FEATURE

Ian Mulheirn
Formerly HM Treasury

The impact of the 2006 National Minimum Wage rise on employment

SUMMARY

The National Minimum Wage (NMW) has risen ahead of earnings since its introduction in 1999. In October 2006, with coverage at its highest level ever, the adult NMW increased by 5.9 per cent, to £5.35 per hour, the largest proportionate rise since 2004. While previous assessments have failed to find any clear evidence of a negative impact on employment, it is of interest to see whether more recent NMW rises have reduced employment.

This analysis evaluates the impact of the October 2006 rise in the NMW on the rate at which people leave employment. Results show no evidence of an increased job exit rate among those directly affected by the rise. Indeed, there is some indication of a positive effect on job retention for men. The analysis therefore provides no reason to think that the minimum wage rise in October 2006 caused employees to be made redundant.

Since 1999, when it was introduced at £3.60 per hour, the National Minimum Wage (NMW) for adults aged 22 and over has risen by around 49 per cent, to £5.35, in October 2006. Had it grown in line with average earnings, NMW would, by then, have been £4.88.[1]

Assessments of the impact of the introduction of the minimum wage have found no firm evidence that it had a negative effect on participation (Metcalf 2007). Since 1999 however, the NMW has risen as a proportion of mean hourly wages – a measure known as the 'bite' of the minimum wage – from around 45 per cent to over 50 per cent in 2006 (BERR 2006). It may be the case that any labour market impact is more likely to be seen as the bite increases, since this measure gives an estimate of the proportion of workers covered by the NMW (BERR 2006). As a result of the consistent increases in the NMW above average earnings growth over the past nine years, it is reasonable to ask whether the minimum wage is starting to have an adverse effect on employment. Despite being, at 5.9 per cent, only the fourth largest increase over the nine years since the introduction of the minimum wage, it is the largest recent increase by some margin. This historical context provides a compelling reason to evaluate the labour market impact of the 2006 increase in the NMW.

This analysis focuses on the impact of the 2006 minimum wage increase on the probability of leaving employment for those directly affected, and compares it with transitions made by those in a control group. The approach taken here draws on two papers: one by Dickens and Draca (2005), examining the employment effects of the October 2003 increase in the NMW; and another by Stewart (2003), evaluating the introduction, and 2000 and 2001 upratings of the NMW.

The results of this analysis show no evidence of an increased job exit rate among those directly affected by the rise. One explanation for this is that a higher minimum wage may be reducing staff turnover among those affected. Indeed, there is some indication of a positive effect on job retention for men. This analysis therefore provides no reason to think that the minimum wage rise in October 2006 caused employees to be made redundant.

Basic theory of the impact of a minimum wage increase

On the standard assumption of perfectly competitive labour markets, the theoretical story told about the impact of the minimum wage is simple. Where the minimum wage is set above the market clearing wage level, the effect is to increase the supply of labour and simultaneously to reduce demand for labour. The result is to lower employment levels and increase unemployment (which, in the perfectly competitive world, did not exist before the minimum wage was introduced).

The impact of a rise in the minimum wage is qualitatively the same and is represented in **Figure 1**. Increasing NMW from NMW_t to NMW_{t+1} reduces labour demand from Q_t to Q_{t+1} and increases labour supply by P_t to P_{t+1}. The result is that unemployment rises from U_t to U_{t+1}.

There are many reasons why this simple view of the world may not give the whole story. For example, labour is not only supplied in discrete units of employees, but also as an amount of time per worker. It is possible that a rise in the minimum wage could leave the number of employees unchanged while reducing the number of hours worked (Stewart and Swaffield 2008). Indeed, the relative effect of a higher minimum wage on variable costs (hourly wage rates) as opposed to fixed costs of employment (for example, hiring and training) implies that, other things being equal, firms might even employ more people for fewer hours per person than before.

