

The Employment and Distributional Impacts of Nationwide Minimum Wage Changes

Giulia Giupponi, *Bocconi University, Institute for Fiscal Studies (IFS), and Centre for Economic Policy Research (CEPR)*

Robert Joyce, *IFS*

Attila Lindner, *University College London (UCL), IFS, HUN-REN CERS Institute of Economics, and CEPR*

Tom Waters, *IFS and UCL*

Thomas Wernham, *IFS and UCL*

Xiaowei Xu, *IFS and UCL*

We assess the impact of nationwide minimum wages on employment throughout the whole wage distribution by exploiting geographical variation in the level of wages. We find a substantial increase in wages at the bottom of the wage distribution, while we detect a small, statistically insignificant negative effect on employment. Combining the estimated change in the wage distribution with a tax and benefit micro-simulation model, we show that the minimum wage generates considerable proportional income gains up to the middle of the household income distribution.

We are grateful to the Low Pay Commission (grant CR20017) and the Open Research Area program (grant ES/X008339/1) for funding and to the Economic and

Submitted September 30, 2022; Accepted October 24, 2023.

Journal of Labor Economics, volume 42, number S1, April 2024.

© 2024 The University of Chicago. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits reuse of the work with attribution. Published by The University of Chicago Press in association with The Society of Labor Economists and The National Opinion Research Center. <https://doi.org/10.1086/728471>

I. 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 are 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 United Kingdom, 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 of the employment effects of minimum wages have been conducted (for reviews, see Neumark and Wascher 2008; Belman and Wolfson 2014; Dube 2019a). However, there are relatively fewer studies of the effects of minimum wages across the wage distribution and how those translate into effects on household incomes. Notable recent exceptions include Cengiz et al. (2019) and Dube (2019b) studying the impact on wages and household incomes in the US context, where the level of minimum wages is relatively low.

This paper proposes a new empirical methodology to estimate the impacts of the minimum wage on employment and wages in a context in which

Social Research Council for cofunding through the Centre for the Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (ES/T014334/1) and through the UK Research and Innovation (UKRI) Strategic Priorities Fund under grant ES/ W010453/1. 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 program (grant 949995). We thank Pietro Biroli, Kirill Borusyak, Alex Bryson, Tim Butcher, Arin Dube, Thomas Le Barbanchon, and Rebecca Riley as well as participants in seminars at the Low Pay Commission, the Institute for Employment Research (IAB)/Labor and Socio-Economic Research Center (LASER) workshop, the Milan Labour Lunch Seminar, the KOF Swiss Economic Institute at Eidgenössische Technische Hochschule (ETH) Zurich, the European Labor Symposium for Early Career Economists, CESifo, the Institute of Labor Economics (IZA) Workshop on Labor Market Institutions, and the Fondazione Rodolfo De Benedetti (fRDB) Annual Workshop for helpful comments. We also thank Anna Becker, Jonathan Cribb, and Emiliano Sandri for help with data and background analyses. This work was produced using data from the Office for National Statistics (ONS). Data from the Households Below Average Income dataset and Family Resources Survey were made available by the Department for Work and Pensions (DWP). We use research datasets that 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. Contact the corresponding author, Attila Lindner, at a.lindner@ucl.ac.uk. Information concerning access to the data used in this paper is available as supplemental material online.

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). Similar to Card (1992), we exploit differences in wage levels between areas, which, at least at the lower end of the wage distribution, are likely to arise from variation in the general price level (i.e., living costs), local aggregate productivity, or local amenities. In addition, instead of simply calculating the employment change for specific subgroups—such as teens in Card (1992)—we trace out employment changes throughout the whole frequency distribution of wages, as in Harasztosi and Lindner (2019) and Cengiz et al. (2019).

To identify the effect of the policy over the wage distribution, we compare trends in employment between groups that would 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. More specifically, we apply a difference-in-differences strategy where we compare the employment change in a low-wage region to the employment change in a high-wage region for workers with similar skills. Since individuals living in higher-wage regions are less exposed to the minimum wage, we can use their employment change as a counterfactual for similarly skilled individuals living in lower-wage regions. By applying this logic, we can estimate the impacts of the minimum wage on the number of jobs throughout the whole wage distribution.

We apply this new methodology to study the impacts of the introduction of the National Living Wage (NLW) in the United Kingdom, as well as 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 international frontier (see fig. A1; figs. A1–A8, C1, D1–D5, E1, F1–F6 are available online). We use data on wages from the Annual Survey of Hours and Earnings (ASHE), a high-quality employer survey on earnings and hours of employees in the United Kingdom, along with employment data from the Annual Population Survey (APS).

We calculate the change in employment and wages by adding up “missing jobs” just below the new minimum wage and “excess jobs” at and slightly above it, in the spirit of Cengiz et al. (2019). We find that over the 2016–19 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 disemployment effects. We estimate an own-wage elasticity of employment of -0.20 (standard error, 0.32), which is in line with many estimates in the literature and corroborates findings of previous work in the United Kingdom using other methods (Dube 2019a). 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 minimum wage itself.

In the second part of the paper, we use the estimated change in the frequency distribution of wages to assess the distributional impact of the policy. To do that, we use the most detailed microsimulation model of the UK tax-transfer system (Waters 2017) to estimate the effects of the minimum wage on net household income, using high-quality household survey data from the Family Resources Survey (FRS). The relationship between minimum wages and household income is complicated. First, it depends on the location of minimum wage workers in the household income distribution and the share of household income that their earnings make up. Second, 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 find that the largest income gains of the minimum wage in cash terms go to the middle of the household income distribution, where households with at least one minimum wage worker are most likely to be located and where marginal tax rates are substantially lower than at the bottom of the distribution. At the same time, in proportional terms the impacts are similar for the poorest households and middle-income households. Our results, therefore, highlight that the minimum wage provides considerable benefits up to the middle of the household income distribution, with effects fading out quickly in the top half.

This finding is in contrast to reduced-form estimates from the United States (Dube 2019b), which find strong effects on posttax incomes at the bottom of the distribution but not in the middle. We document a number of reasons for the difference in distributional impacts in the United Kingdom and the United States. First, minimum wage workers in the United Kingdom are concentrated in the middle of the household income distribution, while in the United States they are predominately located toward the bottom of the distribution. Second, in the United Kingdom, minimum wage workers at the bottom of the distribution gain less from minimum wage increases because they work fewer hours and face higher marginal tax rates than those further up the distribution. Third, individuals at the bottom of the income distribution in the United Kingdom get less of their income from employee work and a much higher share of their earnings from self-employment, which is not covered by the minimum wage.

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. First, 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 United Kingdom, where empirical strategies commonly employed often have limited statistical power (Brewer, Crossley, and Zilio 2019). Second, 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.¹ Third, 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 pioneered by Card (1992). 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. We highlight the empirical relevance of these advantages by comparing estimates using our approach and the traditional regional variation approach and by providing a step-by-step bridging of the two.

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 analyze 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—microsimulations and reduced-form econometric approaches. In particular, we retain a key benefit of microsimulations, which is the ability to easily run counterfactuals. At the same time, we incorporate key advantages of reduced-form methods, since—by integrating the rich information on labor market impacts from the frequency distribution approach—we can account for nonmechanical effects of minimum wages, in particular employment effects and wage spillovers induced by the policy.

Our paper contributes to a vast literature on the employment effects of the minimum wage. Most of the US literature exploits methods based on state-level variation (Card and Krueger 1994, 2000; Neumark and Wascher 2008; 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 subnational variation in

¹ See app. B in Cengiz et al. (2019).

minimum wage rates; see, for instance, Stewart (2002), Dolton, Bondibene, and Wadsworth (2012), Dolton, Bondibene, and Stops (2015), and Dube (2019a) for the United Kingdom; Dustmann et al. (2022) for Germany; and Portugal and Cardoso (2006) for Portugal. Appendix B (apps. A–F are available online) provides a classification of studies of the employment effects of the minimum wage in European countries by method used for identification.

Our paper is also related to a smaller body of work examining the distributional effects of the minimum wage, including impacts on the wage distribution (DiNardo, Fortin, and Lemieux 1996; Lee 1999; Autor, Manning, and Smith 2016; Cengiz et al. 2019) and on the household income distribution (Dube 2019b).

The remainder of the paper is structured as follows. Section II describes the institutional context. Section III details the methodology and data used for the estimation of the effects on labor market outcomes, and section IV illustrates the related empirical results. The methodology used to simulate impacts on household incomes is described in section V, and the simulation results are reported in section VI. Section VII concludes.

