

**The Impact of the Introduction of the UK
Minimum Wage on the Employment
Probabilities of Low Wage Workers, Mark B.
Stewart, April 2001**

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

Mark B. Stewart
University of Warwick

April 2001

* Thanks to the Low Pay Commission and the ESRC for financial support, to Peter Elias and Mary Gregory for helpful advice on the NES data and to Ewa Kmietowicz for helpful comments. The LFS and BHPS data were provided via the Data Archive. The NES data were provided to the Institute for Employment Research 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 individual-level longitudinal data from three contrasting datasets (LFS, BHPS and NES) 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 had to be raised to comply with the new minimum. A “quasi-experiment” approach is adopted and a difference-in-differences estimator used based on position in the wage distribution. The evidence suggests zero, or if anything small positive, employment effects for adult men, young men and young women. The conclusion for adult women is slightly less clear cut. In the BHPS the effect is insignificant (and positive). In the LFS it is negative but insignificant and not robust. In the NES it is negative and approaching significance. However the t-ratio is weak for a sample size of over 300,000 and the case for a negative effect even for this group is not convincing.

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

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 research literature on minimum wages, particularly on their effects and particularly for the United States. A recent review of the literature 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 and the French results in Abowd et al. (2000)). 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

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

economy for most of the last century.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 in the UK 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 and a development rate of £3.20 per hour for adults in the first 6 months of a new job with accredited training. 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 can be viewed as providing an important “quasi-experiment”. Since its introduction followed a period without any minimum, the UK case 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. In addition, it is estimated that in excess of 1.5 million workers (6.4% of employees) were entitled to a wage increase as a result of its introduction (Low Pay Commission, 2000). As the Low Pay Commission point out, “this was a major intervention in the labour market”. This intervention therefore provides the opportunity to investigate the effect on employment of significant wage increases for a large group of 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 employment prospects of 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

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

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 next section lays out the identification and estimation strategy used to investigate the minimum wage impact on individual employment probabilities. Three datasets are used in the empirical analysis: the Labour Force Survey, the British Household Panel Survey and the New Earnings Survey. Their advantages and disadvantages are discussed in Section 3. Results for the basic model specified in Section 2, and for various modifications designed to examine the robustness of the findings, are presented in Section 4 based on the Labour Force Survey, Section 5 based on the British Household Panel Survey and Section 6 based on the New Earnings Survey. Section 7 presents conclusions from the analysis.

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.³ 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.

³ See for example Meyer (1995) and Angrist and Kreuger (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. Blundell et al. (1998) provide an illustration of this type of estimator in a different context.

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. 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 the question, what would their employment position have been if the minimum wage had not been introduced? The objective is to find a suitable comparison to enable one to address a “what if” question of this type.

To be more precise about the estimation method, define e_{0it} to be the employment status of individual i in time period t in the absence of a minimum wage (= 1 if employed, = 0 if not employed) and e_{1it} to be that in the presence of a minimum wage. Thus only one of these is ever observed for a given individual i in a given time period t . 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. Classify employees into a number of groups indexed by g . Then for a given group g in time period t there is direct information on the employment rate in the absence of a minimum wage, $E[e_{0it} | g, t]$, only for $t < t^*$ and direct information on the employment rate in the presence of a minimum wage, $E[e_{1it} | g, t]$, only for $t \geq t^*$. The objective is to estimate the counterfactual $E[e_{0it} | g, t, t \geq t^*]$, i.e. what the employment rate in group g would have been if the minimum wage had not been introduced. This is done using comparisons across g . Suppose that:

$$E[e_{0it} | g, t] = \alpha_g + \gamma_t$$

where the first component is fixed over time and the second component is common across groups. This assumes that in the absence of a minimum wage the difference in the employment rate between two groups is the same in each time period, or equivalently that the growth in employment over time is the same for each group. This is the key identifying assumption in the simple difference-in-differences estimator and will be returned to below.

