

Working Paper 2000-1 / Document de travail 2000-1

The Employment Costs of Downward Nominal-Wage Rigidity

by

Jean Farès and Seamus Hogan

Bank of Canada



Banque du Canada

ISSN 1192-5434
ISBN 0-662-28551-4

Printed in Canada on recycled paper

Bank of Canada Working Paper 2000-1

January 2000

The Employment Costs of Downward Nominal-Wage Rigidity

by

Jean Farès

Research Department
Bank of Canada, Ottawa, Canada
jfares@bank-banque-canada.ca

Seamus Hogan

Health Canada, Ottawa, Canada
Seamus_Hogan@hc-sc.gc.ca

The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements	iv
Abstract/Résumé	v
1. Introduction	1
2. The contract data	2
3. Approaches for estimating the employment costs of low inflation	4
3.1 The histogram approach	5
3.2 The microdata approach	7
3.3 A problem of estimation bias	8
4. Controlling for the bias	9
5. Results	11
5.1 Exogenous variables	11
5.2 The wage equation	11
5.3 The employment equation	15
6. Sensitivity checks	17
6.1 Censoring of the dependant variable	17
6.2 Definition of a wage change	18
6.3 Specification of the wage-freeze term	19
6.4 Aggregate data	22
7. Conclusion	22
Figures	25
References	29

Acknowledgements

We would like to thank Allan Crawford, David Longworth, Tiff Macklem, and Gerald Stuber for their helpful comments. Further suggestions were also received at seminar presentations at the Bank of Canada and the Canadian Economic Association annual meetings.

Abstract

In this paper, we use firm-level wage and employment data to address whether there is evidence of downward nominal-wage rigidity, and whether that rigidity is associated with a reduction in employment. We describe an estimation bias that can result when estimating reduced-form wage and employment equations and suggest a way of controlling for that bias. The adjusted point estimates suggest that nominal-wage rigidity is associated with wages being lower than would have been in the absence of rigidity, rather than higher. Our estimates also suggest that the rigidity tends to be associated with higher rather than lower employment. This conclusion is tempered by the low statistical significance of the estimated coefficients, but our results do suggest the importance of controlling for the bias we identify.

JEL classifications: C33, J23, J31

Bank classification: Labour markets

Résumé

Les auteurs ont recours à des données microéconomiques se rapportant aux salaires et à l'emploi pour tenter d'établir si les salaires nominaux sont rigides à la baisse et si cette rigidité se traduit par une réduction de l'emploi. Ils décrivent le biais lié à l'estimation d'équations de forme réduite relatives au salaire et à l'emploi et proposent une façon d'en tenir compte. Une fois le biais corrigé, les estimations ponctuelles qu'ils obtiennent donnent à penser que la rigidité des salaires nominaux s'accompagne d'un niveau plus faible, et non plus élevé, des salaires et qu'elle se traduit par une augmentation au lieu d'une diminution de l'emploi. La validité de cette conclusion est tempérée par le faible degré de signification statistique des coefficients estimés; il reste que les résultats obtenus font ressortir qu'il importe de tenir compte du biais décelé par les auteurs.

Classifications JEL : C33, J23, J31

Classification de la Banque : Marchés du travail

1. Introduction

The theory that institutional constraints exist that make it difficult for firms to cut nominal wages has a long history in macroeconomics. Keynes, in the *General Theory*, used downward nominal-wage rigidity as a reason why expansionary monetary policy might be more successful than relying on general deflation to end a recession. Tobin (1972) explained how Keynes' idea, when applied to an economy facing heterogeneous rather than aggregate shocks, could generate a long-run trade-off between inflation and unemployment. With the United States and Canada now maintaining low inflation rates, a number of economists have started to reconsider the Tobin hypothesis. Foremost among these are Akerlof, Dickens, and Perry (1996) and Fortin (1996).

The paper by Akerlof, Dickens, and Perry (hereafter, ADP) represents the most comprehensive attempt to date to fully model the Tobin hypothesis. In the ADP model, low inflation generates higher unemployment for the following reasons:¹

1. Even in the absence of aggregate shocks, individual firms are constantly being hit by heterogeneous shocks that lead to increases in the real wages paid at some firms and to decreases at others.
2. Workers resist reductions in their real wages brought about by nominal-wage cuts to a greater extent than they resist real-wage cuts brought about by inflation. Therefore, when inflation is low, a number of firms are constrained by downward nominal-wage rigidity from setting the real wages they would have set if inflation had been higher.
3. As a result, employment is lower at those constrained firms than would have been the case had inflation been higher.
4. Any spillover effects on the wages and employment levels at firms that are not bound by downward nominal-wage rigidity do not fully counteract the direct negative employment effect.

Much of the empirical work examining the Tobin hypothesis has sought to establish the prevalence of downward nominal-wage rigidity in the economy. Comparatively little work has looked at whether that rigidity affects employment. Simpson, Cameron, and Hum (1998) estimate a reduced-form employment equation using industry-level data and find a negative coefficient on the wage-freeze variable. They thus conclude that there is a negative effect on employment from downward nominal-wage rigidity.

1. See Hogan (1997; 1998) for a fuller description and critique of these steps in ADP's analysis.

In this paper, we use firm-level data from the Canadian manufacturing sector to address both the evidence for downward nominal-wage rigidity and its employment effect. Our longitudinal data set provides both the wage and employment changes over the duration of the contract for a large number of collective-bargaining contracts over a period of 21 years. Our approach is to estimate wage-change and employment-change equations to see whether wage freezes are associated with wage and employment changes that are higher or lower than would be expected when controlling for other variables. Our point estimates suggest that nominal-wage rigidity tends to lower rather than raise nominal wages at firms bound by nominal-wage rigidity. Contrary to Simpson, Cameron, and Hum (1998), this is associated with an increase rather than a decrease in employment at those firms. We believe that the difference between our results and those in Simpson, Cameron, and Hum reflects, in part, our control for the estimating bias in these regressions.

The paper proceeds as follows. We describe our contract data in Section 2. In Section 3, we discuss several methods proposed in the literature to address questions about downward nominal-wage rigidity and its employment effects. In Section 4, we outline the approach used in this paper, before presenting our results in Section 5. Section 6 considers the robustness of our results to alternative specifications. Section 7 contains some concluding remarks.

2. The contract data

Our contract data come from the wage-settlements file covering collective-bargaining agreements in the Canadian private union sector over the period 1978–97. The data set includes contracts with 500 or more workers. The available information for each contract includes its starting and ending dates, the base wage in each month of the contract, and the number of employees covered by the agreement at the start of the contract. The base wage is typically the wage paid to the lowest-skill group covered by the collective-bargaining agreement.

Each of these settlements corresponds to a ‘bargaining unit.’ Typically, there is one bargaining unit for each firm, but there are some cases with more than one bargaining unit at a single firm and others with a single bargaining unit covering several firms. However, because of the high degree of overlap, we use the simpler term “firm” to refer to bargaining units in the remainder of this paper.

