working paper

05 08

The Incidence of Nominal and Real Wage Rigidities in Great Britain: 1978–1998

by Richard D. Barwell and Mark E. Schweitzer

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 05-08 September 2005

The Incidence of Nominal and Real Wage Rigidities in Great Britain: 1978–1998

By Richard D. Barwell and Mark E. Schweitzer

This paper analyses the extent of rigidities in wage setting in Great Britain over the 1980s and 1990s. Our estimation strategy, which generalizes the work of Altonji and Devereux (2000), models the notional wage growth distribution – the distribution of nominal wage growth that would occur in the absence of rigidities in pay – while allowing for the presence of measurement error in the data. The model then allows for the possibility that the nominal wage growth of a fraction of the workforce may be subject to a nominal or real downward rigidity. Our model suggests that real rigidities in wage setting are more prevalent than nominal rigidities, although the incidence of these real wage rigidities has fallen gradually over time. If firms cannot cut real wages in response to negative demand shocks they may resort to laying-off workers. Our results support this micro-foundation of the wage-unemployment Phillips curve: workers who are more likely to be protected from wage cuts are also more likely to lose their jobs.

Richard E. Barwell is at the Bank of England. Mark E. Schweitzer is at the Federal Reserve Bank of Cleveland and may be reached at Mark.E.Schweitzer@clev.frb.org or (216) 579-2014. The reported results are based on a confidential data source. The authors would like to thank Paul Beaudry, Ian Bond, William Dickens, Bruce Fallick, Guillermo Felices, Lorenz Goette, Erica Groshen, Mike Joyce, Mike Lettau, Steve Nickell, Jonathan Thomas and fellow members of the International Wage Flexibility Project for their helpful comments. They also thank Saeed Zaman for excellent research assistance.

1 Introduction

Research into the causes and consequences of both nominal and real wage rigidities has a long and distinguished tradition in economics. Macroeconomists have traditionally appealed to wage rigidities as an explanation for the failure of wages to respond to (changes in) the unemployment rate (see, for example, Layard et. al. (1991)). More recently, rigidities in the labour market have been found to dramatically improve the ability of the monetary models of the New Neoclassical synthesis literature to reproduce plausible time series profiles for output, inflation and the real wage over the business cycle (Christiano, Eichenbaum and Evans (2005)).

In this paper, we search for evidence of downward wage rigidities in panel microdata covering Great Britain. If these rigidities reduce the responsiveness of wages to a deterioration in macroeconomic conditions, then at the micro level we should expect to find evidence of 'too many' workers who report wage freezes, and 'too few' wage cuts. Although attention has almost exclusively[1] focused on the presence of nominal rigidities in panel microdata, we might expect that real wages are also sticky. This paper explores the incidence of both nominal and real, or 'cost of living', rigidities in wages.

Of course there is no reason to believe that the incidence of these wage rigidities should be static over time. A decline in the inflation rate is likely to increase the fraction of the workforce who are at risk of a nominal wage cut, so the incidence of nominal wage freezes is likely to rise. Conversely, structural changes in the labour market such as the decline in trade union power, may have reduced the mechanisms which protect workers from either nominal or real pay cuts. A key aim of this paper

[1] See Schweitzer (2003) for an exception: using the non-parametric methodology proposed by Kahn (1997), he is able to identify distortions in the nominal wage growth distribution consistent with real rigidities in wage setting in the data used in this paper.

is to establish whether there has indeed been a shift in the degree of rigidity in wages at the micro level – which can shed some light on whether the British labour market has indeed become more flexible at the macro level.

The starting point for our analysis is the model of Altonji and Devereux (2000) which allows for the presence of downward nominal wage rigidity and measurement error in wage setting. This paper applies an extended version of that model[2] which allows us to jointly estimate the incidence of nominal and real wage rigidities in panel microdata and thereby their impact on the nominal wage growth distribution. Our model allows that workers may fall under one of three regimes – where the growth in their nominal wage is bounded from below either by zero (what we call the *nominal rigidity regime*) or their expectation of the inflation rate (what we call the *real rigidity regime*), or where the growth in their nominal wage is unconstrained (what we call the *flexible* or *unconstrained regime*). Using annual data on wages taken from the New Earnings Survey Panel Dataset (NESPD) we investigate the trends in these regimes over the period: 1978 to 1998.

One reason macroeconomists are interested in wage rigidities is that they can potentially offer an explanation for the behaviour of unemployment over the business cycle: if firms cannot cut wages in the face of a demand shock, they may have to resort to layoffs. Our model generates an estimate of the degree to which rigidities drive a wedge between the wage that reflects productivity, and the wage the firm actually pays. This should help explain the pattern of job destruction at the micro level. We investigate this hypothesis at the end of the paper.

The structure of the paper is as follows. Section 2 summarises the data on which this research is based. Section 3 gives an overview of the underlying macro-

[2] The econometric model used in this paper was developed with fellow members of the International Wage Flexibility Project (IWFP).

environment over the sample period before highlighting the key features of the nominal wage growth distribution which we believe any reasonable model must reproduce. Section 4 describes how we have modified Altonji and Devereux's approach to allow for the presence of both real and nominal rigidities in wage setting, and the idiosyncrasies of our estimation. Section 5 outlines the key results of our model. Section 6 investigates whether the pattern of job destruction in the dataset is indeed consistent with the incidence of wage rigidities our model describes, and hence whether our results support the link between rigidities and business cycle movements in unemployment described in macro-theory.

2 The Data

The data on which this research is based is the New Earnings Survey Panel Dataset (NESPD); for more details on this dataset Schweitzer (2003). The NESPD contains information in panel form on the earnings and employment conditions of a random one percent sample of the workforce since 1975. The data is collected from employers, who are legally obliged to provide information on the various components of pay, hours worked for a specified week in April, together with data on their employees age, sex, occupation and job and whether their pay was affected by a collective agreement, wage council or absence. For the purposes of this paper we define the wage as the gross hourly pay exclusive of overtime of *full-time* workers who have been in the same job for a year or more and whose pay was unaffected by absence.

3 A Preview: A Picture of Nominal Wage Growth in Great Britain

We characterise the (nominal) wage growth distribution using a simple 50 bin histogram where individuals are collected into bins of one percentage point in width (where all outliers are collected into the outside bins). In Figure 1 we illustrate the distribution of nominal wage growth for the year to April 1995 which highlights particularly well the key features of the distribution. For reference, we highlight the points in the distribution which correspond to zero nominal wage growth and the inflation rate over the preceding year, in the month to February (which was the most recent observation on retail price inflation available to agents on the 1st April).

Three key features of these distributions are worthy of particular comment. First, the size of the first and last bins of the histograms – which suggest either that large pay cuts and increases are common, or that data is polluted by measurement error. Second, the 'spike' in the distribution which corresponds to a nominal pay freeze (zero pay growth). Third, the large number of observations clustered around the inflation rate. Ultimately, the final arbiter of our model will be its ability to reproduce these features at a given point in time and their variation in time series. We now turn to discuss each of these three features.

In the year to April 1995 more than one in ten of our sample enjoyed an increase in their nominal wage in excess of 35% or a cut in excess of 15%. Variation in the rate of nominal pay growth across workers on this scale is a somewhat surprising feature of our data. While there are a number of plausible explanations for significant variation in hourly wages from one year to the next *in the absence of a change in job* it is likely that some of these observations reflect measurement error.

Smith (2000) finds that measurement error can explain a significant proportion of both the nominal rigidity and the incidence of wage cuts. But her analysis was based on survey data, while we have access to payroll data that employers are legally obliged to provide, which according to Nickell and Quintini (2003) "ensures a high degree of accuracy".[3] However, there remains the possibility that mistakes in compiling and inputting data on hourly wages will introduce errors in our data. And any undetected changes in job description – especially promotions and demotions – are likely to lead to large absolute wage changes. While the probability of each individual *wage level* observation being mismeasured in the panel is in all likelihood low, the reality is that when these errors occur they may well result in large errors in our estimates of *wage growth*.

