

NOMINAL WAGE RIGIDITY IN THE UNITED KINGDOM

*Jennifer C. Smith**

This paper studies the degree of downward rigidity in nominal wages in the United Kingdom using micro-data. Around 9% of employees who remain in the same job from one year to the next have zero pay growth. But on investigating the causes of rigidity we find that up to nine-tenths can be attributed to 'symmetric' causes (such as contracts and menu costs) or to error. Thus only 1% of workers have pay that may be downwardly rigid. This suggests asymmetric, downward rigidity is not large enough to have serious macroeconomic consequences. The labour market provides almost no evidence to support a positive inflation target.

The idea that nominal wages might be rigid has a long history, stretching back at least as far as Keynes (1936). In addition to any intrinsic interest in the behaviour of the price of labour, the issue has important implications for the macroeconomy. Downwardly rigid nominal wages could provide a rationale for a positive optimal rate of inflation and hence for the pursuance of a positive inflation target. As Tobin (1972) argued, if wages are rigid downwards, negative demand or positive supply shocks leave real wages higher than marginal product, resulting in unemployment. Higher inflation could alleviate this unemployment by reducing the real wage. Thus the cost of downward nominal rigidity – which can also be interpreted as the cost of very low inflation where downward rigidity exists – is unemployment.

Downward rigidity implies asymmetry in wage changes: nominal wages are more likely to rise than to fall. Symmetric rigidities, whereby nominal wages are sticky in both directions, can also have detrimental macroeconomic implications. Unemployment is a feature of many models of long-term or staggered contracts (see, for example, Mankiw and Romer, 1991). However, symmetric rigidities do not have the same implications for the optimal rate of inflation: the effects of symmetric rigidities cannot be alleviated by a higher inflation rate.

In this paper, nominal rigidity is taken to mean a 0% nominal pay change. Clearly wages are in some sense 'rigid' if they do not move as much as they 'should'. It is almost impossible in practice to calculate how much pay 'should' change, although a simple conception of this warranted growth might be related to the change in the individual's productivity. A 0% nominal wage change has potentially special features, however. Psychological studies, which focus on downward rigidity, suggest that zero is a barrier of special importance (Kahneman *et al.*, 1986; Blinder and Choi, 1990). Menu costs also impart

* I would like to thank Mark Stewart, Tony Yates, seminar participants at University College London, Warwick, Strathclyde, the CEPR Unemployment Dynamics Workshop and two anonymous referees for helpful comments. The data used in this paper were made available through the UK Data Archive. The data were originally collected by the ESRC Research Centre on Micro-social Change at the University of Essex, now incorporated within the Institute for Social and Economic Research. Neither the original collectors of the data nor the Archive bear any responsibility for the analyses or interpretations presented here. All remaining errors are my own.

special significance to 'no change'. For these reasons macroeconomic theories incorporating nominal rigidity focus on the 0% point.

The main questions addressed in this paper are straightforward: how extensive is nominal rigidity in general in the UK labour market? In particular, are nominal wages *downwardly* rigid, or is the labour market of the 1990s sufficiently flexible that workers are willing to take cuts in their nominal pay?

It is only over recent years that the reduction in UK inflation to low levels has rendered at all sensible the study of the extent of nominal rigidity in this country. While inflation is in double figures, we can be pretty sure that very few prices, be they of goods or labour, will remain constant for a sufficiently long time to give rise to any costs. Retail price inflation remained below 4% during 1992–6, the period of the current study. Any research attempting to investigate the cost of zero inflation is liable to be criticised on the basis of the 'Lucas critique': microeconomic behaviour may alter when the macroeconomic environment changes. This problem is presumably lessened, the closer are the actual data to the hypothetical situation.

We use the first six waves of the British Household Panel Study (BHPS), the only UK panel survey of individuals, to investigate the extent of rigidity in pay. No comparable work has previously been conducted for the United Kingdom. Five papers have recently studied the issue using similar data from the United States (Akerlof *et al.*, 1996; Card and Hyslop, 1996; Kahn, 1997; Lebow *et al.*, 1995; McLaughlin, 1994).¹ There is complete disagreement among these authors as to the scale of nominal rigidity. McLaughlin concludes that the evidence suggests surprising flexibility in the labour market. Lebow *et al.* judge rigidity to be small in itself and in its consequences. Card and Hyslop find greater rigidity, but even this is not thought large enough to have major macroeconomic consequences. Kahn finds significant downward rigidity for wage earners, but none for salaried workers by the mid-1980s. Akerlof *et al.* dismiss apparent flexibility as merely reflecting measurement error, and maintain that nominal rigidity is important.

UK data on nominal pay growth present a striking picture of a large 'spike' at zero. We examine various possible 'explanations' for this observed rigidity and for what we argue is surprisingly substantial flexibility. The explanations we look to include rounding, long-term contracts and measurement error. Truly rounded pay (resulting not from reporting errors but, for example, from menu or computational costs) and long-term contracts are behavioural features of the labour market that result in symmetric rigidity – but any rigidity that can be attributed to these causes can be subtracted from downward rigidity. Rounding error means that observed rigidity overstates the problem, and this component can also be discounted. Other measurement error is typically thought to reduce observed rigidity, as it would if it were classical. We use data on whether the payslip was checked – a unique feature of the BHPS –

¹ Recent papers by Christofides and Leung (1998), Crawford and Harrison (1998) and Fortin (1996) study the issue using Canadian union contract data. Christofides and Stengos (1998) use these data and some individual data similar to US work and the present study.

to analyse the effect of measurement error. If the payslip was checked in both years, no measurement error is made. Contrary to expectation, we find rigidity higher in the subsample that is subject to measurement error. We hypothesise that people who make errors have an idea of what their pay is that remains fixed over a significant period of time. A part of the observed rigidity can thus be attributed to measurement error. After taking into account these factors, we calculate that only around 1% of workers have pay that might be downwardly rigid. This is a much smaller fraction than is apparent from an initial simple inspection of the data, and is lower than has been found in the United States. Little justification for a positive inflation target can be drawn from such a small amount of downward rigidity.

Some of the results in this paper may seem difficult to believe – the quite common occurrence of nominal pay cuts, for example. It may well be that the difficulty in believing them stems not from the weight of contradictory evidence, but rather from conventional wisdom that has survived because of the previous lack of evidence either way.

The paper is organised as follows. Data are described in Section 1. Section 2 presents the basic facts concerning the extent of nominal rigidity and flexibility in the UK labour market in recent years. Possible ‘explanations’ for rigidity and flexibility are examined in Section 3. Section 4 concludes.

