

**The Impact of the Introduction of the UK Minimum Wage on the
Employment Probabilities of Low Wage Workers ***

Mark B. Stewart
University of Warwick

June 2003

* The author is grateful to the ESRC (under award R000223440) and the Low Pay Commission for financial support, and to two anonymous referees, Wiji Arulampalam, Orley Ashenfelter, Richard Dickens, Andy Dickerson, Peter Elias, Francis Green, Mary Gregory, Dan Hamermesh, Clare Kelly, Ewa Kmietowicz, Francis Kramarz, Alan Krueger, Gauthier Lanot, Alan Manning, Pedro Martins, Robin Naylor, Ian Walker and participants at the CEP Conference on Minimum Wages in the 21st Century, the Royal Economics Society Annual Conference, the Applied Econometrics Association Conference on Wage Econometrics, The European Association of Labour Economists Annual Conference, and the Warwick Labour Economics Workshop for helpful comments. The LFS and BHPS data were provided via the Data Archive. The NES data were provided by the Office for National Statistics.

Address for correspondence: Economics Department, University of Warwick, Coventry CV4 7AL, UK.
E-mail: Mark.Stewart@warwick.ac.uk.

The Impact of the Introduction of the UK Minimum Wage on the Employment Probabilities of Low Wage Workers

Abstract

This paper uses longitudinal data from three contrasting datasets (matched Labour Force Surveys, the British Household Panel Survey and matched New Earnings Surveys) to estimate the impact of the introduction of the UK minimum wage (in April 1999) on the probability of subsequent employment among those whose wages would have needed to be raised to comply with the minimum. A difference-in-differences estimator is used, based on position in the wage distribution. No significant adverse employment effects are found for any of the four demographic groups considered (adult and youth, men and women) or in any of the three datasets used.

Keywords: Minimum wage, employment determination, labour demand, difference-in-differences estimator.

JEL classifications: J38, J23.

Mark B. Stewart
Economics Department
University of Warwick
Coventry CV4 7AL
UK

Mark.Stewart@warwick.ac.uk

1. Introduction

The employment effect of a minimum wage remains one of the most fiercely contested policy questions in economics. In the standard textbook theoretical model of the labour market the introduction of a minimum wage leads perfectly competitive employers to cut employment. The magnitude of the aggregate reduction in employment then depends on the wage rises required to comply with the minimum and on the slope of the labour demand schedule at the relevant point. In contrast a range of monopsony, efficiency wage and search models have been suggested in which a decline in employment may not result and employment may even increase.¹

There is a vast literature on minimum wages, particularly on their effects and particularly for the United States. A recent review is given by Brown (1999). A consensus seemed to have emerged by the 1980s that the effect of minimum wages on employment in the United States was negative although probably fairly small (see for example Brown et al. (1982) for a review). Most of the evidence was based on time series estimation and much of it on teenage employment, where the effects were felt to be largest. Research findings in the 1990s have blown this consensus apart. On the one hand a growing body of research finds zero or positive employment effects (e.g. Card and Krueger (1994, 1995, 2000) for the US, the US results in Abowd et al. (2000) and Machin and Manning (1994) and Dickens et al. (1999) for the UK). On the other hand there is also a body of recent research that finds significant (both statistically and numerically) negative effects (e.g. Kim and Taylor (1995), Currie and Fallick (1996), Burkhauser et al. (2000), Neumark and Wascher (2000) and Neumark et al. (2000) for the US, the French results in Abowd et al. (2000) and Machin et al. (2002) for the UK). Thus the employment effect of minimum wages remains a highly contentious issue.

A new minimum wage was introduced in the UK in 1999 after a number of years with no minimum. The UK had statutory wage floors in many low wage sectors of the economy for most of the last century (although never an economy-wide minimum

¹ See Card and Krueger (1995, chapter 11), Dickens et al. (1999, section 2) and Brown (1999, section 2) *inter alia*.

wage).² The Wages Council system, introduced in 1909, reached its height in terms of coverage in the 1950s, was allowed to wither on the vine in the 1980s and was finally abolished in 1993. There then followed a period without any statutory minimum (except in agriculture) until the introduction of the new minimum wage in 1999 following a change of government. Based on the recommendations from the new Low Pay Commission (LPC, 1998), a minimum wage was introduced on 1 April 1999. The adult rate was set at £3.60 per hour, with a lower youth rate of £3.00 per hour for those aged 18-21 inclusive.³ The youth rate subsequently rose to £3.20/hour in June 2000 and the adult rate to £3.70/hour in October 2000.

There is obviously a need to estimate the impact of the introduction of the minimum wage in a number of dimensions, including employment, as part of the policy evaluation. In addition in the context of the international minimum wage debate the recent UK experience provides an important “quasi-experiment”. Since its introduction followed a period without any minimum, the UK case gives a cleaner test-bed than those for the US and France based on increases to an existing minimum. It allows direct examination of the crucial link between an individual’s position in the wage distribution and subsequent employment probabilities in the absence of a minimum wage and then examination of any post-intervention change in the relationship. As the Low Pay Commission (2000) point out, “this was a major intervention in the labour market”. They and the ONS have estimated that around 1.3 million employees (5.5%) were entitled to a wage increase as a result of its introduction (Low Pay Commission, 2001). It therefore provides the opportunity to investigate the effect on employment of significant wage increases for a large group of workers. Additional advantages of the test to be conducted are that the evidence on the introduction of the UK minimum wage suggests both a very high compliance rate and a lack of spillover effects onto the wages of higher paid workers.

This paper uses individual-level longitudinal data from a number of different sources to estimate the impact of the introduction of the minimum wage on the probability of

² See Low Pay Commission (1998, Appendix 5) and Metcalf (1999) on the history of UK wage floors.

³ There was also a development rate of £3.20 per hour for adults in the first 6 months of a new job with accredited training.

remaining in employment for low wage workers whose wages would have had to be raised to comply with the new minimum – the group directly affected. The paper estimates a model of the individual employment to non-employment transition probability as a function of the individual's initial position in the wage distribution, building on the models used by Linneman (1982) and more recently Abowd et al. (2000). The estimates in this paper refer to the economy as a whole and provide a useful companion to the more tightly focused estimates for the UK residential care homes sector provided by Machin et al. (2002). As they point out about their study, “one must be careful not to extrapolate from studies of one sector (especially the kind of sector we study) to conclusions about the economy as a whole”. This paper looks instead at estimates for the economy as a whole.

The next section lays out the identification and estimation strategy used to investigate the minimum wage impact on individual conditional employment probabilities. Three datasets are used in the empirical analysis: matched Labour Force Surveys, the British Household Panel Survey and the New Earnings Survey panel. Their advantages and disadvantages are discussed in Section 3. Before examining the impact on employment, it is necessary to demonstrate a differential impact on wages. This is done in Section 4. The results for employment are then presented in Section 5. Section 6 presents a test and further examination of the fundamental identifying assumption. Section 7 addresses the issue of measurement error in the wage variable. Further evaluation and examination of the robustness of the estimates is presented in Section 8 and finally conclusions in Section 9.

2. Estimation Strategy

This paper estimates the effects of the introduction of the minimum wage on the employment prospects of those affected. The central feature of the methodology employed is the use of individual-level longitudinal data to compare the employment experience of individual workers whose pay would have had to be increased to comply with the new minimum with that of a similar group who were not directly affected. The introduction of the minimum wage is viewed in this methodology as

what has come to be known as a “quasi-experiment” and a difference-in-differences estimator adopted to estimate its effect.⁴ The approach is an intuitively appealing one given the sharp change in wages brought about at the bottom of the wage distribution by the introduction of the minimum wage.

The starting point of the approach is that, other things equal, one would expect the group of workers whose wages had to be raised to comply with the new minimum (i.e. those initially below the minimum) to be more affected than a group from higher up the wage distribution. A direct comparison of the two groups will not be appropriate to identify any causal effect since, even in the absence of a minimum wage, those at the bottom of the wage distribution have lower subsequent employment probabilities.⁵ This makes the difference-in-differences approach a natural one to take. The difference between the two groups in a period affected by the minimum wage can be compared with the equivalent difference in an earlier period when no minimum wage was in place.

For those directly affected by the introduction of the minimum wage one wants to ask what their employment position would have been if the minimum wage had not been introduced. The objective is to find a suitable comparison to enable one to address a counterfactual question of this type. The approach used here defines the groups for comparison in terms of segments of the real wage distribution in the initial period.⁶ The “comparison” group used contains those slightly above the minimum and is designed to be above the minimum but as close as possible to the “treatment” group in wage terms to make their behaviour, including labour supply, as similar as possible.

⁴ See for example Meyer (1995) and Angrist and Krueger (1999) for a fuller discussion of this type of approach. Heckman and Robb (1985) provide an extensive discussion of estimators for “interventions” of this type. The approach has been used extensively in evaluating the impact of training programs, Ashenfelter (1978) being an influential early example.

⁵ See Stewart (2000) for evidence on this for Britain and an econometric analysis of the relationship.

⁶ This approach has been used, *inter alia*, by Linnerman (1982), Currie and Fallick (1996), Abowd *et al.* (2000) and Neumark *et al.* (2000). Studies using this approach and this type of control group are reviewed by Card and Krueger (1995, Chapter 7, pp. 223-231) and Brown (1999, section 7, pp. 2139-2142). A control group from further up the wage distribution is also used at the establishment level by Katz and Krueger (1992) and Card and Krueger (1994).

