. Finally, studies based on age discontinuities usually adopt regression discontinuity or difference-in-differences strategies comparing workers in age groups subject to different minimum wage levels.

Appendix C. Minimum wages and household income the UK and the US

This section provides some context on minimum wage workers and their location in the household income distribution in the UK and the US. The evidence can help explain differences in the distributional impacts of minimum wages in the two countries, as documented in this paper and in Dube (2019*b*). All analysis is at the individual level, restricted to individuals aged 25-64. For the UK, we use data from the Family Resources Survey for 2015; for the US, from the March Current Population Survey (CPS) and from the Annual Social and Economic Supplement of the Current Population Survey, also for the year 2015. Household income and earnings have been equivalised taking the number and age of household members into account using OECD equivalence scales. Household earnings and income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). An individual is defined as a minimum wage worker if they have a wage lower than 1.25 times the state-specific minimum wage. All results have been weighted using person-level survey weights. Marginal tax rates have been simulated using TAXBEN for the UK tax system (Waters, 2017), and TAXSIM for the US federal and state tax systems (Feenberg and Coutts, 1993).[22]

Figure C1. SHARE OF INDIVIDUALS AGED 25-64 WITH A MINIMUM WAGE WORKER IN THEIR HOUSEHOLD, BY HOUSEHOLD INCOME DECILE

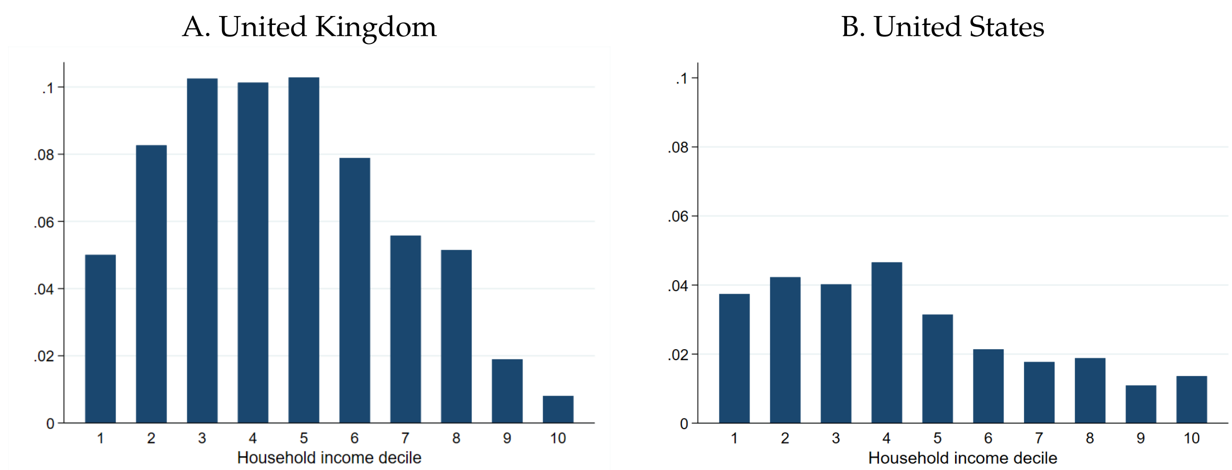

Notes: The graphs report the share of individuals aged 25-64 with a minimum wage worker in their household, by household income decile. Household income has been equivalised taking the number and age of household members into account using OECD equivalence scales. Household income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). An individual is defined as a minimum wage worker if they have a wage lower than 1.25 times the state-specific minimum wage. Panel A refers to the UK and panel B to the US. See text in Appendix C for more details on data sources, sample restrictions and variable definitions.

[22] For TAXSIM, see `https://users.nber.org/~taxsim/`.