In areas where a firm is the only employer in a local area, monopsony power can be held by employers over employees. Monopsony power is an example of where the perfect competition assumption is inadequate and under certain conditions the minimum wage can actually increase employment. For these reasons, the theoretical effect of a minimum wage increase on employment levels is ambiguous.[2]

Impact of the NMW on wages in 2006 and 2004

In order to be able to attribute changes in employment transitions to the minimum wage rise, there must be evidence of a clear impact of the NMW rise on the wages of some employees over the time period of its introduction. If this impact is absent or weak, there is almost nobody affected by the minimum wage rise, and hence no employment impact can be expected to result from it.

Figure 2 uses data from the five-quarter longitudinal Labour Force Survey (LFS) to show the number of employees reporting being paid an hourly wage below the incoming, higher NMW levels for both October 2006 and the earlier increase in October 2004. It shows a substantial reduction in the number of employees in jobs that pay below the incoming NMW, over the period of its introduction.

It is important to note that the chart shows only those people who report an actual hourly wage. It excludes many people who may be paid on a weekly or monthly basis, and therefore do not report their hourly wage, but who are, nonetheless, being paid less than the impending minimum. This means that the estimate of just under 150,000 people being paid below £5.35 in April 2006, for example, is much lower than the actual number of people in the country on wages below that level. The Office for National Statistics uses data from the Annual Survey of Hours and Earnings (ASHE) to estimate the number of jobs that pay below the NMW. In spring 2006, for example, 274,000 jobs filled by employees aged 22 and over paid below the 2005–06 NMW level. The reason for using only the directly reported hourly wage rates in the longitudinal LFS for this analysis is discussed further in **Box 1**.

From Figure 2 it is evident that the number of people reporting hourly pay rates below the impending NMW trended steadily down, from around 150,000 in September 2005 to 110,000 one year later. The pre-rise group (individuals whose employment status was observed in the first and second grey zones in the figure) clearly saw little or no impact of the impending rise on their wages.

In the two months after September 2006, the number dropped very sharply, by 65 per cent, to under 40,000. The group straddling the rise (whose employment status was observed in the second and third grey zones) clearly saw a substantial impact from the rise in a very short timescale. So while there may have been some employer anticipation of the NMW rise, there was clearly a substantial and sudden impact of the minimum wage rise on pay rates, making it amenable to non-experimental evaluation, as discussed below.

The October 2004 rise to £4.85 per hour represented a 7.7 per cent increase, while the increase to the minimum in October 2006, a 5.9 per cent rise. Figure 2 shows that the impact of the later rise was at least as large, in terms of the number of people directly affected, as that of the 2004 rise. This gives cause for concern that there may be an accompanying employment effect of the 2006 rise, despite its being more modest in proportionate terms than that of 2004.

Methodology

In assessing the impact of the minimum wage on employment, the group of interest is those who are directly affected by wage regulation. It is reasonable to assume that this consists of people who are paid (or who would be paid) less than the incoming minimum, since these are the people whose

Figure 1

A rise in the minimum wage under perfect competition

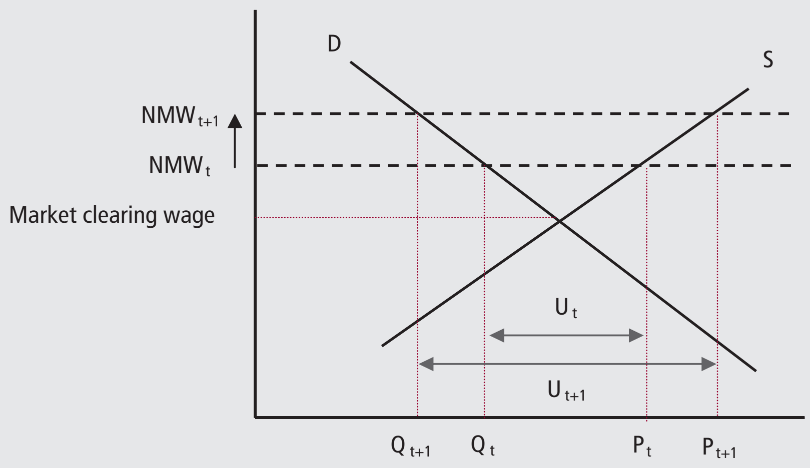

Figure 2

Number of employees in main jobs paid below new NMW rate: by months before/after rate increase