Define e_{it} to be the employment status of individual i in time period t ($= 1$ if employed, $= 0$ if not employed). Suppose that a minimum wage is introduced at a point in time, t^* , and that for observations prior to t^* no minimum wage is in place. The objective is to estimate the counterfactual employment probabilities after t^* as if the minimum wage had not been introduced. Consider two groups of individuals. Suppose that those in group $g = 1$ (the “treatment” group) are directly affected by the introduction of the minimum wage because their wages are initially below the minimum, while those in group $g = 2$ (the “comparison” group) have wages already slightly above the minimum. The following linear model could then be specified for the conditional probability of remaining employed:

$$\Pr[e_{it+1} = 1 \mid e_{it} = 1; g, t] = \alpha_g + \gamma_t + \theta g_{1it} d_{t+1} \quad (1)$$

where the first component is group-specific and fixed over time, the second component captures macro effects and is common across groups, g_1 is a dummy variable for the treatment group and d_{t+1} is a dummy variable taking the value 1 if the new minimum wage was in place at $t+1$. This is a simple difference-in-differences estimator given by double differencing sample means. Consider the case of two time periods: before (period 1) and after (period 2) the introduction of the minimum wage. Then under specification (1), the OLS estimator of θ is given by differencing across these two groups and two time periods: $\{[P_{12} - P_{11}] - [P_{22} - P_{21}]\}$, where P_{gt} is the estimated conditional employment probability for those in group g in period t .

This specification assumes that in the absence of a minimum wage the difference in conditional employment probabilities between the two groups is the same in each time period, or equivalently that the growth in the conditional employment probability over time is the same for each group. This is the first key identifying assumption and will be returned to below (particularly in Section 6). The potential problem with this assumption is that even in the absence of the minimum wage introduction, employment transition rates may evolve differently in the different groups. The second key identifying assumption is that the minimum wage does not alter employment probabilities in the comparison group ($g = 2$). There are two threats to

this assumption.⁷ First, there may be wage spillovers. Those paid slightly above the minimum wage may receive a pay boost to preserve differentials. Second, there may be substitution between groups as a result of the minimum wage introduction. (These two potential effects would be expected to act in opposition to one another.)

The specification used in the paper modifies equation (1) in a number of ways. First, a vector of individual characteristics, x , is added to equation (1), to sweep up any differences in characteristics between the “treatment” and “comparison” groups that are not picked up by the additive group and time effects. Second it is extended to multiple wage groups. In the specification used in the paper, the first group ($g = 1$) contains those directly affected, i.e. those with real wage (adjusted to April 1999 terms) below the appropriate (age-specific) minimum. The second group ($g = 2$) is the “comparison” group and contains those between the minimum and some point slightly above the minimum. The remaining group covers the rest of the wage distribution.⁸ Third, the simple difference-in-differences estimator assumes that the relationship between $\Pr[e_{it+1} = 1 | e_{it} = 1]$ and the wage is captured by the wage group dummies. The actual relationship may not correspond to this step-function specification and a more flexible form may be required. A polynomial in the real wage is included among the variables in the x -vector (as used by Abowd et al., 2000) to capture this.⁹

Incorporating these generalizations above and adopting a logit model specification, the estimated model for the employment transition probability takes the form:

$$\Pr[e_{it+1} = 1 | e_{it} = 1] = \Lambda \{x_{it}'\beta + \alpha_1 g_{1it} + \alpha_3 g_{3it} + \gamma_0 d_{t+1} + \theta^* g_{1it} d_{t+1} + \phi g_{3it} d_{t+1} + \gamma_t\} \quad (2)$$

where $g_{1it} = 1$ if $w_{it} < m_i$ and $= 0$ else, w_{it} is the real wage of individual i in year t and m_i is the value of the minimum appropriate to individual i , $g_{3it} = 1$ if $w_{it} \geq m_i(1+c)$ and $= 0$ else, where the constant c defines the width of the comparison wage group, and Λ is the logit transformation. Thus group 2 [$m_i \leq w_{it} < m_i(1+c)$] acts as the comparison

⁷ Meyer (1995) discusses in more detail the likely “threats” to an identification strategy of this type.

⁸ This third group can also be further subdivided, see Neumark et al. (2000).

⁹ Neumark et al. (2000) employ a similar approach, but use straight line segments for the function of the real wage. The results from using this specification here are very similar to those presented.

group. The probability difference-in-differences is estimated as the marginal effect corresponding to θ^* .

The question addressed is whether an individual whose wage would have had to be increased to comply with the new minimum, has a higher probability of losing their job than a *comparable* person in the wage group just above the new minimum. The methodology is a natural one. It looks at changes in employment status spanning the introduction of the minimum wage compared with changes prior to its introduction and it looks at the difference in this difference between a group directly affected and a group not directly affected.

The difference-in-differences estimator uses a binary indicator for the “treatment” group (those initially paid below the incoming minimum wage), together with its interaction with the post-minimum indicator, to measure the impact on the employment probability. Implicitly any effect on an individual’s employment probability is assumed to be the same irrespective of how far below the minimum the person’s wage initially was, i.e. by how much it would have had to be increased to comply with the new minimum. Any effect is taken to be the same for someone who would have needed, say, a 50p increase in their hourly wage as someone just 5p below the minimum.

An alternative estimator uses a “wage gap” variable, as was done for example by Currie and Fallick (1996). This is constructed as the gap between the individual’s wage at time t and the relevant minimum wage in place at time $t+1$. More explicitly

$$gap_{it} = \begin{cases} m_{t+1} - w_{it} & \text{if } m_{t+1} > w_{it} \\ 0 & \text{else} \end{cases}$$

The model used here is specified in the same way as for the difference-in-differences estimator, but with the dummy variable g_{lit} replaced by this “gap” variable. Equivalently this can be viewed as replacing g_{lit} by the product term $(m_{t+1} - w_{it})g_{lit}$ in equation (2) above.

The potential advantage of this estimation method relative to the difference-in-differences estimator is that precision may be gained by distinguishing between those whose wage needed to be raised a lot and those for whom only a small increase was required. Counter to this, the “gap” estimator may be more susceptible to the problems of measurement error at the very bottom of the distribution.

This “wage gap” variable can also be viewed as a (negative) linear spline term with node at £3.60. This might be taken to imply something of an internal inconsistency in this model, with the range of the wage distribution below £3.60 being represented by a spline, while that above £3.60 is modelled by dummies. This can be rectified by representing the entire wage range by a spline, with the set of nodes including m_{t+1} and $m_{t+1}(1+c)$ for comparability with the difference-in-differences estimator. Both versions of the “wage gap” estimator are considered below and give very similar estimates.

3. Data

Results for the models outlined above are presented for three different datasets: the Labour Force Survey (LFS), the British Household Panel Survey (BHPS) and the New Earnings Survey (NES). Each has advantages and disadvantages for the task at hand. Thus each can be regarded as providing important checks on the results produced by each of the other datasets.

Suitable datasets for the estimation of the model specified above require a number of features. First, they must have a panel or matched cross-section element of at least two time periods: the model estimates the probability of employment at time $t+1$ as a function of the wage (and other factors) at time t . Second, they must provide employment status information for the second time period (time $t+1$) and information on the individual’s hourly rate of pay at time t . Third, they need to be part of a series: there must be observations for which the time interval (t to $t+1$) straddles the introduction of the minimum wage in April 1999 and other observations for which the entire interval falls before April 1999.

Fourth, they must provide information on other factors that influence the conditional probability of an individual being employed at time $t+1$, given employed at time t , to construct suitable control variables. Finally, they must provide reasonably large samples of individuals. The construction of the difference-in-differences estimator requires adequate cell sizes of individuals with real hourly wage rates below the April 1999 minimum and of individuals in the “just above the minimum” control group, both in periods providing panel intervals that straddle April 1999 and in intervals entirely prior to April 1999.

The three datasets listed above meet this fairly demanding set of criteria to varying degrees and have contrasting strengths and weaknesses. Using the combination of all three different datasets therefore provides the broadest possible evaluation of the impact of the introduction of the minimum wage on employment in the context of the estimation strategy outlined above.

3.1 Matched Labour Force Survey data

The models outlined above can be estimated using matched LFS data. The LFS is a quarterly survey with individuals remaining in the sample for up to 5 quarterly waves. For estimation of the equation specified above, LFS data can be used only from 1997 quarter 1 onwards, when earnings questions were added to the wave 1 questionnaire (prior to this earnings questions had only been asked of the (outgoing) wave 5 respondents). The dataset constructed for the empirical analysis in this paper uses matched data from 13 quarterly LFSs and covers the period from March 1997 to March 2000. The estimation procedure uses wage and characteristic information from wave 1 and employment status information from wave 5 (12 months later). In the current context the advantages of the LFS are that it provides a better representation of low earnings workers than the NES and a much larger sample than the BHPS.

3.2 British Household Panel Survey data

Results are also presented based on waves 4 to 9 inclusive of the BHPS. Wave 4 was conducted in the autumn of 1994 and therefore is after the abolition of the Wages Councils. Wave 9 was conducted in the autumn of 1999 and therefore enables

examination of conditional employment probabilities after the introduction of the minimum wage. This six-year panel provides five years' worth of data to estimate the model for the probability of employment at wave $t+1$ conditional on position in the wage distribution (and other characteristics) at wave t . Of these matched waves, the last provides the probability of employment after the introduction of the minimum wage (specifically autumn 1999) given the pre-minimum wage (autumn 1998) and the other four give a control sample for periods when there was no wage floor. The analysis starts in wave 4 to permit this role as a control sample.