1. Data

We use the first six waves of the British Household Panel Study, containing the results of interviews with around 10,000 individuals each year during 1991–6. There are several issues to consider in deciding on a measure of annual pay change. The basic notion that drives our choice is the need to create a ‘like-for-like’ pay measure: we want to know that nominal rigidity is not apparently absent due to task changes which merit a different pay rate, for example. Thus we want to restrict attention to those who have not changed jobs. Unless otherwise specified, all data in this paper refer to ‘stayers’. This excludes both intra- and inter-employer job changers (promotions and grade changes are counted as job changes).²

The individual-level data we consider contain information on total pay but not specifically on basic pay, and nor do they distinguish workers paid on an hourly basis from others. Use of total pay will reduce the apparent extent of rigidity: basic wages could be rigid while there could be flexibility in other components, especially bonuses and overtime. We do have information about whether workers receive bonuses, whether they work paid overtime, and total hours worked, so we can study pay growth when there are no hours changes

² Stayers are defined as those whose job has not changed since 1 September the previous survey year. This excludes all those who changed job between the previous and current interview, but also excludes a very small number of people who were interviewed later in the previous survey period and changed job between 1 September that year and their interview date.

and get close to a measure of the increase in basic wages. We investigate the impact of these factors in Section 2.

The BHPS gives us raw data on pay over a given period of time and derived measures of monthly pay. We study all employees: the reported pay of the self-employed is considered too unreliable. We choose to study gross pay measures, since the use of net measures would confuse pay changes with changes in taxes and allowances, pension contributions, union dues and other factors. We analyse workers' stated *usual* gross pay, rather than their latest pay, because latest pay might be distorted by unusual bonuses, overtime payments and so on.³ Our preferred pay measure is calculated as usual gross pay divided by the number of weeks in the usual pay period. We omit cases where the BHPS imputed data for gross usual pay or calculated gross values from stated net pay in either year. Each technique will induce measurement error of classical form which will tend to reduce misleadingly the observed relative frequency of rigid pay.⁴ We also discard large outlying observations which are clearly due to error.⁵

We choose to focus on pay per week, rather than pay per hour. An hourly pay measure can be created by dividing pay by the 'normal' number of hours worked per month (which correspond to 'usual' pay). Dividing pay by reported hours induces additional measurement error, as hours are reported with error. Measurement error is an important issue when studying the extent of rigidity: measurement error might be mistaken for flexibility. Errors in hours are likely to be greater than errors in pay: the former are more likely to be guessed at or approximated by workers.⁶ (In many cases there is no written record of hours worked, whereas payslips record pay and pay period.) On these grounds, our preferred measure is growth in gross usual weekly pay. We calculate percentage changes as the difference in pay divided by lagged pay to accord with the method used in actual pay settlements. We have just under 9,500 observations on annual pay growth during 1992–6.

The BHPS has a major advantage over other datasets (including the US Panel Study of Income Dynamics (PSID) and Current Population Survey (CPS)) in that it allows employees to check their payslip in reporting their pay level and records whether such a check was made. The advantage this confers

³ The disadvantage with the usual pay measure is its subjectivity. Any measurement error thereby induced needs to be balanced against the measurement error that would result from the exceptional items included in latest pay. We consider the effects of the former likely substantially less than the latter. We investigate the impact of measurement error in Section 3.1

⁴ These affect over 4,500 observations. The effect on estimated rigidity is clear: with these observations sample rigidity is 6.2%; excluding them rigidity rises to 9.0%.

⁵ Errors include, for example, recording pay as referring to a week when comparison with other years' pay and the fact that no job change took place makes it clear the pay refers to a month. We also exclude individuals who record the pay period as less than a week as investigation suggests this reflects error. In addition, we exclude those who switch status between full and part time. Simple hours changes might also generate spurious flexibility in weekly pay (we investigate this in Section 2) but status switches are found to lead to the largest changes. These changes affect just over 100 cases per wave. Qualitative and in many cases quantitative conclusions are unaltered because large errors affect only the extreme tails of the distribution.

⁶ For the similar PSID, validation studies by Duncan and Hill (1985) and Bound *et al.* (1989) concluded that large errors were introduced by dividing reported earnings by reported hours worked.

is that for the paylips-seen subsample we know there is minimal measurement error in pay. We use this information in Section 3 to examine the impact of measurement and rounding error on apparent rigidity and flexibility.

2. Evidence from the Distribution of Individuals' Nominal Pay Growth

Downward nominal rigidity will cause censoring in the nominal wage growth distribution at zero. Observations will be 'missing' from the distribution below zero, having been swept up into a zero 'spike', and the distribution will consequently be asymmetric. The reasons lying behind a reluctance to accept or award nominal pay cuts are usually thought to be psychological. Kahneman *et al.* (1986) and Blinder and Choi (1990) discuss and provide evidence that people have less aversion to cuts in real wages through inflation exceeding zero-or-above nominal pay growth than through cuts in nominal pay. This may involve money illusion. Keynes (1936) famously pointed to the importance of relative wages in motivating this behaviour (see Yates (1998) for a discussion of reasons for downward nominal rigidity).

The mean of the wage growth distribution will shift with inflation and productivity growth. The lower these are, the greater the proportion of workers that would be affected by downward rigidity, and hence the larger the spike and the more asymmetric the distribution. We begin by investigating the data to see whether they accord with these predictions.

Fig. 1 shows histograms of annual growth in nominal usual gross weekly pay for UK employees who remained in the same job, using pooled data and for individual years.⁷ A substantial spike at zero nominal pay change is clearly visible each year. The zero bin contains between 11% and 15% of non-job changers. However, only some of these workers had rigid pay; others experienced very small pay changes. In fact, on average 9.0% of non-job changers experienced exactly zero annual growth in their usual gross weekly pay during 1992–6 (see Table 1).⁸ The proportion rigid varies inversely with the location of the distribution: the higher is the median, the lower is rigidity (but the relationship with inflation is less clear). Our initial examination of the data appears consistent with substantial downward nominal rigidity. In most years there is a perceptible absence of pay changes in the bin above the zero bar, and in some years observations seem to be lacking just below zero, suggesting in addition symmetric causes of rigidity.

The frequency of nominal pay rigidity in the United Kingdom appears on the face of it remarkably similar to that found in the United States. On average over 1971–88, 8% of job stayers were found to have had zero nominal pay change during 1971–88 (Kahn (1997) and Lebow *et al.* (1995), using data from the PSID). A similar 7% of stayers had zero nominal pay growth over

⁷ 50 bins are used, which are chosen so as to exclude a roughly equal fraction (around 5%) of workers at either end of the distribution, to enable the centre of the distribution to be seen clearly.

⁸ Data on hourly pay show lower rigidity and more frequent cuts (4.4% and 30% respectively over 1992–6) which is likely due to measurement error in hours as discussed above.