II. Institutional Context

We analyze the distributional consequences of minimum wages in the context of the United Kingdom, which has increased the minimum wage substantially for most adults since 2016. The United Kingdom 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 United Kingdom as the “National Living Wage”—although 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% 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 thereafter—in real terms (fig. 1A) and as a percentage of the median wage (fig. 1B). 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 consumer price index (CPI)-adjusted real terms. Overall, the 17% real-term increase in the minimum wage applying to those aged ≥ 25 between 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




FIG. 1.—Real minimum wage rate and minimum wage bite, 1999–2019. *A* reports the CPI-adjusted level of the United Kingdom adult minimum wage applying to workers aged 25 and over from April 1999 to April 2019. *B* reports the adult minimum wage bite (i.e., the adult minimum wage rate as a percentage of the median wage). The vertical dashed line corresponds to the NLW introduction on April 1, 2016.

prior 16-year period since the United Kingdom’s minimum wage was introduced in 1999.²

III. Employment and Wage Effects of the Minimum Wage: Methodology and Data

A. Combining the Regional Variation and Frequency Distribution Approaches

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 summarizing those approaches and illustrating how we combine features of both. We then describe our methodology in detail.

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 formalized with a statistical model, where for any two time periods employment changes in location r are modeled as a function of the bite of the minimum wage in that region:

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

$$\Delta E_{rt} = \alpha \text{BITE}_{rt-1} + \gamma_t + \mu_{rt}, \quad (1)$$

where ΔE_{rt} is the change in the employment rate (employment over the regional working age population N_{rt}) in region r between time $t - 1$ and t , BITE_{rt-1} is 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 region r at time $t - 1$, γ_t is the time effect, and μ_{rt} is an error term. The key identifying assumption is a common trends assumption that underlying employment trends across regions are unrelated to the 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 labor market, statistical power can be low. One can think of the problem as being one of a weak first stage: BITE is typically associated only with very small changes in average wages, 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—among teenagers, for example (Card 1992)—although naturally this limits external validity. Another alternative is to further segment the population by demographics, such as sex, age, and skill level, to create additional variation in bite (Stewart 2002; Manning 2016; Dube 2019a). This introduces additional potential problems: namely, the estimated employment effects can be misleading if the policy impact varies across demographic groups.³ Furthermore, when variation across demographic groups is applied, it is unclear whether the estimated employment effects reflect the impact of the policy on certain demographic groups or the overall impact on low-wage jobs.

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

³ With treatment effect heterogeneity, the estimated employment effects can be positive even if the employment change is negative for all groups. To see that, consider a setting with four demographic groups: low- and high-skilled women and men. Suppose that the employment elasticity with respect to the minimum wage is substantially larger for women, but the difference in exposure between low- and high-skilled women is small. At the same time, the employment elasticity for men is small, but the exposure difference between low- and high-skilled men is large. In this case, the standard difference-in-differences approach could yield a positive employment elasticity, as it would pick up that the large exposure differences between men lead to small employment changes, while the small exposure differences for women lead to large negative employment changes.

effects, for example because of firms' desire to maintain pay differentials between different occupations or between supervisory and nonsupervisory 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 a "missing" mass below and an "excess" mass at or above the new minimum. This implicitly defines the total employment effect, which is the difference between the missing and excess masses. Using this framework, the impacts of the minimum wage on the wage distribution and employment are captured jointly in a fully integrated way.

Explicitly disaggregating estimated employment changes by wage bin brings several advantages. First, by focusing on employment changes at the bottom of the wage distribution, we can 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. This could considerably improve statistical precision. Second, by estimating the impacts on employment in every wage bin, a kind of falsification check is automatically produced: significant estimated effects on the number of jobs far up the wage distribution could suggest that the identification strategy may be conflating impacts of the minimum wage with other differences between treatment and control groups. The additional falsification test is especially valuable given the notoriously fierce ongoing debates on the employment impacts of minimum wages in the literature. Finally, estimating effects wage bin by wage bin paints a richer picture of the policy'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 United States 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 United Kingdom, as 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 wage 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 (2019a), Schmitz (2019), and Dustmann et al. (2022).

Nesting the regional and frequency distribution approaches.—Similar to the regional variation approach, our method exploits differences in wage levels between areas, which, at least at the lower end of the wage distribution, are likely to arise from variation in the general price level (i.e., living costs), local aggregate productivity, or local amenities. 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. In practice, workers belonging to the same group—which we will label “skill level,” as it could be thought of broadly as a skill group—will receive higher nominal wages in high-wage than in low-wage areas. This allows us to use trends in the number of jobs in high-wage areas as counterfactuals for trends in the number of jobs in lower-wage areas, effectively matching wage bins across areas that are equivalent in real terms but—due to cost-of-living differences (or other price differences)—differentially exposed to the national minimum wage. Since in practice no region is entirely unaffected by a national minimum wage policy, our approach shares the characteristic of other regional variation approaches of identifying the relative effect of the minimum wage on employment in lower-wage areas compared with higher-wage ones. Retrieving an absolute effect 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. For any employment change within any wage bin observed in a lower-wage area, we can net off an estimated counterfactual change, which is identified from what happens in high-wage areas to workers of the same skill level but with different nominal wages. Aggregating across low-wage regions yields the estimated impacts of the change in the minimum wage on the frequency distribution of wages in those regions (in the relative sense described above). The identifying assumption is that absent minimum wage changes, employment changes for each skill level would evolve in the same way across lower- and higher-wage regions.

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 extensions of it. Those approaches implicitly assume that the population of workers living in areas less affected by the minimum wage are a good control group for the population of workers living in more affected areas. Our approach relies on a weaker assumption, as it compares only narrowly defined subsets of workers with similar skill levels living in different areas.

In addition to that, our approach differs from the traditional approaches that combine regional with demographic variation in exposure to the minimum wage, since we do not exploit variation across skill groups. The underlying idea of those traditional approaches is to exploit identifying variation

coming from skill-level differences across demographic groups on top of variation across locations. Our approach does not use such skill-level variation for identification per se. Rather, it separately identifies the effect of the policy on workers of different skill levels. As a result, our approach transparently shows which parts of the skill distribution drive the estimated changes in employment. In appendix C, we provide a step-by-step mapping from our methodology to the regional variation approach.⁴

B. Methodology

To implement our approach, we first need to define a skill level for each worker in our sample. A skill level identifies workers who would earn the same wage if they lived in the same place at the same point in time. This requires purging wages of place and time effects so that we can use those transformed wage levels as indicators of skills. Then we use a difference-in-differences-style framework to examine differential trends in employment by skill level across regions that are more and less affected by the minimum wage. This delivers our estimates of the impacts of the minimum wage. We describe both steps in turn below.

Purging wages of place and time effects.—We have individual wage observations from across the United Kingdom and across our sample period, and we want to assign each of those observations a skill level. This requires purging wages of both place and time effects. For example, our data include individuals in Hull in 2016 and individuals in London in 2017. To know whether an individual from the first group has the same skill level as an individual from the second group, we need to know what wage the Hull-2016 observation would earn (i) if they were in London rather than Hull and (ii) after applying expected wage growth over the 2016–17 period.

To identify place effects, we run a Mincer-style regression of raw individual log wages $\ln w_{it}^*$ on location effects, year effects, and individual controls. This is estimated using only pre-policy-reform data, pooled from 2012 to 2014, to avoid any confounding effect of the minimum wage increase. Our regression specification is

$$\ln w_{it}^* = \ln \delta_{r(i,t)} + \theta_t + X'_{it}\beta + \nu_{it}, \quad (2)$$

where X is a vector of individual and firm characteristics, $\ln \delta_{r(i,t)}$ is the logarithm of the location-specific relative pay premia, θ_t is the year effect, and ν_{it} is an error term. Covariates include gender interacted with full-time/part-time status and age, one-digit occupation, one-digit industry, and a dummy for being in a graduate job (based on the four-digit Standard Occupational Classification code). In some specifications, we also include person effects

⁴ We invite interested readers to go through secs. III.B–IV.A before turning to app. C.