Most of the settlements are multi-year contracts. (The sample mean for contract duration is about 28 months.) Multi-year contracts specify in advance a wage change for each year of the contract. In these cases, it is not clear what is the best way to define the wage change applying to each contract. Crawford and Harrison (1998) discuss three definitions of a wage change that have

been used in studies of wage rigidity. These are the “lifetime,” “first-year,” and “year-over-year” definitions. The lifetime definition takes the average annual wage increase over the lifetime of the contract as the wage change for that observation. The first-year definition uses the wage change applying in the first year of the contract, and the year-over-year definition treats each year of a multi-year contract as a separate observation.

The choice of definition of a wage change greatly affects the proportion of contracts that show wage freezes (that is, zero wage change). This is illustrated in Table 1, which shows the number of wage freezes in the private sector in different periods according to two of the three definitions of the wage change. There are a number of contracts that specify no wage change in the first year of the contract but allow for wage increases in subsequent years. These are considered wage freezes under the first-year definition but not under the lifetime definition. In this paper, we adopt the lifetime definition of wage changes because we have employment data only for the start of each contract. By matching consecutive contracts for individual bargaining units, we are able to construct the employment change over the lifetime of the contract but not for each year. Thus, we are forced to use a lifetime definition of employment changes. Since we want to match the wage-change and employment-change data, we adopt the same definition for wage changes. It is important to note, however, that the lifetime definition may tend to understate the extent of nominal-wage rigidity and its effects on employment. This is because contracts with wage freezes in the first year but not every year of a multi-year contract may produce a temporary employment effect for the period of the wage freeze that we are unable to observe in our data. For this reason, it will sometimes be useful to restrict our attention to one-year contracts in which all three definitions of a wage freeze coincide. For this reason, we also report in Table 1 the percentage of wage freezes for one-year contracts only.

Table 1: Percentage of union contracts with wage freezes^a

Wage-change definition		1978–1982	1983–1991	1992–1997	Total
Lifetime		0.26	3.97	7.6	3.73
First-year		0.79	15.36	30.54	14.70
One-year contracts only		0.54	31.03	45.57	20.70
Number of contracts	Total	1137	1888	907	3945
	One-year	185	145	79	414

a. Private sector wage settlements, 1978–1997

Figure 1 shows histograms of nominal-wage changes in the contract data over the periods 1978–82, 1983–91, and 1992–97, corresponding to times of high, medium, and low inflation, respectively, and for the entire sample. These histograms illustrate four points: first, there is a pronounced spike at zero wage change in the low-inflation and medium-inflation periods; second, nominal-wage cuts are relatively infrequent even in periods of low inflation; third, small positive wage increases are also infrequent; and fourth, there is a marked reduction in the variance of wage changes from the high-inflation to the low-inflation period. The first two points suggest a *prima facie* case that there is significant downward nominal-wage rigidity affecting these data. The third, however, gives some reason for caution: it is not surprising that there are few wage cuts during the high-inflation and medium-inflation periods, and the tighter distribution around a still-positive median in the low-inflation period suggests that one would not necessarily expect a large number of wage cuts in that period either. The absence of small positive wage changes suggests that there is at least some upward as well as downward rigidity of nominal wages contributing to the spike at zero wage changes. The challenge in searching for evidence of downward nominal-wage rigidity is to ascertain how much of this spike is the result of left rather than right censoring of the underlying distribution of wage changes.

Figure 2 shows histograms of contract-to-contract employment changes (at an annual rate) over the same time periods. The employment-change distribution also has a large spike at zero. This could be due to several reasons: the rounding-off in reporting employment levels; bargaining unit coverage might not be as variable as actual employment; or, it may reflect that it takes a major shock for firms to adjust their employment levels.² One clear difference between the employment-change and the wage-changes histograms is the absence of any visually significant changes in the distribution across inflation periods. There is no indication that the changes in the wage-change distribution have led to any corresponding changes in the employment-adjustment distribution. This suggests that any effects of nominal-wage rigidity on employment, whether positive or negative, have not been concentrated at any point of the employment-change distribution. Therefore, the comparison of the distribution across inflation regimes will not reveal much about the employment effects of nominal-wage rigidity.

3. Approaches for estimating the employment costs of low inflation

A number of recent papers have considered the issue of downward nominal-wage rigidity and its associated employment effects, using data sets of wage and employment changes similar to the one

2. In Section 6, the econometric implications of these measurement errors are discussed in detail.

described in the previous section. Most of this literature has focused purely on whether the histograms of wage changes provide evidence for downward nominal-wage rigidity and has ignored the question of whether that rigidity is associated with a reduction in employment. In this section, we first outline the histogram approach, and then describe a more general approach that also addresses the question of employment effects.

3.1 The histogram approach

When using distributions of observed wage changes to infer whether there is downward rigidity in nominal wages, the main problem is that we are not looking simply to reject a null hypothesis that wages are fully flexible. We also need to be able to support the alternative hypothesis that wages are more rigid downwards than upwards. The large spikes at zero wage change in our data set and in the similar data sets considered in this literature make it clear that there is a large amount of nominal-wage rigidity. (Given the different rates of inflation and productivity growth in different years and in productivity growth across firms, the spike cannot reflect underlying economic fundamentals.) It is important, however, whether that spike arises from symmetrical rigidity—in which small underlying positive or negative wage changes are rounded off to zero—or from purely downward rigidity. In the former case, removing the rigidity through inflation would not lower real wages at firms that were constrained by rigidity; however, it would in the latter case.

It is reasonable to expect at least some symmetrical rigidity of nominal wages, due either to menu-cost effects or to the fact that, when the underlying fundamentals suggest only a small change in the wage, zero wage change is a focal point for both the level change and the rate of change. Firms and workers typically bargain over much more than just the base wage in a contract. It is reasonable to imagine that, faced with a wage change that would only be a few cents, they choose the focal point of zero and then concentrate negotiations on other areas. This presumption is supported by the fact noted earlier that there are very few small wage increases in our data set.

In the literature, the main approach used to infer to what extent the distribution of observed nominal-wage changes is the result of downward rather than symmetrical rigidity is to compare the distribution of *actual* wage changes to a hypothetical distribution of *notional* wage changes that would occur if there were no nominal-wage rigidity. The papers using this approach need to impose some maintained assumptions on the nature of the distribution to identify the counterfactual distribution of notional wage changes. Card and Hyslop (1996) assume that the notional distribution of wage changes is symmetric and that the upper half of the distribution is unaffected by nominal rigidities. They can then use the upper half of the observed distribution to infer the shape of the lower half of the notional distribution. Using wages from two different U.S. data sources (PSID

and CPS), they found that about 10 per cent of workers experienced nominal-wage rigidity in high-inflation periods; this fraction became 15 per cent when inflation decreased to 5 per cent. The empirical results show that a 1 per cent increase in the inflation rate reduces the fraction of workers experiencing nominal-wage rigidity by about 0.8 per cent. Downward rigidity exerts a small but measurable effect on average real-wage growth, with a bigger effect in low-inflation years.