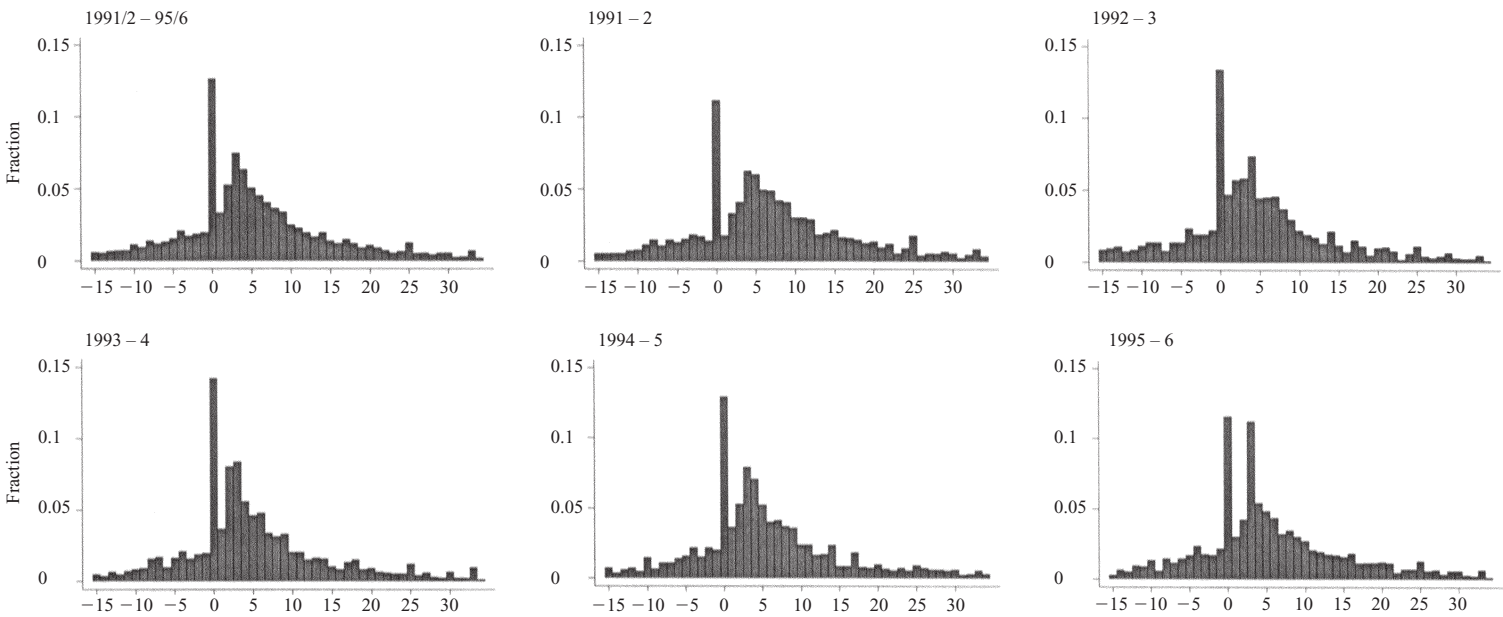


Fig. 1. Annual Nominal Pay Growth

Table 1
The Extent of Nominal Rigidity and Downward Flexibility

Year	% Rigid	% Cuts	Median nominal pay growth (%)	Inflation (%)	Productivity growth (%)
1991–2	8.1	20.9	5.9	4.0	1.7
1992–3	9.7	25.1	3.4	1.9	3.5
1993–4	10.0	22.9	3.2	2.2	3.3
1994–5	9.4	22.5	3.8	3.3	2.5
1995–6	7.8	23.4	3.7	2.6	1.8

Note: % Rigid = Proportion of non-job changers with zero annual growth in nominal usual gross weekly pay. % Cuts = Proportion of non-job changers with negative annual growth in nominal usual gross weekly pay. Source of price and productivity data: NSO. Inflation is October to September average of annual RPI growth (since pay changes may occur at any time in the year between interviews, which are conducted from September). Productivity growth is Q4 to Q3 average of annual growth in whole economy productivity.

1976–86 (McLaughlin (1994), using the same data source). But these averages mask substantial variation over time (and with inflation): in the United States, the proportion of (hourly-paid) workers experiencing zero wage change varied from about 7% in a 10%-inflation environment to over 15% when inflation fell to 5% (Card and Hyslop (1996), using the PSID and the CPS).⁹ Hence UK rigidity is noticeably lower than that in the United States at comparable inflation rates.

We would expect wage stickiness to decline as the time period we consider lengthens. This could be due to nominal contracts coming to an end (see Section 3.3), and also to the fact that as the pay growth distribution shifts upward (as it will with cumulatively higher inflation and productivity growth over time) a smaller proportion of workers have ‘warranted’ pay changes less than zero,¹⁰ so censoring would affect a smaller fraction of the cumulative distribution. Fig. 2 shows pay changes over two to five year periods.¹¹ The picture is as expected. Over a two-year horizon, the spike at zero is still quite noticeable, and 4% of non-job changers had zero pay growth – less than half the proportion rigid over one year. Over three years the spike is still visible and 2.5% of non-movers had nominally rigid pay. Over four years less than 2% had no pay change. Over a five-year horizon the spike is much less striking: only 0.9% of workers who remained in the same job throughout this period had rigid pay. A comparison of BHPS and PSID data reveals that UK nominal rigidity over the longer term remains about half that in the United States for comparably low rates of inflation (see Lebow *et al.*, 1995, and Card and Hyslop, 1996).

Nominal pay cuts are very frequent in the United Kingdom: on average 23% of workers suffered nominal cuts in their usual weekly gross pay (for the same

⁹ Rigidity is typically found to be higher for hourly-paid than for salaried workers.

¹⁰ ‘Warranted’ pay changes are those that maintain equality of real pay with the value of the worker’s marginal product.

¹¹ The histograms use the same end-points as Fig. 1 for ease of comparison.