FIG. 2.—Estimated wage premia. The graph reports estimates of $\ln \delta_r$ from regression equation (2) (solid circles) using data from 2012 to 2014. The capped vertical bars show 95% confidence intervals based on robust standard errors clustered at the TTWA level. Estimates in black refer to regions in the bottom nine deciles of the distribution of wage premia (treatment group), and those in gray refer to regions in the top decile of the distribution (control group).

in the regressions. To account for the left censoring at the minimum wage, we use a Tobit as our benchmark specification. Estimates of (the logarithm) of location effects ($\ln \delta_r$)—which we call regional “wage premia”—are shown in figure 2.⁵

We implement the rest of the analysis using wages that are adjusted by average wage inflation over time.⁶ From here on, our notation w_{it}^* refers

⁵ Our estimation of the regional wage premia raises two main concerns. The first one is that regression eq. (2) does not account for unobserved differences in skill levels across locations. In sec. D.2 of app. D, we provide alternative estimates of regional wage premia based on a two-way fixed effects estimation including individual fixed effects in the regression. A second concern for the identification of regional wage premia is sorting of workers across regions based on a (person-region) match component of wages. The presence of sorting would imply that our regression eq. (2) is misspecified and our estimate of $\ln \delta_r$ biased. We formalize this issue and test its empirical relevance in sec. D.2 of app. D.

⁶ Part of the growth in average wages can be driven by changes in the minimum wage itself. To make sure that this is not what drives our estimates, in our benchmark specification we control for the direct impact of the minimum wage on average wage growth by estimating time effects using the following regression: $\ln w_{it}^* = \gamma_r + \beta \text{GAP}_{r(i,t)} + \tau_t + \varepsilon_{it}$, where $\ln w_{it}^*$ is the raw log hourly wage of individual i in

to raw wages and w_{it} to wages adjusted for time effects. Unless otherwise stated, our discussion will always refer to wages adjusted for time effects.

Using our estimates of place effects $\ln \delta_r$, we can obtain time- and place-adjusted wages:

$$\exp(\ln w_{it} - \ln \delta_{r(i,t)}) = \frac{w_{it}}{\delta_{r(i,t)}}. \quad (3)$$

As noted above, we refer to this as “skill levels.” Workers with the same skill level would earn the same amount if they lived in the same place at the same time.

Estimating the effect of the minimum wage on the frequency distribution of wages.—Our basic strategy is as follows. For any wage level in the low-wage (treated) locations, we use the change in employment in high-wage (control) locations for workers of the same skill level as counterfactual employment change. This allows us to trace out the impact of the minimum wage across the wage distribution.

To build intuition, consider first a situation with two periods and two regions: region H is the control (high-wage) region, and region L is the treated (low-wage) region. We normalize the wage premium in the treated region to $\delta_L = 1$; in the control region, the wage premium is $\delta_H > 1$. We are interested in the impact of the minimum wage below a given wage level $w = c$. Let $\Delta E_L(c)$ denote the change in total employment in the low-wage region at wages below c , as a share of the population. To identify the impact of the minimum wage, we need to find the counterfactual employment rate change at this wage level in the absence of the minimum wage, $\Delta E_L(c)^{CF}$. We do this by finding the change in the employment rate of similarly skilled workers in the high-wage region. A worker who in the low-wage region earns c would earn $c\delta_H$ in the high-wage region. We therefore use the employment rate change below $c\delta_H$ in the high-wage region as the counterfactual for the employment rate change below c in the low-wage region, and we calculate the percent change in employment as follows:

$$\text{DiD}_{E(c)} = \frac{\Delta E_L(c) - \Delta E_L(c)^{CF}}{\text{EPOP}_{t-1}} = \frac{\Delta E_L(c) - \Delta E_H(c\delta_H)}{\text{EPOP}_{t-1}}, \quad (4)$$

where $\Delta E_L(c) - \Delta E_L(c)^{CF}$ is the difference-in-differences employment rate change below c and $\text{EPOP}_{t-1} = \lim_{c \rightarrow \infty} E_{t-1}(c)$ is the employment to population rate in the baseline period $t - 1$. Notice that by dividing by the

location r and year t , $\text{GAP}_{r(i,t)}$ is the mechanical increase in average wages that the higher minimum wage would induce for workers in location r in year t relative to $t - 1$, τ_t is the time trend that we want to extract, and ε_{it} is an error term. Table A1 (tables A1, B1, C1, C2, D1 are available online) reports the estimated coefficients β and τ_t for different year pairs in our sample. In sec. IV.C, we assess the robustness of our results to estimating wage trends without controlling for $\text{GAP}_{r(i,t)}$. This makes virtually no difference to our results.

employment to population rate, we express employment changes as a share of the prereform (national) employment.

In this simple two-region example, we can trace out the impacts of the minimum wage across the entire wage distribution by estimating equation (4) for every level of c . For this difference-in-differences strategy to identify the causal effect of the minimum wage on employment, we need to assume that—absent changes in the minimum wage—employment rates would evolve in the same way in the treated (low-wage) and control (high-wage) region for every level of c . We assess the validity of this common trends assumption in section D.2 of appendix D.

Furthermore, the minimum wage should ideally not be binding in the control region. In practice, though, there will be some jobs affected by the minimum wage even in the control (high-wage) region. Our methodology does not allow us to identify the impact of the minimum wage on the skill levels corresponding to those jobs. Indeed, the lack of a control group for workers that are affected by the policy in both low- and high-wage regions implies that the effect of the policy on those worker cannot be calculated without some extrapolation. While this is an important limitation of our approach, it is a general feature of all empirical designs studying the impact of nationwide minimum wage changes.⁷

We now move away from the two-region example to our setting, where we have multiple locations in each of the treatment and control regions, and multiple time periods. Let H and L now refer to the set of high-wage (control) and low-wage (treated) locations, respectively. Let $\Delta E_{rt}(c)$ denote the change in employment rate below wage c for some treated locations $r \in L$ between time $t - 1$ and t . To create $\Delta E_{rt}(c)^{CF}$, we calculate the following (population-weighted) average over control locations of the change in the employment rate below the corresponding skill level:

$$\Delta E_{rt}(c)^{CF} = \frac{\sum_{r' \in H} \Delta E_{r't} \left(\frac{\partial_{r'} c}{\partial_r} \right) N_{r't-1}}{\sum_{r' \in H} N_{r't-1}}, \quad (5)$$

where $N_{r't-1}$ is the population in control location r' at time $t - 1$.

For each wage level, we then take all treated locations and average over the difference between actual and counterfactual employment rate changes, weighted by the pretreatment population in that location. The average effect of the minimum wage below wage c in treated locations at time t is therefore

$$\text{DiD}_{E,(c)} = \frac{1}{\text{EPOP}_{t-1}} \frac{\sum_{r' \in L} (\Delta E_{r't}(c) - \Delta E_{r't}(c)^{CF}) N_{r't-1}}{\sum_{r' \in L} N_{r't-1}}. \quad (6)$$

⁷ To alleviate this concern, in our empirical implementation we use very high wage locations as the control group, as we explain in more detail below.

Empirical implementation.—In practice, we implement the approach described above by calculating employment rate changes in discrete wage bins of £ x rather than calculating the cumulative distribution function at different wage levels. Let e_{krt} denote the employment density in wage bin k , which runs from k to $k + x$:

$$e_{krt} = E_{rt}(k + x) - E_{rt}(k). \quad (7)$$

For every wage bin in every treated location, it is possible to find the range of skills exactly corresponding to e_{krt} in every control location. By applying the same logic as in equation (6), this leads to the following estimator of the employment change in wage bin k in treated locations:

$$\begin{aligned} \text{DiD}_{e_{kt}} &= \frac{\sum_{r' \in L} (\Delta E_{r't}(k + x) - \Delta E_{r't}(k + x)^{\text{CF}}) N_{r't-1}}{\text{EPOP}_{t-1} \sum_{r' \in L} N_{r't-1}} \\ &\quad - \frac{\sum_{r' \in L} (\Delta E_{r't}(k) - \Delta E_{r't}(k)^{\text{CF}}) N_{r't-1}}{\text{EPOP}_{t-1} \sum_{r' \in L} N_{r't-1}} \\ &= \frac{1}{\text{EPOP}_{t-1}} \frac{\sum_{r' \in L} (\Delta e_{kr't} - \Delta e_{kr't}^{\text{CF}}) N_{r't-1}}{\sum_{r' \in L} N_{r't-1}}, \end{aligned} \quad (8)$$

where

$$\begin{aligned} \Delta e_{kr't}^{\text{CF}} &= \Delta E_{rt}(k + x)^{\text{CF}} - \Delta E_{rt}(k)^{\text{CF}} \\ &= \frac{\sum_{r' \in H} \left[\Delta E_{r't} \left(\frac{\delta_{r'}(k+x)}{\delta_r} \right) - \Delta E_{r't} \left(\frac{\delta_{r'}k}{\delta_r} \right) \right] N_{r't-1}}{\sum_{r' \in H} N_{r't-1}}. \end{aligned} \quad (9)$$

In essence, our proposed estimator for the wage bin-specific employment rate change (Δe_{kt}) is the (population-weighted) average difference between the actual and the counterfactual employment rate change ($\Delta e_{krt} - \Delta e_{krt}^{\text{CF}}$) across low-wage locations.