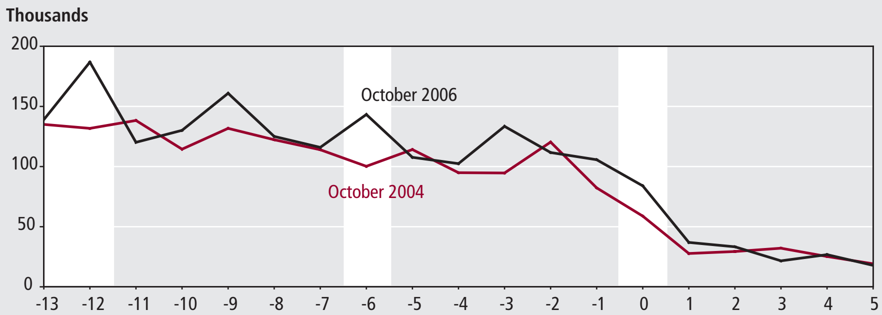

Box 1
Data

The analysis presented in this article makes use of the longitudinal LFS. The LFS is a large continuous quarterly survey of households, covering around 53,000 private households in the UK. The questionnaire covers a range of issues, including labour market behaviour and household characteristics. Participants are interviewed for five consecutive quarters. Tracking their responses over time in this way allows analysis of their transitions in and out of employment.

ASHE is the preferred source for earnings estimates. The data are considered more reliable than the LFS because they are provided by employers from the payroll. However, ASHE does not contain data on the broader range of individual characteristics covered by the LFS, which in this case allows a richer regression model for the evaluation. Secondly, and more importantly, ASHE is a panel data set that observes individuals once a year in April. This analysis compares the changes in employment status of individuals over a time period where no minimum wage increase took place, to those changes over a time period where such an increase did occur. Since the NMW has risen in each year following its introduction, the annual nature of the ASHE panel does not lend itself to this approach because it is impossible to find a comparison group of people who did not experience a minimum wage change. For this reason, the quarterly nature of the longitudinal LFS is more appropriate.

There are a few limitations to using the longitudinal LFS in this way. First, LFS earnings data are only collected in the first and fifth waves of interviewing, substantially reducing the sample to two-fifths of its normal size. This problem is compounded by the fact that many people choose not to answer income questions. Finally, high attrition rates (the gradual decrease in the workforce by natural wastage) across the waves further reduce the number of observations when using the longitudinal data sets for analysis.[3]

The extent to which sample attrition may distort the results of the analysis if certain groups systematically drop out of the survey while others do not needs to be considered. If, by contrast, observations are missing at random, then sample attrition will not adversely affect the results. It is difficult to check which of these situations prevails. However, looking at the distribution of characteristics in one representative quarter of the LFS and comparing it with respondents that stay in for five quarters used in this analysis gives some clues as to its reliability. On a range of variables such as the marital status, ethnicity, country of birth and qualifications of respondents, there does not appear to be any substantial difference in the composition of the two samples. This implies that the attrition rate in the panel sample does not jeopardise the results.

In the LFS, the derived variable ‘hourpay’ has been shown to be unreliable for measuring the extent of low pay (Dickens and Manning 2004) since it is calculated from reports of weekly gross income and usual hours, rather than being directly reported. There are good reasons, therefore, to believe that the directly reported hourly pay rate variable, ‘hrrate’, is more accurate than the derived wage variable for measuring the extent of low pay (Skinner *et al* 2002).

For this reason, and in line with the approach of other literature in this field (Dickens and Draca 2005), the analysis relies on the directly reported hourly pay variable. Using the variable ‘hrrate’ tends to limit the sample and exclude a large group of potentially affected people, since many people are not paid by the hour and therefore do not report an hourly rate.

At the low pay levels considered here, however, the need for accuracy in the wage variable, in order to prevent contamination of the treatment and control groups with people whose wages are inaccurately recorded, outweighs concerns over sample size. Aside from reducing the sample size, this does not pose a major problem for the evaluation, since capturing the full number of people paid below the incoming NMW is not necessary. Rather, changes in the probability of people leaving work, rather than the absolute numbers, reveal the impact of the NMW rise on employment.