The main advantages of the BHPS are the proper panel nature of the data, the coverage of the complete earnings distribution (in contrast to the NES), and the consultation of individual's payslips wherever possible. The main disadvantage, relative to both the LFS and the NES, is the much smaller sample size.

3.3 New Earnings Survey panel data

The NES, conducted in April of each year, surveys employees with a particular final two digits to their National Insurance number so long as they are in employment, providing a 1% random sample of employees in employment. It can therefore be used to provide data to estimate the model outlined in section 2 above by matching people across years. Non-employment can be inferred from an individual's absence from the survey in a particular year, although not observed directly.

There are also other reasons why an individual may be missing from the survey in a particular year. Primary among these is that the NES excludes most of those whose weekly earnings falls below the PAYE deduction threshold (particularly in small organisations). Thus, for example, someone in a low-paying job who reduces their hours may as a result fall below the PAYE threshold and not appear in the survey in the subsequent year and be incorrectly classified as not in employment.¹⁰ Those already part-time (predominantly women) are likely to be most affected by this. The analysis conducted below therefore excludes women employed part-time at time t .

¹⁰ At the time of the 1999 survey the PAYE tax threshold was £83.37 per week. Thus someone on exactly the minimum wage of £3.60/hour would have to work 23 hours/week or fewer to fall below the PAYE tax threshold.

Other potential misclassifications result from the fact that there is a gap (of about a month) between the drawing of the sample for a particular year and the reference week in which the survey is conducted. Those unemployed or out of the labour force when the sample is located, but who enter employment before the survey date, will be excluded. Any employees who are with one employer when the sample is located and who have moved to another by the survey date may also be excluded if they cannot be traced.

The NES under-represents the low paid in year t as a result of these exclusions. Its great strengths are the likely accuracy of the wage data it provides, much of it direct from computerised payroll records, and the enormous samples, giving satisfactory cell sizes for the four key groups needed in the construction of the difference-in-differences estimator. As with the BHPS, and for the reasons discussed there, NES data for the years 1994 to 1999 inclusive are used to estimate the model. Since the survey is conducted in April of each year, the NES estimates will only capture immediate short-run effects.

4. The effect on wages

Evaluating the impact of the introduction of the minimum wage on employment by using a difference-in-differences estimator based on position in the wage distribution, as is done in this paper, requires that those in the directly affected group received a bigger wage boost than those above the minimum and that this difference shows up significantly in the wage variable used. Before applying the estimators outlined in Section 2 to models of the probability of subsequent employment, it is necessary to demonstrate that the introduction of the minimum wage caused wages in the directly affected group to rise significantly more than they would have done in its absence. Equivalent estimators to those to be used in the next section for employment are applied here to wage growth.

Looking first at the evidence from the matched LFS data, two alternative wage measures are used. The standard LFS wage variable is that recommended by the Office for National Statistics (ONS). It is constructed as gross pay last time paid (converted to a weekly basis) divided by the number of paid hours usually worked (both referring to main job only). In this construction, there is a potential mismatch between the hours measure used and the number of hours worked in the period covered by the reported gross pay. Information is also collected in the LFS on hours actually worked in the reference week (in the main job). The reference week does not necessarily correspond exactly to the pay period either. However it will do so for employees paid weekly at the end of each week. For this group using actual hours will provide a more accurate wage measure than that based on usual hours. Low paid employees, who are the focus of the estimation methods used in this paper, are more likely to be weekly paid and for them the wage based on actual hours is likely to be preferable to that based on usual hours. While the wage constructed using usual hours may provide a more accurate measure for those higher up the wage distribution, this suggests that the wage variable based on actual hours will be a more accurate measure for those with low wages, who are the focus of attention here. However since the wage based on usual hours is the ONS-recommended measure, the two are used in parallel in the results presented.¹¹

Before looking at the estimators to be applied in the next section, Figure 1 shows Nadaraya-Watson nonparametric kernel regression plots of the relationship between the percentage wage growth and the initial wage for observations entirely prior to the introduction of the minimum wage and for periods that span its introduction.¹² Separate regressions above and below the minimum are used to avoid smoothing across the minimum. Panel (a) plots the relationships for the wage variable based on “usual” hours. The difference in wage growth between “before” periods and periods spanning the introduction of the minimum wage is greatest for those with an initial wage below £3.60/hour. There is some evidence of a gap between the two kernel regression plots for those in the comparison group, but it is much less than that for

¹¹ For the construction of the “treatment” and “comparison” groups, both wage measures are converted to real terms (April 1999 prices) using the Retail Price Index (all items).

those below £3.60. Panel (b) of Figure 1 shows the corresponding kernel regression plots when the wage variable based on using “actual” hours is used. The conclusion is similar.

The significance of this difference between the group differences can be evaluated by applying the estimators outlined in Section 2, and used in the next section on the employment model, to an equation for the percentage wage growth (with the sample restricted to those who remain in employment). Results of this for all three data sources, and in the case of the LFS for both wage measures, are given in Table 1. All three datasets give rise to a small number of extreme outliers for wage growth.¹³ These regressions (except where stated otherwise) are therefore after suitable minor trimming.¹⁴ The first two columns of Table 1 give estimates for the two wage measures in the LFS. The remaining two columns give corresponding estimates for the NES and BHPS. The first row of the table provides raw difference-in-differences estimates for wage growth. All are significantly greater than zero, i.e. the difference between the annual rate of wage growth in the “treatment” and “comparison” groups is significantly higher for periods that straddle the introduction of the minimum wage than for periods that do not. The estimated wage effect is slightly larger for the LFS wage variable based on actual hours and the BHPS than for the other LFS wage variable and the NES, but all are of the same order of magnitude.

These raw difference-in-differences estimates distinguish between before and after the introduction of the minimum wage, but do not distinguish by time within these two phases. The next estimates include a complete set of month-by-month dummies (year dummies only in the case of the NES). The addition of these dummies causes only slight changes in the estimated effects in each of the datasets.

An alternative to trimming to take care of the extreme outliers in wage growth is to use robust regression on the full samples. These estimates are presented in the third

¹² This figure is restricted to the sample of adults, i.e. those to whom the minimum wage rate of £3.60 applies, to make clearer the positioning of the “treatment” and “comparison” groups.

¹³ For example in the LFS and using the first wage measure, the range is from observations where w_t is more than 350 times w_{t+1} to ones where w_{t+1} is more than 400 times w_t .

¹⁴ For each dataset the top and bottom 1% in wage growth terms are removed.

row of Table 1.¹⁵ The conclusions are similar. In all cases the estimated effect is significantly positive. The earlier findings are also strengthened by the fact that robust regression on the full samples gives very similar estimates to those on the trimmed samples. The final row of Table 1 gives “wage gap” estimates. Again the estimated effect is strongly significantly positive in all four columns.

5. Effects on the probability of remaining in employment

5.1 Estimates using matched LFS data

The sample is restricted to those aged at least 18 but below 60 who were employees at the time of the first wave interview, but excludes those who were full-time students or on government schemes. The dependent variable in the model to be estimated is the individual’s employment status at time $t+1$. The employment category is defined to include both employees and self-employed. The base group includes both the unemployed and those who have become economically inactive. The robustness of the results to variation in this and other variable definitions are considered in Section 8 below.

Table 2 presents the results for four demographic groups - male and female adults and youths - with the age division set at 22 to match that in the level of the minimum wage. Results are presented for both the wage variables discussed in the previous section. The first block of the table gives the raw (i.e. without control variables) linear difference-in-differences estimates. The corresponding absolute “robust” t -ratios from the regression are given in parentheses. The difference-in-differences estimate is insignificantly different from zero for both wage variables for all four demographic groups. Seven of the eight estimates are also positive, indicating a positive impact of the minimum wage on the probability of remaining in employment for the group directly affected.¹⁶

¹⁵ These estimates down-weight outliers. Iterative estimation is used with weights based on absolute residuals. First Huber weights are used to convergence, then biweights to final convergence.

¹⁶ It should be noted that the estimates presented in this paper are of effects on the probability of employment at $t+1$ conditional on employment at t , i.e. of the probability of remaining in employment. They do not include any impact on flows from non-employment into employment.

A logit form is estimated for the full model with control variables. Controls are included for the variables listed at the foot of Table 2. Separate models are estimated for men and women and the main effects (and the intercept) are allowed to vary within these across the two age groups. The second block of Table 2 presents estimates of the “marginal effect” of the dummy variable of interest, derived from the logit coefficient estimate and evaluated at the sample proportion (or equivalently the sample means of the explanatory variables). These are probability difference-in-differences and can be interpreted as the effect of the introduction of the minimum wage on the probability of remaining in employment. The table gives estimates calculated by the standard partial derivative adjustment, scaling the coefficient by $p(1-p)$, where p is the sample proportion.¹⁷ The absolute value of the robust asymptotic t-ratio of the logit model coefficient is given in parentheses in each case.

The full model difference-in-differences estimates are insignificantly different from zero for both wage variables for all four demographic groups. For the wage variable based on actual hours, the difference-in-differences estimate is also positive for all four demographic groups. For the wage based on usual hours it is positive for three of the groups. Only for adult women and using this latter wage variable is there a negative (although insignificant) effect. The absolute t-ratio in this case is less than 1.