That there exist nominal rigidities in wage setting in the UK over this period appears almost incontrovertible. We can clearly observe a spike in the nominal wage growth distribution in the bin of the histogram corresponding to (near) zero wage growth $-0.5\% \leq \dot{w} < 0.5\%$ over the year 1994/5. The spike is also apparent in years of double digit inflation, though smaller in size. But it is the extent to which there is a *disproportionate* mass of observations in this bin of the histogram that is key to identifying the rigidity.

We think of downward real wage rigidities as any mechanism which attempts to protect workers from a cut in their ex-post real wage by protecting them from cuts in their *expected* real wage. But we do not expect to see a spike in the nominal wage growth distribution at some expectation of inflation, akin to that we see at zero on account of nominal rigidities. Expectations of the increase in the cost of living will vary across agents in the economy, not least because the timing of wage bargains will vary so agents have access to different information sets. As a result, we would expect real wage rigidities to deliver a clustering of observations over an interval in the nominal wage distribution (rather than at a precise point), where the width of that interval is determined by how diffuse the distribution of inflation expectations is (at

[3] In the validation exercises, Card and Hyslop (1997) and Bound *et. al.* (1989), payroll data of this kind is typically taken as the truth against which the survey responses are compared.

the time workers agree their settlement). Given the behaviour of inflation over the period it is likely that the dispersion of these expectations will also have varied significantly over the period. .

4 The Econometric Model

The econometric model which we employ in this paper to explain the distribution of nominal wage growth across workers in the NESPD extends the model presented by Altonji and Devereux (2000) to account for real wage rigidities. We allow that workers may fall under one of three regimes. Wages may be flexible or *unconstrained*, in which case their wage is unprotected and set equal to the *notional* wage on all occasions. Alternatively, wages may be subject to *nominal* rigidities, where we assume that the nominal wage growth is bound from below by zero. Finally, wages may be subject to *real* rigidities, where we assume that nominal wage growth is bound from below by an expectation of inflation. See the Appendix for a derivation of the specific likelihood function used in these estimates.

The essence of this paper is the estimation of the time series variation in the fractions of the workforce who are protected from nominal and real pay cuts through the distortions they create in the nominal wage growth distribution. As well as potentially increasing the mass in the extreme tails of the distribution, errors in the data are also likely to smooth out any distortions in the underlying distribution of nominal pay growth. It therefore follows that if the incidence of errors in the data varies from year to year then this may lead to a corresponding variation in our estimates of the fraction of the workforce who lie in each of the regimes[4]. We

[4] So, for example in the case of the nominal regime, for a given spike in the observed nominal pay growth distribution, the greater the incidence of errors in the data the greater the underlying spike must have been in the actual distribution, and therefore

therefore allow for a small yet significant degree of measurement error, which is constant across the sample period[5]. The NESPD has had only minor methodology changes and steady large sample making this assumption reasonable.

Our model assumes that depending on which regime workers fall in, zero or an expectation of the inflation rate forms a lower bound on the actual wage growth distribution, so that only wage changes above these thresholds are set according to the notional wage relationship. We use this assumption to identify the notional relationship. Focusing on the right-hand tail of the nominal wage growth distribution, we simulate the conditional mean notional wage change for our set of covariates using a quantile regression at the 75th percentile.

If the model is estimated year by year it is likely that the marginal effect of a given covariate on the relative propensity to fall in each of the regimes to vary from year to year; however the concept of wage rigidities suggest a longer-term departure from fully flexible wages. It is probable that most industries have at least one year in which the vast majority of individuals received substantial nominal wage increases. We therefore estimate the window over a five-year time frame as it is less likely that an industry or group will repeatedly do much better than average five years in a row.

5 Results

For each of the windows of data over which we estimate the model we obtain a set of parameters which describe the key behavioural relationships in the model: the rate of notional wage growth, the variability in notional wage growth and the relative propensities for wage growth to be bound from below by zero or an expectation of the

---

by implication the greater the proportion of the sample who must be in the nominal regime.

[5] We ‘assume’ that 89 % of observations in the NESPD are correctly measured. This value was chosen via a grid search: it was the value of q which gave the highest average value of the likelihood when the model was estimated over successive five year windows of the sample period (see later).

inflation rate, together with a series of constants which specify the mean and standard deviation of the lower bound for wage growth in the real regime and an estimate of the standard deviation of measurement error in the NESPD. In order to set these results in perspective, we shall first review the results obtained when the model is estimated over a particular five year windows (the five years of wage growth centred on the year to April 1995) to build up an explanation of the pattern of wage growth in that period, before turning to a more comprehensive review of our results over the period as a whole.

Our model begins with a description of the underlying or notional wage growth relationship. This is effectively identified outside as in input to the likelihood function using quantile regression techniques to pin down a conditional quantile function of the notional wage growth distribution. This distribution reflects a simple wage growth equation where we include the age of each individual and its square and a series of occupational and industrial dummy variables. The coefficients on the age terms imply a standard lifetime earnings profile, with the rate of wage growth falling monotonically with age. Of the remaining controls, five industry dummies are significant, while only one occupational dummy is.

We find that other things equal, the standard deviation of notional wage growth declines with age, and this relationship is statistically significant, consistent with younger workers doing more ‘job shopping’. Moreover, we find that the majority of our industrial controls are significant, and all the occupational controls are, along with our control for whether an individual was covered by a union agreement. Combining these results implies that the estimates of the notional wage distribution will be a mixture of several distinct normals, which would allow the notional to potentially deviate from symmetry, unimodality, and other restrictions implied by normality.

Our model allows for the possibility that only a fraction of the workforce will have their wages protected. In table 1 we illustrate the estimated fraction of different sub-samples of the workforce who fall in each of the three regimes.

The key result illustrated in the table above is that real rigidities are far more prevalent than nominal rigidities. We estimate that in this period a little over two in five workers in the sample are members of the real rigidity regime. Of the remainder, three out of every four fall in the unconstrained regime so only about 15 % of workers are found to fall in the nominal rigidity regime. In terms of which individuals are members of the real regime it is very much a case of the usual suspects: older workers, those in the public sector, and those covered by union agreements have the highest proportion in the real regime.

Our model also produces an estimate of the mean and standard deviation of the expectation of inflation which forms the lower bound on wage growth in the real regime. Our results imply that the mean of the lower bound on wage growth in the real regime in the year to April 1995 was 2.6 %. This figure should be set against a rate of RPI inflation in the year to February 1994 (the latest observation on retail inflation available to wage bargainers) of 2.4%, and a rate of RPI inflation of 3.3% in the year to April 1995 (the period over which wage growth is measured in the NESPD). This result corroborates our assumption that the lower bound on the wage growth in the real regime is an approximate peg with the cost of living. The standard deviation of this lower bound on wage growth for individuals in the real regime is about 0.6 percentage points in 1994.

Finally, we also estimate of the standard deviation of measurement error in the NESPD, given the model's estimate of 89% of the observations are correctly measured. The standard deviation of measurement error in nominal wage growth is estimated to be about 34 percentage points. This confirms our suspicion that when they occur, errors are likely to lead to extreme wage growth observations.

One way to illustrate how well our model performs is by simulating the actual nominal wage growth distribution for the year 1994/5 using the estimated parameters. The top left hand corner of figure 2 illustrates the actual distribution of nominal wage growth (within a specified range[6]) and the bottom right hand corner we illustrate our result. Generally, our model is able to reproduce most characteristics of the actual distribution. The simulation with an estimate of the distribution of predicted mean notional wage growth for groups as previously discussed, which is tightly concentrated. Once the stochastic component representing individual-level variation in notional wage growth is included, the distribution becomes more disperse than the raw data. In the middle row of the figure we illustrate the simulated wage growth distributions for individuals in each of the three regimes. Since the rigidity regimes can only generate wage cuts through measurement error it is indeed the case that the mass in the left hand tail of the actual distribution is accounted for primarily by the unconstrained regime. In terms of the clusters of observations in the distribution, the nominal regime generates the clear spike at zero as we might expect. Perhaps more interestingly, the real regime illustrates a more diffuse clustering of observations,

[6] We concentrate on those wage growth observations that fall in the range $-14.5\% \leq \hat{w} < 35.5\%$, collecting all observations outside this range into the extreme bins of the histogram. Although this does not allow us to compare how well the model fits the *density* of extreme wage growth observations, it still allows us to compare how well the model fits the *frequency* of these observations.

which reflects the non-trivial standard deviation of the lower bound we highlighted earlier.