Due to the ambiguous effects of the adult minimum on younger workers through substitution effects, the analysis here is conducted for adults aged 22 and over, only.

wages would need to rise in order to comply with the law. This group is referred to as the ‘treatment’ group.

One way to estimate whether the 2006 NMW rise had an impact on employment of the treatment group would be to examine employment rates for low-wage people, before and after the rise. There are two fundamental difficulties with this approach. The first relates to the problems with using an employment rate, and the second to the shortcomings of the simple before and after approach.

Employment transitions rather than employment rates

It is impossible to define an employment rate for people of a given wage level since one cannot know the wage level at which people out of work might join the labour market. One way around this problem is to look instead at the probability of people, initially employed (at time t) being out of work some time later (at time t+1). This is, in other words, the probability of making a transition out of employment. If this probability changes between a group of people who experienced the minimum wage rise between periods t and t+1, and a group for whom no minimum wage increase occurred in that time period, then that change might be attributable to the rise in NMW.

The time period between time t and t+1 must be less than a year, so that it is possible to identify a group whose employment transition probability was unaffected by the increase in the NMW. But the interval must also be large enough to take account of the fact that different employers may increase their wages at different times, over a period of months, around the time of the October NMW change. For this analysis, therefore, each group’s employment status is observed at six-month intervals.

For this analysis, two groups of people are identified:

- the ‘pre-rise’ group contains those who did not experience a minimum wage increase. It is constructed from people who were first interviewed between

November 2005 and March 2006. This group's employment status is then observed again six months later (at their third interview) between May and September 2006, just before the October rise becomes statutory
- the 'post-rise' group straddles the October NMW rise. This group's first wave interviews are conducted between May and September 2006, and their employment status is observed six months later between November 2006 and March 2007

Reported hourly pay rates from the longitudinal LFS (see Box 1) are used to assign people into the group affected by the NMW increase and the comparison group who were unaffected. Employment status observed in waves one and three for each individual is then used to create a variable to indicate each person's employment status after six months, given that they were employed in wave one. This employment transition indicator is the dependent variable in the regression model used. The variable takes a value of 1 if an individual was employed at time t and workless at time t+1, six months later, and 0 if they were employed at both t and t+1.

If employers raise wages to comply with the new minimum well in advance of the October deadline, those changes might affect the wages and employment chances of people classified as being interviewed in the pre-rise time period. It would therefore contaminate the 'before' group with wage changes expected to happen 'after' October 2006.

It is likely that some employer anticipation of the new rate does occur, particularly in large companies that settle their wage levels for all employees at the same time of year. To the extent that this happens, estimates of the probability of employment transition will be biased downwards. However, as Figure 2 shows, many employers only adjust their wages to comply with the new minimum close to the October deadline.

Difference-in-difference evaluation

The second problem is that, even using such a transition variable, the before and after comparison of transition probability alone is unreliable as an evaluation technique. It implicitly assumes that, in the absence of a change in the minimum wage, the probability of transition would have been the same as before. This is a very strong assumption to make, not least because seasonal labour demand patterns may cause transition probabilities to vary through the year.

To avoid making such an implausible assumption, the trend in transition probabilities for a similar but unaffected (control) group can be taken as the counterfactual for the (treatment) group who were affected by the change. This approach requires far weaker assumptions than those relied upon for a simple before and after analysis. Taking the difference between the before and after changes of treatment and control groups, is known as difference-in-difference (DiD) evaluation.

In an experimental setting, the minimum wage rise would apply to some people, but not all employees. A low-wage control group could then be monitored to demonstrate what would have happened to low-paid workers who were affected by the rise had they not been. Since the minimum wage is a legal requirement on all employers, however, no such experimental control group exists. It is therefore necessary to look elsewhere for an appropriate substitute. An obvious candidate group is those employees who were paid at, or just above, the new minimum before the rise. How can one know whether the behaviour of this control group adequately reflects what would have happened to the treatment group in the absence of the NMW rise?