We implement the above estimator using the following (population weighted) regression specification:

$$\frac{\Delta e_{krt} - \Delta e_{krt}^{\text{CF}}}{\text{EPOP}_{t-1}} = \sum_{f=\underline{F}-x}^{\bar{F}} \alpha_f \mathbb{I}[k = f] + \eta_{krt} \quad \text{for } r \in L, \quad (10)$$

where $\mathbb{I}[k = f]$ is an indicator function taking the value 1 if wage bin k corresponds to values between f and $f + \text{£}x$ above the new minimum wage at time t and is equal to 0 otherwise. We center the wage bin indicators around the postreform minimum wage so that the changes in the distribution of wages are easy to visualize. We aggregate effects at the tails of the wage distribution: we sum up all changes in wage bins below $\underline{F}$, the postreform minimum wage, and all changes in wage bins above $\bar{F}$, which we set at £15 above

the new minimum wage. The α_f coefficients show the estimated bin-by-bin change in the employment rate relative to the new minimum wage, expressed as a share of the national employment rate. In our benchmark specification, we implement regression equation (10) by pooling different years, which is why our α_f coefficients are not indexed by t .⁸

Approximation.—Calculating the true counterfactual based on equation (9) is possible but computationally intensive. This is because for each wage bin in each treated region there is a separate combination of wage bins in control regions covering workers of the same skill level. As a result, we approximate the employment rate change in equation (9) in the following way. First, we estimate the exact counterfactual employment change for each wage bin for a reference region. Then we approximate the counterfactual employment change for workers in wage bin k and location r by taking an average of the counterfactual employment changes calculated in the first step among workers in the reference region who are of the same skill level as wage bin k in r . We provide more details on the approximation in section D.1 of appendix D.

Parametrization of benchmark specification and bootstrapping.—For our main estimates we set $x = 0.25$, partitioning the wage distribution into wage bins k of £0.25 width. We define high-wage locations (the control group) as those with wage premia in the top decile of the distribution of $\ln \delta_{r(i,t)}$ and low-wage locations (treated group) as those with premia in the bottom nine deciles. In section IV.C, we assess the sensitivity of our results to different bin widths and definitions of control locations.

For statistical inference, we use a bootstrap procedure (with 100 replications). To allow for clustering at the local level, we randomly draw locations (not individuals) with replacement. The sample is then composed of workers in those randomly drawn locations, with duplicates of those in locations that were drawn multiple times. We bootstrap all steps in our methodology, from the estimation of wage trends and local wage premia, to that of counterfactual employment rate changes and the α_f coefficients.

Calculating the employment effect.—As in Cengiz et al. (2019), the set of α_f 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 $\Delta b = \sum_{f=F-x}^0 \alpha_f$ and the excess mass above it as $\Delta a = \sum_{f=0}^{\bar{F}} \alpha_f$. Since the α_f coefficients identify changes in employment as a percentage of the pretreatment national employment rate, the missing and excess mass can be interpreted analogously. Their sum, which we define as $\Delta e = \Delta a + \Delta b$, represents the total percent change in the employment rate due to the minimum wage. For our baseline estimates, we set $\bar{F}$ equal to the minimum wage (NLW) plus £5, meaning that any employment changes occurring more than £5 above the

⁸ This step corresponds to further averaging eq. (6) over time periods, which would lead to $\sum_t \sum_{r \in L} (\Delta E_{rt}(k) - \Delta E_{rt}(k)^{CF}) N_{rt-1} / \sum_t \sum_{r \in L} N_{rt-1}$.

new minimum are not assumed to result from the minimum wage change and do not contribute to Δe . We show the sensitivity of our results to alternative choices of $\bar{F}$, and we also routinely report an estimate of $\Delta e_{\text{total}} = \sum_{-\infty}^{\infty} \alpha_f$, which aggregates the estimated α_f over the entire support of the wage distribution. This is never far from our central estimate of Δe , which is reassuring evidence in favor of our identifying assumptions: it implies that employment rate changes within given skill levels were very similar between treatment and control regions whenever we look beyond the lower portion of the wage distribution.

A conceptual difference between our α_f 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 United Kingdom does not provide geographic variation in minimum wage policy, so there are no geographic areas that are completely “untreated” that can be used as controls to identify absolute effects. One can, however, recover absolute effects across the whole economy with some extrapolation, as we explain below.

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 α_f 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 $\sum_{f=\bar{F}-x}^{\bar{F}} \alpha_f$), divided by the share of the workforce earning below the new minimum wage in the year before treatment ($\bar{b}_{-1}$), across the whole population (i.e., both high- and low-wage regions):

$$\% \Delta e - \% \Delta e^{\text{CF}} \approx \frac{\Delta e - \Delta e^{\text{CF}}}{\bar{b}_{-1}} = \frac{\sum_{f=\bar{F}-x}^{\bar{F}} \alpha_f}{\bar{b}_{-1}}. \quad (11)$$

We also use the estimated α_f coefficients 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 average wage of affected workers. We then divide that by prepolicy average wages among affected workers, as illustrated in the following formula:

$$\% \Delta w - \% \Delta w^{\text{CF}} \approx \frac{\frac{\overline{wb}_{-1} + \sum_{f=\bar{F}-x}^{\bar{F}} (f + \overline{\text{MW}}) \alpha_f}{\bar{b}_{-1} + \sum_{f=\bar{F}-x}^{\bar{F}} \alpha_f}}{\overline{wb}_{-1} / \bar{b}_{-1}} - 1. \quad (12)$$

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 prepolicy period, the average wage is computed as the ratio of the preperiod wage bill

among those paid less than the new minimum $\overline{wb}_{-1}$ divided by the share of the workforce earning below the new minimum $\bar{b}_{-1}$.⁹ This is the denominator in formula (12).

To understand how the proportional relative effect of the minimum wage on the 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 prepolicy wage bill $\overline{wb}_{-1}$ and the wage bill increase generated by the minimum wage in low-wage regions relative to high-wage ones $\sum_{f=E-x}^{\bar{F}} (f + \overline{MW})\alpha_f$, where $\overline{MW}$ is the average wage in the bin where the minimum wage falls in the postperiod. This is then divided by the sum of the prepolicy number of workers paid below the new minimum plus the relative increase in the number of affected workers in low- versus high-wage regions ($\bar{b}_{-1} + \sum_{f=E-x}^{\bar{F}} \alpha_f$). 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).

Having estimated the wage elasticity of employment, we can also calculate the absolute effect of the minimum wage on employment. To do this, we first estimate the absolute wage effect of the minimum wage by comparing the wage distribution before and after a minimum wage increase (uprating the earlier year using the τ_t from the specification illustrated in n. 6). Second, we multiply the absolute wage effect by the own-wage elasticity of employment to get an estimate of the absolute change in employment.

C. Data and Sample Construction

Our primary data source for the analysis of the impacts of the NLW on employment, wages, and hours is ASHE for the years 2010–19. A large-scale employer-completed survey of earnings and hours of employees in the United Kingdom, 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.

In our empirical implementation, a region r is defined as a travel-to-work area (TTWA). TTWAs are statistically defined geographic units that are constructed by the United Kingdom's Office for National Statistics, based on commuting flows, to approximate local labor markets. They identify self-contained areas in which most people both live and work. Since ASHE is (weighted to be) representative at the national level but not at the local level, we rescale employment counts in ASHE to match employment counts in the locally representative APS. We also use APS to get the working age

⁹ To compute $\overline{wb}_{-1}$, we deflate wages in each region with the reference being the average wage premium in the low-wage regions. In other words, $\overline{wb}_{-1}$ is in the price (wage) of the low-wage regions. This makes it consistent with our estimated α_f .

population in each TTWA. We group TTWAs with fewer than 200 observations in ASHE with their nearest neighboring 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 137 geographic areas. We check the sensitivity of our results to different degrees of aggregation.