Suppose initially that the minimum wage has a constant effect on the employment probability for those in group $g = 1$ and no effect on those in group $g = 2$:

$$\begin{aligned} E[e_{1it} | g = 1, t] &= E[e_{0it} | g = 1, t] + \theta \\ E[e_{1it} | g = 2, t] &= E[e_{0it} | g = 2, t]. \end{aligned}$$

Consider two time periods, t_1 when no minimum wage was in place and t_2 when one was. So $t_1 < t^* \leq t_2$. Then differencing across these two groups and across these two time periods gives θ :

$$\begin{aligned} &\{E[e_{it} | g = 1, t = t_2] - E[e_{it} | g = 2, t = t_2]\} \\ &- \{E[e_{it} | g = 1, t = t_1] - E[e_{it} | g = 2, t = t_1]\} = \theta \end{aligned}$$

Thus the simple, or raw, difference-in-differences estimator is got by double differencing sample means.

Alternatively, under the above assumptions, the employment status of individuals in all groups and all time periods can be written:

$$e_{it} = \alpha_g + \gamma_t + \theta D_{it} + \varepsilon_{it} \quad (1)$$

where $D_{it} = 1$ if individual i is affected by the minimum wage, i.e. if individual i is in group $g = 1$ and if $t \geq t^*$, $D_{it} = 0$ otherwise, and where $E[\varepsilon_{it} | g, t] = 0$. Thus the raw difference-in-differences estimator is also given by a regression using micro data pooled across groups and time periods with additive group and time dummies plus an

interaction term between the “ $g = 1$ ” dummy and another dummy variable for all time periods with the minimum wage in place. If there are more than two groups, additional interaction terms of this type for groups $g > 2$ will also be required.

This paper defines the groups indexed by g in terms of segments of the real wage distribution in the preceding period.⁴ 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.⁵

The simple difference-in-differences specification can be extended to produce a “regression adjusted” difference-in-differences estimator by adding a vector of individual characteristics, x_{it} , that are thought to affect the probability of employment as control variables to equation (1) to give:

$$e_{it} = x_{it}'\beta + \alpha_g + \gamma_t + \theta D_{it} + \varepsilon_{it} \quad (2)$$

The objective in adding these control variables to the equation is to sweep up any differences in characteristics between the “affected” or “treatment” group ($g = 1$) and the “comparison” or “control” group ($g = 2$) not picked up by the additive group and time effects.

The first key identifying assumption is still that interaction terms are zero in the absence of the minimum wage. This is now after controlling for differences in observable characteristics. The problem here 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 group $g = 2$. There are two threats to

⁴ This approach has been used, *inter alia*, by Linneman (1982), Currie and Fallick (1996), Abowd *et al.* (2000) and Neumark *et al.* (2000).

⁵ This third group can also be further subdivided, see Neumark *et al.* (2000).

this assumption.⁶ 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.) These possible threats need to be kept in mind when interpreting results from this approach and when the robustness of results to modifications in the definition of the “control” group are investigated.

The empirical model used analyses the probability of being employed in the subsequent period as a function of the individual’s wage group in the current period, and therefore for those currently employed. It is therefore a model of the employment transition probability. This requires a slight adjustment to the specification above. In addition since a probability is being analysed, there are certain well known disadvantages to the use of a linear specification. and a logit model is adopted. The estimated model takes the form:

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

where $g_{1it} = 1$ if $w_{it} < m_i$ and $= 0$ else, where 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 , where $g_{3it} = 1$ if $w_{it} \geq m_i(1+c)$ and $= 0$ else, where the constant c defines the width of the comparison wage group, d_{t+1} is a binary variable taking the value 1 if the new minimum wage was in place at time $t+1$, x_{it} is a vector of other factors that influence the probability of remaining in employment, γ_t are time effects for the remaining time periods and Λ is the logit transformation, i.e. the CDF of the logistic distribution, $\Lambda(z) = [1 + e^{-z}]^{-1}$. Thus group 2 [$m_i \leq w_{it} < m_i(1+c)$] acts as the comparison group and θ is still the parameter of interest giving us the difference-in-differences estimator. Thus 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.