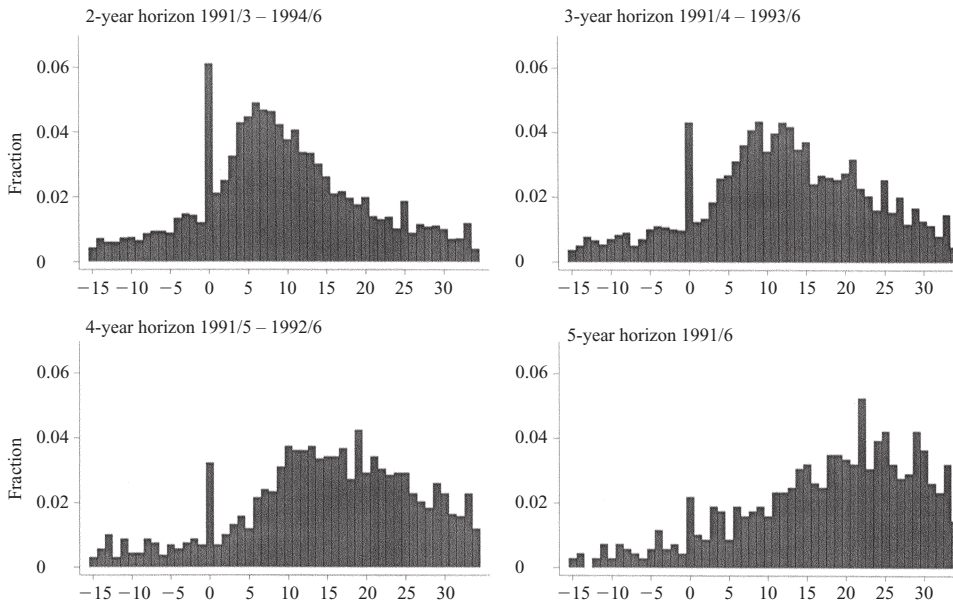


Fig. 2. *Nominal Pay Growth over Two to Five Year Horizons*

job) over a one-year span during 1992–6 (see Table 1). The median annual cut was 7.6%. Even over longer periods, a large proportion of workers suffered cuts in their nominal level of pay, even though they remained in the same job.

This degree of downward flexibility is surprisingly high compared to the United States: there, on average over 1976–86, 17% of those who did not change jobs experienced nominal pay cuts over a one-year period (McLaughlin (1994), using the PSID). Lebow *et al.* (1985, also PSID) and Card and Hyslop (1996, CPS data) find similar proportions of 18% and 15–20% respectively. However, in the United States cuts tend to be larger: these nominal cuts averaged nearly 12% (McLaughlin, 1994). There is no clear support from UK data for a negative relationship between the proportion workers taking pay cuts and the location of the pay growth distribution, as we would have expected if downward rigidity were prevalent, although there is some US evidence for this (see Card and Hyslop, 1996).

In the next section we suggest that apparent rigidity is largely not downward rigidity. Instead, it reflects rounding (both ‘true’ and in error), long-term contracts and measurement error. Together with the substantial downward flexibility we have noted, we suggest this means the UK labour market is characterised by downward rigidity to a far lesser extent than has typically been suggested. But before explaining away apparent rigidity, we will examine apparent flexibility further.

First, because our data refer to total pay they include bonuses and overtime. These non-basic wage components are likely to be more variable than the basic wage itself. Indeed, the receipt of overtime pay serves to significantly raise the variance of pay growth, and increases the frequency of nominal cuts from 21%

to 28%. But bonuses do not affect flexibility: the proportion of those who receive bonuses taking cuts is the same as those who do not (p-value 0.19). These figures indicate that bonuses and overtime pay do not explain all flexibility: a substantial proportion of workers whose pay consists primarily of the basic wage experience cuts (21% of those in receipt of neither bonuses nor overtime pay).

Second, we have focused on weekly pay rather than pay per hour to avoid the measurement error problems described in Section 1. But changes in hours could misleadingly result in weekly pay changes although hourly pay was constant, raising apparent flexibility and reducing rigidity. In order to assess the extent of this we can look at pay growth for those whose hours did not change, which minimises the measurement error problem.¹² Rigidity is higher for these workers (11% compared to 8% for those whose hours change). But although 5 percentage points fewer non-hours changers experienced cuts than hours changers, a substantial 20% did face reductions in their nominal pay.

What flexibility is left once we eliminate hours changers and those in receipt of bonuses or overtime pay? 19% of non-job changers with no bonus or overtime payments nor hours changes experienced nominal cuts. Although this is 4 percentage points lower than all stayers, it is clear that downward flexibility remains high and is largely not attributable to changes in hours or non-basic-pay components.

3. Explaining Rigidity and Flexibility

In this section we examine in some detail three possible ‘explanations’ for the observed pay growth distribution: measurement error, rounding and long-term contracts. The first is generally thought to raise variation in observed pay growth, while the last two are typically regarded as likely to lead to rigid pay. It is important to look to these types of explanation because we can discount rigidity or flexibility due to these factors as affecting the costs of very low inflation, although they may nevertheless be of macroeconomic significance since contracts and truly rounded pay reflecting menu costs will affect the economy’s response to shocks.

3.1. *Measurement Error*

It is widely thought that measurement error will increase the dispersion of pay growth, raising apparent flexibility and reducing the observed extent of rigidity. We will show that the former statement is true but not the latter: although true flexibility is lower than suggested by error-ridden pay data, measurement error also leads the data to overstate rigidity.

Widely varying views have been put forward as to the importance of measurement error in comparable datasets. ‘The empirical patterns of wage variability

¹² Hours include usual and overtime hours since pay includes overtime pay. Around 40% of the sample had no change in hours from one year to the next.

do not appear to be driven by measurement error. The wage data are fairly clean, and wide dispersion of the wage growth distribution survives econometric corrections for measurement error' (McLaughlin, 1994, p. 385). McLaughlin estimates that allowing for measurement error could reduce the standard deviation (of real wage growth) from 14.2 to 9.5. He compresses the data to conform to this revised figure, preserving the mean, and finds that the frequency of nominal wage cuts falls 5 percentage points to 12%, and their average size from 12% to 8%. Lebow *et al.* (1996) agree that measurement error should be relatively small. They refer to the PSID validation study (relating to a sample of workers employed by a single large firm) analysed by Duncan and Hill (1985) and Bound *et al.* (1989). Weekly earnings were found to be reported reasonably accurately, and errors in the change in earnings are not much larger than those in levels, due to a negative correlation between the errors and true earnings. Similar findings have been reported for the CPS by Bound and Krueger (1991).

In contrast, Akerlof *et al.* (1996) argue that measurement error is responsible for virtually all negative wage changes in individual survey datasets. They suggest the true distribution is characterised by asymmetry entailing almost total censoring at zero. If a normal measurement error is added to such a censored distribution, many negative values would be generated. Merely reducing their frequency by compressing the variance, as in McLaughlin, does not in these circumstances recover the true distribution. Akerlof *et al.* perform an interesting exercise based on telephone survey evidence they collected, which includes only 2.7% cuts and a massive 45% rigid. They add an error which is normally distributed with standard deviation 0.167 (which is their estimate of the standard deviation of the difference between employee- and employer-reported wages from the 1977 CPS validation study) and unit autocorrelation to 55.8% of the distribution, and assume that 44.2% make no error. This generates a distribution with fatter tails than the PSID, implying that a 'clean' PSID would reveal even fewer cuts than their own survey. Akerlof *et al.*'s evidence has been cited by US monetary policymakers as helping justify a positive inflation rate.