Finally in the bottom left and bottom centre panels of the figure we illustrate the distortion in the distribution created by each of these rigidity regimes, by focusing on the density of actual wage growth observations in each regime who are swept up to the spike. Quite apart from the fact that more individuals are in the real rather than the nominal regime, these figures illustrate why the aggregate wage sweep is larger for the former rather than the latter regime..

So how well does our model fit the underlying characteristics of the data and how important is the inclusion of the real regime? Figure 3 illustrates the fit of the model. The blue line is the difference between the fraction of the sample attributed to each histogram bin in the data versus the full model. While the performs reasonably well, it is clear that our model allocates insufficient mass at both the nominal spike at zero and the cluster around the expectation of inflation. The under-prediction in the interval $-0.5\% \leq \dot{w} < 0.5\%$ is equivalent to a little over one per cent of the NES sample, and the under-prediction in the interval $2.5\% \leq \dot{w} < 3.5\%$ is equivalent to a little over three per cent of the NES sample. The importance of allowing for the fact that a significant fraction of the workforce have their real wages protected is immediately obvious from the alternative simulation shown by the red line in Figure 3. When we switch off the real regime in our model, the fit of the model deteriorates dramatically.

We now turn to summarise the time series variation in our results. Taking the two decades of data as a whole, one can identify the following broad trends in the size of the three regimes (see Figure 4). First, there was a gradual decline in the share of the real rigidity regime. Second, the share of the nominal rigidity regime was broadly constant. Third, there was therefore an increase in the share of the workforce who were protected from neither real nor nominal wage cuts. Between 1977/8 and 1997/8 the fraction of the sample in the unconstrained regime rose by over 25 percentage points. In this sense, our results do indicate that the British labour market has become more flexible. Furthermore, simulations based on holding the parameters of our model fixed at their estimated values using the first window of data reveals that little of this change is associated with the changing characteristics of the workforce.[7] .

We also find that throughout the sample period, the estimate of lower bound on wage growth is consistent with the underlying rate of retail price inflation. In Figure 5 we graph the implied mean lower bound on nominal wage growth against the latest observation on the rate of RPI inflation at the beginning of the year (RPI inflation to February) and the rate of RPI inflation over the year (RPI inflation to April). There was considerable volatility from month to month in the rate of inflation at the start of our sample period so would might expect agents to have been uncertain over the likely path for inflation in the future. In the model increased uncertainty over inflation should be reflected in the large standard deviation of the real regime lower bound.. The estimated standard deviation of the lower bound falls by a factor of two between the 1970s and the mid 1980s and then again by a factor of two by the end of the sample period (Figure 6). Finally, as we move further towards the low and stable inflation environment of the late 1990s we find that our two alternative observations on the rate of RPI inflation generally fall inside the approximate 95% confidence interval around our estimate of the lower bound on wage growth. We think this is strong evidence that the lower bound on wage growth primarily reflects expected inflation.

---

[7] Due to space limitations we do not show these simulations. They are available in Bank of England Working Paper.

Over time the fraction of the workforce whose wage growth is constrained on account of protection from real wage cuts is found to decline (see Figure 7). However, as the rate of inflation falls over the sample period and the notional wage growth distribution shifts to the left, the fraction of the workforce swept up to the spike at zero increases somewhat. Once again given our estimates we can calculate the distortion in the wage growth distribution generated by these regimes. Our results suggest that over the 1990s the cumulative effect of sweeping up all real notional wage cuts in this regime raises average wage growth by around two percentage points (see Figure 8). The contribution of the nominal rigidity regime is significantly lower– the cumulative impact of sweeping up all nominal wage cuts always adds less than fifty basis points to the average rage of wage growth.

6 Consequences: the link between Rigidities and Unemployment

Macroeconomists are primarily interested in wage rigidities because they have the potential to explain the behaviour of unemployment over the business cycle. A key feature of the model presented in this paper is that some workers are protected from nominal or real cuts in pay. Laying-off workers is a costly process so we might expect a zone of inaction where small deviations of wages from the level implied by productivity (as described by the notional wage relationship) do not lead firms to start laying-off workers. But any reasonable model of labour demand ought to predict that the probability of job destruction should be increasing in the size of wedge between the actual wage and the notional wage.

Our results indicated that the notional wage growth distribution is varied enough for that notional wage changes below zero or the inflation rate are never rare and sometimes quite common and that workers vary in their likelihood of being protected from pay cuts. It follows then that there is considerable variation in the size of the expected wage sweeps (the difference between notional and realized wage growth). If wage rigidities play an important role in fluctuations in unemployment over the business cycle then the variation in the size of these wage sweeps at the individual level should help explain the pattern of job destruction. That is the final hypothesis we investigate in this paper. To be clear, we are not directly investigating a link between the incidence of wage rigidities and the unemployment rate, but rather the link between the incidence of wage rigidities and the gross outflow from employment (job destruction). But inflow shocks are central to understanding the behaviour of unemployment over the business cycle (see Burgess and Turon (2005)).

We wish to test whether the size of the expected wage sweep helps explain the pattern of job destruction in the data. We therefore need to calculate what the notional wage and the probability of falling in each of the regimes would have been for individuals who lose their job, which then allows us to calculate what their expected sweep would have been. It goes without saying that we can only estimate these hypothetical expected sweeps using data on the sub-sample of individuals who keep their job from one year to the next, but these estimates can be applied to both workers keeping and losing their current job.

Given our data source – annual observations from payroll data – we can only capture job destruction at an annual frequency. We experiment with two measures of job destruction: our favoured definition which focuses on those individuals who are in the NESPD (and therefore employed) in year *t-1* but are absent in year *t*, and a looser definition which allows the individual to be in the NESPD in year *t*, so long as they report changing their job. Our results, which are presented in the table 2, indicate that the expected wage sweep does help explain the pattern of job destruction. We include simple demographic characteristics of the individuals (age, age squared, and sex) that might alter the flow rates in the cross-sections independently. Having a larger expected sweep is statistically associated with both a higher probability of not being employed in the next year and with some change in jobs between the two years.

The expected level of sweep is influenced by a variety of factors, but clearly the probability of being in the rigid regimes and the location of the notional wage distribution both change substantially from year to year. These and other factor causes the expected wage sweep to vary substantially from year to year, with average expected sweep peaking in 1990 at 0.4 percentage points, but dropping to less than a basis point in 1980. While this is certainly variation appropriate for identifying the effects of wages sweeps on employment flows one of the advantages of our estimates is that they allow the same effects to identified cross-sectionally. Table 2 also repeats the panel regressions with a full set of year dummy variables which absorb all of the time-series variation in the data and the estimates of effects of expected wage sweep on subsequent status rise slightly. The standard errors in these estimates remain very small.

Of course we are really interested in the economic significance of the expected wage sweep variable – not its statistical significance in our regressions. With many zeros, the average sweep in the data is 0.0024, so on average having no wage sweep would lower the flow by around ½ a percentage point in both equations for the representative person. However, in 1990, when actual wage growth slowed and the wage floor in the real regime increased markedly, the average sweep rises to 0.0043, increasing the predicted flow rate by a further 0.4 percentage points. By 1996, wage growth is relative strong (5%) compared to the wage floor in the real regime (3%) and the fraction of the workforce who are protected from real pay cuts has also fallen. As a result, the predicted outflow rate falls by a percentage point between 1990 and 1996. These do not directly translate in different unemployment rates, mostly because we do not observe hiring rates, but the underlying mechanism for rigidities to create unemployment is certainly evident.