The third block of the table gives the corresponding “wage gap” estimates. The results presented are for the model with the binary “treatment” indicator replaced by the “gap” variable defined in Section 2. The additional use of a linear spline above the minimum produces similar estimates and the same qualitative conclusions. Looking first at the wage based on usual hours, the estimated effects in the logit model with control variables included are all smaller (with same sign) than the corresponding dummy variable (i.e. difference-in-differences) estimates. The absolute t-ratios are also smaller, with the exception of young men, and in all cases indicate insignificance. The overall conclusion is the same when the wage gap construction is

¹⁷ An alternative method evaluates the difference in predicted probabilities. This is given by $\Lambda(\theta^* + \Lambda^{-1}(p)) - p$, where θ^* is the appropriate logit coefficient in the specification above and Λ is the logit transformation. This gives very similar estimates to those in Table 2 in all cases.

used as when the dummy variable construction is used. When the wage based on actual hours is used, the estimated effect is insignificant for three of the demographic groups, but significantly positive for young men. With both wage measures there is no evidence of any significant adverse effect.

The interpretation of the estimated effect is of course slightly different when this variable is used. It now measures the effect on the employment probability per unit “gap” (i.e. per £ below the minimum). For adult women, for example, a wage “gap” of £1 (i.e. a wage that needed to be raised from £2.60 to £3.60 to comply) is estimated to have reduced the employment probability by half a percentage point (but insignificantly different from zero). Alternatively the estimates can be interpreted in terms of the implied elasticities. These are also given in the table.

5.2 British Household Panel Survey estimates

Table 3 presents the results from estimating the same models using data from alternative datasets. The first alternative uses the BHPS. As far as possible the specification adopted is equivalent to that used in the analysis of the LFS data in Section 4.1 above. The sample is restricted to original sample members (and excludes the ECHP sub-sample) who provided full interviews and who were aged at least 18 but below 60 at time t . It is restricted to those who were employees at t , but excludes full-time students and those on government training schemes. The dependent variable, the individual’s employment status at time $t+1$, is constructed on the basis of whether or not the individual had a job in the week before the interview. The base group includes both the unemployed and those who were economically inactive. The wage variable is usual gross pay (converted to a weekly basis) divided by paid hours usually worked “in a normal week” and converted to real terms. (Both pay and hours are for the individual’s main job only.)

The raw (i.e. without control variables) linear difference-in-differences estimates based on the BHPS are all insignificantly different from zero for each of the four demographic groups, as in those based on the LFS data. They are also positive for both male groups and for young women. The estimated effect for adult women is negative (although insignificantly different from zero) and of a similar magnitude to

that for the LFS data. As in the LFS analysis described in the previous section, a logit form is estimated for the full model with control variables. Similar control variables are also used as far as possible and are listed in the notes to Table 3. The logit coefficient estimate is converted to a “marginal effect”, i.e. a probability difference-in-differences, as in the LFS analysis, and the absolute value of the “robust” t-ratio of the logit coefficient is given in parentheses below the difference-in-differences estimate.

For both male and female youth groups the samples for the below-minimum group and/or the comparison group post-introduction are either empty or too small once observations with missing values for control variables are excluded. The with-controls difference-in-differences estimates are therefore given only for the two adult groups. They are positive but insignificantly different from zero for both these groups. For women the difference-in-differences estimate switches from negative to positive when the control variables are added, in contrast to the results for the LFS data, although the estimates are insignificantly different from zero in both cases. “Wage gap” estimates are also given in Table 3. They too are insignificantly different from zero for both groups. In summary, there is no evidence of significant negative employment effects in the BHPS data.

5.3 New Earnings Survey estimates

Table 3 also presents the results from estimating the same models using data from the NES. As far as possible the specification adopted is equivalent to that used in the analysis of the LFS and BHPS data. However the NES provides much less information on individual characteristics than the other two datasets and hence the model is able to control for far fewer control variables. The sample is restricted to those aged at least 18 but below 60 and excludes those who held more than one job and those who worked less than one hour in total in the reference week for year t .

The wage variable is the basic hourly rate, constructed as total gross weekly earnings for the reference week (converted to a weekly basis if necessary) less overtime earnings, divided by normal basic hours (defined to be “basic hours for the employee in a normal week, excluding meal breaks and overtime”) and converted to real terms. Despite the enormous sample size, the NES raw difference-in-differences estimates

are all insignificantly different from zero for each of the four demographic groups, as in those based on the LFS and BHPS datasets.

As in the LFS and BHPS analyses, a logit form is estimated for the full model with control variables. Those used are listed in the notes to Table 3. As with the LFS and BHPS analyses, the logit coefficient estimate is converted to a “marginal effect”. The full model difference-in-differences estimates are insignificantly different from zero for all four demographic groups, with an absolute t-ratio of less than 1 in all cases. They are also positive for the three groups other than adult women. The corresponding “wage gap” estimates are also given in Table 3. They are also insignificantly different from zero for all four demographic groups. The positive estimates for the two male groups and for young women are similar to the difference-in-differences estimates. The negative effect for adult women is much smaller (in absolute value) than the difference-in-differences estimate and with a much smaller absolute t-ratio.

6. Examination of the fundamental identifying assumption

To add to the credibility of the results presented in the previous section, it is important to examine the reasonableness of the underlying methodological assumptions made. A fundamental identifying assumption underlying the difference-in-differences approach is that there are no interactions between the wage group dummies and the time effects in the absence of the minimum wage. One of the strengths of the use of the UK introduction as a test bed is that it provides a period with no minimum wage in place. This allows the assumption to be tested directly by estimating the model on the pre-minimum period and testing for the presence of significant interactions between group and time variables.

To conduct the test, interactions are included between the full set of month-by-month dummy variables and the wage group dummies. The joint significance of the interactions between the monthly dummies and the treatment group indicator is then examined. The test statistics reported here are based on the LFS data. In the model

without control variables this test gives test statistics with p-values of 0.39 and 0.32 for men and women respectively using the wage variable based on usual hours and 0.21 and 0.47 using the wage variable based on actual hours. In the model with control variables added, the test statistics have p-values of 0.42 and 0.39 for men and women respectively using the wage variable based on usual hours and 0.18 and 0.43 using the wage variable based on actual hours. Overall the assumption of zero interaction terms in the absence of the minimum wage is well supported by the data.

It is also informative to take this a step further by looking at the individual month-by-month interaction coefficients. After estimation on the full time period used in Section 5, the post-minimum interaction coefficients can be assessed relative to the distribution of the pre-minimum coefficients. These are plotted in Figure 2 (slightly smoothed), standardized relative to their distribution in the phase prior to the introduction of the minimum wage. For the employment probability models for both men and women the post-minimum interaction coefficients do not look out of line with what we would expect on the basis of the distribution of pre-minimum coefficients and the hypothesis of no employment effects.¹⁸ A corresponding plot of the interaction coefficients in the parallel wage growth model (as described in Section 4) is also given in Figure 2 for comparison. The significant gap between pre- and post-minimum levels can be seen. Thus in employment terms there does not seem to be anything unusual about the months immediately following the introduction of the minimum wage compared with the months preceding it, and this is in contrast with the wage growth interaction terms which are higher post-minimum than pre-minimum.

7. Measurement error in the wage variable

A potential problem with the LFS and BHPS estimates presented above is that of measurement error in the constructed wage variable. This could lead to misclassification of individuals into wage groups and thereby to a dilution of the

¹⁸ An anonymous referee has pointed out that this can be viewed as an empirical Bayes approach.

estimated effect on employment. This section attempts to measure the magnitude of this potential bias and correct the estimates for it.

The case of measurement error in a binary explanatory variable is slightly different to the classical measurement error problem. It has been examined in detail by Aigner (1973) and Bollinger (1996) among others. Important differences from the classical case are that the measurement error (the difference between the measured and true values of the binary variable) does not have zero mean and is not uncorrelated with the true value. The results from the classical case therefore do not apply in their standard form. The results in the Aigner and Bollinger papers are however useful here.

To apply their results here, note that the raw difference-in-differences estimator can, in its simplest form (two wage groups and two time periods), be expressed as the differences between the slope coefficients from two regressions of the employment indicator on the single dummy variable g_t , one for each of the two time periods (“before” and “after”). Aigner and Bollinger both give expressions for the probability limit of the least squares estimator of the slope coefficient in such a regression in terms of the misclassification probabilities.¹⁹ The expressions are equivalent. Aigner’s is in terms of misclassification probabilities conditional on the measured value, Bollinger’s in terms of misclassification probabilities conditional on the true value. The probability limit of the least squares estimator is proportional to the true value of the slope coefficient (and the asymptotic bias is towards zero, as in the classical case). The expressions of Aigner and Bollinger can therefore be used to produce a consistent “corrected” estimator (by inversion) if consistent estimates of the conditional misclassification probabilities can be produced. This is the procedure used here.

From March 1999 onwards the LFS has included a direct question on the basic hourly rate (for those paid by the hour). This provides a far more accurate measure of the appropriate hourly rate than the constructed hourly wage measures described and used

¹⁹ Equation (11) in Aigner (1973) and equation (10) in Bollinger (1996), the latter with a slight misprint.

above. Unfortunately this starting date prevents the variable from being used for the estimation of the effect of the introduction of the minimum wage presented above, being only one month before the minimum was introduced. However comparison of this variable with the constructed wage used in the analysis in this paper can be used to estimate the extent of measurement error and the magnitude of the asymptotic bias in the difference-in-differences estimates of the effect on employment and then to construct “corrected” difference-in-differences estimates that are consistent.

The procedure adopted uses the March 1999 data (the only month for which responses to the direct question are available prior to the introduction of the minimum wage) to provide sample estimates of the two misclassification probabilities.²⁰ These probabilities are then assumed to apply in all periods used in the estimation, both before and after the introduction.²¹ Using this procedure the corrected raw difference-in-differences estimate for adult women is adjusted from the value of -0.006 in Table 1 to -0.009. Asymptotic standard errors are constructed using the “delta method” (as used in Bollinger, 1996). This gives an absolute asymptotic t-ratio that is unchanged from that in Table 1 at 0.32. The main qualitative conclusions are unaltered by correction of the estimator for measurement error.