Kahn (1997) also uses U.S. data from the PSID. She assumes that the shape of the notional distribution around its median is the same in periods of low inflation and high inflation. This allows her to use the distribution in times of high inflation—when nominal rigidity does not bind—to infer the shape in times of low inflation. Kahn observes a significant fraction of household heads with no nominal-wage or salary change. She finds that salaried workers do not receive pay cuts less frequently than would be expected, especially after 1982, but for non-salaried workers, she finds evidence for downward rigidity.

The identifying assumptions used in these papers are problematic. There is no reason to assume *a priori* that the underlying forces affecting wage changes are symmetric. Indeed, an important reason for wage changes to differ among industries is the different rates of labour-productivity growth. One would expect the inter-industry distribution of underlying productivity growth to be positively skewed, since there is no upper bound to how fast productivity can grow while negative technological progress is not common. Furthermore, there is a limit to how far firms can reduce their workers' real wages without facing a labour-supply constraint. No such constraint exists for wage increases. If the notional distribution of wage changes is positively skewed, then reflecting the actual distribution about the median to obtain an estimate of the notional distribution will tend to overstate the number of nominal-wage cuts we would expect to see in the data. Thus, the extent of downward rigidity would also be overstated.

The assumption that the mean level of inflation does not affect the shape of the wage-change distribution around its median may also lead to overestimating the degree of downward nominal-wage rigidity in the data. As many authors have pointed out, theory predicts that the variance of the wage-change distribution should be lower when inflation is lower (see Golub [1993] for a recent survey of these papers). Using the distribution from a high-inflation era to infer the shape of the notional distribution when inflation is low will therefore tend to overpredict the extent to which we should observe wage reductions.

The histogram approach described here does not directly address the issue of whether the nominal rigidity generates an adverse effect on employment. One way of extending the analysis would be to calculate the extent to which nominal wages have been increased due to downward rigidity and then to put that estimate into a labour-demand equation to calculate the resulting

employment effect. There are two reasons why this might overstate the employment cost of any downward nominal-wage rigidity. First, if downward nominal-wage rigidity is perceived as a temporary constraint, firms may be reluctant to reduce their workforce by as much as they would if the increase in their real wage was perceived as permanent. Therefore, misleading conclusions could be reached if the labour demand does not adequately capture short-run dynamics. Second, employment levels depend not only on the number of job slots created by firms but also on the vacancy rate for those job slots. By increasing job-acceptance rates and reducing quit rates, an increase in the wage can have a negative effect on the vacancy rate and hence a positive effect on employment. Although it is unlikely that such a supply-side effect would dominate the demand-side effect on employment of an increase in wages, it will imply that the short-run elasticity of employment with respect to wages is less than the elasticity of labour demand.

The microdata approach, described in the next subsection, provides an alternative means of identifying whether wage freezes are associated with downward or symmetric rigidity. The approach also allows direct estimation of employment effects of that rigidity.

3.2 The microdata approach

We use the term “microdata approach” to describe the use of contract-level information such as location, industry, etc. to try to infer the expected wage change in each contract. In this approach, one tries to *explain* the distribution of wage changes using the additional information rather than imposing a particular form on the hypothetical underlying distribution.

Crawford and Harrison (1998) estimate the distribution of notional wage changes using a proportional-hazard model in which a contract-specific variable is used to estimate changes in the moments of the notional distribution over time. This work suggests that the overall effect on wages from nominal rigidity is very small, but did not address the question of employment effects resulting from nominal-wage rigidity.

One way of using the microdata approach to identify both the direction of any nominal-wage rigidity and its associated employment effects is to estimate reduced-form wage and employment equations directly, and then to ask whether the presence of a wage freeze has any explanatory power. A pure reduced-form model along these lines would take the form

$$\Delta W_{it} = \beta' X_{it} + \omega D0_{it} + \varepsilon_{it}, \quad (1)$$

$$\Delta E_{it} = \theta' X_{it} + \delta D0_{it} + \mu_{it}, \quad (2)$$

where ΔW_{it} and ΔE_{it} are the change in wage and employment, respectively, at firm i in period t ; X_{it} is a vector of exogenous explanatory variables; and DO_{it} is a dummy variable that equals one when $\Delta W_{it} = 0$. If the vector of exogenous variables contains a lot of explanatory power, then a positive coefficient on the wage-freeze variable in the wage equation would indicate that wage freezes bring about an increase in wages on average (that is, that observed rigidity in nominal wages is not purely downward). Alternatively, a negative coefficient in the employment equation would indicate that rigidity is associated with a reduction in employment.

Simpson, Cameron, and Hum (1998) estimate an employment equation similar to Equation (2). They use an industry-level data set in which the left-hand-side variable in Equation (2) is the average annual employment change in each industry in each period; and the wage-freeze variable, rather than being a dummy variable, is the proportion of firms in the industry with wage freezes. They find a negative coefficient on the wage-freeze variable and thus conclude that there is a negative effect on employment from downward nominal-wage rigidity. A problem with the reduced-form approach, however, may have biased this coefficient downwards. We describe this problem in the remainder of this section before returning to see how it might have affected Simpson, Cameron, and Hum's results.

3.3 A problem of estimation bias

The sign of the coefficient of the wage-freeze variable in Equations (1) and (2) can be taken only as an indication of the direction in which wages are affected by nominal-wage rigidity and of the associated employment effects if the explanatory power of the exogenous variables is very high. If the explanatory power of X is low, we can expect the wage-freeze variable to be correlated with the error term, thus biasing its estimated coefficient. To see the direction of this bias, note that—even in the absence of downward nominal-wage rigidity—the average nominal-wage increase in any year would be positive due to the combined effect of labour-productivity growth and inflation. Observations with wage freezes, then, represent unusually low wage changes. Therefore, the wage-freeze variable would be negatively correlated with the error term in Equation (1), thus biasing down the coefficient.

With the employment equation, the direction of the bias is less clear as it depends on the correlation between wage changes and employment changes. If the wage changes in the data were largely the result of wage-bargaining shocks that pushed firms along their labour-demand curves, then we would see a negative correlation between wage changes and employment changes and the wage-freeze dummy would be positively correlated with the error term. If, on the other hand, wage changes were largely the result of shocks to labour demand, then we would see a positive

correlation with employment changes and a negative correlation between the wage-freeze variable and the error term. In this case, the coefficient on the wage-freeze variable would be biased downwards, thus biasing the results in favour of finding evidence of an employment cost due to nominal rigidity.