To examine the role of the various factors that underpin these expected wage sweeps in greater depth we now focus on the variation in the implied job destruction rates for a representative pair of workers in cross-section and time-series. Consider two thirty five year old men: one is a plant operative in the manufacturing sector; the other is a manager in the financial sector. For much of the sample period, finance managers enjoyed faster nominal wage growth, and are therefore less likely to be at risk of real or nominal pay cuts. Changes in the probability of these workers falling under the real or nominal rigidity regimes therefore has no impact on the expected wage sweep - and therefore the predicted probability of job destruction – for a typical finance manager in most years. But changes in the probability of being protected from real or nominal pay cuts matter for the exit rate of the typical plant operative (green line in Figure 9). If his probability of being protected from pay cuts increases – to match those of an otherwise identical public sector worker – the probability of the typical plant operative losing his job can rise by as much as a percentage point, in years where his expected wage falls below the lower bound in the real regime (pink line). And for plant operatives with below average notional wage growth (blue line) – for whom the downward wage rigidities are more likely to bind – outflow rates are affected over even more of the sample period, and by a much larger amount (red line).

The impact on outflow rates of increasing the probability that a plant operative is protected from pay cuts varies from year to year. That reflects two factors: first, the change in the probability of falling under these regimes varies from year to year;

second, and more importantly for our purposes, the rate of notional wage growth relative to the lower bounds in these regimes for our workers varies from year to year.

Another key determinant of the expected wage sweep is the level of the lower bound in the real regime relative to the rate of notional wage growth. Other things equal, the higher the lower bound the higher the fraction of the workforce who are potentially protected from 'real' pay cuts. Consider the impact of increasing the lower bound in the real regime for our typical finance manager (green line in Figure 10) by one standard deviation (pink line). Over most of the sample period median notional wage growth is still high enough to ensure that our typical finance manager is not affected by this 'error'. But for those finance managers further down the notional wage growth distribution (blue line), there is a large impact on predicted outflows over most of the sample period (red line). But by the mid 1990s there is so little variation in the lower bound of the real regime (Figure 6) that the impact on outflow rates is marginal. These results consider just a handful of the cases explored in the regressions shown in Table 2. However, they do indicate that variation in the exposure of workers to wage rigidities, and in particular real wage rigidities, can alter their prospects of remaining employed in their current job.

7 Conclusions

Casual inspection of the distribution of nominal wage growth across a representative sample of British workers reveals three striking features: a clear and now well remarked spike in the distribution at zero, a clear clustering of observations around the prevailing rate of retail price inflation and for want of a better description 'fat tails'. In this paper we have argued that the first two of these features of the wage growth distribution are the product of rigidities in wage setting which protect workers from cuts in their nominal and (expected) real wage respectively. The final feature of the distribution is largely a by-product of measurement error of various kinds in earnings data. We estimate a model capable of including all of these features on a rich source of data on pay growth in Great Britain, and find that both nominal and real rigidities are needed to fit the pattern of wage changes seen in this data.

However, the frequencies of these rigidities are not constant over the sample period.. The fraction of the workforce whose wages are unconstrained has increased since the late 1970s. Conversely, the fraction of the workforce whose nominal wage growth is bound from below by an expectation of inflation has fallen; with the fraction for whom zero is the lower bound remaining broadly constant. On this basis we argue that the labour market has indeed become more flexible over this period.

Our model also recognised the fact that when workers and firms agree nominal increases in hourly wages for the forthcoming year they do not know with certainty what the rate of inflation in the cost of living will be. On this our results are quite intuitive: the mean of the lower bound tracks somewhat between observations on inflation available to agents when the wage is agreed and inflation out-turns over the following year. In other words, our results suggest that agents are not naïve, in that they are able to predict the direction if not the extent of changes in inflation over the near future. Perhaps more interestingly, we find that the standard deviation of this lower bound falls sharply over time, which likely reflects the gradual improvement in the credibility of monetary policy focusing on a fixed inflation target after 1992.

Finally, the estimated rigidities lead to potentially meaningful deviations away from the unconstrained notional wage. And where wages deviate from the notional the probability of lay-offs might be expected to increase. Our results confirm this hypothesis: rigidities in wage setting lead to economically significant variation in the job destruction rate at the micro level.

References

Altonji, J and Devereux, P (2000), ‘Is there Nominal Wage Rigidity? Evidence from Panel Data’, *Research in Labor Economics*, Vol. 19, pp. 313-350.

Bewley, T (1999), *Why Wages Don’t Fall During a Recession*, Harvard University Press.

Bound, J, Brown, C, Duncan, G and W Rodgers (1989), ‘Measurement Error in Cross-Sectional and Longitudinal Labor Market Surveys: Results from Two Validation Studies’, *NBER Working Paper no.* 2884.

Burgess, S and Turon, H (2005), ‘Unemployment Dynamics in Britain’, *Economic Journal*, 115, pages 423-448.

Card, D and Hyslop, D (1997), ‘Does Inflation “Grease the Wheels of the Labor Market?”’, in Romer, C and Romer, D (*eds.*) *Reducing Inflation: Motivation and Strategy*, University of Chicago Press.

Christiano, L, Eichenbaum, M and Evans, C (2005), ‘Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy’, *Journal of Political Economy*, 113, pages 1-45.

Dunlop, J (1938), ‘The Movement of Real and Money Wage Rates’, *Economic Journal*, 48, pages 413-434.

Dunlop, J (1939), ‘Trends in the “Rigidity” of English Wage Rates’, *Review of Economic Studies*, 6, pages 189-199.

Fehr, E and Goette, L (**2005),** ‘Robustness and Real Consequences of Downward Nominal Wage Rigidity’, *Journal of Monetary Economics*, forthcoming.

Howitt, P (2002), ‘Looking Inside the Labor Market: A Review Article’, *Journal of Economic Literature*, 40, pages 125-138.

Kahn, S (1997), ‘Evidence of Nominal Wage Stickiness from Microdata’, *American Economic Review*, 87, pages 993-1008.

Koenker, R and Hallock, K (2001), ‘Quantile Regression’, *Journal of Economic Perspectives*, 15, pages 143-156.

Layard, R, Nickell, S and R Jackman (1991), *Unemployment: Macroeconomic Performance and the Labour Market*, Oxford University Press.

McLaughlin, K (1994), ‘Rigid Wages?’, *Journal of Monetary Economics*, 34, pages 383-414.

Nickell, S and Quintini, G (2003), ‘Nominal Rigidity and the Rate of Inflation’, *Economic Journal*, 113, pages 762-781.

Office for National Statistics (2000), 'NESPD: The Information Pack', *mimeo*.

Schweitzer, M (2003), 'Wage Flexibility in Britain: Some Micro and Macro Evidence', Chapter 9 in Fagan, G, Mongelli, F and Morgan, J (*eds*), *Institutions and Wage Formation in the New Europe*, Edward Elgar.

Smith, J (2000), 'Nominal Wage Rigidity in the United Kingdom', *Economic Journal*, 110, pages C176-C195.