Uniquely, the BHPS gives interviewees a chance to check their payslip when reporting their pay. The survey records whether the payslip was checked, and even whether it was the latest slip or an earlier one. (In the vast majority of cases – 84% – it was the latest payslip that was seen.) Measurement error is likely to be minimal in cases where the payslips were seen. Table 2 describes 'true' flexibility among those whose payslips were seen (see also Fig. 3).¹³ 7% fewer of the payslips-seen report cuts than the payslip-unseen. Allowing for measurement error in the full sample reduces the frequency of pay cuts by over 5 percentage points (23% of the full-sample figure). Nevertheless, a surprisingly large 18% of employees truly experience nominal pay cuts.

¹³ Because we focus on pay growth, payslips must have been checked both this interview and the previous in order to eliminate measurement error. 'Payslips seen' refers to the latest or an earlier payslip having been checked in both interviews. 'Payslips not seen' includes all whose payslip was not checked in at least one of these interviews.

Table 2
The Influence of Measurement Error, 1992–96

	Proportion of sample (%)	% Cuts	% Rigid
Payslip not seen	70	25.2	10.4
Payslips seen	30	17.8	5.6
Percentage in full sample possibly due to measurement error	n/a	23	38

Note: Payslips-seen subsample is measurement-error-free. Payslip-not-seen subsample may be subject to measurement error. Calculation of proportions in full sample possibly due to measurement error is based on the possibility that all cuts and rigidity among payslip-not-seen subsample in excess of payslips-seen subsample might be due to measurement error.




Fig. 3. *Nominal Pay Growth Among Those Whose Payslips Were Not Seen Compared to the Payslips-seen*

Despite exhibiting less flexibility, the measurement-error-free subsample does not feature a higher degree of rigidity. At 5.6%, rigidity in nominal pay growth is almost 50% lower for the payslips-seen subsample than for the payslip-unseen (Table 2). This is initially surprising: it might be expected that eliminating measurement error would have most effect in the tails of the pay growth distribution, reducing the variance of the sample and lowering apparent pay flexibility, but would thereby increase the apparent likelihood of rigidity. Adding a classical measurement error to the true pay distribution generates higher dispersion in pay growth and lower measured rigidity.

The actual error process may be complex. Let the measurement error in pay levels be u_{it} , Δw_{it} be observed pay growth and Δw_{it}^* be true 'latent' pay growth. Then

$$\Delta w_{it} = \Delta w_{it}^* + \Delta u_{it} \quad (1)$$

and the proportion observed rigid is

$$\Pr(\Delta w_{it} = 0) = \Pr(\Delta w_{it}^* = a) \times \Pr(\Delta u_{it} = -a | \Delta w_{it}^* = a) \quad (2)$$

where a is any value of latent pay growth.

There are several possible explanations for why we observe $\Pr(\Delta w_{it} = 0) > \Pr(\Delta w_{it}^* = 0)$. Considering the special case of $\Delta w_{it}^* = 0$ is misleading:

from expression (2) this implies $\Pr(\Delta u_{it} = 0 | \Delta w_{it}^* = 0) > 1$, which is obviously impossible. The observed data could be generated by a specific form of error process, whereby if a worker reports her pay with error in one year, she is likely to report the same pay the next year. This could be explained by the worker having an inflexible idea of their pay, a perception of their pay that does not change over time. In particular, we can from expression (1) write $\Pr(\Delta w_{it} = 0) = \Pr(\Delta u_{it} = -\Delta w_{it}^*)$: the error changes to exactly offset actual pay growth.¹⁴ If some workers make this type of reporting error it could generate the higher spike at zero among the payslip-unseen.¹⁵

To conclude, substantial downward flexibility remains after we eliminate measurement error. This is contrary to previous claims in the literature. Furthermore, the measurement error process serves to exaggerate rigidity. This is also contrary to what has previously been assumed, and casts some doubt on evidence that has been used to support a positive optimal inflation rate in the United States. Over 1992–6, rigidity among the payslips-seen of 5.6% compared to 9.0% for all stayers indicates that if measurement error were eliminated, rigidity would fall by 35–40% (3.4 p.p.). Overall, then, our investigation of measurement error has suggested that the UK labour market exhibits remarkably high nominal flexibility, and we have already knocked down the zero spike to 5.6%.

3.2. Rounding

Some of the spike at zero might be caused by workers rounding very small pay changes to zero. This could occur from above as well as below: there appears no *a priori* reason to consider rounding from one side more likely. Rounding might arise through error, or pay might be truly rounded. In the latter case, the worker's productivity might just happen to merit pay of a round number, or rounding might reflect some sort of cost or bounded rationality. This would have a similar effect to menu costs, in the sense that it draws observations from either side of the rounded value.

In order to assess the extent of rounding, we need to turn to the 'raw' pay data: the amount stated by the individual in the interview. We construct a variable that tells us whether the raw pay that underlies usual gross weekly pay is rounded or not. The relevant raw pay figure is latest gross pay if that was usual and usual pay if the latest payment was unusual. Because respondents are asked to answer to the nearest £1, pay is already rounded to this order. Respondents are allowed to choose pay covering whatever period they choose:

¹⁴ There is some evidence from the United States that error and wages are negatively correlated (see Duncan and Hill, 1985, and Bound *et al.*, 1989).

¹⁵ There is also the possibility of sample-selection bias. The 30% of the sample who can and are willing to locate their payslip are better-organised workers, for example (which may be correlated with higher productivity and pay growth). Nevertheless, we do not investigate this selection issue further here, as we believe that the error process we have described correctly describes the formation of beliefs among workers who do not monitor their actual pay level. Measurement error and selection bias are further examined in Smith (1999).

the smallest time interval is a week, the largest is 52 weeks.¹⁶ Obviously we would expect rounding to take correspondingly different orders of magnitude. We investigate the extent of rounding to the nearest £5, £10, £50, £100, £500 and £1000. We are concerned to find out whether rounding of the pay level in any given year is associated with rigidity over the following year.

Table 3 suggests that rounding is quite common: nearly half the sample rounded pay to the nearest £10.¹⁷ Rounding is more common among those whose pay was rigid than in the full sample. Rigidity among rounders is likewise more common than in the full sample: nearly 11% of respondents stating pay rounded to the nearest £50 reported the same pay level the following year, whereas 9% of the full sample reported rigid pay. Rigidity is typically more likely, the higher the degree of initial rounding.

Any rigidity due to rounding is not due to downward nominal rigidity, and hence we can eliminate it as a source of costs of low inflation. Nevertheless, there may be macroeconomic costs associated with rigidity due to rounding as long as this rounding reflects menu costs or bounded rationality. Neither rounding error nor productivity that simply happens to coincide with a rounded value for pay cause costs.