It is possible that the results found by Simpson, Cameron, and Hum are affected by this bias. To control for the possible effect of labour-demand shocks, they include an output variable in their set of exogenous variables. Unfortunately, this will not deal with the problem, as output can be an indicator of both labour-demand shocks and wage shocks. That is, output could fall because of a fall in demand for that industry's product and thus in the derived demand for labour, or because a positive wage shock has induced firms to reduce employment and hence output. Although output may capture labour-demand shocks, it may also capture the effect of downward nominal-wage rigidity. Thus it would not increase the ability of the wage-freeze variable to represent a causal relationship from downward nominal-wage rigidity to employment. In other words, including industry-level output in a regression of industry-level employment will not control for labour-demand shocks, but will simply change the equation from being about employment to being about labour productivity.

In the remainder of this paper, we consider an alternative way of dealing with this estimation bias in the wage and employment equations.

4. Controlling for the bias

We can correct for the estimation bias in the employment equation simply by adding ΔW_{it} as a right-hand-side variable to give

$$\Delta E_{it} = \theta' X_{it} + \delta DO_{it} + \gamma \Delta W_{it} + \mu_{it}. \quad (3)$$

In this form, we are asking if wage freezes have any explanatory power for employment over and above that contained in the wage change. We give no structural interpretation to the coefficient on the wage-change variable. Equation (3) is still a reduced-form equation in which any correlation between wage and employment changes could be the result of either labour-demand shifts or wage-bargaining shocks such as might result from changes in the local unemployment rate. In this form, however, we do give a structural interpretation to the wage-freeze dummy. The maintained assumption here is that an effect on the wage due to nominal-wage rigidity constitutes a pure wage shock. As long as the coefficient on ΔW_{it} reflects, at least in part, labour-demand shocks that have

not been captured in X_{it} , then the wage-freeze dummy will have isolated an effect on employment from wage shocks that is not captured in ΔW_{it} .

For the wage equation, we obviously cannot control for the correlation between the wage-freeze dummy and the error term simply by adding ΔW_{it} as a right-hand-side variable. For this equation, then, we adopt the following approach.

First, we estimate the wage equation without a dummy variable for wage freezes:

$$\Delta W_{it} = \beta' X_{it} + \varepsilon_{it}. \quad (4)$$

Second, we use the estimated model to find the average forecast error for the observations where $\Delta W_{it} = 0$. Specifically, the estimated residuals are

$$\hat{\varepsilon}_{it} = \Delta W_{it} - \hat{\beta}' X_{it}, \quad (5)$$

and let $\psi = E(\hat{\varepsilon}_{it} | (\Delta W_{it} = 0))$. We estimate ψ as the average residual for the observations with wage freezes:

$$\hat{\psi} = \text{Average}(\hat{\varepsilon}_{it} | (\Delta W_{it} = 0)).$$

For the reasons given above, we would expect that this average forecast error to be negative.

Finally, we use the residuals from the regression using observations that are not wage freezes to extrapolate an expected residual when $\Delta W_{it} = 0$. To do this, we run the following regression:

$$\hat{\varepsilon}_{it} = \alpha + \lambda \Delta W_{it} + \eta_{it}. \quad (6)$$

Again, unless the explanatory power of X is extraordinarily high, we would expect the estimated slope of this regression line to be positive and the estimated intercept, $\hat{\alpha}$, to be negative.³ This estimated intercept constitutes the model's prediction for the estimated residual when $\Delta W_{it} = 0$. Under the null hypothesis that there is no nominal rigidity censoring the wage-change distribution, the following equality would hold:

$$\psi = \alpha. \quad (7)$$

3. Note that estimating Equation (6) is not subject to the classical generated regressor problem, given that the generated variable is the dependent variable in this regression. However, we might have introduced a measurement error problem that would have an effect on the estimated standard errors but not on the estimated coefficients.

If downward nominal-wage rigidity is increasing nominal wages on average, then $\psi > \alpha$. Our test then for downward nominal-wage rigidity is to compare $\hat{\psi}$ with $\hat{\alpha}$.

5. Results

5.1 Exogenous variables

Our vector of explanatory variables, \mathbf{X} , includes five regional unemployment rates⁴ to capture the labour-market conditions at the signing of a new contract; a set of unrestricted-year dummy variables to capture the effect of inflation and the business cycle on nominal-wage and employment growth; a set of regional dummy variables to control for structural differences in regional unemployment rates; and the growth in industry-level output to capture shocks to firm-specific labour-demand curves. The industry output data were taken from various Cansim series produced by Statistics Canada on *Gross Domestic Product by Industry*. These data are classified by 1980 standard industrial codes (SICs) at the three-digit industry level. Because these data are available only for the manufacturing sector, we restrict our wage and employment data to that sector.

We also tried including the growth in industry-specific input prices as another industry-level variable that might capture shocks to firm-specific labour-demand curves. We found, however, that this variable was highly correlated with industry output. Thus, it did not improve the explanatory power of our regressions by the criterion of the adjusted R-squared. Also, it entered both our wage and employment regressions with the opposite sign from what we would expect if it were truly capturing labour-demand shocks. Accordingly, we chose to omit this variable from our regressions. This in no way changed the qualitative results we report below.

5.2 The wage equation

For the wage equation, we ran the following regression:

$$\Delta W_{it} = \beta^c + \beta^q Q_{it-1} + \beta^u u_{it} + \sum_{\tau=1}^T \beta^{y\tau} DY_{it}^{\tau} + \sum_{\alpha=1}^A \beta^{r\alpha} DR_i^{\alpha} + \varepsilon_{it}. \quad (8)$$

For each observation, ΔW_{it} is the wage change for a contract negotiated at firm i in period t , expressed as an average annual percentage rate of change from the previous contract.⁵ Q_{it-1} is the

4. The regions are: the Atlantic provinces, Quebec, Ontario, the Prairies, and British Columbia.

5. This ex post wage-change measure includes all COLA calculations when a contract has a COLA clause.

average annual percentage growth in industry output in the industry to which firm i belongs over the period of the previous contract. Finally, u_{it} is the regional unemployment rate⁶ for firm i in period t , and DY_{it}^τ and DR_i^α are the year and regional dummy variables, respectively.

DY_{it}^τ is not a binary dummy variable that takes the value of 1 in the year the contract is settled and 0 otherwise. Rather, each year is weighted according to the proportion of the contract that applied to that year. Formally, the dummy variable for Year τ is the number of months in that year that the contract settled in period t was in force, divided by the total number of months in the contract. The regional unemployment rate is the unemployment rate applying in the region of firm i at the start of the contract. In the case of multi-regional contracts, the national unemployment rate was used instead. It is intended to capture shocks to worker bargaining power. It should have a negative effect on wage growth. Q_{it-1} is intended to capture effects on firm-level labour demand over the period of the previous contract. If wages are correlated with firm profitability, then lagged output should have a positive effect on wage growth.

Although our right-hand-side variables are disaggregated only to the industry-region level, the different contract lengths and settlement dates across firms allow firm-level variation in these variables. For instance, for two firms in the same industry that both negotiated new contracts in period t but whose previous contracts were of different length, Q_{it-1} will not necessarily take the same value as it will be averaging over a different number of years. The year dummy variables similarly vary according to the settlement date and contract length of each contract.