Figure 1 :The Distribution of Nominal Wage Growth over the Year to April 1995

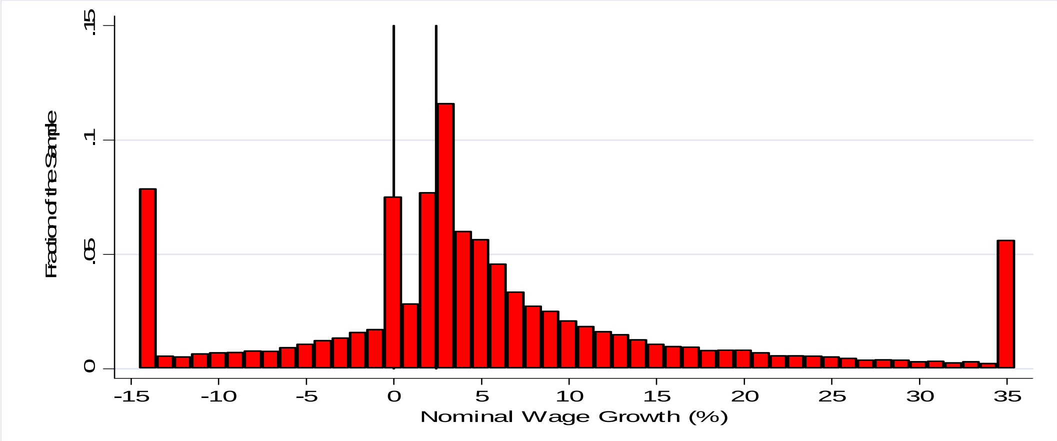

Figure 2 : Illustration of Model Predictions : NESPD 89% Measurement Error Free

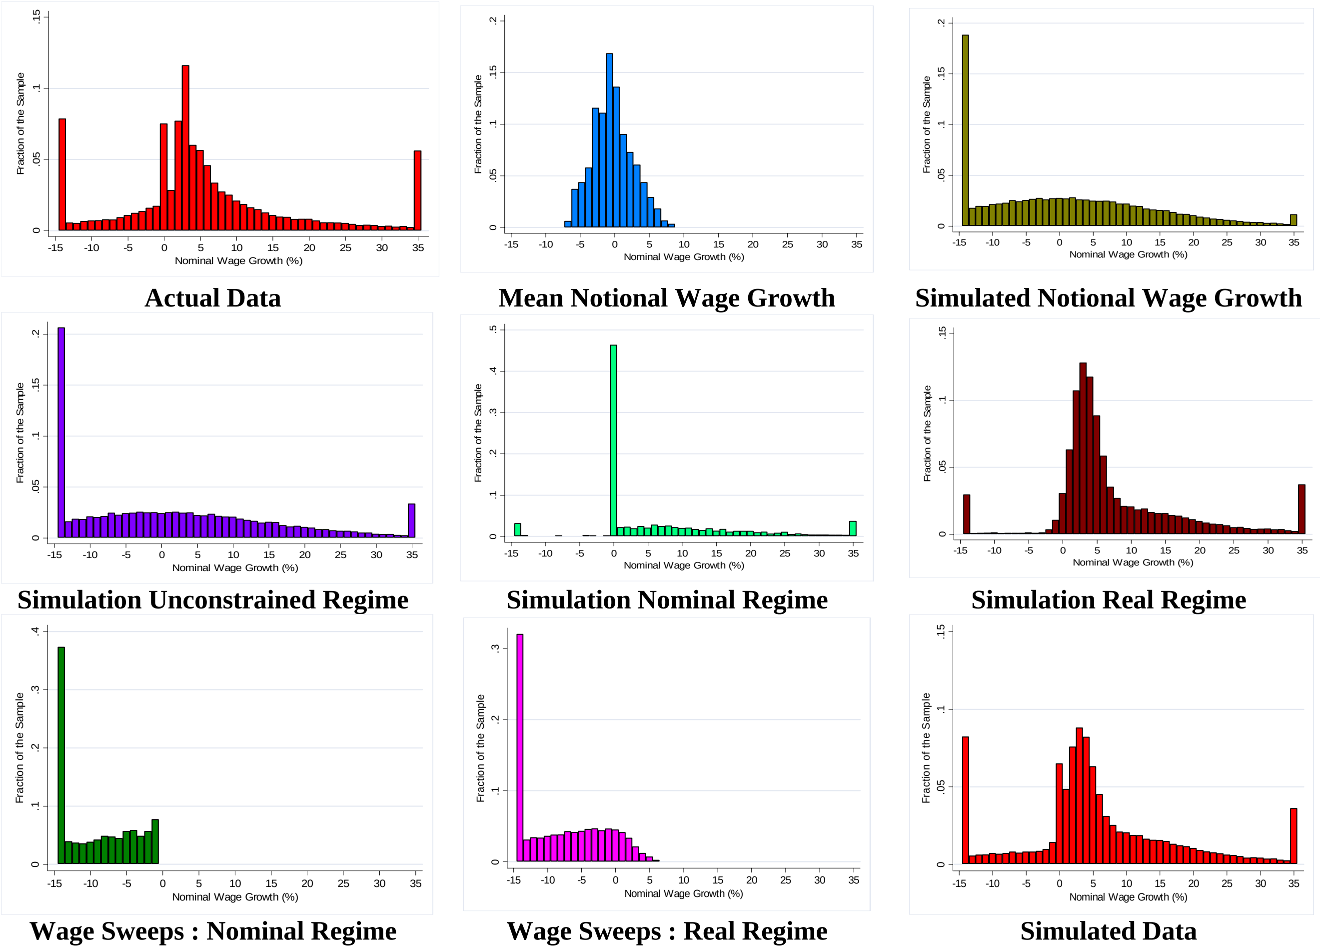

Figure 3: The Fit of the model with and without real rigidities: Actual minus Frequencies

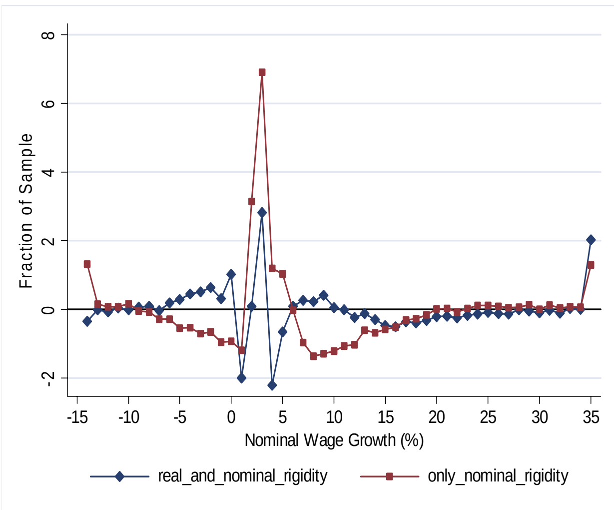

Figure 4: Size of the Three Wage Setting Regimes: *R* ,*N* ,*U* (%)

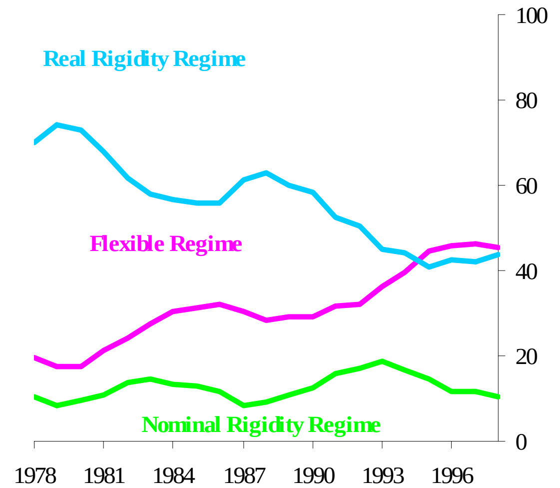

Figure 5: The lower bound in the real regime (%)

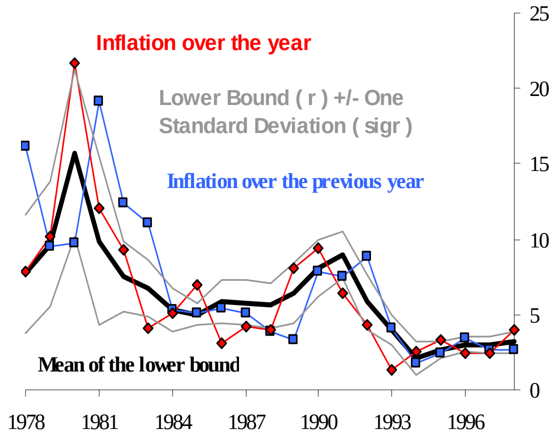

Figure 6: Variation in the lower bound in the real regime

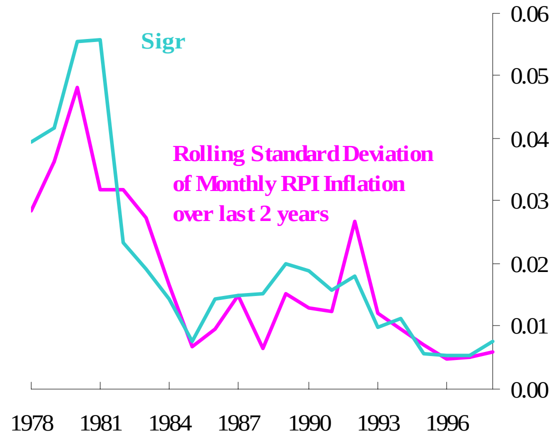

Figure 7 : The incidence of real and nominal wage freezes (%)