IV. Employment and Wage Effects of the Minimum Wage: Results

A. 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 (10). Each circle represents our estimate of employment changes—averaged over the four minimum wage increases from 2015 to 2019—in each wage bin relative to the level of the new NLW in each of the years we consider. For ease of visualization, we plot wage bins of £1 width: these are linear combinations of the £0.25 width wage bins (α_f coefficients) estimated in regression equation (10).

The figure shows a clear and significant drop in jobs just below the NLW, indicating 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 5.44% (standard error, 0.22%) of total employment in the previous year (Δb). We also find a large increase in the number of jobs at, or within £1 of, the new minimum wage (approximately 4.5% of pretreatment employment) and at wages slightly higher than the NLW, with spillovers stretching up to around £2 above it. This is around the twentieth percentile of hourly wages, which is broadly consistent with evidence of wage spillovers from minimum wages found previously (Autor, Manning, and Smith 2016; Cengiz et al. 2019; Harkness and Avram 2019). Our point estimates also indicate some small and statistically insignificant spillovers up to around £5 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 -0.11% (standard error, 0.16%) of pretreatment employment—a very small decline that is not statistically significant.

We calculate the own-wage employment elasticity using the formula described in section III.B. Our central estimate is -0.20 (standard error, 0.32)—a small effect, which is in line with several other estimates in the literature (Dube 2019a), including previous studies in the United Kingdom (Stewart 2004; Dube 2019a; Manning 2021). The 95% confidence interval allows us to rule out an employment elasticity that is more negative than -0.83 . As a result, we can rule out the large employment effects that would be implied by some other estimates in the minimum wage literature (Dube 2019a).




FIG. 3.—Impact of the minimum wage on the wage distribution: baseline estimates for workers aged 25–64. The graph reports linear combinations of estimates of the coefficients α_i from regression equation (10) 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 circle represents our estimate of employment changes, averaged over the four minimum wage increases from 2015 to 2019, in each £1 wage bin relative to the level of the new NLW in each of the years we consider. The £1 wage bins are linear combinations of the £0.25 wage bins (α_i coefficients) estimated in regression equation (10). Employment rate changes are normalized by the baseline national employment rate, so that the sum of the effects across all wage bins can be interpreted as the total percent change in employment arising from the change in the minimum wage. Estimated effects in wage bins below the new NLW, as well as those in wage bins more than £15 above the NLW, are aggregated in one single-point estimate. The gray line shows the running total of employment changes up to that point in the distribution. The vertical bars underlying the circles and the shaded area around the gray line show the bootstrapped 95% confidence intervals associated with 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 Δtotal (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 subcomponents—the percent change in affected employment and affected wages—are also reported. See section III.B for further details on these statistics.

It is also worth highlighting that here and throughout we report cluster-bootstrapped standard errors, where we allow for clustering at the regional level. Since we bootstrap all estimation steps in our analysis—including the estimation of wage trends and regional wage premia, on top of counterfactual wage changes and the α_f coefficients—our standard errors are more conservative than in similar studies in the literature. Indeed, estimates of the minimum wage bite are typically not bootstrapped in studies based on the regional-bite approach. Panel A of table 1 reports our main estimates, with bootstrapped standard errors and robust standard errors clustered at the TTWA level. The less conservative approach leads to a substantially more precise estimate, which allows us to rule out even modest negative employment effects of the policy.

The gray line in figure 3 shows the running total of employment changes up to that point in the distribution. For example, at £5 above the NLW, the gray 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 running sum is consistent with an employment effect that is close to zero, as is our estimate of the effect over the entire wage distribution ($\Delta_{\text{total}} = 0.25\%$; standard error, 0.31%). 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 low-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 D.2 of appendix D.

Figure A2 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 show up only in 2018 or 2019. The employment change up to £5 of the NLW is estimated at -0.42% , which is in fact very close to four times the estimated average effect from pooling the four consecutive-year periods. This estimate is more imprecise, however (standard error, 0.68%), because it uses much less data than our central estimates, which effectively pool the results from four different minimum wage increases.

B. Identification Tests

We summarize here the set of assumptions underlying our methodology and the identification tests that we run to assess their validity, and we refer the reader to section D.2 of appendix D for a more detailed discussion.

Table 1
Impact of the Minimum Wage on the Wage Distribution: Robustness Checks

	Δb (%)		$\Delta a + \Delta b$ (%)		Elasticity	
	Est (1)	SE (2)	Est (3)	SE (4)	Est (5)	SE (6)
A. Main Specification						
Main specification bootstrapped SE	-5.44	.22	-.11	.16	-.20	.32
Main specification clustered SE	-5.44	.17	-.11	.11	-.20	.14
B. Bin Width						
£0.10	-5.47	.28	-.10	.16	-.22	.41
£0.50	-5.51	.24	-.18	.20	-.26	.34
C. Threshold for Grouping TTWAs						
100 observations	-5.48	.21	-.07	.16	-.13	.35
400 observations	-5.32	.29	-.05	.19	-.10	.40
D. Cutoff for Calculation of Δa						
$\tilde{F} = 4$	-5.44	.22	-.18	.16	-.38	.36
$\tilde{F} = 6$	-5.44	.22	-.01	.15	-.02	.29
E. Estimation of Wage Premia and Wage Growth						
Bottom half of distribution	-4.91	.36	-.16	.16	-.39	.60
AKM (including person effects)	-5.15	.42	-.20	.26	-.48	4.08
Grouped AKM (including person effects)	-6.87	1.09	-.23	.46	-.36	.75
No industry/occupation controls	-5.61	.28	-.20	.18	-.37	.37
Only full-time workers	-5.50	.22	-.13	.15	-.24	.31
OLS estimation	-5.40	.23	-.09	.16	-.17	.32
Alternative adjustment to time effects	-5.07	.23	-.07	.15	-.15	.37
F. Treatment and Control Regions						
Bottom two deciles as treatment	-6.86	.36	-.15	.22	-.22	.34
Top two deciles as control	-5.60	.26	-.15	.14	-.28	.31
Excluding London	-5.36	.30	-.10	.23	-.19	.50

NOTE.—This 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 parameterizations of regression eq. (10). Estimates are averaged over the four minimum wage increases from 2015 to 2019. Columns 1, 3, and 5 report our central estimates; cols. 2, 4, and 6 report bootstrapped standard errors. See sec. III.B for further details on these statistics. Panel A reports baseline estimates from fig. 3 with bootstrapped standard errors and analogous estimates with robust standard errors clustered at the TTWA level. Panel B shows robustness to the choice of wage bin width (x), where we rerun the analysis bins 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 neighboring travel to work areas to 100 and 400 observations instead of 200. Panel D varies the $\tilde{F}$ cutoff for the calculation of Δa to £4 and £6, rather than £5. Panel E shows robustness to changes in the specification used for the estimation of wage premia and wage growth. In one variant, we estimate wage premia from regression eq. (2) using only the bottom half of the wage distribution in each region. In a second, we estimate wage premia using an AKM regression (including person and location effects) rather than a Mincerian regression. In a third variation, we estimate wage premia using an AKM regression (including person and location effects) on grouped TTWAs. In a fourth one, we drop industry and occupation controls from our Mincerian specification in regression eq. (2). In a fifth one, we estimate regression eq. (2) on full-time workers only. In a sixth one, we use an OLS rather than a Tobit model. In a seventh one, we estimate the specification in n. 6 without the GAP control. Panel F shows estimates for different definitions of treatment and control regions. Instead of comparing the bottom nine deciles of regional wage premia to the top decile, we compare (i) regions in the first two deciles (treated) to the top decile (control) and (ii) regions in the bottom eight deciles (treated) to the top two deciles (control). We also run the main specification excluding London from the set of control regions.

Stability of regional wage premia and correlation with NLW bite.—Our definition of skill levels is based on local wage premia estimated in the prepolicy period. This definition rests on the assumption that our estimated $\ln \delta_i$ are stable over time and are not affected by the NLW introduction. Both conditions are shown to be supported by the data (fig. D1A, D1B). In addition, we document that our definition of treatment and control regions is indeed capturing differential exposure to the NLW (fig. D1C).