⁶ Meyer (1995) lists and discusses in more detail the likely “threats” to an identification strategy of this

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. This type of model, with a control group from further up the wage distribution, has been used on individual-level data by several important studies in the minimum wages literature, including Linnerman (1982), Currie and Fallick (1996), Abowd et al. (2000) and Neumark et al. (2000).⁷

An important defect of the simple difference-in-differences approach in the current context is that those with lower pay have a lower probability of remaining employed even in the absence of a minimum wage. There is thus a relationship between $\Pr[e_{it+1} = 1 \mid e_{it} = 1]$ and w_{it} even before the introduction of a minimum wage.⁸ Abowd et al. (2000) control for this effect by including a polynomial in the real wage among the variables in x_{it} . Neumark et al. (2000) employ a similar model, but use straight line segments for the function of the real wage.

3. Data

Results for the model 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 matched cross-section or panel element of at least two time periods: the model estimates the probability of employment at time $t+1$ as a

type.

⁷ 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 Card and Krueger (1994).

⁸ See Stewart (2000) for evidence on this for Britain and an econometric analysis of the relationship.

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.

This is a fairly demanding set of criteria. The three datasets listed above meet them to varying degrees and with 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 UK Labour Force Survey data

The model 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 data from 11 quarterly LFSs: from 1997 quarter 1 (March – May) to 1999 quarter 3 (September – November). The estimates therefore capture employment effects up to 8 months after the introduction of the minimum wage. 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 much better representation of low earnings workers than the NES and that it provides a much larger sample than the BHPS. The LFS is also conducted quarterly, while the other two surveys are only conducted annually. The LFS also has a number of disadvantages. Compared with the BHPS, the LFS has only a limited panel dimension, makes extensive use of proxy responses, and does not consult respondent's payslips (something which is done in the BHPS). Compared with the NES, the LFS gives serious concerns about measurement error in the constructed hourly pay rates, particularly due to that in its hours variable.

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 is the latest available. It 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 estimates will capture employment effects up to 9 months after the introduction of the minimum wage.

The advantages of the BHPS are the proper panel nature of the data, the coverage of the complete earnings distribution (in contrast to the NES), the extensive range of information on individual characteristics relative to the NES, the lack of proxy respondents, 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. The additional disadvantage relative to the NES is the potentially greater measurement error problems with the wage variable.

3.3 New Earnings Survey panel data

The NES is conducted in April of each year. It surveys employees with a particular final two digits to their National Insurance number so long as they are in employment. It can therefore 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.

In practice 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 but working full-time who switches to a slightly higher paying part-time job 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. 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 lag between the drawing of the sample for a particular year and the reference week in which the survey is conducted. First, those unemployed or out of the labour force when the sample is located, but who enter employment before the survey date are excluded. Second, those with one employer when the sample is located who have moved to another by the survey date and cannot be traced are also excluded. Thus there is some uncertainty about an individual's exact economic status for years when they do not appear in the NES. Certain groups such as domestic service workers and piece-rate homeworkers are also excluded from the survey.

The NES also systematically under-represents the low paid in year t as a result of the exclusions just described. It also provides more limited information on personal characteristics than the other two datasets from which to construct control variables (since it is provided from employer records). 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.

The NES estimates will only capture immediate short-run effects, since the survey is conducted in April of each year and is currently only available up to 1999. In fact, the list of employees and their employers to whom NES questionnaires are to be sent is taken from Inland Revenue PAYE records in mid March (about a month before the survey pay period). Thus those who are not employees at this date will not be in the survey and will be classified as non-employed by the construction here. To this extent the NES estimates will capture instances of employers deciding to lay off employees before the April 1 start date rather than increase their pay rate to the minimum wage. Those in employment in mid March, but laid off before the survey date (which is defined to be the pay period that includes April 14) will also be classified as not in employment in the analysis presented below. (Most low-wage workers are paid weekly. For them the pay period will have been the week Monday 12 April to Friday 16 April.) The NES estimates will therefore also capture those who are laid off immediately after the introduction of the minimum wage. However this means that the NES estimates presented below will capture only very short-run effects.

4. Labour Force Survey results

4.1 Results for the basic model

The sample is restricted to those interviewed at both the first and fifth quarterly waves and who were aged at least 18 but below 60 at the time of the first wave interview. It is also restricted to those 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 (plus very small numbers of unpaid family workers and those on government employment and training programs). 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 4.2 below.