We used three different techniques to estimate Equation (8). Our base technique is OLS estimated using the entire data set. This is appropriate, however, only under the null hypothesis of no nominal rigidity. If there is some censoring of wage changes at zero, OLS will produce biased coefficients. To deal with the case of censored data, we tried two other techniques: a Tobit model in which it is assumed that the distribution of wage changes is left-censored due to downward nominal-wage rigidity; and OLS using only those observations for which the wage change is not equal to zero, so that the estimates of the parameters are not contaminated by the presence of censored data.

The Tobit model is defined as

$$\Delta W_{it} = \begin{cases} \beta' X_{it} + \varepsilon_{it} & \text{if } \beta' X_{it} + \varepsilon_{it} > 0 \\ 0 & \text{Otherwise} \end{cases}$$

6. In this paper, we do not distinguish between long-term and short-term unemployed. An extension could explicitly add in these different measures that may have a differential impact on wages as indicated by the wage determination literature.

where the residuals, ε_{it} , are assumed to be independently and normally distributed with mean zero and a common variance, σ^2 . This model imposes the assumption that negative wage changes are never observed. Although this is too strong, there are only four observations in our sample where $\Delta W_{it} < 0$, so the assumption of complete censoring is not strongly contrary to the data. We removed those four observations from the data when running the Tobit model.

In the Tobit model, the mass of observations at zero are assumed to be purely the result of left censoring. That is, it allows observations at zero to be the result of downward rigidity, but does not admit the possibility of symmetrical rigidity in which small positive wage changes are also censored to zero. If wages are as likely to be censored from the right as well as the left, the Tobit model will tend to underestimate the value of the residuals for observations where $\Delta W_{it} = 0$, leading it to bias down its estimate of $\beta'X_{it}$.

Table 2 presents the results of the wage-change estimation, which are very similar across the different estimation approaches. Lagged industry output growth and regional unemployment enter with the expected sign and are statistically significant. Omitted from Table 2 are the year and regional dummy variables. In each case, the variables were jointly significant.

Table 2: Estimates of the wage-change regression^a

	OLS (full sample)	Tobit (excl. $\Delta W < 0$)	OLS (excl. $\Delta W = 0$)
Δ Industry output (-1)	0.029 (.024)	0.028 (.034)	0.025 (.053)
Unemployment rate	-0.26 (.002)	-0.22 (.010)	-0.34 (.000)
\bar{R}^2	0.64	0.18	0.64
Sample size	1214	1192	1154
p values given in parentheses for two-sided tests			

a. Source: Wage settlements from the manufacturing sector. All regressions include year and regional dummy variables.

Table 3 shows the average residual when $\Delta W_{it} = 0$, $\hat{\psi}$, and the predicted residual based on the pattern of residuals for observations where ΔW_{it} does not equal zero, $\hat{\alpha}$. $\hat{\psi}$ represents the naive estimate of whether nominal rigidity results in wages being higher or lower than they would

otherwise be, conditional on the explanatory variables, X . As expected, all three estimates for this error are negative and statistically significant, indicating that nominal rigidity has lowered wages instead of increasing them. The predicted forecast error $\hat{\alpha}$, however, shows why this is a biased estimate. The average forecast errors for observations where ΔW_{it} is small but not zero are negative. As a result, $\hat{\alpha}$, found by extrapolating the forecast errors when $\Delta W_{it} > 0$, is also negative.

Table 3: Actual vs. expected wage-change residual

	OLS (full sample)	Tobit (excl. $\Delta W < 0$)	OLS (excl. $\Delta W = 0$)
$\hat{\psi}$	-3.24 (0.000)	-3.09 (0.010)	-3.72 (0.000)
$\hat{\alpha}$	-1.96 (0.000)	-1.93 (0.000)	-2.18 (0.000)
$\hat{\lambda}$	0.34 (0.000)	0.33 (0.000)	0.35 (0.000)
$\hat{\psi} - \hat{\alpha}$	-1.28 (0.184)	-1.16 (0.183)	-1.54 (0.117)
p values given in parentheses for one-sided tests			

The key result in Table 3, then, is not that $\hat{\psi} < 0$, but rather that $\hat{\psi} < \hat{\alpha}$. This implies that, even when controlling for the fact that we expect a negative residual when using the model to explain wage freezes, the model still finds that wage freezes involve lower-than-expected wage changes. This difference between the average and predicted forecast errors is not statistically significant at the 10 per cent level in any of the three cases. However, the sign of the difference gives a point estimate that wage freezes tend to result more from positive wage changes being censored down to zero than from negative wage changes being censored up.

As noted earlier, if there is right as well as left censoring, then the Tobit estimation—which assumes only left censoring—will produce estimates of βX_{it} that are biased downwards, and thus estimates of ψ that are biased upwards. This could explain why the gap between the average actual and the predicted forecast error is less with this estimation than with the other two.

We also tried estimating $\hat{\alpha}$ and $\hat{\lambda}$ in Equation (6) by taking observations for ΔW_{it} only close to zero, because those residuals should provide a better extrapolated prediction for the expected residual at $\Delta W_{it} = 0$. The qualitative results of Table 3 were unaffected by this. As a final check,

Figure 3 allows an eyeball estimation of $\hat{\alpha}$. This graph plots ΔW versus the residuals from the full-sample OLS estimation of the wage equation and shows the estimated regression line as well as a spline through the medians for each value of ΔW_{it} . At $\Delta W = 0$, the regression line is clearly above the spline. (i.e., $\hat{\psi} < \hat{\alpha}$, the opposite result predicted by downward nominal-wage rigidity.)

5.3 The employment equation

For the employment equation, we estimated the following equation under various restrictions:

$$\Delta E_{it} = \theta^c + \theta^{q_1} Q_{it} + \theta^{q_2} Q_{it-1} + \sum_{\tau=1}^T \theta^{y_\tau} DY_{it}^\tau + \sum_{\alpha=1}^A \theta^{r_\alpha} DR_i^\alpha + \delta D0_{it} + \gamma \Delta W_{it} + \mu_{it}, \quad (9)$$

where ΔE_{it} is the average annual percentage change in employment over the duration of the contract, and the other variables are as defined in Equation (8). As we described in Section 4, the wage-change is included in this reduced-form equation to deal with a problem in which ongoing labour-demand shocks tend to affect ΔW_{it} and ΔE_{it} in the same direction, thus biasing down the estimate of δ due to the correlation between wage freezes and ΔW_{it} . If labour-demand shocks rather than wage-bargaining shocks are the predominate cause of wage changes, then the coefficient on ΔW_{it} will be positive. We also include Q_{it} and Q_{it-1} in some regressions. Recall that Q_{it} is the average annual growth rate in industry output in the industry to which firm i belongs over the period of the contract signed in period t . We include these variables as alternative ways of controlling for labour-demand shocks that affect both ΔW_{it} and ΔE_{it} .