Since the asymptotic bias in the uncorrected estimate is towards zero, the corrected raw difference-in-differences estimates for the other three demographic groups all remain positive and increase in value. For example for adult men the corrected estimate increases from 0.020 to 0.038 and the asymptotic t-ratio falls slightly from 0.68 to 0.65. If the misclassification rates used above are regarded as upper bounds, then the bounds on the true employment effect based on the derivations in Bollinger (1996) are equal to the estimates in Table 2 on one end and the corrected estimates in this and the previous paragraph on the other.

²⁰ Aigner’s expression for the probability limit is used, since it leads to an easier construction of standard errors. Thus misclassification probabilities conditional on the measured value are constructed. Note that the “corrected” estimate based on inverting Bollinger’s expression is identical.

²¹ A more complex procedure was also examined where the March 1999 data was used to estimate the relative frequencies of the wage groups using the direct hourly rate variable. Artificial groups were then constructed on the data from April 1999 onwards (i.e. after the introduction of the minimum wage)

8. Further evaluation and examination of the robustness of the estimates

To further add to the credibility of the results presented in Section 5, the robustness of the estimates to modifications in the underlying specification needs to be examined. As well as the fundamental assumption examined in the previous section, the difference-in-differences methodology also depends crucially on the one hand on the distinction between the periods before and after the intervention and on the other hand on the distinction between the “treatment” group and the untreated “comparison” group. Both are considered in more depth in various ways in this section. A range of other robustness checks is also presented.

8.1 The “before” and “after” distinction

The implementation of the difference-in-differences methodology in the previous section compares observations for periods which straddle the introduction of the minimum wage with observations for periods entirely prior to the introduction. It therefore assumes that there is a “pre-minimum wage” phase up to the end of March 1999 and a “post-minimum wage” phase from the beginning of April 1999, with a clear distinction between them. For example, individual observations of the employment status in March 1999 (as a function of the individual's real wage in March 1998) are prior to, and hence treated as unaffected by, the introduction. If instead a significant part of the necessary wage increases to comply with the minimum took place prior to the legally required implementation date of April 1st, then the “before” and “after” contrast would be weakened. The methodology requires that the wage adjustment took place at the legal due date for a sufficient number of employees.

The new direct question on the basic hourly rate of pay used in Section 7, since it is not available prior to March 1999, cannot be used as the wage measure in the model to estimate the effect of the introduction of the minimum wage. However since it provides a more accurate measure of the basic wage (for the hourly paid), it is useful for addressing the issue of the timing of implementation. Figure 3 shows the distribution of the basic hourly wage rate (in 10p bands) for hourly paid adult

to match these relative frequencies and the misclassification probabilities estimated using these. This procedure produced almost identical corrected estimates.

employees in the LFS in March 1999 (upper graph) and in April/May 1999 (lower graph). Caution should be exercised due to the much reduced samples. However, the lower end of the wage distribution shows a sharp change between the two months. The April/May 1999 distribution shows a pronounced “spike” at the adult minimum wage of £3.60, while that for March 1999 does not. The 10p band starting at £3.60 quadruples in relative frequency, with a commensurate fall for the bands below this.²²

This new direct question on the basic hourly rate of pay also allows an examination of the issue of compliance. Employment effects could also be absent due to non-compliance with the new minimum. The Low Pay Commission (2000) concluded that the array of evidence collected by them “suggests that the vast majority of employers are complying” with the new minimum (page 113), but are not specific in quantitative terms. Ashenfleter and Smith (1979) propose a simple measure of the compliance rate: the fraction of employees earning the minimum wage or less who earn exactly the minimum. Using this measure, the LFS basic hourly rate question gives a compliance rate for adults of 84% in May 1999 rising to 92% by November 1999.²³

Returning to the “before” and “after” distinction, an alternative approach to the issue is to exclude from the analysis a period of time immediately prior to the April 1st legal implementation date, creating a “neutral zone” between the pre- and post-introduction periods. The first row of each half of Table 4 gives the difference-in-differences estimates for the full model when employment status observations in March 1999 are excluded. The estimates change very little and all remain insignificantly different from zero. The next row excludes the 3-month period January – March 1999. Again the difference-in-differences estimates change very little. All are insignificant and the only negative effect, that for adult women using the wage based on usual hours, weakens slightly. Excluding one or three months to create a “neutral zone” between the “before” and “after” periods has little effect on the results.

²² The equivalent histogram for youths, not shown because based on a rather small sample, shows spikes at both the youth and adult minima. Relatively low usage of the youth minimum was also found by Katz and Krueger (1992) for the US fast-food industry.

²³ These calculations include an adjustment for those estimated to qualify for the “developmental rate” or for the “accommodation offset”, although both of these have very little effect on the estimated rate.

8.2 Possible anticipatory action

Another possible problem is that employers could have made employment responses before the introduction of the minimum in April 1999. As long as any response occurred in or after April 1998, it will be covered by the 12-month moving window that straddles the introduction of the minimum wage and so will be captured by the current methodology. Whether the results are sensitive to the exact positioning of this switch between the pre-introduction and straddling-introduction periods can be examined by moving this switch date back slightly. This is done in the next block of each half of Table 4 for 1-month and 3-month shifts. This alters the estimated effects very little. They remain insignificantly different from zero in all cases. For the one combination that produced a negative effect in the base specification (adult women and the wage based on usual hours), the estimated effect and coefficient's t-ratio both decline in absolute value.

8.3 Definition of the "comparison" group

The difference-in-differences methodology depends on there being a meaningful difference between the "treatment" group and the untreated "comparison" group. This distinction may be threatened by spill-overs or by measurement error in the wage variable. In fact Dickens and Manning (2001) suggest that "the spill-over effect has been non-existent" (page 11) for the UK experience (in contrast to evidence for the US experience). Their analysis finds that "the national minimum wage has had virtually no impact on the pay of workers not directly affected" (page 3). Stuttard and Jenkins (2001) conclude similarly. However it seems worthwhile investigating the sensitivity of the results to the definition of the comparison group.

There is a trade-off when selecting the comparison group. On the one hand, a fairly narrow definition is desirable to keep the comparison group as similar to the directly affected group (in terms of unobservables) as possible. Moving the comparison group further up the wage distribution makes the possibility of interactions more of a threat and increases the burden of adjustment through the real wage polynomial. On the other hand, widening the definition of the comparison group or moving it further away from the "treatment" group lessens the problem of misclassification due to

measurement error and makes the underlying wage rates less similar, giving the test more leverage. Moving the comparison group further up the wage distribution reduces the impact of the “threats” to the identification strategy discussed in Section 2: both wage spillovers and substitution between the groups become less likely. An additional benefit of widening the definition is that it increases the size of the comparison group and (other things equal) the precision of estimation.

Difference-in-differences estimates for a series of alternative comparison groups are given in the next block of each half of Table 4. In all cases the estimates remain insignificantly different from zero for both wage variables and all four demographic groups. Extending the upper limit of the comparison group from the minimum + 10% to the minimum + 15% or + 20% using the wage based on usual hours moves the difference-in-differences estimate for adult women slightly closer to zero and reduces its significance. The estimates for the other three groups remain positive. For the wage based on actual hours the estimate is negative for adult women, but small and insignificant. The other two comparison groups examined introduce a gap between the minimum wage and the lower limit of the comparison group. The estimated effect for adult men is now negative for three out of four cases, but never significantly so. The estimates for the other three demographic groups are very similar to above. The results based on the use of modified comparison groups support the conclusions drawn in Section 5 on the basis of the original definition.

8.4 Definition of the “treatment” group

Some of those in the “treatment” group required only a small wage rise to bring them up to the incoming minimum. One way to sharpen the experiment is to focus the “treatment” group on those well below the minimum prior to its introduction. (An alternative to this is to use the “wage gap” estimator as was done in Section 5.) This modification also addresses another potential problem related to the age profile of wages. Someone on £3.50 per hour six months before introduction, who was still on the upward part of the age profile, might be expected to have been earning £3.60 by April 1999 even in the absence of the minimum wage. The likelihood of this for someone in the “treatment” group is reduced by this re-definition of the “treatment” group.

One potential disadvantage of this modification is that it may reduce the similarity in unobservables between the “treatment” and “comparison” groups. Another is that cell sizes for the “treatment” group are reduced, which will reduce the precision of estimation. The results of reducing the “treatment” group in this way are given in the final block of each half of Table 4. Two reductions are considered: to £3.30 and to £3.00 (with equal reductions for youths). The “comparison” group is kept unchanged. The results change very little as a result of either of these modifications.

8.5 Other robustness checks on the LFS results

Table 5 presents the results of a series of other examinations of the robustness of the findings. The first issue addressed concerns overtime hours and the payment of an overtime premium. The LFS analysis in section 5, in common with much of the literature, uses an average hourly earnings measure – averaged across all hours. The 1999 BHPS is a useful source of information on overtime premia. Of those paid by the hour for whom overtime is available 40% receive no overtime premium, i.e. their rate of pay for overtime hours is equal to their basic rate. Of those who receive a premium, time-and-a-quarter and time-and-a-half are dominant modal values and the median premium is a quarter. Estimates on the LFS using a wage variable constructed assuming a premium of 25% for all overtime hours are given in Table 5. The difference-in-differences estimate for adult women and usual hours falls slightly (in absolute terms) and its statistical significance declines. The corresponding estimate for adult men is now negative, but still insignificant. The estimates based on actual hours all remain positive. Allowing for an overtime premium does not produce evidence of a significant employment effect.