The wage variable is constructed as gross pay in the most recent pay period, converted to a weekly basis and then divided by hours per week. Both pay and hours are for the individual's main job only. The hours variable is total usual hours, including paid overtime and excluding any unpaid overtime. This construction is that recommended by the Office for National Statistics (ONS) and used by the Low Pay Commission in their first report (LPC, 1998). The wage is converted to real terms (April 1999 prices) using the Retail Price Index (all items). Potential modifications to various aspects of the construction of the wage variable are also considered in section 4.2 below.

The top half of Table 1 presents sample raw probabilities of being in employment in year $t+1$ (i.e. at the time of interview for wave 5) broken down by the wage group the individual was in at t (i.e. at the time of interview for wave 1, roughly 12 months previously) and by whether or not the minimum wage was introduced between these two dates. Employment probabilities are given 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. The first two rows of the table give these probabilities for those whose real wage (in April 1999 prices) at time t was below the April 1999 minimum wage appropriate to their age. The first of the two rows is for observations where this period straddled the introduction of the minimum wage, i.e. it gives employment probabilities in April 1999 or later for those with a real wage below the minimum one year previously. For this group wages would have needed to be raised to comply with the minimum wage and by more than the rate of inflation. The next row gives the corresponding probabilities for earlier periods (i.e. for periods with both observation points prior to the introduction of the minimum wage). The next row then

gives the differences in these probabilities between periods which saw the introduction of the minimum wage and periods which did not. For all four demographic groups the probability of being in employment at $t+1$ is smaller after the introduction of the minimum wage than for periods without a minimum.

The next three rows of the table present the equivalent statistics for those in the "control group", defined as those whose real wage (in April 1999 prices) at the time of interview at t was between the appropriate April 1999 minimum and the minimum plus 10%. For this group the probability of being in employment at $t+1$ is smaller after the introduction of the minimum wage than for periods without a minimum for men and for young women, but is higher for adult women.

The raw (i.e. without control variables) linear difference-in-differences estimator is then given by subtraction. The results are given in the next row of the table. The results can be equivalently obtained by a linear regression as outlined in Section 2. The corresponding absolute "robust" t-ratios from the regression are given in parentheses. The difference-in-differences estimate is insignificantly different from zero for all four demographic groups. For both adult men and young women the estimate is also positive, implying a positive impact of the minimum wage on the employment probability for the group directly affected.

It is worth pointing out, by way of additional explanation, that this raw difference-in-differences estimator can also be viewed (and calculated) the other way round, i.e. by considering differences in employment probabilities between those directly affected (i.e. with real wage below the minimum) and those in the "control group" and then taking the difference in this between the period that saw the introduction of the minimum wage and earlier periods that did not. Calculating it this way round, for the period when the minimum wage was introduced the difference in probability between the directly affected group and the "control group" was negative for all four demographic groups: -.021, -.043, -.041 and -.058. Similarly for the earlier periods without the introduction of the minimum wage the difference in probability was again negative for all four groups: -.036, -.035, -.022 and -.080. Taking the differences

between the corresponding probability differences in these two sets of numbers again gives the raw linear difference-in-differences estimates given in Table 1.

A logit form is estimated for the full model with control variables. Similar control variables to Abowd et al. (2000) are used, with some additions. Controls are included for age completed full-time education, a set of highest educational qualification indicators, labour market experience (a quartic), length of tenure with current employer (a quadratic), part-time status, marital status, ethnic status, an indicator if the job at time t is not permanent (fixed term contracts, seasonal, agency temping, or casual work), a public sector indicator, whether the individual has a health problem or disability (lasting more than 12 months) which limits the kind of paid work they can do, regional dummies, and year and month dummies. Following Abowd et al. (2000) a cubic in the real hourly wage is also included in the x -vector to control for the relationship in the absence of a minimum wage as discussed in Section 2 above. 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.