Accounting for unobserved individual heterogeneity in the estimation of wage premia.—One limitation of our Mincerian specification in equation (2) is that it does not account for unobserved individual heterogeneity. We can deal with this issue by including person effects in the regression and estimating regional wage premia in an Abowd-Kramarz-Margolis (AKM)–style specification. The regional wage premia estimated using the AKM specification and the Mincerian specification in equation (2) are highly correlated (fig. D2A). To minimize the risk of limited-mobility bias, we also derive an alternative measure of the AKM premia by grouping TTWAs into 30 groups. The grouped AKM premia are also highly correlated with the Mincerian regional wage premia (fig. D2B). In section IV.C, 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.

Sorting bias in the estimation of wage premia.—A potential concern for the identification of regional wage premia based on estimating regression equation (2) is sorting of workers across regions based on an idiosyncratic match component of wages. In the presence of sorting, the regression equation (2) would be misspecified and our estimate of $\ln \delta_i$ would be biased. In section D.2 of appendix D, we test for and find no indication of the presence of sorting bias.

Validity of common trends assumption.—Our difference-in-differences strategy rests on a common trends assumption that absent changes in the minimum wage, employment rates in each skill level would evolve in the same way across lower- and higher-wage regions. As we already noted in section IV.A, 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. Figure D4 reproduces estimates from regression equation (10) using pre-NLW years 2011–15 and taking the 2016–19 NLW rates as placebo minima in each year (from 2012), respectively. This placebo test yields a close-to-zero employment effect throughout the wage distribution, corroborating the assumption of no differential trends.

Impact of the minimum wage in the control group.—Even though ideally we would like the minimum wage not to be binding in the control regions, in practice some jobs will be affected by the minimum wage even in the

control (high-wage) regions. As we already noted in section III.A, this implies that our estimates can identify the relative effect of the minimum wage in lower-wage versus higher-wage regions. This is an important limitation of our approach that is a general feature of empirical designs studying nationwide minimum wage changes. Nevertheless, as we document in figure D5, the minimum wage did not have an impact on total employment in the control regions, suggesting that our relative estimates are unlikely to differ substantially from the overall impact of the policy.

C. Robustness Checks

Robustness to parametrization.—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 parameterizations. For reference, our headline estimates are reported in panel A of table 1.

Panel B shows robustness to the choice of wage bin width (x), where we rerun the analysis using bins of £0.10 or £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 neighboring TTWAs to 100 and 400 observations instead of 200. Panel D varies the $\tilde{F}$ cutoff for the calculation of Δa to £4 and £6 rather than £5. In panel E, we show robustness to changes in the specification used for the estimation of wage premia and wage growth. In one variant, we estimate wage premia from regression equation (2) using only the bottom half of the wage distribution in each region.¹⁰ In a second variant, we estimate wage premia using an AKM regression on movers across TTWAs 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. In a fourth one, we drop industry and occupation controls from our Mincerian specification in regression equation (2). In a fifth one, we estimate regression equation (2) on full-time workers only. In a sixth one, we use an ordinary least squares (OLS) rather than a Tobit model. In a seventh one, we estimate the specification in footnote 6 without the GAP control. Panel F shows estimates for different definitions of treatment and control regions. We start by restricting the set of treated regions to those in the bottom two deciles of the regional wage premia distribution. We then alter the definition of both treatment and control regions by comparing regions

¹⁰ To define the subgroup of workers with wages in the bottom half of the wage distribution, we estimate regression eq. (2) using observations over the entire wage distribution and use the estimated coefficients—with the exception of the estimated location effects $\ln \delta_r$ —to predict individual wages. We then estimate our new $\ln \delta_r$ from regression eq. (2) on the bottom half of the distribution of predicted wages.

in the bottom eight deciles to regions in the top two deciles. We also run the main specification excluding London from the set of control regions.¹¹

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 the one using only regions in the bottom two deciles as treated. 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 (e.g., 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 lead to imprecise estimates. This is because our sample size is too small to obtain a precise enough estimate of the wage premia.

Effects on 16–64-year-olds.—All of the estimates we illustrated so far are based on the sample of individuals aged 25–64, that is, the age group that the NLW legally applies to. Yet there are reasons to believe that people under the age of 25 could be affected 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 to fairness concerns (Giupponi and Machin 2023). As this would effectively represent an increase in labor cost for those under 25, one might see impacts on employment in that age group as a result. Alternatively, to the extent that the NLW makes those under 25 cheaper to employ than older workers, labor 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 affected by minimum wage policy.

To jointly capture this wide range of factors, we apply our approach to examine effects on individuals aged 16–64. Figure A3 reports estimates of employment changes around the NLW as a share of the pretreatment employment rate among those 16–64, using local wage premia from regression equation (2) also estimated on the 16–64 population.¹² Our estimates of the fall in employment below the NLW (5.36%; standard error, 0.21%) and of

¹¹ 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 one. Therefore, this robustness check is mainly informative for the elasticity it delivers.

¹² Wages are deflated using time trends estimated on the 16–64-year-old population too. See n. 6 for details.

the total employment effect up to £5 above the NLW (−0.06%; standard 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 (standard error, 0.34). These results suggest that the wages of those under 25 were positively affected 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 2023). Our estimates also suggest that the overall employment effect of the NLW for those under 25 was either broadly neutral or positive.

D. Heterogeneity Analysis

In this section, we study heterogeneity in the effects that the minimum wage has on wages and employment by gender and age. To this end, we use equation (10) to estimate the effect of the NLW on the frequency distribution of wages for each subgroup, normalized to the pretreatment employment rate of that subgroup. Panel B of table 2 shows the estimates separately by gender. The fall in employment below the NLW is more pronounced for women than for men, as is the corresponding rise at (or just above) it. This is expected given that women are more likely to be receiving low wages. The point estimate of the total employment change up to £5 of the NLW is slightly positive for men (0.22%; standard error, 0.18%) and slightly negative for

Table 2
Impact of the Minimum Wage on the Wage Distribution: Heterogeneous Effects

	Δb (%)		$\Delta a + \Delta b$ (%)		Elasticity	
	Est (1)	SE (2)	Est (3)	SE (4)	Est (5)	SE (6)
A. Main Specification						
Main specification	−5.44	.22	−.11	.16	−.20	.32
B. Gender						
Women	−7.39	.31	−.47	.23	−.69	.38
Men	−3.55	.19	.22	.18	.61	.48
C. Age Group						
Young (25–34)	−6.40	.28	−.66	.56	−1.54	2.89
Middle age (35–54)	−4.77	.25	−.07	.19	−.15	.52
Old (55–64)	−6.59	.43	−.25	.72	−.35	.99

NOTE.—This 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 demographic groups over which regression eq. (10) is estimated. The frequency distribution of wages for each demographic subgroup is normalized to the pretreatment employment rate of that subgroup. Estimates are averaged over the four minimum wage increases from 2015 to 2019. Columns 1, 3, and 5 report our central estimates; cols. 2, 4, and 6 report bootstrapped standard errors. See sec. III.B for further details on these statistics. Panel A reports baseline estimates from fig. 3. Panel B shows heterogeneity by gender, and panel C shows that by age group.

women (-0.47% ; standard error, 0.23%). The negative effect for women is just statistically significant at the 95% level. Results by different age groups among the ≥ 25 population are shown in panel C of table 2. Estimated effects on employment for 35–54- and 55–64-year-olds are small and not statistically significant. Effects are somewhat more negative but highly imprecisely estimated for 25–34-year-olds.

V. Effects on the Household Income Distribution: Methodology and Data

A. 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 labor market outcomes, is a central policy question yet a difficult one to answer (Dube 2019b). 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 benefit system. Hours of work determine how a change in wages will translate into a change in earnings, although the relationship is complicated by the fact that minimum wage increases may itself cause changes in hours worked and in the likelihood of remaining employed. Moreover, the impact of the minimum wage on the household income distribution is sensitive to whom individuals affected by the minimum wage live with, for two reasons. First, households with more affected earners will be more affected by changes to wages than households with only one. Second, 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 United Kingdom has an individually assessed system of income and earnings taxation and a system of cash transfers that 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 that, in the United Kingdom, are really just cash transfers by another name. While the United Kingdom 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 United States, where the distributional impacts of minimum wages on net incomes have been studied previously (Dube 2019b). This context is important 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; conversely, taxes rise when earnings increase. Thus, the tax and benefit system shapes 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 disemployment

effects will be partially mitigated by tax decreases and benefit increases. Figure A4 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 increases than for workers in higher net household income deciles.