We can estimate of the proportion of rigidity due to rounding under the assumption that for those who round, rounding causes all rigidity in excess of the proportion rigid for those who did not round. The proportion of rigidity due to rounding is given by

$$\Pr(\text{rigid because round}) = \left[\begin{array}{c} \Pr(\text{rigid}|\text{round}) \\ -\Pr(\text{rigid}|\text{not round}) \end{array} \right] \times \Pr(\text{rigid and round}) \quad (3)$$

where $\Pr(\text{rigid}|\text{round}) = \text{Freq}(\text{rigid \& rounded})/\text{Freq}(\text{rounded})$ and similarly for $\Pr(\text{rigid}|\text{not round})$, and $\Pr(\text{rigid and round}) = \text{Freq}(\text{rigid \& rounded})/\text{Freq}(\text{pay growth observed})$. The term in brackets is ‘excess rigidity’ among those who round compared to those who do not round: the proportion of those who round who had zero pay growth over the following year, less the proportion of those who did not round who had zero pay growth. This excess rigidity applies only to ‘rigid rounders’, so we multiply by the proportion of these in the full sample to assess what proportion of the full sample is rigid due to rounding.

To each degree of rounding, this proportion of rigidity due to rounding is given in column 3 of Table 3. ‘Excess rigidity’ is quite large: for example, the difference between the proportions of rounders to £50 and non-rounders-to-£50 who had rigid pay was 3.7 percentage points on average over 1992–6 (11.6% compared to 7.9%). When multiplied by 36.9% – the proportion of those with rigid pay that rounded to £50 – this gives percentage points of rigidity due to rounding to £50: 1.4 percentage points. This constitutes 15% of

¹⁶ We discard observations where reported pay period is less than a week because inspection of the data reveals these are likely to reflect errors.

¹⁷ The data in Table 3 are cumulative, in the sense that, for example, figures for ‘rounded to £5’ include pay rounded to any number divisible by 5. Cumulative data allow flexibility in the choice of what degree of rounding is rounding.

Table 3
Rounding and Nominal Rigidity, 1992–6

Pay rounded to nearest ...	Proportion of sample (%) (1)	% Rigid (2)	% of rigidity due to rounding (3)	% Rigid adjusted for rigidity due to rounding (4)	% liable to rounding error (5)	% of rigidity due to truly rounded pay (6)	% Rigid adjusted for rigidity due to rounding error (7)
Not rounded	44	6.9	n/a	n/a	n/a	n/a	n/a
£5	56	11.2	27.0	6.6	64	9.6	7.4
£10	47	8.6	16.3	7.5	70	4.9	7.9
£50	29	10.7	15.1	7.6	76	3.6	7.9
£100	22	11.1	11.8	7.9	77	2.7	8.2
£500	8	14.6	5.7	8.5	80	1.2	8.6
£1000	5	12.2	2.5	8.8	80	0.5	8.8

Note: Rounding is cumulative – for example, everyone who rounds to £10 rounds to £5. Rounding refers to base year of percentage changes. Figures in columns (2), (4) and (7) can be compared to the full-sample 9.0% rigid.

total rigidity (column 3).¹⁸ Column 4 reports the percentage rigid after subtracting those whose rigidity is due to rounding. On the basis of rounding to £50, 1.4 percentage points of rigidity can be discounted as causing costs of very low inflation, leaving rigidity of 7.6%.

Some of those who round do not do it often; for these individuals, rounded pay may well be true pay. Unusually, we can estimate what proportion of rounded pay is ‘truly’ rounded, and what proportion is mis-reporting by workers. Truly rounded pay includes both the coincidence of ‘warranted’ pay (reflecting the value of the employee) with a rounded value and rounding which is the optimal result of, for example, computational costs or bounded rationality. To estimate the proportion truly rounded, we use information on whether the payslip was checked in reporting the pay figure. Rounding error by those whose payslips were checked will be minimal. We assume all rounding by the payslips-seen subsample reflects truly rounded pay. In contrast, rounding by those whose payslips were not seen includes both truly rounded pay and rounding error. The difference between the extent of rounding in these two subsamples gives us the extent of rounding error. Rounding is as expected lower among the payslips-seen: only 11% of those who checked their payslips rounded to £50, for example, compared to 36% of those whose payslips were not seen.¹⁹ In all categories, the degree of rounding is significantly higher among the payslip-unseen (with p-values on a test of equality less than 0.001).

Given that a worker has stated a rounded pay figure, the probability of that worker having made a rounding error is given by

¹⁸ Since more people round to a lower degree (column 1), the proportion of rigidity that is associated with rounding tends to fall as the degree of rounding rises. 27% of zero pay growth during 1992–6 was associated with pay figures rounded to £5, but only 2.5% was connected to pay rounded to £1000 (column 3).

¹⁹ As before, ‘payslips seen’ refers to the latest or an earlier payslip having been checked in the current and the previous interview. ‘Payslip unseen’ includes all whose payslip was not checked in at least one of these interviews.

$$\Pr(\text{error in rounding}) = \Pr(\text{unseen}|\text{round}) - \Pr(\text{seen}|\text{round}) \quad (4)$$

where $\Pr(\text{unseen}|\text{round}) = \text{Freq}(\text{rounded \& unseen})/\text{Freq}(\text{round})$ and similarly for $\Pr(\text{seen}|\text{round})$. This is ‘excess rounding’ among the unseen: the proportion of rounders whose payslips are not checked, less the proportion of rounders whose are. This difference is the proportion of rounders who are likely to be making rounding errors. The likelihood of error is large: the majority of workers do not check their payslip, and this majority is higher for rounders than for non-rounders. 45% of non-rounders checked in both years; among rounders to £50 only 12% checked in both years. 76% of rounders to £50 might be making rounding error (88%–12%).²⁰

We can use this estimate of rounding error to decompose the proportion of rigidity due to rounding into that due to rounding error – ‘false’ rigidity – and that resulting from truly rounded pay. The proportion of rigidity associated with truly rounded pay is

$$[1 - \Pr(\text{error in rounding})] \times \Pr(\text{rigid because round}). \quad (5)$$

This is the probability of *not* making an error, given rounded pay (i.e. 1 minus expression (4)), multiplied by the proportion of rigidity associated with rounding (expression (3)). Only a very small proportion of annual rigidity is associated with truly rounded pay: for example, only 0.3% of workers (3.6% of overall rigidity of 9.0%) have rigid pay because their pay is truly stable at a value rounded to £50 from one year to the next (see column 6 of Table 3).²¹ On average over 1992–6, we conclude that truly rounded pay, reflecting menu cost-type explanations – explanations drawing observations from either side of zero – could account for just over 0.3 percentage points of rigidity.

Finally, column 7 of Table 3 shows the extent of rigidity after we have eliminated those cases where apparent rigidity can be put down to rounding error. The proportion of those whose pay is rigid making a rounding error is simply the product of expressions (4) and (3) (columns 5 and 3 respectively). Expression (3) gives the proportion of rigidity specifically due to rounding; expression (4) gives the likelihood of error among those who round. The proportion rigid adjusted for rounding error is (column 7) then given by

$$\Pr(\text{rigid}) \times \{1 - [\Pr(\text{error in rounding}) \times \Pr(\text{rigid because round})]\}. \quad (6)$$

On the basis of rounding to £50, error accounts for around 1 percentage point of rigidity in annual pay growth.