Figure C2. CORRESPONDENCE BETWEEN HOUSEHOLD EARNINGS DECILE AND HOUSEHOLD INCOME DECILE ONTO HH INCOME DECILE AMONG INDIVIDUALS AGED 25-64

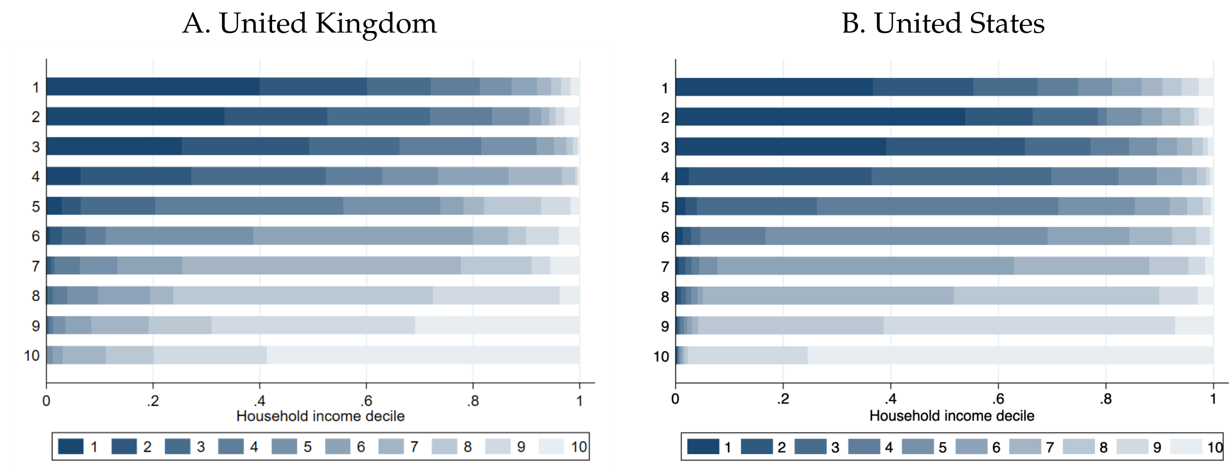

Notes: The graphs show the proportion of individuals aged 25-64 in each household income decile (horizontal axis) by household earnings decile (vertical axis). Household income and earnings have been equivalised taking the number and age of household members into account using OECD equivalence scales. Household earnings and income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). Panel A refers to the UK and panel B to the US. See text in Appendix C for more details on data sources, sample restrictions and variable definitions.

Figure C3. MEDIAN HOURS WORKED AMONG MINIMUM WAGE WORKERS AGED 25-64, BY HOUSEHOLD INCOME DECILE

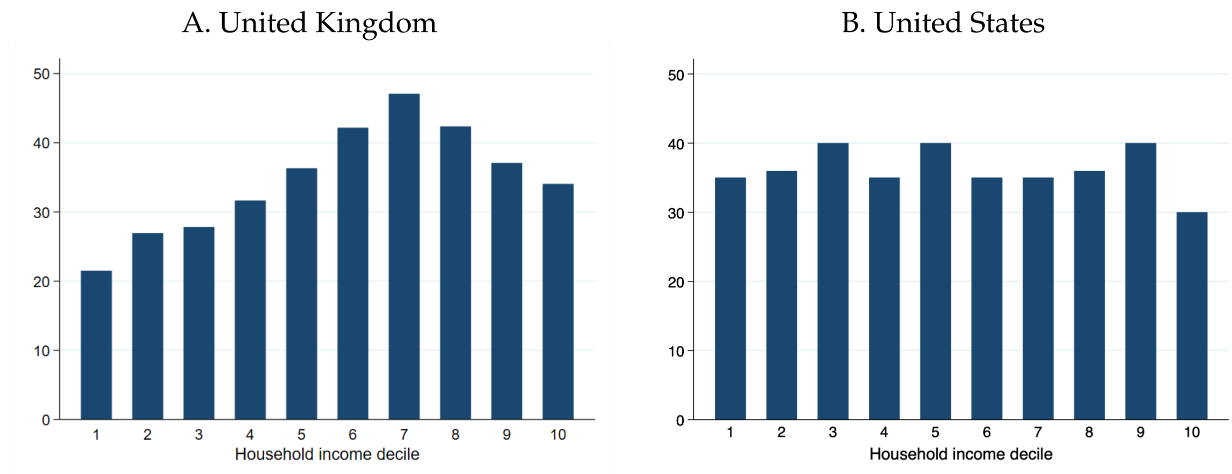

Notes: The graphs show median hours worked among minimum wage workers aged 25-64, by household income decile. Household income has been equivalised taking the number and age of household members into account using OECD equivalence scales. Household income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). An individual is defined as a minimum wage worker if they have a wage lower than 1.25 times the state-specific minimum wage. Panel A refers to the UK and panel B to the US. See text in Appendix C for more details on data sources, sample restrictions and variable definitions.