Two key assumptions about the control group are necessary for DiD analysis to be appropriate:

- the first required for the difference-in-difference technique to yield accurate results is the 'common trend' assumption. For NMW analysis, this requires that, in the absence of the minimum wage rise, the transition probabilities of both the treatment and control groups would have moved together over time. This assumption means that the DiD methodology, unlike a before and after approach, strips out any macroeconomic trends or seasonal fluctuations that might affect transition probabilities, but which are unrelated to the NMW change itself, and
- the second required for the application of DiD in this case, is that the wage changes experienced by the treatment group should have no impact on the employment transitions of the control group. This is known as the assumption of 'no spillover effects'

Together, the common trend and spillover assumptions allow DiD to isolate the impact of the NMW rise on transitions of the treatment group.

Choosing a control group for which these assumptions hold is therefore central to the analysis. Unfortunately, the demands of each assumption tend to pull in opposite directions. For example, an attempt to minimise any spillover effects would, in this case, involve choosing a control group further away from the treatment group in wage terms (for example, a control group paid between one-and-a-half and twice the NMW). But these employees are less likely to be in similar employment to those in the treatment group. They are therefore less likely to face the same seasonal or macroeconomic fluctuations. The result is that the common trend assumption becomes harder to sustain: cleaners, for example, face a very different labour market from investment bankers, and therefore do not face common employment trends.

In view of the need for a credible common trend assumption, it is sensible to limit the wage range of the control group. In line with the literature, the control group for the central results of this analysis is taken as people earning between the new NMW level of £5.35 per hour and 10 per cent higher than the incoming level, or £5.89 per hour. Implicitly, it is therefore assumed that there is no spillover effect of the rising NMW on this control group. This assumption is supported by the findings of Dickens and Manning (2003). To ensure that the transition rates for the treatment group are compared with those for people in this wage band only, the model is adjusted to control for people with wage rates higher than NMW + 10 per cent.

To avoid reporting error contaminating the treatment group, the model is adjusted to exclude from it those people who report being paid less than the previous year's minimum wage (£5.05). This is done on the grounds that these observations are either erroneous, or evidence of non-complying employers. Either way, these employees would not be expected to benefit from the NMW increase.

To control for the changing composition of the sample over time, a vector of individual characteristics is included in the model. A wide range of individual characteristics are included in this control vector including, but not restricted to, age group, whether in full or part-time work, sex, qualification level and whether in temporary employment. A full list of control variables is available in the technical note.

However, using a binary indicator to show whether a person is directly affected is a relatively blunt tool to analyse the impact of the NMW rise. Theory suggests that those paid substantially below the new minimum would be more likely to be laid off in the face of a 30p NMW increase than those just below the new minimum rate. Therefore, one might expect employment effects to be concentrated among those people facing the biggest wage rises, and to be proportional to the wage increase necessary to comply with the incoming minimum.

It is possible to exploit this idea by including a measure of the 'wage gap' for each treated individual, measured as NMW_{t+1} minus $wage_{it}$. This variable takes a (positive) value only for people in the treatment group, creating a continuous variable that measures the extent of the necessary pay rise for the employer to comply with the new minimum for that individual. Obviously, the maximum wage rise required by the NMW increase is 30p, from £5.05 to £5.35. Using this continuous wage gap variable, in place of the treatment binary, it is possible to measure the marginal effect on employment transitions for people facing a notional £1 rise in wages.

Results

Across all adults, the results indicate that the minimum wage rise had no impact on the probability of leaving work of those directly affected. Among males, the results suggest that the change caused the probability of leaving employment to fall by 4 percentage points over the period of the rise, for the treated group. This compares to a pre-rise exit rate of 11.7 per cent for the treatment and control groups. None of the other results are significantly different from zero at the usual levels, however.

As discussed above, the possibility of spillover effects of the minimum wage rise on employees at or just above £5.35 is something that makes choosing these people as a control group problematic. As a robustness test, **Table 1** shows the results of using people paid between 10 and 20 per cent above the new NMW as the control group. Here, all results are insignificantly different from zero and similar to those for the lower-wage control group.