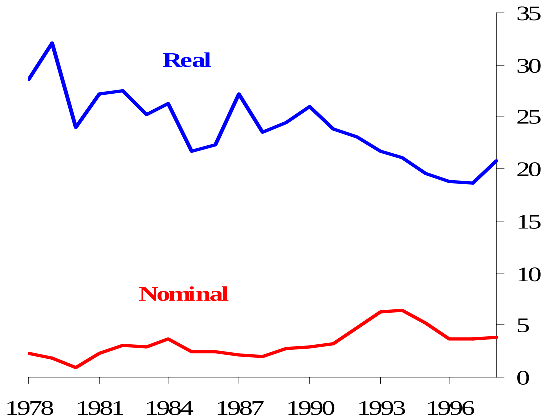

Figure 8: The incidence of real and nominal wage sweeps (pp)

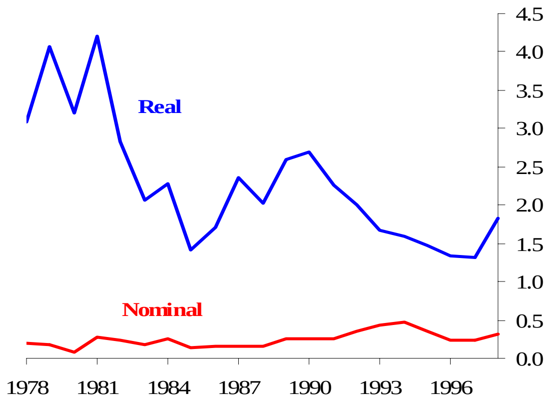

Figure 9: Impact of changes in the probability of being protected from pay cuts: plant operatives

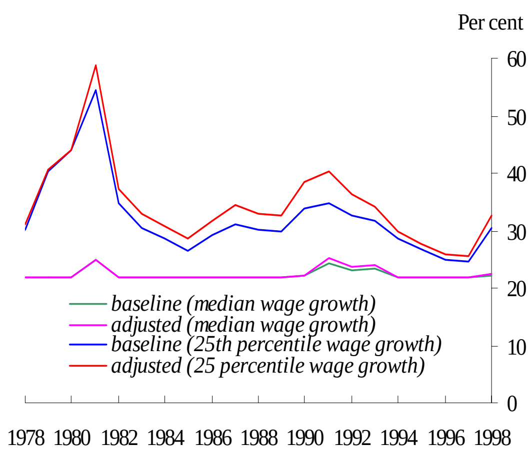

Figure 10 : Impact of changes in the lower bound in the real regime: finance managers

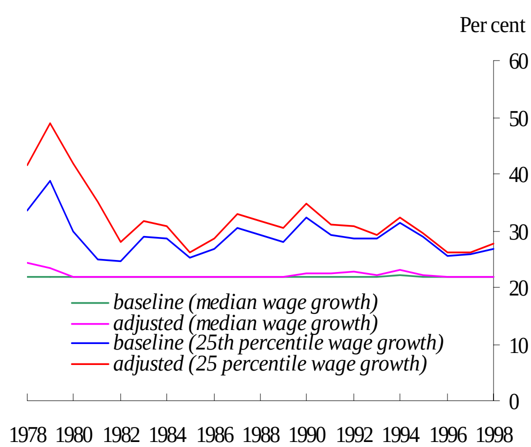

Table 1 : Variation in Membership of Regime by Characteristics (%)

Variable	Real Regime	Nominal Regime	Unconstrained
young	36.4	16.5	47.1
old	44.3	13.9	41.8
public sector	53.0	5.9	41.1
private sector	35.6	18.2	46.2
white collar	40.8	15.7	43.5
blue collar	41.2	11.8	47.0
union coverage	52.2	5.5	42.3
no union coverage	35.7	18.6	45.7
Sample Average	**41.0**	**14.4**	**44.6**

Table 2: The Role of Wage Rigidities in Job Destruction

	Not employed	**Not employed**	**Not in same job**	**Not in same job**
Sweep	12.17 (0.33)	14.33 (0.34)	8.39 (0.31)	9.47 (0.32)
Age	-0.06 (0.001)	-0.05 (0.001)	-0.07 (0.001)	-0.07 (0.001)
$(Age/100)^2$	5.87 (0.19)	4.83 (0.19)	6.78 (0.18)	6.01 (0.18)
Male	-0.13 (0.003)	-0.14 (0.003)	-0.12 (0.003)	-0.13 (0.003)
Constant	0.20 (0.03)	0.04 (0.03)	1.08 (0.02)	0.08 (0.02)
Time Dummies	NO	YES	NO	YES
Log-likelihood	-1053078.6	-1049249.4	-1199016.5	-1194786.3
Prob. $\chi^2>0$	0.00	0.00	0.00	0.00
Observations	1,927,262	1,927,262	1,927,262	1,927,262

Appendix: The Estimated Likelihood Function

The econometric model we use to explain the nominal wage growth distribution extends work by Altonji and Devereux (2000) [hereafter AD]. To simplify matters we shall begin by reviewing the AD model, which should also help to highlight the contribution made in this paper. At the heart of the AD model is Akerlof *et. al.* (1996)'s concept of the *notional* wage: the wage that the firm would choose to pay the worker in the absence of the costs imposed upon it by the (unspecified) nominal rigidity *in that period.*[8] These costs create an incentive for the firm to depart from the notional wage and protect a worker from an otherwise desirable cut in their wage.

Preliminaries: A Review of the AD Model

Let us define the logarithm of the notional wage: $\tilde{w}_{it}$ for individual i at time t as a function of a set of explanatory variables: x_{it} and a time dummy: α_t to capture cyclical effects as follows:

$$\tilde{w}_{it} = \alpha_t + x_{it}\beta + \varepsilon_{it} \qquad \varepsilon_{it} \sim N\left(0, \sigma_{\varepsilon}^2\right) \qquad \textbf{(1)}$$

In the baseline AD model an individual's notional wage varies systematically according to her age, her occupation, the industry in which she is employed, and so on. Equation (**1**) captures the fact that the average rate of pay growth is likely to vary across industries. But it is also likely that pay growth is more volatile in some industries than others. That could reflect the fact that product demand is more volatile in certain industries (perhaps through greater exposure to international competition), or that the use of bonus payments, merit pay and other more volatile elements of pay

[8] In the A-D model firms are not assumed to be myopic; they are aware that when they set wages today they may be stuck with paying them next year as well.

varies across industries (and occupations).[9] The most straightforward way to allow for this variation in the volatility of pay growth across different groups in the panel is to endogenise the parameter: σ_ε. In particular we allow for the following relationship:

$$\sigma_{\varepsilon,it} = \sigma_{\varepsilon 0} + s_{it}.\varpi + \varepsilon_\sigma \qquad \textbf{(2)}$$

where the standard deviation of the notional wage growth regression depends on some vector of parameters: s_{it}. Allowing for this heteroscedasticity does not complicate our estimation greatly, but it does offer the potential for us to generate a notional wage growth distribution which is less Gaussian in appearance. In other words by extending the model in this direction we allow that the aggregate nominal wage growth distribution may be a mixture of many Gaussian distributions, each with different means and variances. As a result, there is no restriction on the aggregate distribution to be either uni-modal or symmetric.

Let us assume that whenever the notional wage is greater than the actual wage in the previous year, then the firm pays the notional wage. However, where the notional wage this year is lower than the actual wage last year, firms may choose to leave the actual wage unchanged, up to some threshold value: $-\xi$. Thereafter, firms will cut wages; although the firm may protect the worker from the full extent of the notional wage cut by some factor: ψ. All workers are potentially covered by the nominal rigidity, so the model will not predict any wage growth in the interval $(-\xi, 0)$, and if

[9] Of course firms use of these more volatile elements of pay is in itself endogenous – they do so either because the nature of production allows them, say, to identify effort with greater accuracy (and reward it accordingly), or because the variability of conditions in the product market – from which their demand for labour is derived – require it.