Although we follow Simpson, Cameron, and Hum in using output to control for labour-demand shocks, we do not face the same problem that, for a constant level of labour productivity, output cannot distinguish between demand shocks and wage shocks as a source of employment changes. This is because our output variables are taken at the industry level while the employment change and wage-freeze dummy are firm-level variables. Given that there is not a high degree of intra-industry correlation across firms in DO_{ij} , this difference in the level of aggregation will imply that changes in output will largely reflect shocks that are industry-wide whereas the wage-freeze variable will be firm-specific.

Table 4 shows the results from the OLS estimation of the employment-change regression. In Column I, we do not try to control for labour-demand shocks that occur prior to the start of the contract. The comparison with Column II illustrates the bias that can result from this. The first thing to note here is the positive and statistically significant coefficient on the wage-change variable in Column II. This suggests that the variables in X have failed to fully capture the effects of labour-demand shocks that affect employment and the wage rate in the same direction. Given this positive

correlation between ΔW_{it} and ΔE_{it} , the feature highlighted in Table 3 (wage equation tends to over-predict when $\Delta W_{it} = 0$) is likely to extend to the employment equation. This is confirmed by comparing the coefficient on the wage dummy in the first two columns in Table 4. In Column I, the coefficient is negative, which would provide a point estimate that nominal-wage rigidity is associated with employment loss. When controlling for the fact that low wage increases are associated with low employment gains, however, the coefficient on the wage-freeze dummy changes signs. Although the coefficient is not statistically significant in either case, the magnitude of the difference between the two regressions is large. In Columns III and IV of Table 4, we include the change in industry output for the period of the previous contract to see if that can control for demand shocks in place of ΔW_{it} . The comparison of the coefficient on the wage-freeze variable between Columns I and III suggests that this variable is picking up some of the effect of labour-demand shocks in jointly affecting ΔW_{it} and ΔE_{it} . However, the positive and statistically significant coefficient in Column IV when ΔW_{it} is included as a right-hand-side variable suggests that there is still enough correlation between ΔW_{it} and ΔE_{it} , to severely bias down the coefficient on the wage-freeze variable. This is confirmed by comparing that coefficient between Columns III and IV.

Table 4: Employment adjustment 1978–1997^a

	I	II	III	IV
Wage change (one-sided)	—	0.26 (0.003)	—	0.22 (0.046)
Wage freeze (one-sided)	-0.88 (0.316)	0.11 (0.426)	0.61 (0.326)	1.41 (0.246)
Δ Industry output (two-sided)	-0.005 (0.929)	-0.005 (0.927)	0.000 (0.989)	0.008 (0.907)
Δ Industry output (-1) (two-sided)	—	—	0.03 (0.585)	0.03 (0.585)
\bar{R}^2	0.021	0.025	0.018	0.020
Sample size	1557	1557	1192	1192
p values given in parentheses for the type of test specified under the variable name				

a. All regressions include year and regional dummy variables.

Overall, the results in Table 4 show two things: first, there is a positive correlation between wage changes and employment changes that cannot be controlled for using output-change

variables; second, when both types of demand-shock controls are included, the point estimates suggest that nominal-wage rigidity is associated with increases rather than decreases in employment. The explanatory power of the regressions, however, is very low and so the positive coefficients on the wage-freeze variable are not statistically significant. In the next section, we consider some alternative ways of running our employment equation to check the robustness of our results.

6. Sensitivity checks

The lack of explanatory power in our employment equations probably results in part from both measurement error in the dependent variable and the lack of firm-level explanatory variables. Although there is little that we can do to address these problems with the data set used in this paper, there are some sensitivity checks we can apply to check the robustness of the results.

6.1 Censoring of the dependant variable

We noted in Section 2 that there is a large spike at zero in the employment-change distribution that may be due to the definition of the employment coverage⁷ and/or to round-off error in reporting of employment levels. In this case, the spike would be indicative of measurement error in the dependant variable that would not bias the estimated coefficient but would increase the variance in the error term of the regression and hence increase the standard errors around the estimated coefficients. Although both round-off error and costly adjustment can add measurement error at all points along the employment-change distribution, the error is likely to be more prevalent in those observations with a reported change of zero. We therefore tried estimating our employment equations, omitting all observations for which there was zero employment change. This implied removing about one-third of the data, including close to half of the observations with wage freezes. The results are reported in Table 5. Although there was some improvement in the overall fit of the equation (as measured by the adjusted R-squared), the standard errors on the wage-freeze term increased. This was probably the result of there being only 21 and 19 observations with wage freezes in Columns I and II and Columns III and IV, respectively. The estimated coefficient on the wage-freeze term was reduced in each specification. This implies that the employment changes reported as zero in our data set were more likely to have been higher than expected than were the non-zero employment changes. To the extent that this reflects a structural fact rather than random

7. Employment covers in some instances a bargaining unit, not all of whose members may be actually employed.

variation in a small sample, it suggests that employment-adjustment costs impose greater downward rigidity of employment than upward rigidity for firms bound by nominal-wage rigidity. This would be consistent with increases in trend employment.

Table 5: Employment adjustment excluding zero changes

	I	II	III	IV
Wage change (one-sided)	—	0.37 (0.003)	—	0.35 (0.035)
Wage freeze (one-sided)	-2.28 (0.239)	-0.96 (0.336)	-0.06 (0.443)	1.08 (0.366)
Δ Industry output (two-sided)	0.015 (0.845)	0.014 (0.856)	0.034 (0.737)	0.042 (0.682)
Δ Industry output (-1) (two-sided)	—	—	0.074 (0.379)	0.074 (0.376)
\bar{R}^2	0.038	0.045	0.035	0.037
Sample size	1046	1046	803	803
p values given in parentheses for the type of test specified under the variable name				

6.2 Definition of a wage change

One of the problems when analyzing wage-settlements data to look for evidence of nominal-wage rigidity is to define exactly what is meant by a wage change, and hence which contracts have experienced wage changes of zero. There are a number of reasons why the definition of a wage change used in this paper may overstate or understate the true extent of nominal rigidity in the firms covered by the data set.

First, we have information only on the base wage at each firm. The base wage is typically the wage paid to the lowest-skill group covered by the collective-bargaining agreement. An implicit assumption for our analysis is that variation over time in intra-contract wage differentials is small enough to be safely ignored so that the evolution of the base wage is representative of the pattern of wages at each firm in general.

Second, even for the workers in the skill group covered by the base wage, that wage does not describe other bargained-over aspects of worker compensation such as benefits, guarantees of job security, etc. It is possible that the presence of nominal-wage rigidity (either downward or

symmetrical) is a relatively benign constraint as it simply transfers the margin of bargaining from wages to some non-wage form of compensation. However, any overestimate of the measured wage rigidity will likely be offset by a downward bias to its coefficient, thus not affecting the magnitude of the total estimated employment effect, although it could possibly lower the power of the test to reject the null hypothesis of no employment effects.