Another important factor determining the distributional consequences of the NLW is the mapping from the individual wage distribution to the household income distribution, which is influenced by hours of work, the incomes of other household members, and the tax and benefit system. We illustrate this relationship in figure A5. 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). While 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 more than 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 fig. A6).¹³

Previous studies have often used simulation approaches in order to estimate the impacts of minimum wage increases on household incomes (Sabia and Burkhauser 2010; Brewer and Agostini 2017). 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 microsimulation tool is used to account for interactions between earnings and the tax and transfer system, arriving at a more accurate estimate of impacts on net (i.e., after taxes and transfers) income. This is particularly important in institutional settings, where income-related transfers, especially those for working households, are widespread, as in the United Kingdom.

Microsimulation 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

¹³ Working households are defined as households with at least one member with positive earnings from employment.

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 2019b). 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 noncompliance 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 captured only the mechanical effects of minimum wage increases or have had to introduce further assumptions 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 in app. E), this means we can address all of 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 disemployment 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 (v) we use a tax-benefit microsimulator to calculate net household incomes. We describe these steps in more detail below.¹⁴

¹⁴ A different approach in the spirit of Dube (2019b) would directly assess the impact of the minimum wage on household incomes by comparing households in high- and low-paid regions. We do not do this for two reasons. First, the data we use on household incomes are too small a sample to allow us to split them into

B. Data and Sample Construction

Our main data source is the FRS, an annual cross-sectional survey of around 20,000 households that forms the basis of the United Kingdom'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 2016 prices.¹⁵ The national minimum wage was constant over that period, at the same level observed in the 2015 ASHE data used as the 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.

C. 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 few workers at precisely the minimum wage, just as one would expect if there is measurement error in the derived hourly wage (Skinner et al. 2002). Figure E1 reports the hourly wage distribution in the FRS (October 2014 to September 2015) and ASHE (April 2015). A comparison of the two distributions highlights the presence of measurement error in the FRS. In appendix E, we provide a detailed discussion of the challenges that measurement error in hourly wages poses for our simulation and the approach we adopt to correct for it.

D. Simulating Impacts on Household Income

Imputing employment effects.—Our central estimates from the frequency distribution analysis implied a small, although not statistically significant, disemployment effect. To simulate the effect of this, we randomly select the applicable fraction of workers who earn at or below the new minimum wage and set their earnings to zero. This assumes that a worker who would have earned just £0.01 less than the new minimum wage is as likely to lose their job as a worker who would have been on the prereform minimum wage.

fine geographic areas. Second, even if we had access to a larger sample, our approach has some distinctive merits relative to the direct approach. In practice, a large number of factors affect benefit entitlements and thus household incomes. Since minimum wages will have only a relatively small impact on most households' income, small changes in the tax code could bias the estimates substantially. Differential trends across areas in family composition or housing costs could have large effects on household income estimates while having a limited impact on our labor market estimates.

¹⁵ To be consistent with the frequency distribution analysis, we use τ_{2016} from the specification in n. 6 to uprate earnings. For other financial variables we use official price indices, such as average rents.

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 spill-over effects (causing some wages to increase beyond the NLW). The first step is to calculate the postpolicy cumulative distribution function of wages that 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.¹⁶ 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-reranking assumption, meaning that we assume 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’s. 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-described 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 microsimulation to account for the knock-on effects of earnings changes on taxes paid and transfers received, accounting for all of the relevant demographic and economic characteristics of the household. We do this using TAXBEN, the IFS tax-benefit microsimulator, which is the most detailed microsimulation model of the UK tax-transfer system (Waters 2017). We use the parameters of the 2016–17 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

¹⁶ 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_i from regression eq. (10) by the ratio of the overall absolute employment effect to the overall relative employment effect. See “Calculating the employment elasticity” in sec. III.B for details on how we 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. In addition, 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 plus £0.25. We base the distribution of the latter on the observed hourly wages in the bin around the October 2014 minimum wage in the base data.

appear to have been entitled based on their characteristics, we assume that they continue not to take up that benefit in our simulation.¹⁷ In a relatively small number of cases, households gain entitlement to a transfer as a result of the simulated impacts of the minimum wage on labor 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.¹⁸ We then randomly assign take-up using that household-specific probability.¹⁹

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 United Kingdom as a “benefit unit” (or more commonly in the United States, 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.

VI. Effects on the Household Income Distribution: Results

We now turn to the effects of the NLW on household incomes. We simulate the impact on net household incomes of the employment and wage effects of a £1 increase in the NLW, based on the employment and wage effects estimated from our baseline frequency distribution specification for those aged 25–64 shown in figure 3.²⁰

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.²¹

¹⁷ The exception is that we assume full take-up of child benefits, since child benefit take-up rates are more than 95%.

¹⁸ We classify families by four categories: couples and singles, with and without children.

¹⁹ A caveat is that self-reported take-up in survey data tends to imply lower overall benefit 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 ran the analysis under the assumption of full take-up. The key conclusions are unchanged.

²⁰ The FRS data we use cover the period when the NMW was £6.50. We simulate an increase to £7.50.

²¹ We assign households to income deciles based on their household equivalized income, using the Organization for Economic Cooperation and Development (OECD)–modified equivalence scale, before the introduction of the NLW. We compute changes in household incomes, with the household as the unit of analysis.




FIG. 4.—Impact of the minimum wage on household incomes: decomposition by income source. The graph reports the simulated distributional effect of a £1 increase in the NLW on household income. The vertical bars separately show the cash effect on net income (income after taxes and benefits in dark gray) and net tax payments (taxes minus benefits in light gray). 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 on the basis of pre-NLW income in this sample. Income is equalized and net of taxes and benefits.

The bars separately show the effect on net income (income after taxes and benefits) and net tax payments (taxes minus benefits). The two together sum 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 (rightward-most estimates).

On average, a £1 NLW increase raises net household incomes across households with someone aged between 25 and 64 by 0.31%. Around a third of the increase in pretax earnings is offset by reduced income-related benefits or higher taxes. Those clawbacks are even higher, reaching almost half of the total increase in pretax 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.²² This limits the

²² The bottom decile includes a significant number of households that are not in receipt of benefits—perhaps because they are not entitled by virtue of having

effect of the NLW on poorer households' income. The proportional effect of the NLW is broadly flat across the bottom half or so of the distribution. 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 of that seen in the second.

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 £30.68 per week, not accounting for spillovers or disemployment effects. After accounting for taxes and reductions in means-tested benefits, this leads to an increase in net household income of £21.60 per week. Average household income among minimum wage workers is £585 per week, meaning that the increase in income resulting from the minimum wage is 3.7% on average. But only 3% of working-age households contain a minimum wage worker. Even in the third and fourth deciles, where minimum wage workers are most common, only 5% 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, 4% have a minimum wage worker. The highest concentration of minimum wage workers is in the lowest decile, where the number rises to 8.5%. As can be seen in figure 5, the poorest 30% of working households each see proportional gains of around 1%. Effects then steadily decline as one moves further up the distribution.

One of the arguments for the NLW applying only to workers aged 25 and over was that it would be better targeted to minimum wage workers in poor households. Indeed, a teenager paid the minimum is more likely to be in a richer household than an older worker also paid at the minimum. We inspect this argument by running a mechanical simulation in which we raise the wages of all employees aged 16–64 up to the NLW and compare it to the same mechanical simulation for those 25–64. The results reported in figure A7 show that the gains in income in the 16–64 scenario are almost twice as large as in the 25–64 one, as one would expect. In addition, it does seem that the 16–64 scenario is less progressive than the 25–64 one. For example, the gains to the second and third deciles are around 70% to 75%

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.




FIG. 5.—Impact of the minimum wage on household incomes of working households: decomposition by income source. The graph reports the simulated distributional effect of a £1 increase in the NLW on household income. The vertical bars separately show the cash effect on net income (income after taxes and benefits in dark gray) and net tax payments (taxes minus benefits in light gray). 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 on the basis of pre-NLW income in this sample. Income is equivalized and net of taxes and benefits.

higher, whereas gains in the middle are 100% higher and gains for the ninth and tenth deciles are more than 200% higher. The 25-year-old age restriction thus seems to improve the distributional targeting of the policy.

Thus far we have been analyzing effects at the household level. An alternative approach is to assume that families or benefit units are the unit of income sharing and analyze effects at the family level, as discussed in section V.D. 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 income. Figure A8 shows the average effect of NLW increases on family incomes (among all families, not just those in work). In cash terms, the patterns are weaker than 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.