To summarise our findings relating to rounding: apparent rigidity of 9% contains almost 1.5% attributable to rounding. This can be excluded from possible downward nominal rigidity as it reflects either error or symmetric factors. Hence we have chopped the zero spike down to around 7.5%. Not all

²⁰ As might be expected, the higher the degree of rounding, the more likely is that rounding to reflect error. Of those who round to £5, 64% are liable to rounding error, whereas 80% of those who round to £1000 may well be stating inaccurate pay figures (column 5 of Table 3).

²¹ The proportion of rigidity due to truly rounded pay falls as the degree of rounding rises because rounding to higher degrees is more likely to reflect error (column 5) and is more scarce so the proportion of rigidity associated with rounding to higher degrees is lower (column 3).

the 1.5% chopped off is costless in terms of unemployment, though: we have shown that almost one-third of this represents possibly costly symmetric, menu cost-type factors.

3.3. 'Long-term' Contracts

Is rigidity simply due to the prevalence of long-term contracts? Pay negotiations in the United Kingdom typically occur at annual intervals, and pay awards are also often made every 12 months even for workers not covered by collective bargaining. We have already seen that rigidity declines over time (Fig. 2). The decline in rigidity among non-job changers from 9% over a one-year horizon to 4% over two years could be caused by annual contracts coming to an end, or multi-stage long-term contracts specifying a nominal change after a certain period of time. Contracting models suggest wages will be nominally rigid for a certain length of time, between contract negotiations (see Fischer, 1977, and Taylor, 1980). Perhaps rigidity is simply due to the fact that a proportion of workers' pay is observed at a frequency lower than their contract duration. BHPS interviews are typically conducted at annual frequency, but the time between interviews ranges from 3 to 19 months. Table 4 shows that consecutive pay data for 56% of interviewees is recorded at a sub-annual frequency.²² Is lack of pay change due to annual contracts for some of these workers responsible for apparent rigidity? This is important to the extent that it might 'explain' rigidity, and because if most rigidity will disappear after a year, its consequences may not be all that detrimental in terms of costly low inflation, although the likelihood remains that rigidity due to such contracts will exacerbate the effects of shocks. Furthermore, if long-term contracts are responsible for rigidity, this has implications for which workers are affected. Long-term fixed nominal contracts could apply to workers anywhere in the pay growth distribution: they are a symmetric explanation for rigidity. Hence any rigidity that can be explained by long-term contracts can be eliminated as not due to downward rigidity.

Table 4 records how rigidity varies with the number of months between interviews. There is no monotonic decline in rigidity as time between interviews rises, although zero pay change is most common for those whose pay was last recorded 9 or fewer months earlier, and least common for those whose pay was last recorded more than 14 months ago (these differences are significant, as shown by the p-values in the final column). The delay in the decline in rigidity after 12 months could be due to delayed settlements: evidence from collective bargaining suggests that negotiations are often protracted beyond the official settlement date (resulting in backdated pay awards).

The decline from 12.5% to 4.0% rigidity around the annual interval is very

²² There is a large proportion who are interviewed at '11 month' intervals. Explanations include our strict definition of months (because there is often a settlement date which remains constant from year to year we require 365 days to have elapsed before we classify the number of months between interviews as 12) and BHPS survey behaviour (fieldwork was completed earlier in 1996 and 1992 than in the previous years).

Table 4
*Extent of Nominal Rigidity due to 'Long-term'
 Contracts, 1992–96*

Months between interviews	Proportion of sample (%)	% Rigid	p-value for H_0 : % rigid is same as 12 months†
≤9	4.6	12.5	0.007*
10	12.5	8.4	0.632
11	38.9	9.3	0.178
12	32.5	8.7	n/a
13	8.4	8.4	0.405
14	1.4	8.9	0.526
≥15	1.7	4.0	0.017*
Any	100.0	9.0	n/a

Note: † p-values for one-tailed tests. Alternative hypotheses: ≤9, 10, 11 months: greater than 12; 13,14, ≥15 months: less than 12.

* Null hypothesis can be rejected at 5% level.

large, suggesting the prevalence of annual contracts may be responsible for a substantial part of observed rigidity. After 14 months, rigidity is the same as that recorded over a two-year horizon. We calculate the influence of annual contracts as the proportional difference between average rigidity over a one-year horizon and that remaining after 14 months, i.e. $(9.0-4.0)/9.0$, which is 56%. This proportion of rigidity may adversely affect the economy's response to shocks, but it does not contribute to downward rigidity. We conclude that the zero spike can be knocked down to half its apparent size on the basis of this explanation.

3.4. *Overlap between the Explanations*

To summarise this section's investigation of the three explanations for rigidity and flexibility: rounding accounts for around 15% of rigidity on average over 1992–6; measurement error accounts for 35–40%; and annual contracting accounts for around 50% (see Fig. 4).²³ If all these explanations were independent, we would have accounted for all (or more than all!) of rigidity.

Measurement error and long-term contracts are independent. But we have seen that rounding and measurement error are not independent. Excluding measurement error reduces the proportion of rigidity accounted for by rounding to just under 4% (this is rigidity due to pay that remains truly rounded from one year to the next). The data do not reveal any overlap between rounding and long-term contracts: the proportions stating rounded pay levels are the same for those interviewed at a more-than-14-month interval as they

²³ In Fig. 4 the proportion of zero-pay-growth observations accounted for by each explanation is removed. Note that the histograms include wage growth of between -0.5% and 0.5% in the 'zero' bin. Values other than zero form a substantial proportion: almost three-quarters of the height of the zero bin in the 'All explanations' chart is accounted for by non-zero pay growth. This explains the apparent discrepancy between the height (about 4%) of the bin and remaining rigidity (around 1%).