Finally, as noted in Section 2, the definition of a wage change matters a lot. Our data set allows us to construct the employment change over the lifetime of the contract but not for each individual year of multi-year contracts. Because we needed to match the employment-change and wage-change data in our employment regressions, we adopted an equivalent lifetime definition of the wage change. This will, however, tend to understate the extent of nominal-wage rigidity in the data and the associated employment effects, as employment may have been temporarily affected by nominal rigidity in the first year or two of a contract before returning to the unconstrained level in the final year. Although our data show the wages in each year of the contract, they do not provide a way of observing temporary employment effects that get reversed before the end of the contract. This problem would bias down the absolute magnitude of the estimated effect of nominal rigidity on employment but not bias its sign. As a check, we tried running our employment equations including only those observations containing one-year contracts. These are reported in Table 6. With this subset of the data, there is a dramatic increase in the point estimates of the employment effect from nominal-wage rigidity, even in the specifications where the wage-change term is not included to control for labour-demand shocks. One interpretation of this is that wage freezes in one-year contracts frequently represent neither downward nor symmetrical rigidity, but simply the failure to negotiate a new contract during the 12 months following the expiration of the previous contract. Given the general tendency for nominal wages to rise, such a failure to negotiate a new contract would constitute a very strong form of upward rigidity on nominal wages, which is consistent with the strongly positive point estimates for the change in employment. This result illustrates why the very high proportion of wage freezes in the data under the first-year definition of a wage-change almost certainly overstates the extent of downward nominal-wage rigidity; the presence of a wage-freeze in the first year of a multi-year contract will frequently represent negotiation delays rather than genuine wage rigidity.

6.3 Specification of the wage-freeze term

The positive coefficient on the wage-freeze dummy variable in our employment equations is interpreted to mean that, *on average*, employment was higher than would be expected in contracts with zero wage change, given the other information about the contract. This is consistent with the

Table 6: Employment adjustment using one-year contracts

	I	II	III	IV
ΔW_{it} (one-sided)	—	0.37 (0.149)	—	0.18 (0.373)
$D0_{it}$ (one-sided)	6.01 (0.207)	8.52 (0.270)	8.90 (0.182)	9.90 (0.136)
Δ Industry output (two-sided)	-0.229 (0.356)	-0.206 (0.400)	-0.356 (0.322)	-0.34 (0.326)
Δ Industry output (-1) (two-sided)	—	—	0.02 (0.949)	0.02 (0.942)
\bar{R}^2	0.12	0.13	0.14	0.14
Sample size	162	162	125	125
p values given in parentheses for the type of test specified under the variable name				

fact that, in our wage equation, we found wage freezes to result in wages that were lower on average than would be expected. This suggests an alternative way of modelling the wage-freeze term that could use some of the information masked by these averages. Specifically, we replace the wage-freeze dummy variable in our employment equations with a wage-shock term equal to the difference between the actual wage change and the wage change predicted by our wage equation. The wage-shock term is an estimate of how much higher the wage change is because of nominal-wage rigidity.

We report the results of these regressions in Table 7. Because of the way the wage-shock variable is defined, the sign on its coefficient is negative. The wage-shock term is negative on average (since our wage equation tended to overpredict wage changes that were zero). Thus, the negative coefficient on the wage-shock term in Table 7 is consistent with our previous point estimates that wage freezes have been associated with greater-than-expected employment changes.

One way of interpreting these results is that nominal-wage rigidity provides a natural experiment that can help identify the labour-demand curve. The maintained assumption of our analysis has been that the wage-change variable captures systematic demand and wage-bargaining shocks (other than those caused by nominal-wage rigidity) that are not captured by the other

Table 7: Employment adjustment using wage-shock variable

	I	II	III	IV
ΔW_{it} (one-sided)	—	0.25 (0.032)	—	0.25 (0.033)
<i>Wage shock</i> (one-sided)	-0.29 (0.274)	-0.60 (0.123)	-0.27 (0.271)	-0.58 (0.130)
Δ Industry output (two-sided)	0.003 (0.956)	0.015 (0.824)	0.004 (0.952)	0.016 (0.821)
Δ Industry output (-1) (two-sided)	—	—	0.02 (0.615)	0.02 (0.647)
\bar{R}^2	0.019	0.021	0.019	0.021
Sample size	1192	1192	1192	1192
p values are given in parentheses for the type of test specified under the variable name				

covariates, and that nominal-wage rigidity then represents an exogenous wage shock. Under this assumption, the coefficient on the wage-shock variable in Columns II and IV of Table 7 can be interpreted as the elasticity of labour demand. We would not want to push this interpretation too far, however. Our methodology for estimating the wage shock understates the variability in this variable, as the predicted wage can capture the variability only in the explanatory variables of our wage equation and not in the error term. This would tend to bias up the estimate of the (absolute) elasticity. On the other hand, as we noted above, the existence of wage freezes may overstate the degree of wage rigidity as there are other margins such as performance pay on which employers can adjust compensation. Therefore, the measured variation in the wage-shock term might overstate the true variation in the exogenous wage change caused by wage rigidity, and thus bias down the estimate of the elasticity.

Our reason for using the wage-shock term, however, is not to estimate the elasticity of demand but rather to check the robustness of our results. The higher significance level of the wage-shock term compared with the wage-freeze dummy variable in Table 4 suggests that there is some useful information contained in our wage equation. Also, the resulting estimate of a negative elasticity of demand suggests that the underlying assumption—that employment is demand-constrained rather than supply-constrained, so that wage shocks due to nominal-wage rigidity result

in movement along negatively sloped labour-demand curves—is valid. Therefore, although neither our point estimate of a negative effect on wages on average from nominal rigidity nor our point estimate of a positive effect on employment is statistically significant, we do have some confidence that the two results are independent results. That is, if we treat the wage-settlements data base as a population, we have some confidence that the lower-than-expected wage changes associated with wage freezes in that population are directly related to the higher-than-expected employment changes via a labour-demand relationship. If we think of the data base as a sample, we cannot be sure that the point estimates are not simply the result of there being an unusually high degree of upward rather than downward rigidity in the sample. These results do suggest, however, that we should be very wary of using histogram analysis of this data base to conclude that there is substantial downward nominal-wage rigidity in the economy, and that *ipso facto*, there is a permanent negative effect on employment from low inflation.