The results are also robust to modification of the definition of employment used in the construction of the dependent variable. Unpaid family workers and those on government employment and training programmes are treated as employed in the basic specification. Table 5 gives the results if these groups are instead treated as non-employed. Very little changes as a result.

A very broad definition of non-employment is used in the basic specification, to include all those who are economically inactive (as well as those who are unemployed). This includes those who are not seeking work and would not accept it if offered. The three largest groups among the economically inactive in the data all fall into this latter category. They are those who do not wish to work because they are staying at home to look after a family, because they are retired, and because they are long-term sick or disabled. The next row of each half of Table 5 gives the results when this group (economically inactive, not seeking work and do not wish to work) are excluded – on the grounds that they left employment for other reasons. All estimates remain insignificantly different from zero.

Using an alternative price deflator to construct the real wage (RPI excluding mortgage interest payments in place of the all-items RPI) changes the estimates very little. Neither does trimming the wage distribution. Due to the method of construction, there are a small number of very small or large estimated wage rates. The first modification trims around 0.2% of individuals from each end of the distribution, the second modification trims 0.5%. Neither changes the difference-in-differences estimates by much for any of the four demographic groups or either wage variable.

Overstatement of hours can lead to some individuals being wrongly classified as below the minimum and dilute the effect on the “treatment” group. Another modification examined excludes those with reported usual hours per week greater than 60 (about 2% of the sample). This too has very little effect on the estimates. Nor does use of weights in the estimation to compensate for differential non-response.

The final modifications examined concern the use of proxy responses in the LFS. Of the sample used for estimation in the basic specification here, around a quarter are proxy responses. The vast majority of these are responses by the individual’s spouse or partner. ONS has developed “adjustments” for hourly earnings for proxy responses. Its “lower adjustment” increases hourly earnings by 1% for spouse/partner proxies and by 10% for other proxies. Its “upper adjustment” increases hourly earnings by 3% for spouse/partner proxies and by 13% for other proxies. These adjustments are rather large for the “other” proxies, whose responses are viewed as

much less reliable. The remaining two modifications presented in Table 5 exclude “other” proxy responses entirely and make the ONS lower and higher adjustments for the spouse/partner proxies.²⁴ The difference-in-differences estimates are very little changed by these modifications.

In summary, this range of modifications confirms the main findings on the basis of the basic specification. The difference-in-differences estimate of the effect of the introduction of the minimum wage on the probability of employment is insignificantly different from zero for all four demographic groups and both wage variables.

8.6 Robustness of the BHPS and NES estimates

A similar range of modifications was also examined for the other two datasets used above – the BHPS and the NES. The findings based on these two datasets are also robust to the various modifications considered. In the case of the BHPS, attention is restricted to the two adult groups for the same reasons as above. For both these groups the conclusions from the results given in Section 4 are robust to: (1) the use of an alternative comparison group, (2) restricting the treatment group to those well below the minimum, (3) allowing for an overtime premium, (4) using an alternative definition of employment, (5) using an alternative price index, (6) trimming the wage distribution, (7) excluding those with very long hours, (8) using sampling weights, and (9) adjusting the wage according to whether or not a payslip was seen by the interviewer.

Turning to the NES, some of the robustness checks conducted on the other datasets are not needed (e.g. an assumption about an overtime premium is not required since separate basic and overtime pay is recorded and a wage rate for basic hours is used in the model) and others are not possible (e.g. no weights are provided). The modifications considered were: (1) the use of an alternative comparison group, (2) restricting the treatment group to those well below the minimum, (3) using an alternative price index, and (4) trimming the wage distribution. For these modifications the conclusions from the results given in Section 4, in particular the

²⁴ Retaining the “other” proxies and making the corresponding ONS adjustments produces very similar results in the current context.

statistical insignificance of the difference-in-differences estimates for all four demographic groups, are robust.

9. Conclusions

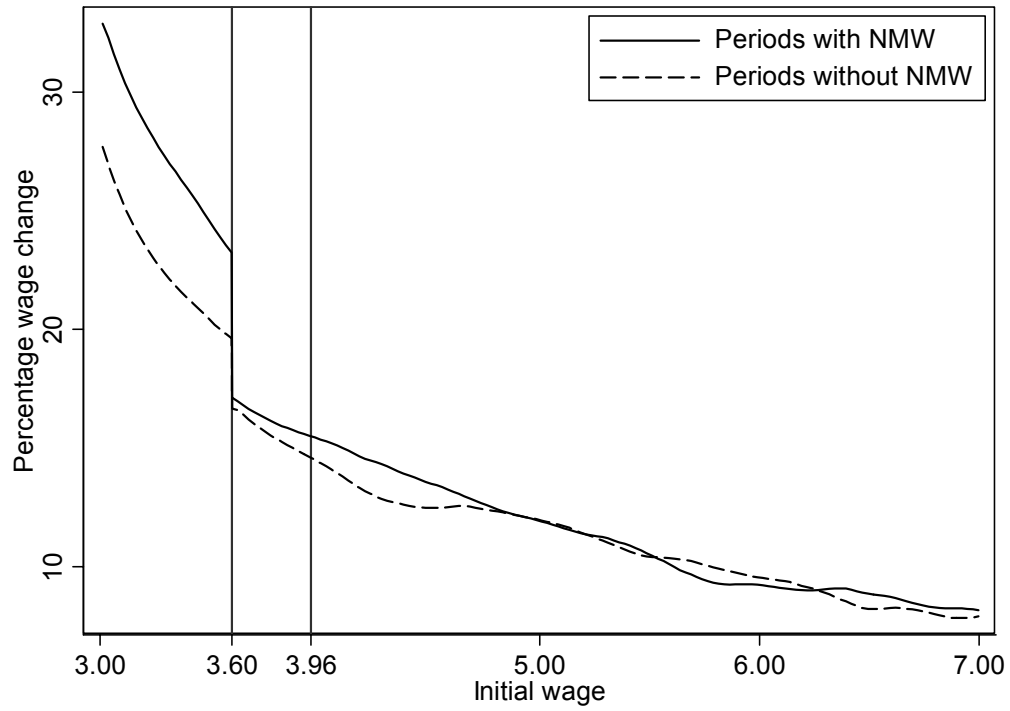
This paper uses individual-level longitudinal data from matched Labour Force Surveys, the British Household Panel Survey and the New Earnings Survey panel to estimate the impact of the introduction of the UK minimum wage in April 1999 on the conditional probability of subsequent employment among those whose wages would have to be raised to comply with the new minimum. A difference-in-differences estimator is employed using individuals from slightly higher up the wage distribution as the comparison group. The estimated impact of the introduction of the minimum wage on the probability of remaining in employment is insignificantly different from zero for all four demographic groups (male and female adults and youths) and all three datasets. This finding is robust to an extensive range of modifications considered.

The estimated effect is also found to be positive (although insignificant) for both male groups and for young women in all cases. The estimated effect is negative (although insignificant) for adult women in the LFS, but this depends on the construction of the hourly wage rate using “usual” hours, and is no longer the case if actual hours in the reference week are used. It is also negative for this group in the NES, although even less significant. In the BHPS the estimated effect is positive for this group (but insignificantly different from zero). These findings too are robust to the modifications considered. To sum up, the evidence presented in this paper indicates no significant adverse employment effects of the introduction of the UK minimum wage in any of the four demographic groups considered or in any of the three datasets examined.


References

- Abowd, John M., Francis Kramarz, David N. Margolis and Thomas Philippon (2000), "A Tail of Two Countries: Minimum Wages and Employment in France and the United States", mimeo, CREST, Paris, September.
- Aigner, Dennis J. (1973), "Regression with a binary independent variable subject to errors of observation", *Journal of Econometrics*, 1, 49-60.
- Angrist, Joshua and Alan Krueger (1999), "Empirical Strategies in Labor Economics", in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Volume 3 (Elsevier).
- Ashenfelter, Orley (1978), "Estimating the Effect of Training Programs on Earnings", *Review of Economics and Statistics*, 60, 47-57.
- Ashenfelter, Orley and Robert S. Smith (1979), "Compliance with the Minimum Wage Law", *Journal of Political Economy*, 87, 333-350.
- Bollinger, Christopher R. (1996), "Bounding mean regressions when a binary regressor is mismeasured", *Journal of Econometrics*, 73, 387-399.
- Brown, Charles (1999), "Minimum Wages, Employment and the Distribution of Income", in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Volume 3 (Elsevier).
- Brown, Charles, Curtis Gilroy and Andrew Kohen (1982), "The Effect of the Minimum Wage on Employment and Unemployment", *Journal of Economic Literature*, 20, 487-528.
- Burkhauser, Richard V., Kenneth A. Couch and David C. Wittenburg (2000), "A Reassessment of the New Economics of the Minimum Wage Literature with Monthly Data from the Current Population Survey", *Journal of Labor Economics*, 18, 653-680.
- 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*, 772-793.
- Card, David and Alan B. Krueger (1995), *Myth and Measurement: The New Economics of the Minimum Wage*, Princeton NJ: Princeton University Press.
- Card, David and Alan B. Krueger (2000), "Minimum Wages and Employment: A Case Study of the Fast-food Industry in New Jersey and Pennsylvania: Reply", *American Economic Review*, 90 (December), 1397-1420.
- Currie, Janet and Bruce C. Fallick (1996), "The Minimum Wage and the Employment of Youth: Evidence from the NLSY", *Journal of Human Resources*, 31, 404-428.
- Dickens, Richard, Stephen Machin and Alan Manning (1999), "The Effects of Minimum Wages on Employment: Theory and Evidence from Britain", *Journal of Labor Economics*, 17, 1-22.
- Dickens, Richard and Alan Manning (2001), "The national minimum wage and wage inequality", mimeo, Centre for Economic Performance, LSE, November.