Figure C4. MEDIAN MARGINAL TAX RATE AMONG MINIMUM WAGE WORKERS AGED 25-64, BY HOUSEHOLD INCOME DECILE

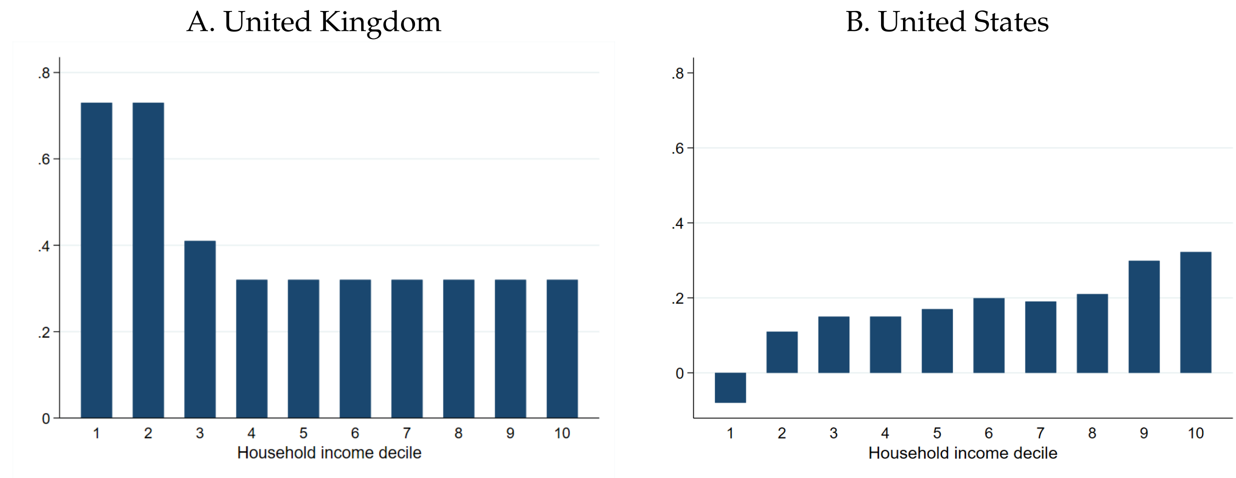

Notes: The graphs show median marginal tax rates among minimum wage workers aged 25-64, by household income decile. Household income has been equivalised taking the number and age of household members into account using OECD equivalence scales. Household income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). An individual is defined as a minimum wage worker if they have a wage lower than 1.25 times the state-specific minimum wage. Marginal tax rates have been simulated using TAXBEN for the UK and TAXSIM for the US. Panel A refers to the UK and panel B to the US. See text in Appendix C for more details on data sources, sample restrictions and variable definitions.