6.4 Aggregate data

Simpson, Cameron, and Hum (1998) use the same data set as that used here to estimate an employment equation, but apply a slightly different methodology. Specifically, they aggregate the firm-level data into industry and time-period groups, and then use the average employment growth in each industry-period as their dependant variable. Their wage-freeze variable is the proportion of firms in the industry-period that experienced wage freezes. Simpson, Cameron, and Hum find that the coefficient on wage freezes is negative and statistically significant. As we noted above, however, we believe that their estimated coefficient on wage freezes is biased downwards because they do not adequately control for labour-demand shocks. Faruqui (1999) has applied the methodology outlined in this paper to the industry-aggregated data of Simpson, Cameron, and Hum. The use of aggregate data solves a lot of the problems with measurement error in the data. Accordingly, the explanatory power of the regressions in his study is higher than with the disaggregated data, but the basic conclusions remain the same: The coefficient on the wage-change variable is positive, indicating the presence of demand shocks not captured by other variables. Thus a downward bias on the wage-freeze variable could be present in regressions where the wage change is omitted. Faruqui found, however, that the sign of the coefficient on the wage-freeze variable is highly sensitive to the specification and method of aggregating the data.

7. Conclusion

This paper evaluates evidence of the employment effect of downward nominal-wage rigidity, using firm-level data from the Canadian manufacturing sector. Most papers looking for such evidence

have used the counterfactual-distribution approach to show the extent of wage-change censoring that might be due to downward rather than symmetrical rigidity. In some applications, however, this approach requires restrictive assumptions on the wage-change distribution. We estimate the effects of downward nominal rigidity by directly estimating reduced-form wage and employment equations.


For the wage equation, we used a two-stage procedure that controls for the fact that wage freezes represent unusually low wage changes and so create a correlation between a wage-freeze dummy variable and the error term in the equation, thus biasing the results against finding evidence for downward nominal-wage rigidity. The results of this two-step estimation, using a variety of estimation approaches, suggest that the wage-change equation is not underpredicting wage changes at the zero wage-change rate, and thus there is no evidence of downward nominal-wage rigidity.

For the employment equation, we show that there is a positive correlation between wage changes and employment changes in the data, due to the apparent presence of unexplained labour-demand shocks. As a result, there is again a negative correlation between the wage-freeze dummy and the error term in the equation that biases the results in favour of finding evidence of an employment cost to downward rigidity. When we control for this by adding wage-changes as a right-hand-side variable, our point estimates change from indicating an employment cost from nominal rigidity to an employment gain.


Unfortunately, due to the paucity of firm-level data matched to our wage and employment data, our tests are not very powerful. As a result, none of our point estimates for whether nominal rigidity has increased or decreased wages or whether it has generated increases or decreases in employment at those firms constrained by the rigidity is statistically significant. Our main conclusion, therefore, is methodological: Our results show clearly that, when we control for the bias that arises from wage freezes representing unusually low wage changes, we generate a substantial change in the point estimates indicating the direction of nominal rigidity and its associated employment effects. Future work using reduced-form wage and employment equations needs to take that bias into account.

Figure 1: Wage histograms


Low inflation (1992–1997)



Moderate inflation (1983–1991)



High inflation (1978–1982)



Full sample






Figure 2: Employment histograms


Low inflation (1992–1997)



Moderate nflation (1983–1991)



High inflation (1978–1982)



Full sample





Figure 3: Residual and actual wage changes

References

- Akerlof, G., W. Dickens, and G. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 1: 1–59.
- Card, D. and D. Hyslop. 1996. "Does Inflation 'Grease the Wheels of the Labor Market'?" NBER Working Paper No. 5538.
- Crawford A., and A. Harrison. 1998. "Testing for Downward Rigidity in Nominal Wage Rates." In *Price Stability, Inflation Targets, and Monetary Policy*, 179–218. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.
- Faruqui U. 1999. "Employment Effects of Nominal Wage Rigidity: An Examination using Wage-Settlements Data." Bank of Canada Working Paper. Forthcoming.
- Fortin, P. 1996. "The Great Canadian Slump." *Canadian Journal of Economics* 29: 761–787.
- Golub, J. 1993. "Inflation, Inflation Uncertainty, and Relative Price Variability: A Survey." Federal Reserve Bank of Kansas City, Working Paper No. RWP 93-15.
- Hogan, S. 1997. "What Does Downward Nominal-Wage Rigidity Imply for Monetary Policy?" Bank of Canada Working Paper No. 97-13.
- . 1998. "What Does Downward Nominal-Wage Rigidity Imply for Monetary Policy?" *Canadian Public Policy* 24: 513–525.
- Kahn, S. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review* 87: 993–1008.
- Simpson, W., N. Cameron, and D. Hum. 1998. "Is Hypoinflation Good Policy?" *Canadian Public Policy* 24: 291–308.
- Tobin, J. 1972. "Inflation and Unemployment." *American Economic Review* 62: 1–18.

Bank of Canada Working Papers

Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. *Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.*

1999

99-20	The Expectations Hypothesis for the Longer End of the Term Structure: Some Evidence for Canada	R. Lange
99-19	Pricing Interest Rate Derivatives in a Non-Parametric Two-Factor Term-Structure Model	J. Knight, F. Li, and M. Yuan
99-18	Estimating One-Factor Models of Short-Term Interest Rates	D. Mc Manus and D. Watt
99-17	Canada's Exchange Rate Regime and North American Economic Integration: The Role of Risk-Sharing Mechanisms	Z. Antia, R. Djoudad, and P. St-Amant
99-16	Optimal Currency Areas: A Review of the Recent Literature	R. Lafrance and P. St-Amant
99-15	The Information Content of Interest Rate Futures Options	D. Mc Manus
99-14	The U.S. Capacity Utilization Rate: A New Estimation Approach	R. Lalonde
99-13	Indicator Models of Core Inflation for Canada	R. Dion
99-12	Why Canada Needs a Flexible Exchange Rate	J. Murray
99-11	Liquidity of the Government of Canada Securities Market: Stylized Facts and Some Market Microstructure Comparisons to the United States	T. Gravelle
99-10	Real Effects of Collapsing Exchange Rate Regimes: An Application to Mexico	P. Osakwe and L. Schembri
99-9	Measuring Potential Output within a State-Space Framework	M. Kichian
99-8	Monetary Rules When Economic Behaviour Changes	R. Amano, D. Coletti, and T. Macklem
99-7	The Exchange Rate Regime and Canada's Monetary Order	D. Laidler
99-6	Uncovering Inflation Expectations and Risk Premiums from Internationally Integrated Financial Markets	B.S.C. Fung, S. Mitnick, and E. Remolona
99-5	The Quantity of Money and Monetary Policy	D. Laidler
99-4	An Intraday Analysis of the Effectiveness of Foreign Exchange Intervention	N. Beattie and J-F. Fillion
99a-3	Forecasting GDP Growth Using Artificial Neural Networks	G. Tkacz and S. Hu
99-2	Capital Gains and Inflation Taxes in a Life-cycle Model	C. Leung and G.-J. Zhang
99-1	Dynamic Employment and Hours Effects of Government Spending Shocks	M. Yuan and W. Li

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à:

Publications Distribution, Bank of Canada
234 Wellington Street Ottawa, Ontario K1A 0G9

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9

E-mail / Adresse électronique: publications@bank-banque-canada.ca

WWW: <http://www.bank-banque-canada.ca/>