The wage gap estimator results are shown in **Table 2**. These marginal effects show the impact of raising wages by £1 per hour on transition rates among the treated group. None of the results are statistically significant at the usual levels, indicating that no transition rate effect was detected. Again, the sign for the male results is not what might be expected from basic theory.

Discussion

In common with Dickens and Draca (2005), which looked at the impact of the minimum wage increase in 2003, these results show no evidence of an adverse employment impact from the 2006 change. But, unlike those results, there is some suggestion of a positive impact on retention for men. This result should be treated with some caution since there are only 117 observations for the post-rise male treatment group.

One reason for these benign results may relate to the fact that this analysis has only looked at outward transition probabilities: it could be the case that rather than laying off workers in the face of a rising minimum wage, employers choose to hire fewer people. In this case, an employment effect would not show up in analysis of employment exit rates.

It is necessary to consider the reliability of the identifying assumptions of the analysis to assess the validity of the results. First, it may be that the common trend assumption of DiD analysis is untenable, and that the employment transition chances of the treatment and control groups do not move together over time. Second, and perhaps more likely, it may be that positive effects on wages from the increasing minimum reach further up the wage distribution, thus violating the spillover assumption. This might occur if wage differentials, between employees with different levels of experience, are restored by the employer raising everyone's wages by the same proportion.

The common trend assumption is ultimately not testable given national implementation of the NMW rise. However, if those in the control group are in different industries or job types to the treated, the assumption may be unsound. This could be because different sectors (for example, manufacturing versus retail) are subject to different macroeconomic, and therefore employment, trends. Including control variables for industry type might be expected to limit the impact of differential trend, although doing so makes very little difference to the reported results.

Lam *et al* (2006) have suggested some effect of the minimum wage on the pay rates of workers who do not benefit from it directly. Restoration of pay differentials within firms may make the 'no spillover' assumption hard to sustain. If this is the case, the actual impact of the increase on transition rates might be bigger than the one estimated here.

The impact of changes in NMW on the wages of those affected occurs almost immediately (see Figure 2). Much of the work on spillover effects seems to suggest that restoration of pay differentials, by contrast, takes place over a relatively long

Table 1
Difference-in-difference estimates of transition probability

Treatment binary	Control group	All adults	Females	Males
Raw difference-in-difference	NMW to NMW+10 per cent	–0.009 (0.59)	0.010 (0.64)	–0.043* (0.08)
	NMW+10 per cent to NMW+20 per cent	0.005 (0.78)	0.026 (0.22)	–0.037 (0.17)
Control vector	NMW to NMW+10 per cent	–0.017 (0.22)	–0.003 (0.87)	–0.040** (0.02)
	NMW+10 per cent to NMW+20 per cent	–0.005 (0.72)	0.011 (0.52)	–0.035* (0.06)

Notes:

Negative numbers indicate reduced probability of transition; P-values in brackets.
* and ** indicate statistical significance at the 90 per cent and 95 per cent levels, respectively.

Table 2
Wage gap estimates of transition probability

Wage gap	Control group	All Adults	Females	Males
Raw difference-in-difference	NMW to NMW+10 per cent	–0.008 (0.91)	0.054 (0.48)	–0.191 (0.2)
	NMW+10 per cent to NMW+20 per cent	0.043 (0.54)	0.102 (0.17)	–0.138 (0.35)
Control vector	NMW to NMW+10 per cent	–0.043 (0.50)	0.009 (0.90)	–0.219* (0.09)
	NMW+10 per cent to NMW+20 per cent	0.006 (0.92)	0.052 (0.41)	–0.164 (0.20)

Notes:

Negative numbers indicate reduced probability of transition; P-values in brackets.
* indicates statistical significance at the 90 per cent level.