- Heckman, James J. and R. Robb (1985), "Alternative Methods for Evaluating the Impact of Interventions" in J.J. Heckman and B. Singer (eds.), *Longitudinal Analysis of Labor Market Data*, Cambridge: Cambridge University Press.
- Katz, Lawrence F. and Alan B. Krueger (1992), "The effect of the minimum wage on the fast-food industry", *Industrial and Labor Relations Review*, 46, 6-21.
- Kim, Taeil and Lowell J. Taylor (1995), "The Employment Effect in Retail Trade of California's 1988 Minimum Wage Increase", *Journal of Business and Economic Statistics*, 13, 175-182.
- Linnerman, Peter (1982), "The Economic Impacts of Minimum Wage Laws: A New Look at an Old Question", *Journal of Political Economy*, 443-469.
- Low Pay Commission (1998), *The National Minimum Wage*, First Report, Cm. 3976, London: The Stationery Office.
- Low Pay Commission (2000), *The National Minimum Wage: The Story So Far*, Second Report, Cm. 4571, London: The Stationery Office.
- Low Pay Commission (2001), *The National Minimum Wage: Making a Difference: The Next Steps*, Third Report (volume 2), Cm. 5175, London: The Stationery Office.
- Machin, Stephen and Alan Manning (1994), "The Effects of Minimum Wages on Wage Dispersion and Employment: Evidence from UK Wages Councils", *Industrial and Labor Relations Review*, 47, 319-329.
- Machin, Stephen, Alan Manning and Lupin Rahman (2002), "Care Home Workers and the Introduction of the UK National Minimum Wage", mimeo, LSE. Forthcoming in *Journal of the European Economic Association*.
- Metcalf, David (1999), "The Low Pay Commission and the National Minimum Wage", *Economic Journal*, 109, F46-F66.
- Meyer, Bruce (1995), "Natural and Quasi-experiments in Economics", *Journal of Business and Economic Statistics*, 13, 151-161.
- Neumark, David, Mark Schweitzer and William Wascher (2000), "The Effects of Minimum Wages Throughout the Wage Distribution", NBER Working Paper 7519.
- Neumark, David and William Wascher (2000), "Minimum Wages and Employment: A Case Study of the Fast-food Industry in New Jersey and Pennsylvania: Comment", *American Economic Review*, 90 (December), 1362-1396.
- Stewart, Mark B. (2000), "The Inter-related Dynamics of Unemployment and Low Pay", paper presented at SOLE/EALE World Congress, Milan, and Working Paper, University of Warwick.
- Stuttard, Nigel and James Jenkins (2001), "Measuring low pay using the New Earnings Survey and the Labour Force Survey", *Labour Market Trends*, 55-66.



(a) Wage based on “usual” hours



(b) Wage based on “actual” hours

Figure 1: Kernel Regressions for Wage Growth (%): LFS

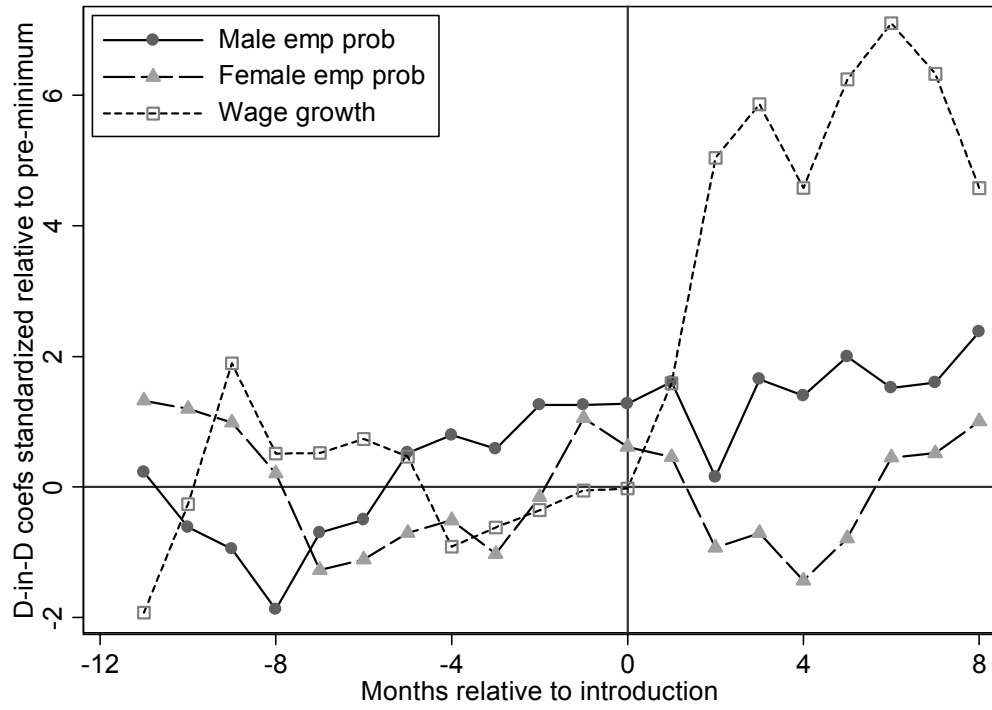


Figure 2: Plot of interaction coefficients

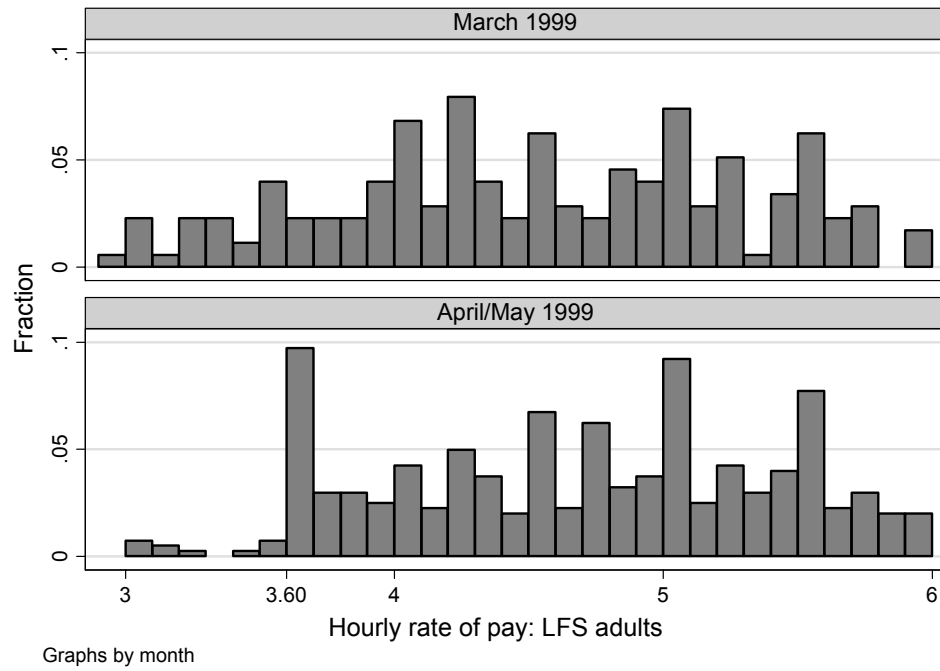


Figure 3: Basic hourly wage rate, hourly paid adult employees, LFS

Table 1
**Difference-in-differences and “wage gap” estimates of the effect of the
introduction of the minimum wage on wage growth.**

Basic specification: Matched Labour Force Survey data.

	LFS (W _U)	LFS (W _A)	NES	BHPS
Raw difference-in-differences	4.633 (2.68)	6.024 (2.95)	4.033 (4.92)	7.494 (2.53)
Full set of time dummies added	4.605 (2.66)	6.100 (2.99)	4.047 (4.93)	7.170 (2.45)
Robust regression	4.887 (3.86)	5.018 (3.00)	4.178 (10.11)	7.422 (3.26)
"Wage gap" estimator	6.899 (4.47)	8.837 (4.81)	10.012 (10.05)	9.062 (3.25)

Notes:

1. W_U = wage based on “usual” hours, W_A = wage based on “actual” hours.
2. Sample sizes: LFS: 44,076; NES: 450,319; BHPS: 14,047.
3. t-ratios in parentheses.

Table 2
**Difference-in-differences and “wage gap” estimates of the effect of the
introduction of the minimum wage on the probability of subsequent
employment.**

Basic specification: Matched Labour Force Survey data.