FIG. 6.—Impact of the minimum wage on household incomes: decomposition by source of response. The graph reports the simulated distributional effect of a £1 increase in the NLW 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 plus spillovers effect accounts for changes in the wage distribution as a result of the NLW, stripping out disemployment effects. The total change incorporates the full set of effects as estimated in figure 3, that is, the mechanical effect, spillovers, and disemployment effects. The graph also shows proportional impacts across all households in the sample (rightward-most crosses). The sample includes households with at least one person aged 25–64. Households are ranked on the basis of pre-NLW income in this sample. Income is equivalized and net of taxes and benefits.

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 V.D, except rather than using our estimated impacts on employment in each wage bin, we simply move those observed earning under the NLW up to the NLW. To incorporate spillovers—but not disemployment effects—we impute a postpolicy frequency distribution of wages based on the prepolicy distribution and the parameter estimates from section IV. However, we make a simple adjustment to those parameter estimates to purge them of the implied disemployment effects and isolate only the marginal effect of the spillovers. Namely, we add the estimated number of displaced workers back into the wage bin that starts at the level of the postpolicy NLW. This number is found by summing over all employment

changes up to £5 above the NLW. We then incorporate the disemployment effect to recover the total estimated effect shown in figure 4.

The spillovers are estimated to have a large effect, approximately half of the direct mechanical change in most deciles. Spillover effects are larger in the bottom six deciles of the distribution, reflecting the fact that workers benefiting from spillover are mostly located in those deciles and higher-wage earners tend to be in higher-income households. This amplifies the distributional effect for middle-income households. The disemployment effects have reasonably similar effects across the distribution, although they are slightly bigger toward the bottom, where more workers are directly affected by the NLW and so are at risk of job loss.

Comparison with evidence on impacts on household income in the United States.—A comparison of the distributional impacts of minimum wages in the United States and the United Kingdom reveals that minimum wage policies have a much more progressive impact in the United States than in the United Kingdom. Our evidence for the United Kingdom indicates that the most significant gains from minimum wage increases go to the middle of the working-age household income distribution, certainly in cash and to some extent in percentage terms (fig. 4). This is different from what has been found for the United States. Using US Current Population Survey data from 1984 to 2013, Dube (2019b) documents substantial and statistically significant positive effects of minimum wage increases on family income after taxes and transfers for percentiles between the seventh and twentieth, declining sharply to around zero by the thirty-fifth percentile.²³ In appendix F, 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 summarize the evidence here and refer the reader to appendix F for more details.

A first reason for the observed differences is that in the United Kingdom, minimum wage workers tend to be concentrated in the middle of the household income distribution, while in the United States, they are predominantly located toward the bottom of the distribution (fig. F1). This seems to be explained by the fact that because of transfers and other sources of household income, there is less of a correspondence between household earnings and household income in the United Kingdom than in the United States (fig. F2).²⁴ A second reason for the discrepancy is that in the United Kingdom (but not the United States), minimum wage workers at the bottom of the income distribution are less likely to gain from minimum wage increases because they work fewer hours (fig. F3) and face higher marginal tax rates (fig. F4) than

²³ See figs. 5 and 6 in Dube (2019b).

²⁴ 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 United Kingdom than in the United States.

those higher up the income distribution. Finally, individuals at the bottom of the income distribution in the United Kingdom get less of their income from work (fig. F5); a much higher share of their total earnings comes from self-employment, which is not covered by the minimum wage (fig. F6).

VII. Conclusion

We have examined the effects that the introduction of the United Kingdom's NLW has had on wages, employment, and household income, covering the period between the introduction of the NLW and the last prepandemic uprating—that is, 2015–19. 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 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 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 for not only employment and spillover effects onto those with higher wages but also 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.

We find that the NLW and its increases up to 2019 had substantial effects on wages toward 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 £2 per hour above it (approximately the twentieth 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 increases 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 prepolicy workforce in lower-wage regions relative to high-wage regions, with a 95% confidence interval spanning -0.5% to $+0.2\%$. Hence, we can rule out large effects with high confidence. The finding of small, negative, and statistically insignificant employment effects

is consistent across alternative specifications. There is some evidence of more negative impacts on employment of women than men. Those under 25 years old were also affected, with large positive spillover effects onto their wages.

Looking at the effects of the minimum wage increases on household incomes, we find that the biggest cash gains go to the middle of the working-age household income distribution. In proportional terms, a roughly similar effect is felt in the bottom half of the distribution, with impacts fading out in the top half. If we look only at households working before the introduction of the NLW, however, both the cash and the proportional impacts are more progressive, with the largest proportional increases at the bottom and steadily declining effects above that. Our results also quantify the distributional impacts from spillovers and disemployment effects. We show that such effects—especially spillovers—play an important role in shaping the distributional implications of the minimum wage.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67, no. 2:251–333.
- Ahlfeldt, Gabriel M., Duncan Roth, and Tobias Seidel. 2018. The regional effects of Germany's national minimum wage. *Economics Letters* 172:127–30.
- 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, no. 1:58–99.
- Belman, Dale, and Paul J. Wolfson. 2014. *What does the minimum wage do?* Kalamazoo, MI: W. E. Upjohn Institute for Employment Research.
- Brewer, Mike, and Paola De Agostini. 2017. *The National Minimum Wage, the National Living Wage and the tax and benefit system*. Colchester: Institute of Social and Economic Research.
- Brewer, Mike, Thomas Crossley, and Federico Zilio. 2019. What do we really know about the employment effects of the UK's National Minimum Wage? IFS Working Paper no. W19/14, Institute for Fiscal Studies, London.
- 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.
- Card, David. 1992. Using regional variation in wages to measure the effects of the federal minimum wage. *ILR Review* 46, no. 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, no. 4:772–93.

- . 2000. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania: Reply. *American Economic Review* 90, no. 5:1397–420.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics* 134, no. 3:1405–54.
- Clemens, Jeffrey, and Michael Wither. 2019. The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers. *Journal of Public Economics* 170:53–67.
- Corlett, Adam. 2021. Improving our understanding of UK poverty will require better data. London: Resolution Foundation Spotlight.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach. *Econometrica* 64, no. 5:1001–44.
- 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.
- 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, no. 1:78–106.
- Dube, Arindrajit. 2019a. Impacts of minimum wages: Review of the international evidence. London: HM Treasury.
- . 2019b. Minimum wages and the distribution of family incomes. *American Economic Journal: Applied Economics* 11, no. 4:268–304.
- Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. Minimum wage effects across state borders: Estimates using contiguous counties. *Review of Economics and Statistics* 92, no. 4:945–64.
- Dube, Arindrajit, and Attila Lindner. 2021. City limits: What do local-area minimum wages do? *Journal of Economic Perspectives* 35, no. 1:27–50.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge. 2022. Reallocation effects of the minimum wage. *Quarterly Journal of Economics* 137, no. 1:267–328.
- Giupponi, Giulia, and Stephen Machin. 2023. Company wage policy in a low-wage labor market. Working paper.
- Harasztsi, Peter, and Attila Lindner. 2019. Who pays for the minimum wage? *American Economic Review* 109, no. 8:2693–727.
- Harkness, Susan E., and Silvia Avram. 2019. The impact of minimum wage upratings on wage growth and the wage distribution: A report prepared for the Low Pay Commission. Low Pay Commission, London.
- Lee, David S. 1999. Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage? *Quarterly Journal of Economics* 114, no. 3:977–1023.

- Manning, Alan. 2016. The elusive employment effect of the minimum wage. LSE CEP Discussion Paper no. dp1428, Centre for Economic Performance, London.
- . 2021. The elusive employment effect of the minimum wage. *Journal of Economic Perspectives* 35, no. 1:3–26.
- 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, no. 3:608–48.
- Neumark, David, and William L. Wascher. 2008. *Minimum wages*. Cambridge, MA: MIT Press.
- 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, no. 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, no. 3:592–623.
- Schmitz, Sebastian. 2019. The effects of Germany's statutory minimum wage on employment and welfare dependency. *German Economic Review* 20, no. 3:330–55.
- 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, suppl.: 653–76.
- Stewart, Mark B. 2002. Estimating the impact of the minimum wage using geographical wage variation. *Oxford Bulletin of Economics and Statistics* 64, suppl.: 583–605.
- . 2004. The employment effects of the National Minimum Wage. *Economic Journal* 114, no. 494:C110–C116.
- Waters, Tom. 2017. TAXBEN: The IFS tax and benefit microsimulation model. IFS User Guides, Institute for Fiscal Studies, London.