Figure 3
Number of employees paid below NMW + 20 per cent: by months before/after NMW change in October 2006

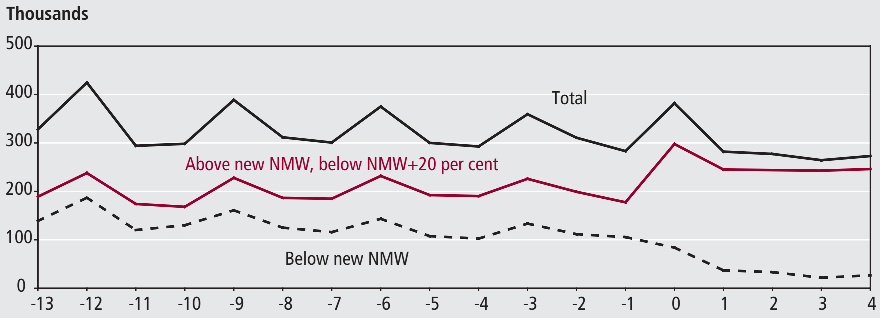

time period. This difference in timing should mean that the 'no spillover' assumption is reasonable given the time periods under observation here. Furthermore, the estimated spillover effects here appear modest relative to the direct impact of the minimum wage increase for those directly affected.

Figure 3 shows the number of people paid less than NMW + 20 per cent. Although the data are noisy, there is no sign of a dramatic drop in that number around the time of the NMW rise. This suggests that spillover effects were not large within three months of the NMW rise. It therefore supports the idea that spillover does not present a huge problem for this type of analysis, although it may reduce the measured impact of the change.

Although it should be treated with some caution due to the small sample sizes involved, the most intriguing aspect of the results is that the NMW rise appears to have improved job retention among men who were directly affected. This may seem counterintuitive, given the standard approach to considering the impact of a minimum wage increase. However, it is also a common observation that employment retention rates improve in firms that pay higher wages.

One study, looking at recruitment and retention in lower-paying labour markets estimated that firms that raised their wages relative to competitors by 10 per cent saw staff turnover rates fall by 15 percentage points to 23 per cent (Brown *et al* 2001). This particular study related to employment paying above the minimum wage, where a pay rise relative to the wages of other firms has the effect of making those jobs clearly more valuable than their employees' alternative, outside, option. Such a change would be expected to reduce turnover.

While the minimum wage cannot in itself make one minimum wage job more attractive than another, it does still raise the value of employment relative to not being in employment. To the extent that being out of work is an employee's alternative option, therefore, an NMW increase would also be anticipated to reduce turnover and therefore transition probabilities. A further explanation for the results for men is that, in increasing the coverage of the single minimum pay rate, the minimum wage increase reduces opportunities for employees to achieve higher wages by moving from job to job, and hence reduces staff turnover.

These results therefore suggest no evidence that the 2006 minimum wage increase had an adverse effect on employees' chances of remaining in work. Indeed, for men, the impact on retention may, in fact, have been positive.

Notes

1 See Low Pay Commission website at www.lowpay.gov.uk/lowpay/index.shtml
2 For a textbook guide on monopsony, see Manning A (2003) Monopsony in Motion, Princeton University Press.
3 It is only possible to match 60 per cent of the people who appear in wave one with their wave five responses one year later.

CONTACT

elmr@ons.gsi.gov.uk

REFERENCES

Brown D, Dickens R, Gregg P, Machin S and Manning A (2001) *Everything Under a Fiver: Recruitment and Retention in Lower Paying Labour Markets*, Joseph Rowntree Foundation.

Department for Business, Enterprise & Regulatory Reform (2006) *Government Evidence to the Low Pay Commission*.

Dickens R and Draca M (2005) *The Employment Effects of the October 2003 Increase in the Minimum Wage*, Low Pay Commission.

Dickens R and Manning A (2003) *The Impact of the National Minimum Wage on the Wage Distribution in a Low-Wage Sector*, Royal Economic Society Annual Conference 2003.

Dickens R and Manning A (2004) *Has the minimum wage reduced UK wage inequality?* Journal Of The Royal Statistical Society, vol. 127(4).

Dickens R and Manning A (2006) *The NMW and Wage Inequality – An Update*, Presentation to the Work and Pensions Economics Group, October.

Lam K, Ormerod C and Vaze P (2006) 'Do company wage policies persist in the face of minimum wages?', *Labour Market Trends* 114(3), pp 69–81.

Metcalf D (2007) 'Why has the British National Minimum Wage had little or no impact on Employment?', CEP Discussion Paper No 781.