ξ is equal to infinity, the model doesn't predict any *genuine* wage cuts at all (although in the presence of measurement error we may observe pay cuts in the data). So if we define the logarithm of the actual wage as w_{it}, then the model can be summarised as follows:

$$
\begin{aligned}
&w_{it} - w_{it-1} = x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} && if \quad 0 \le x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} \\
&w_{it} - w_{it-1} = 0 && if \quad -\xi < x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} < 0 \\
&w_{it} - w_{it-1} = \psi + x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} && if \quad x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} \le -\xi
\end{aligned}
\qquad \textbf{(3)}
$$

Finally, we have to account for the potential distortions that measurement error can create in the observed wage growth distribution. Assume that with probability: q an observation is correctly measured in the data. Further assume that when errors occur, those errors, m_{it}, are independently and normally distributed across individuals and time periods, with mean zero and variance: ${\sigma_m}^2$. We can therefore construct a composite measurement error term: η_{it} as follows:

$$
\begin{aligned}
&\eta_{it} = 0 && with \quad probability \quad q^2 \\
&\eta_{it} = m_{it} \ or \ m_{it-1} \sim N\left(0, {\sigma_m}^2\right) && with \quad probability \quad 2.q.(1-q) \\
&\eta_{it} = m_{it} + m_{it-1} \sim N\left(0, 2{\sigma_m}^2\right) && with \quad probability \quad (1-q)^2
\end{aligned}
\qquad \textbf{(4)}
$$

Given this composite error term we can therefore re-write our model as follows:

$$
\begin{aligned}
&w_{it} - w_{it-1} = x_{it}\beta + \alpha_t + \eta_{it} + \varepsilon_{it} - w_{it-1} && if \quad 0 \le x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} \\
&w_{it} - w_{it-1} = \eta_{it} && if \quad -\xi < x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} < 0 \\
&w_{it} - w_{it-1} = \psi + x_{it}\beta + \alpha_t + \eta_{it} + \varepsilon_{it} - w_{it-1} && if \ x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} \le -\xi
\end{aligned}
\qquad \textbf{(5)}
$$

Incorporating Real Rigidities

The AD model can explain the spike we observe in the nominal wage growth distribution at zero. However the model cannot explain the distortions that real rigidities may deliver deviations further up the nominal wage growth distribution. In order to capture the impact of these real wage rigidities we need to extend the model presented in the previous section.

We allow that workers may fall under one of three regimes. Wages may be flexible or *unconstrained*, in which case their wage is unprotected and set equal to the notional wage on all occasions. Alternatively, wages may be subject to *nominal* rigidities, where we assume that the nominal wage growth is bound from below by zero. Finally, wages may be subject to *real* rigidities, where we assume that nominal wage growth is bound from below by an expectation of inflation. We thus trade complexity in terms of the number of rigidities we allow for in the data with simplicity in how each of these rigidities are modelled – if a worker's wage is protected, then the growth in their wage is bound from below by zero or an expectation of inflation. In the terminology of the AD model, we assume that in each regime ξ tends to infinity.

From the outset it should be stressed we define the incidence of nominal and real rigidity in wage setting not by the fraction of the workforce whose wage growth is constrained (the size of the spike/cluster) but by the fraction of the workforce whose wage growth is *potentially constrained*. Nor is there any simple one for one correspondence between changes in the size of the regimes and changes in the fraction of the workforce whose wage growth is actually constrained. For example, given a sharp fall in the inflation rate the number of workers protected from pay cuts may rise even if the fraction of the workforce in the nominal rigidity regime is falling.

For a worker who falls in the unconstrained regime we know that the change in the log of the actual wage is given by the change in the log of their notional wage. Once again, allowing for the possibility that wage level observations may be mismeasured, we have that :

$$w_{it} - w_{it-1} = x_{it}\beta + \alpha_t + \eta_{it} + \varepsilon_{it} - w_{it-1} \qquad \textbf{(6)}$$

For those workers who fall under the nominal rigidity regime, we assume that irrespective of the notional wage, wages will not be cut in nominal terms from one period to the next. Therefore, we have that:

$$\begin{aligned} w_{it} - w_{it-1} &= x_{it}\beta + \alpha_t + \eta_{it} + \varepsilon_{it} - w_{it-1} && \text{if} \quad 0 \le x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} \\ w_{it} - w_{it-1} &= \eta_{it} && \text{if} \quad x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} < 0 \end{aligned} \qquad \textbf{(7)}$$

For those in the real rigidity regime, we assume that nominal wage growth will never fall below some expected inflation rate: r_i. We allow for the fact that this expectation of the inflation rate may vary across workers. In particular we assume that these expectations of the inflation rate are normally distributed around some mean expectation: r with variance: $\sigma_r^{\ 2}$. We therefore have that :

$$\begin{aligned} w_{it} - w_{it-1} &= x_{it}\beta + \alpha_t + \eta_{it} + \varepsilon_{it} - w_{it-1} && \text{if} \quad r_i \le x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} \\ w_{it} - w_{it-1} &= r_i + \eta_{it} && \text{if} \quad x_{it}\beta + \alpha_t + \varepsilon_{it} - w_{it-1} < r_i \end{aligned} \qquad \textbf{(8)}$$

Finally, we have to specify which regime workers fall into. It is natural to think of modelling the probabilities that each individual will fall under each of the three regimes: y_{ui} , y_{ni} & y_{ri} with a probit equation. But workers must fall under one of these three regimes so one of these three probit regressions would be redundant: an individual must fall under one of the three regimes with probability one! To avoid introducing adding-up constraints into the model we can re-write these three probabilities as two *relative propensities*, as follows:

$$y_{rni} = y_{ri} - y_{ni} = c_i\vartheta + z_{1i} \quad \textbf{(9)}$$
$$y_{rui} = y_{ri} - y_{ui} = c_i\chi + z_{2i} \quad \textbf{(10)}$$

where c_i is a vector of explanatory variables and the z_{ni}'s are standard normal variables. It then follows that the probability of falling in a given regime can be written as a product of a number of cumulative density functions of the standard normally distributed variable. The probabilities of individual i falling in the real ($\boldsymbol{R}$), nominal ($\boldsymbol{N}$) and unconstrained ($\boldsymbol{U}$) regimes are then given as follows:

$$\Pr\{ i = \boldsymbol{R}\} = \Pr\{y_{ri} - y_{ni} > 0\} \times \Pr\{y_{ri} - y_{ui} > 0\} = \Phi\{c_i\vartheta\} \times \Phi\{c_i\chi\} \quad \textbf{(11)}$$
$$\Pr\{ i = \boldsymbol{N}\} = \Pr\{y_{ri} - y_{ni} < 0\} \times \Pr\{y_{ni} - y_{ui} > 0\} = \Phi\{-c_i\vartheta\} \times \Phi\left\{\frac{c_i\chi - c_i\vartheta}{\sqrt{2}}\right\} \quad \textbf{(12)}$$
$$\Pr\{ i = \boldsymbol{U}\} = 1 - \Phi\{c_i\vartheta\} \times \Phi\{c_i\chi\} - \Phi\{-c_i\vartheta\} \times \Phi\left\{\frac{c_i\chi - c_i\vartheta}{\sqrt{2}}\right\} \quad \textbf{(13)}$$

We now briefly summarise our approach to estimating the values of our set of parameters: $\beta, \vartheta, \chi, \sigma_e, \sigma_r, \sigma_m, q, r$, given individual data: x_i, c_i.

Estimation of the Model

The essence of this paper is the estimation of the time series variation in the fractions of the workforce who are protected from nominal and real pay cuts through the distortions they create in the nominal wage growth distribution. In this section we briefly discuss how we overcome some practical difficulties we have to overcome in estimating our model.

Errors in the data will tend to smooth out any spikes in the underlying distribution of nominal pay growth that have been created by rigidities. It follows that if the incidence of errors in the data varies from year to year then this may lead to a corresponding variation in our estimates of the fraction of the workforce who lie in each of the regimes.[10] Of course, any errors in the NESPD ought to be rare, since our data is reported from the payroll records of employers. More importantly, given the size of the NESPD we would not expect the actual incidence of these errors to vary significantly from year to year. We therefore allow for a small yet significant degree of measurement error and constrain the parameter: q to be constant across the sample period. In the results presented in this paper we assume that 89 % of observations in the NESPD are correctly measured.[11]

One of the key aims of this paper is to recover time series data on the fraction of the workforce who fall in each of the regimes: where the (expected) real wage is protected ($\boldsymbol{R}$), where the nominal wage is protected ($\boldsymbol{N}$) and the where the wage is completely unconstrained ($\boldsymbol{U}$). These in turn are determined by the explanatory variables: c_i, and the coefficients: θ, χ of the model. If the model is estimated year by year the marginal effect of a given covariate on the relative propensity to fall in each of the regimes may be found to be extremely erratic over time. Allocation of a particular group of workers into either the nominal or real rigidity regime will essentially depend on there being significant numbers of those workers clustered either around zero or the identified expectation of inflation. It would not be surprising to find that in one particular year the vast majority of individuals in a given industry

[10] For example in the case of the nominal regime, for a given spike in the observed nominal pay growth distribution, the greater the incidence of errors in the data the greater the underlying spike must have been in the actual distribution, and therefore by implication the greater the proportion of the sample who must be in the nominal regime.