	Adult men	Young men	Adult women	Young women
Raw linear difference-in-differences estimates				
Wage based on usual hours	.020 (0.68)	.012 (0.15)	-.006 (0.32)	.067 (0.79)
Wage based on actual hours	.015 (0.49)	.059 (0.72)	.011 (0.60)	.079 (0.86)
Logit difference-in-differences estimates with controls				
Wage based on usual hours	.014 (0.93)	.073 (0.88)	-.010 (0.93)	.119 (1.15)
Wage based on actual hours	.011 (0.74)	.155 (1.36)	.006 (0.51)	.065 (0.78)
Logit “wage gap” estimates with controls				
Wage based on usual hours	.006 (0.77)	.058 (1.22)	-.005 (0.67)	.045 (1.00)
<i>Implied elasticity</i> Wage based on actual hours	.005 (0.83)	.047 (2.70)	-.003 (1.02)	.038 (0.11)
<i>Implied elasticity</i>	.006	.108	.006	-.004

Notes:

1. Based on 13 Quarterly Labour Force Surveys using data from March 1997 to March 2000. Sample size = 54,165, made up of 26,312 men (1,087 young) and 27,853 women (942 young).
2. Logit models contain as control variables: age completed full-time education, highest educational qualification dummies, labour market experience (quartic), length of tenure with current employer (quadratic), part-time, marital status, ethnic status, dummy for job at time t not permanent, public sector, health problem or disability limits kind of work can do, real hourly wage (cubic), regional dummies, year and month dummies.
3. Absolute values of "robust" t-ratios on model coefficients in parentheses.

Table 3

Difference-in-differences and “wage gap” estimates of the effect of the introduction of the minimum wage on the probability of subsequent employment.

Comparisons on LFS, BHPS and NES data.

	Adult men	Young men	Adult women	Young women
Logit difference-in-differences estimates with controls				
LFS estimates with wage based on usual hours	.014 (0.93)	.073 (0.88)	-.010 (0.93)	.119 (1.15)
LFS estimates with wage based on actual hours	.011 (0.74)	.155 (1.36)	.006 (0.51)	.065 (0.78)
BHPS estimates	.044 (1.29)	NA	.013 (0.43)	NA
NES estimates	.005 (0.32)	.001 (0.02)	-.011 (0.73)	.048 (0.79)
Logit “wage gap” estimates with controls				
LFS estimates with wage based on usual hours	.006 (0.77)	.058 (1.22)	-.005 (0.67)	.045 (1.00)
<i>Implied elasticity</i>	.005	.047	-.003	.038
LFS estimates with wage based on actual hours	.007 (0.83)	.138 (2.70)	.008 (1.02)	-.005 (0.11)
<i>Implied elasticity</i>	.006	.108	.006	-.004
BHPS estimates	.029 (1.29)	NA	.026 (1.27)	NA
<i>Implied elasticity</i>	.026		.024	
NES estimates	.004 (0.28)	.003 (0.10)	-.001 (0.10)	.042 (0.68)
<i>Implied elasticity</i>	.005	.004	-.001	.036

Notes:

1. LFS sample and control variables as in Table 2.
2. BHPS estimates based on 6 waves for 1994 (wave 4) to 1999 (wave 9) inclusive. Sample size = 16,796, made up of 7,851 men (561 young) and 8,945 women (518 young). Control variables: age completed full-time education, highest educational qualification dummies, labour market experience (quartic), length of time in current job (quadratic), part-time, marital status, ethnic status, dummy for job at time t not permanent, public sector, health limits kind of work can do, real hourly wage (cubic), regional dummies, year and month dummies.
3. NES estimates based on the New Earnings Surveys for 1994 to 1999 inclusive (April of each year). Sample size = 537,697, made up of 341,957 men (16,440 young) and 195,740 women (12,075 young). Control variables: age (quartic), dummy for more than 12 months in current job, part-time status, real hourly wage (cubic), regional dummies and year dummies.

Table 4
Modifications to basic specification: Difference-in-differences estimates:
Matched Labour Force Survey data.

	Adult men	Young men	Adult women	Young women
<u>Wage based on usual hours</u>				
<u>“Neutral zone” between “before” and “after” periods</u>				
Sample excludes March 1999	.017 (1.14)	.079 (0.93)	-.010 (0.88)	.119 (1.12)
Sample excludes January – March 1999	.022 (1.36)	.058 (0.64)	-.006 (0.51)	NA*
<u>Allow anticipation effects</u>				
Switch date moved back 1 month	.020 (1.30)	.082 (0.98)	-.009 (0.80)	.113 (1.10)
Switch date moved back 3 months	.025 (1.57)	.032 (0.36)	-.001 (0.09)	NA*
<u>Alternative comparison groups</u>				
Minimum to (minimum +15%)	.007 (0.54)	.077 (1.10)	-.008 (0.86)	.094 (1.14)
Minimum to (minimum +20%)	.002 (0.17)	.056 (0.92)	-.008 (0.92)	.112 (1.41)
(Minimum +5%) to (minimum +15%)	-.005 (0.37)	.084 (0.96)	-.005 (0.43)	.108 (0.98)
(Minimum +5%) to (minimum +20%)	-.007 (0.57)	.049 (0.73)	-.006 (0.43)	.108 (1.32)
<u>Only well below minimum in treatment group</u>				
< £3.30 (£2.70 for youths)	.020 (1.29)	.050 (0.57)	-.007 (0.63)	.163 (1.52)
< £3.00 (£2.40 for youths)	.017 (0.98)	.065 (0.70)	-.009 (0.69)	.151 (1.35)
<u>Wage based on actual hours</u>				
<u>“Neutral zone” between “before” and “after” periods</u>				
Sample excludes March 1999	.011 (0.78)	.162 (1.39)	.006 (0.55)	.068 (0.78)
Sample excludes January – March 1999	.015 (0.96)	.088 (0.78)	.008 (0.64)	NA*
<u>Allow anticipation effects</u>				
Switch date moved back 1 month	.012 (0.81)	.164 (1.43)	.006 (0.57)	.065 (0.77)
Switch date moved back 3 months	.016 (1.05)	.039 (0.33)	.007 (0.64)	NA*
<u>Alternative comparison groups</u>				
Minimum to (minimum +15%)	.010 (0.80)	.090 (0.98)	-.005 (0.46)	.050 (0.67)
Minimum to (minimum +20%)	.004 (0.33)	.038 (0.56)	-.004 (0.47)	.051 (0.71)
(Minimum +5%) to (minimum +15%)	.006 (0.41)	.061 (0.51)	-.010 (0.92)	.027 (0.31)
(Minimum +5%) to (minimum +20%)	-.001 (0.08)	.007 (0.10)	-.008 (0.81)	.037 (0.45)
<u>Only well below minimum in treatment group</u>				
< £3.30 (£2.70 for youths)	.015 (0.97)	.151 (1.30)	.010 (0.87)	.089 (1.02)
< £3.00 (£2.40 for youths)	.013 (0.80)	.179 (1.46)	.009 (0.69)	.064 (0.66)

Note: Sample & control variables as Table 2. * Small cells (& cell with employment rate of 1).

Table 5
Modifications to basic specification: Difference-in-differences estimates:
Matched Labour Force Survey data.

	Adult men	Young men	Adult women	Young women
<u>Wage based on usual hours</u>				
Assume overtime premium of 25% for all overtime hours	-.001 (0.04)	.063 (0.81)	-.008 (0.68)	.079 (0.74)
Treat unpaid family workers and those on government programmes as non-employed	.011 (0.75)	.107 (1.30)	-.011 (0.97)	.114 (1.08)
Exclude those who are inactive, not seeking and don't want work	.004 (0.39)	.034 (0.43)	-.009 (1.25)	.006 (0.07)
Wage deflated by RPI excl. mortgage interest payments	.023 (1.50)	.081 (0.95)	-.009 (0.79)	.152 (1.49)
Trim wage distribution at £0.50 & £48.00	.012 (0.81)	.051 (0.62)	-.010 (0.89)	.118 (1.13)
Trim wage distribution by 0.5% at top & bottom	.013 (0.87)	.047 (0.56)	-.009 (0.83)	.121 (1.16)
Exclude those with reported usual hours per week >60	.010 (0.65)	.073 (0.86)	-.008 (0.76)	.132 (1.28)
Use distributed sampling weights	.012 (0.81)	.074 (0.88)	-.010 (0.82)	.097 (0.93)
ONS "lower adjustment" to wage for spouse proxies + exclude "other" proxy responses	.019 (1.18)	.010 (0.07)	-.011 (0.94)	.094 (0.64)
ONS "upper adjustment" to wage for spouse proxies + exclude "other" proxy responses	.026 (1.57)	.040 (0.26)	-.010 (0.87)	.094 (0.64)
<u>Wage based on actual hours</u>				
Assume overtime premium of 25% for all overtime hours	.001 (0.02)	.133 (1.17)	.004 (0.33)	.126 (1.27)
Treat unpaid family workers and those on government programmes as non-employed	.011 (0.76)	.095 (0.88)	.007 (0.59)	.054 (0.62)
Exclude those who are inactive, not seeking and don't want work	.002 (0.15)	.111 (1.08)	.001 (0.17)	.042 (0.54)
Wage deflated by RPI excl. mortgage interest payments	.018 (1.25)	.171 (1.50)	.002 (0.19)	.073 (0.87)
Trim wage distribution at £0.50 & £48.00	.008 (0.58)	.123 (1.09)	.006 (0.55)	.063 (0.74)
Trim wage distribution by 0.5% at top & bottom	.008 (0.54)	.116 (1.02)	.006 (0.55)	.068 (0.80)
Exclude those with reported usual hours per week >60	.008 (0.56)	.154 (1.31)	.007 (0.58)	.077 (0.91)
Use distributed sampling weights	.011 (0.69)	.167 (1.47)	.010 (0.76)	.051 (0.55)
ONS "lower adjustment" to wage for spouse proxies + exclude "other" proxy responses	.012 (0.74)	NA	.005 (0.45)	-.021 (0.17)
ONS "upper adjustment" to wage for spouse proxies + exclude "other" proxy responses	.014 (0.89)	NA	.007 (0.63)	-.021 (0.17)

Note: Sample and control variables as in Table 2.