For the purpose of interpretation, the logit coefficient estimate is converted to a “marginal effect” of the dummy variable of interest evaluated at the sample proportion (or equivalently the sample means of the explanatory variables), i.e. a probability difference-in-differences. This can then be interpreted as the effect of the introduction of the minimum wage on the probability of employment. Two alternative methods are used to calculate this from the logit coefficient. The table gives those calculated by scaling the coefficient by $p(1-p)$, where p is the sample proportion. This is the standard partial derivative adjustment. In all cases the second method, which evaluates the difference in predicted probabilities, gives very similar estimates.⁹ The absolute value of the robust asymptotic t -ratio of the coefficient in the logit model is given in parentheses in each case.

The full model difference-in-differences estimates are insignificantly different from zero for all four demographic groups. For three of the groups the effect is positive.

⁹ The difference in predicted probabilities is given by $\Lambda(\alpha + \Lambda^{-1}(p)) - p$, where α is the appropriate logit coefficient in the specification above and Λ is the logit transformation: $\Lambda(z) = [1 + e^{-z}]^{-1}$.

For young men the effect, negative in raw probability terms, turns positive when the control variables are introduced. Only for adult women is there a negative (although insignificant) effect. The absolute t-ratio is also highest for this group (although only 1.49, giving a p-value of 0.14). However, the sample sizes involved should be kept in mind when evaluating these t-ratios. In very large samples it is often felt that standard 5% or 10% critical values are inappropriate and that the critical values used should rise with the sample size, i.e. we should be more demanding of our test statistics as the sample size increases. The Schwarz criterion¹⁰, sometimes used in this context, implies using a critical value for the female sample of 3.17. The t-ratio for adult women is less than half this value.

For adult women the effect is exactly the same as the raw difference-in-differences estimate, i.e. the introduction of the control variables makes no difference, and it is worth noting that the raw estimate is mainly driven by the increased employment probability after the introduction of the minimum wage for those in the "just above" control group. However not too much should be read into this given the magnitudes of the t-ratios involved. The bottom line is that for all four demographic groups the estimated impact of the introduction of the minimum wage on the probability of remaining in employment is insignificantly different from zero.

4.2 Robustness issues

The implementation of the basic model in the previous sub-section required a number of specification choices to be made, including particularly the detailed definitions of variables. This sub-section presents the results for various modifications to the basic specification to examine the robustness of the previous findings.

The first modification considered concerns the definition of the "slightly above the minimum" control group. In the basic specification, the definition used by Abowd et al. (2000), namely between the minimum and the minimum plus 10%, is adopted. This is quite a narrow range and, particularly for the youth samples, leads to quite

¹⁰ The Schwarz criterion involves comparing the "t-ratios" to a critical value of $\sqrt{[\ln(\text{sample size})]}$.

small cell sizes for the control group. However a fairly narrow definition is required to keep the control group as similar to the directly affected group as possible.¹¹ There is a trade-off here. On the one hand, widening the definition increases the size of the control group and (other things equal) the precision of estimation and also reduces the impact of the “threats” to the identification strategy described in Section 2. On the other hand it weakens the similarity and hence increases the burden of adjustment through the real wage polynomial.

An alternative of the minimum plus 15% was examined. The results for the full model (i.e. with control variables) are given in the second row of Table 2. (The corresponding results for the basic specification are repeated in the first row to facilitate comparison.) The difference-in-differences estimates of the impact of the introduction of the minimum wage alter very little as a result of this modification. The estimated effect is insignificantly different from zero for all four demographic groups and the numerical changes are small.

Another set of modifications investigated concerns the months immediately prior to the April 1999 introduction of the minimum wage. As part of its construction the difference-in-differences estimator compares observations for periods which straddle the introduction of the minimum wage with observations for periods entirely prior to the introduction. 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. However with firms adjusting low wages upwards in the run-up to the minimum wage's introduction, this may not be realistic. The next modification examines the impact of excluding observations shortly prior to the introduction of the minimum, creating a “neutral zone” between the pre- and post-policy-introduction periods.

The third row of Table 2 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

¹¹ As Meyer (1995) points out, “this design is most plausible when the untreated comparison group is very similar to the treatment group so that interactions are less likely”.