[11] This value was chosen via a grid search – it was the value of q which gave the highest average value of the likelihood function.

received relatively large nominal pay increases. If the model is estimated on a year by year basis, we are likely to find that working in that industries in that particular year is linked with membership of the unconstrained regime.

We therefore estimate the window over a longer time frame – pooling five successive observations on wage growth (i.e., six years of wage level observations from the NESPD) – and constrain the marginal impact of most of the covariates to be fixed across the period.

Our model assumes that if workers fall into a rigidity regime, their actual wage growth distribution is bound from below, so that only wage changes above these thresholds are set according to the notional wage relationship. So it may be difficult to identify that notional wage relationship using standard regression techniques which will focus on the conditional mean of nominal wage growth. Instead, we choose to focus on the right-hand tail of the nominal wage growth distribution, since rigidities will not bind in this region.[12] Our approach to identification of the dependence of the notional wage on our set of covariates is therefore to rely on quantile regression techniques.[13]

In the same way that classical regression techniques allow us to identify the relationship between the conditional mean of the dependent variable and a set of covariates, quantile regression techniques allow us to identify the relationship between dependent and independent variables at any given point in the conditional distribution of the dependent variable. More formally, while the classical linear

[12] Of course the further we move into the tails of the wage growth distribution the more likely we are to be capturing reporting errors.
[13] The following brief exposition of these techniques draws heavily on Koenker and Hallock (2001).

regression model derives an estimate of the (linear) conditional mean function by solution of:

$$\underset{\beta \in \Re}{arg\,min} \sum_{i}^{N} \left(y_i - x_i \beta \right)^2 \qquad \textbf{(14)}$$

the linear quantile regression model derives an estimate of the conditional quantile function from the solution of:

$$\underset{\beta \in \Re}{arg\,min} \sum_{i}^{N} \rho_{\tau} \left(y_i - x_i \beta \right) \qquad \textbf{(15)}$$

where $\rho_{\tau}(u)$ is the 'tilted' absolute value function[14] defined as follows

$$\rho_{\tau}(u) = \left(\tau - I\{u < 0\} \right).\, u \qquad \textbf{(16)}$$

In practice we obtain a given conditional quantile of the notional wage growth distribution – where the quantile is chosen such that it is far enough from the median to ensure rigidities do not bind, yet close enough to minimise the loss of efficiency in our estimator – from which we can recover an estimate of the conditional mean of the wage growth distribution.

If, as we have assumed, the errors of these quantile regressions are homoscedastic then for a given estimate of the parameter: σ_{ε},[15] the median (and hence the mean, for

[14] The intuition here is that minimising the sum of asymmetrically weighted absolute residuals – giving different weights to positive and negative residuals – yields the conditional quantile function, in the same way that minimising the sum of absolute residuals yields the conditional median function (Koenker and Hallock (2001) pp.145-6).

[15] Of course, we have allowed the standard deviation of the disturbances in the notional wage relationship to vary systematically across individuals, according to [**4.2**]. However, at each stage in the maximisation process, we obtain an estimate of

a normally distributed error) predicted wage change: μ_{50} is pinned down. We can thus specify the mean notional wage growth of each individual in the panel, so that only the error in the notional wage growth regression is estimated alongside the remaining parameters of our model in the maximisation problem. Of course, identifying the notional wage relationship outside the main maximisation problem has practical advantages as well – in so doing we reduce the size of the parameter space of the likelihood function significantly.

The remaining parameters of our model are then estimated using maximum likelihood techniques. Given the nature of the model, and in particular the number of regimes the model allows and the composite nature of the error term, the actual likelihood function is far from straightforward. In what follows we therefore offer a heuristic description of the likelihood function; more details are provided in Appendix 2 of this paper.

The Likelihood Function

Given our assumptions, the growth in the nominal wages of workers who fall under the unconstrained regime will be normally distributed, with a mean pinned down by their predicted notional wage (given the actual wage last year), and a standard deviation which will depend on their characteristics, the standard deviation of measurement error in the data, together with the probability of measurement error occurring. However, for those workers who fall under either the nominal or real rigidity regime, the probability density function of actual wage growth is truncated, with a potentially large point mass at either zero or the expected inflation rate. When

---

$\sigma_{\varepsilon,it}$ for each individual given their characteristics, from which we can obtain an estimate of their mean notional wage.

constructing the likelihood we therefore distinguish between two types of observation in each of these rigidity regimes: those where the rigidity does not have an impact on wage growth (firms pay the notional wage) and those where it binds (notional wages are swept up to that lower bound). Our likelihood function therefore reflects this collection of observations into five distinct categories:

- Individuals who fall under the unconstrained regime ($\boldsymbol{U}$).
- Individuals who fall under the nominal rigidity regime ($\boldsymbol{N}$), whose wage growth is constrained and swept up to spike at zero.
- Individuals who fall under the nominal rigidity regime ($\boldsymbol{N}$), whose wage growth is unconstrained.
- Individuals who fall under the real rigidity regime ($\boldsymbol{R}$), whose wage growth is constrained, and swept up to the individual specific lower bound: r_i.
- Individuals who fall under the real rigidity regime ($\boldsymbol{R}$), whose wage growth is unconstrained.

Given a set of *observed* rates of nominal wage growth: $D = \{d_i\}$ and an underlying set of unobserved *actual* rates of nominal wage growth: $A = \{a_i\}$ then for a given set of explanatory variables: x_i, c_i, s_i (and suppressing reference to the time period t) the likelihood function can thus be written down as:

$$
\begin{aligned}
L(D|\zeta) \; = \; & \prod_i^N Pr(i \in \boldsymbol{U}|c_i) \times L(d_i | i \in \boldsymbol{U}, x_i, s_i) \\
& + Pr(i \in \boldsymbol{N}|c_i).I\{a_i = 0\}.Pr(a_i = 0 | i \in \boldsymbol{N}, x_i, s_i).L(d_i | i \in \boldsymbol{N}, x_i, s_i) \\
& + Pr(i \in \boldsymbol{N}|c_i).I\{a_i > 0\}.Pr(a_i > 0 | i \in \boldsymbol{N}, x_i, s_i).L(d_i | i \in \boldsymbol{N}, x_i, s_i) \\
& + Pr(i \in \boldsymbol{R}|c_i).I\{a_i = r_i\}.Pr(a_i = r_i | i \in \boldsymbol{R}, x_i, s_i).L(d_i | i \in \boldsymbol{R}, x_i, s_i) \\
& + Pr(i \in \boldsymbol{R}|c_i).I\{a_i > r_i\}.Pr(a_i > r_i | i \in \boldsymbol{R}, x_i, s_i).L(d_i | i \in \boldsymbol{R}, x_i, s_i)
\end{aligned}
\qquad \textbf{(17)}
$$

where $I\{\cdot\}$ denotes the indicator function, and $\zeta = \{\beta, \psi, \vartheta, \chi, \sigma_e, \sigma_r, \sigma_m, r\}$ denotes our vector of model parameters, where β is estimated outside the model. The contribution of each observation to the likelihood is therefore defined by a product of three terms: the probability that each individual falls in a given regime, the probability that the observation is constrained, and the likelihood of the observation, conditional on the regime and whether the observation is constrained or not. Given the assumptions of normality made on the distribution of $\varepsilon_i , m_i , r_i$ the likelihood can then be maximised in the standard fashion.

Federal Reserve Bank
of Cleveland
Research Department
P.O. Box 6387
Cleveland, OH 44101