Figure C5. AVERAGE SHARE OF HOUSEHOLD INCOME DERIVED FROM WORK AMONG INDIVIDUALS AGED 25-64, BY HOUSEHOLD INCOME DECILE

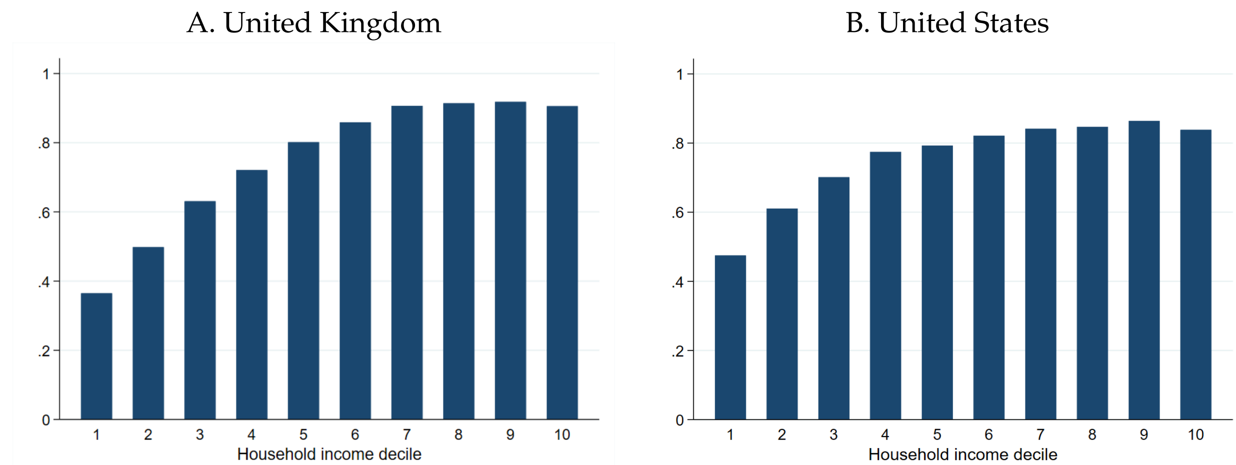

Notes: The graphs show the average share of household income derived from work among individuals aged 25-64, by household income decile. Household income and earnings have been equivalised taking the number and age of household members into account using OECD equivalence scales. Household income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). The share of household earnings divided by household income is censored at one. Panel A refers to the UK and panel B to the US. See text in Appendix C for more details on data sources, sample restrictions and variable definitions.

Figure C6. SHARE OF HOUSEHOLD EARNINGS COMING FROM SELF-EMPLOYMENT AMONG INDIVIDUALS AGED 25-64, BY HOUSEHOLD INCOME DECILE

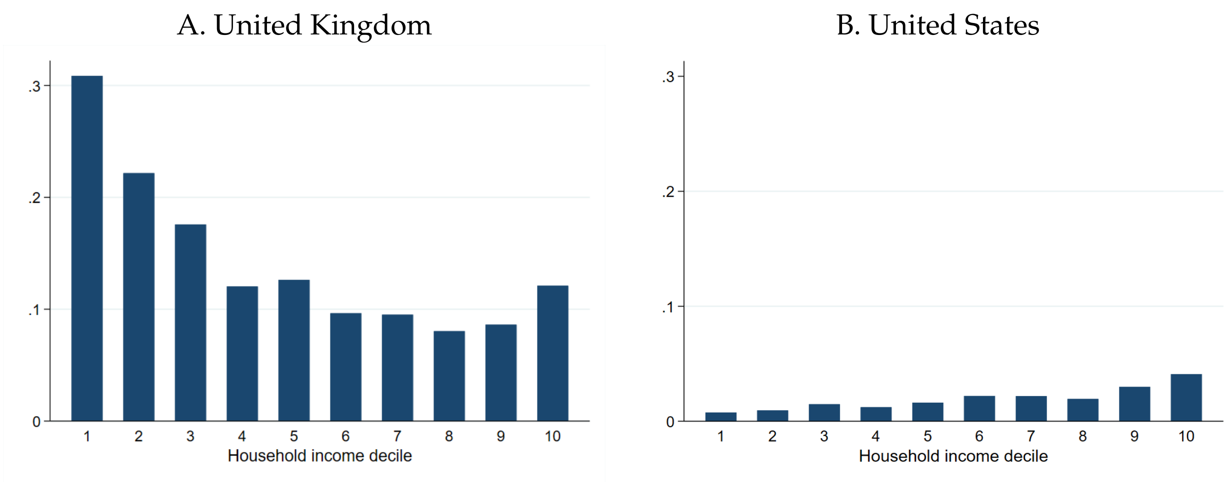

Notes: The graphs show the share of household earnings coming from self-employment among individuals aged 25-64, by household income decile. Household income has been equivalised taking the number and age of household members into account using OECD equivalence scales. Household income deciles are based on the whole population (i.e. an individual in the top decile is in the top decile of all people, not just those aged 25-64). Panel A refers to the UK and panel B to the US. See text in Appendix C for more details on data sources, sample restrictions and variable definitions.