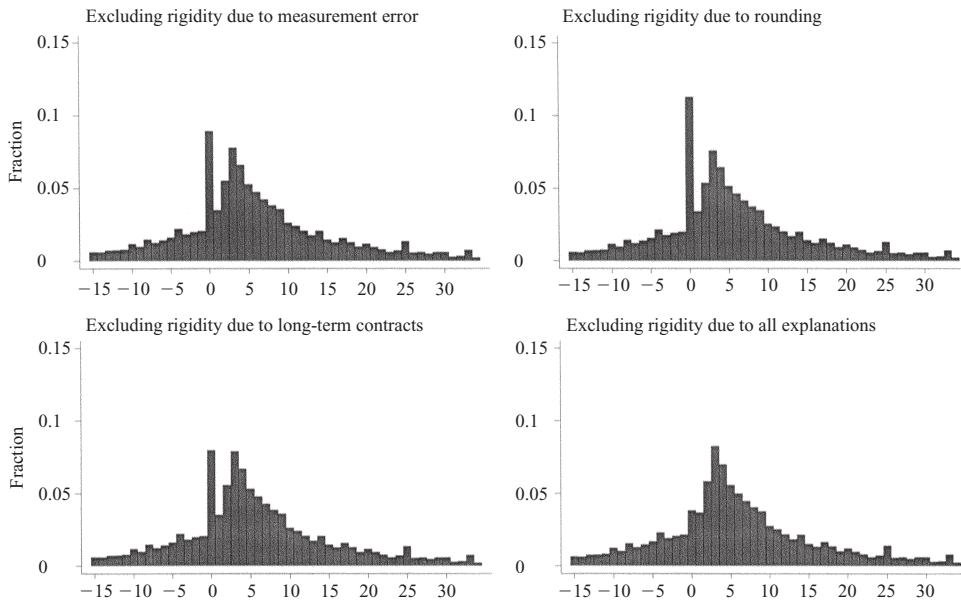


Fig. 4. *Chopping Down the Spike: Accounting for Rigidity in Nominal Pay on the Basis of Measurement Error, Rounding and Long-term Contracts, 1992–6*

Note: Rigidity remaining after accounting for the relevant explanation (note that non-zero bins have been left unaltered). Compare with top left chart in Fig. 1.

are for the full sample. So even taking interdependence into account, our three explanations account for up to 90% of rigidity (50% from contracts, 35% from measurement error and 4% from rounding). Thus only around 1% of stayers have zero pay growth we cannot account for. We conclude there is very little residual rigidity to attribute to downward-rigidity-inducing explanations. Symmetric rigidities, however, have been shown to exist: these include all explanations except measurement error, namely around half observed rigidity, or around 4.5 percentage points, on top of the 1 percentage point we have not accounted for.

4. Conclusion

This paper has investigated the extent of nominal rigidity in the UK labour market in the 1990s, a period when inflation has been sufficiently low that we might expect any tendency toward downward nominal rigidities to be reasonably marked. We find that 9% of non-job changers have zero growth in their usual weekly pay over a one-year horizon during the period of study. This is a similar order of magnitude to that found in the United States on average over the 1970s and 1980s, but is just over half US figures at comparably low rates of inflation.

We examine whether rigidity is due to rounding or ‘long-term’ contracts, and whether apparent flexibility merely reflects measurement error. Any

rigidity due to these factors is not caused by downward rigidity. We estimate that rounding accounts for up around 15% of rigidity and that annual contracting accounts for 50%. Substantial downward flexibility remains after we eliminate measurement error: 18% of non-job changers truly experience nominal pay cuts over a one-year horizon. Furthermore, we find that the measurement error process serves to exaggerate rigidity. Both of these are contrary to what has typically previously been assumed. Measurement error accounts for over 35% of observed rigidity. This result has implications for recent influential work in the United States which has assumed measurement error of classical form, thus inferring higher true rigidity.

Taking into account interdependence between them, our explanations account for up to 90% of rigidity. Only 1% of stayers have zero pay growth we cannot account for. This suggests the UK labour market exhibits surprisingly high nominal flexibility and little downward nominal rigidity. We should not look to the labour market as the source of major unemployment costs of zero inflation that would justify a positive inflation target.

University of Warwick

References

- Akerlof, G. A., Dickens, W. T. and Perry, G. L. (1996). 'The macroeconomics of low inflation.' *Brookings Papers on Economic Activity*, vol. 1, pp. 1–76.
- Bound, J., Brown, C., Duncan, G. J. and Rodgers, W. (1989). 'Measurement error in cross-sectional and longitudinal labor market surveys: results from two validation studies', NBER Working Paper no. 2884.
- Blinder, A. S. and Choi, D. H. (1990). 'A shred of evidence on theories of wage stickiness.' *Quarterly Journal of Economics*, vol. 105, no. 4, pp. 1003–15.
- Bound, J. and Krueger, A. B. (1991) 'The extent of measurement error in longitudinal earnings data: do two wrongs make a right?' *Journal of Labor Economics*, vol. 9, no. 1, pp. 1–24.
- Card, D. and Hyslop, D. (1996). 'Does inflation "grease the wheels of the labor market"?' NBER Working Paper no. 5538.
- Christofides, L. N. and Leung, M. T. (1998). 'Wage adjustment in Canadian contract data: wage rigidity and menu costs.' University of Guelph mimeo.
- Christofides, L. N. and Stengos, T. (1998). 'The symmetry of the wage-change distribution: non-parametric tests using contract and survey data.' University of Guelph mimeo.
- Crawford, A. and Harrison, A. (1998). 'Testing for downward rigidity in nominal wages.' In *Price Stability, Inflation Targets and Monetary Policy*, pp. 179–225. Ottawa: Bank of Canada.
- Duncan, G. J. and Hill, D. H. (1985). 'An investigation of the extent and consequences of measurement error in labor-economic survey data.' *Journal of Labor Economics*, vol. 3, no. 4, pp. 508–32.
- Fischer, S. (1977). 'Long-term contracts, rational expectations and the optimal money supply rule.' *Journal of Political Economy*, vol. 85, no. 1, pp. 191–205.
- Fortin, P. (1996). 'The great Canadian slump.' *Canadian Journal of Economics*, vol. 29, no. 4, pp. 761–87.
- Kahn, S. (1997). 'Evidence of nominal wage stickiness from microdata.' *American Economic Review*, vol. 87, no. 5, pp. 993–1008.
- Kahneman, D., Knetsch, J. L. and Thaler, R. (1986). 'Fairness as a constraint on profit seeking: entitlements in the market.' *American Economic Review*, vol. 76, no. 4, pp. 728–41.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. London: MacMillan.
- Lebow, D. E., Stockton, D. J. and Wascher, W. L. (1995). 'Inflation, nominal wage rigidity and the efficiency of labor markets.' Federal Reserve Board of Governors Finance and Economics Discussion Series no. 95–45.
- Mankiw, N. G. and Romer, D. (1991). *New Keynesian Economics, Volume 1: Imperfect Competition and Sticky Prices*. London: MIT Press.
- McLaughlin, K. J. (1994). 'Rigid wages?' *Journal of Monetary Economics*, vol. 34, no. 3, pp. 383–414.
- Smith, J. C. (1999). 'Restoring wage flexibility: the nature of earnings measurement error in survey data.' University of Warwick mimeo.

- Taylor, J. (1980). 'Aggregate dynamics and staggered contracts.' *Journal of Political Economy*, vol. 88, no. 1, pp. 1–23.
- Tobin, J. (1972). 'Inflation and unemployment.' *American Economic Review*, vol. 62, no. 1, pp. 1–18.
- Yates, A. (1998). 'Downward nominal rigidity and monetary policy.' Bank of England Working Paper no. 82.