Skinner C, Stuttard N, Beissel-Durrant G and Jenkins J (2002) *The measurement of low pay in the UK Labour Force Survey*, Oxford Bulletin of Economic Statistics 64, pp 653–76.

Stewart M (2003) *Modelling the Employment Effects of the Minimum Wage*, Low Pay Commission.

Stewart M and Swaffield J (2008) *The Other Margin: Do Minimum Wages Cause Working Hours Adjustments for Low-Wage Workers?* Economica, 75, pp 148–67.

TECHNICAL NOTE

Difference-in-difference methodology

Formally, the DiD method can be derived as follows. Mean employment in each (treatment and control) group g at time t can be denoted as:

$$\overline{E_t^g} = \alpha_g + \phi_t$$

where α_g is a group-specific effect and ϕ_t is a time effect common across all groups. Assuming that the NMW rise then has a constant effect Ψ on employment for the treated group and no effect on the untreated (no spillover assumption), mean employment rates for each group can be denoted as shown below:

$$\overline{E_t^g} = \alpha_g + \phi_t + \Psi \qquad \text{if g = treatment}$$

$$\overline{E_t^g} = \alpha_g + \phi_t \qquad \text{if g = control}$$

The difference-in-difference estimator is then given by:

$$\Psi = \left(\overline{E}_A^T - \overline{E}_B^T\right) - \left(\overline{E}_A^C - \overline{E}_B^C\right)$$

where $\overline{E}$ is the mean employment outcome for the treatment (T) and control (C) groups in the time period before (B) and after (A) the NMW increase.

As explained above, however, it is not possible to derive employment rates for the treatment and control groups. Rather, it is necessary to consider labour market transitions, and construct the difference-in-difference estimator as:

$$\Psi = \left(\overline{P}_A^T - \overline{P}_B^T\right) - \left(\overline{P}_A^C - \overline{P}_B^C\right)$$

where $\overline{P}$ is the mean probability that someone employed in period 0 remains employed in period 1, that is, $P[E_{it1}=1 \mid E_{it0}=1]$. This is calculated for treatment (T) and control (C) groups, both before (B) and after (A) the minimum wage increase.

An alternative way of finding Ψ, the impact on transition probability of the NMW rise, is to estimate the following regression model:

$$P\left[e_{it+1} = 1 \mid e_{it} = 1\right] = \alpha_g + \phi_t + \Psi g_{it} d_{t+1}$$

where α_g is a group specific effect, ϕ_t is a time effect common to all groups, g_{it} is a dummy variable denoting whether an individual belongs to the treatment group, and d_{t+1} is a dummy taking the value 1 when the new rate of the NMW is in force.

In order to avoid the reporting error contaminating the treatment group, a dummy variable (shown as g_1 below) is introduced to the model. This excludes those people reporting being paid less than the previous year's minimum wage (£5.05) from the treatment group. The treatment group is represented as g_2 and the control group (dummy variable excluded from the model as the baseline) as g_3.

Finally, a vector of individual characteristics, X, is included in the model to control for changing composition of the sample over time. A wide range of individual characteristics are included in this control vector, including sex, a cubic function of age, qualification levels, type of housing tenure, whether in social housing, past employment experience, whether in full or part-time work, whether in public sector employment, whether job is temporary, region of residence, ethnicity, and whether a couple or single. The basic probit model to be estimated therefore becomes:

$$P\left[e_{it+1} = 1 \mid e_{it} = 1\right] = X_{it}\beta + \alpha_1 g_{1it} + \alpha_2 g_{2it} + \alpha_4 g_{4it} + \phi_t + \gamma g_{1it} d_{t+1} + \psi g_{2it} d_{t+1} + \varphi g_{4it} d_{t+1}$$

where X is the vector of individual characteristics, and the impact estimator is the marginal effect captured by ψ.

As in Stewart (2003), the model also includes a polynomial in the wage rate. This controls for the fact that those in lower-paid work are less likely to remain in employment over a given time period than higher-paid people. Controlling for this effect strengthens the claim of a credible common trend assumption.