

IZA DP No. 2799

Real and Nominal Wage Rigidities in Collective Bargaining Agreements

Louis N. Christofides
Paris Nearchou

May 2007

Real and Nominal Wage Rigidities in Collective Bargaining Agreements

Louis N. Christofides

*University of Cyprus, University of Guelph,
CLLRNet, CESifo and IZA*

Paris Nearchou

University of Cyprus

Discussion Paper No. 2799
May 2007

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Real and Nominal Wage Rigidities in Collective Bargaining Agreements^{*}

An earlier study of wage agreements, reached in the Canadian unionized sector between 1976-99, found that wage adjustment is characterized by downward nominal rigidity and significant spikes at zero. We extend this earlier approach to encompass the possibility of real as well as nominal wage rigidity. The addition of real wage rigidity variables enhances earlier results and suggests that real rigidity increases significantly the mass in the histogram bin containing the mean anticipated rate of inflation, as well as in adjacent bins. Downward nominal wage rigidities and spikes at zero remain important.

JEL Classification: J52, J31

Keywords: real, nominal wage rigidities

Corresponding author:

Louis N. Christofides
Department of Economics
University of Cyprus
Kallipoleos 75
P.O. Box 20537
1678 Nicosia
Cyprus
E-mail: louis.christofides@ucy.ac.cy

^{*} We thank M. Legault, Human Resources Development Canada, for the data and the Social Sciences and Humanities Research Council for financial support. The paper was presented by Nearchou at the September 2006 EALE conference in Prague. Nearchou thanks IZA for the opportunity to present his work at the 9th IZA European Summer School in Labor Economics in April 2006.

1 Introduction

Monetary policies in a number of countries have, at least until the current oil price shocks, succeeded in limiting price inflation. A by-product of this success has been concern with the extent to which this inflation record has been achieved at a cost. In a low inflation environment, downward nominal wage rigidity (DNWR) may mean that nominal-wage reductions, called for by bargaining pair-specific productivity shocks, do not occur, thereby compromising the efficiency of the labour market. Indeed, some studies go as far as to look for the unemployment consequences of such low-inflation mechanisms. If inflation greases the wheels of the labor market, then its absence may lead to costs. An expanding literature covering a number of countries takes advantage of the recent periods of low price inflation and attempts to measure the extent and consequences of DNWR.¹ This literature has been further energized by the International Wage Flexibility Project (IWFP), led by William Dickens and Erica Groshen.²

An important concern of studies in this literature should be the extent to which real rigidities can be treated as part and parcel of the more general wage adjustment process. Naturally, the extent to which price inflation and particularly anticipated price inflation feed into nominal-wage adjustment is a subject that goes at least as far back as Friedman (1968). While nominal-wage adjustment is clearly conditioned by price inflation effects, the extent to which downward real wage rigidity (DRWR) exists, its implied impact on the shape of the wage adjustment distribution in the neighborhood of the anticipated rate of inflation, and possible interactions of this process with DNWR are issues that deserve further attention.

A particularly good data set for studying these effects is the Human Resources Development Canada (HRDC) record of the provisions of collective bargaining agreements reached in the Canadian unionised sector. The data is thought to be very accurate because it refers to legally

¹An extensive review of the literature is contained in Christofides and Leung (2003).

²Much more information is provided in the proceedings of the project's Final Conference (June 17-18, 2004).

binding provisions, it covers all industries over all of Canada, and it covers high as well as low inflation periods since 1976. In an earlier paper by Christofides and Leung (2003), the HRDC data were used to examine DNWR and menu cost behaviour in the period 1976-1999 using parametric techniques inspired by Kahn (1997). In this paper, we extend the earlier study to more explicitly encompass DRWR and its interaction with DNWR. A strength of the HRDC data for current purposes is that the diverse inflation experience that it encompasses makes it possible to differentiate DNWR from DRWR processes. The results obtained indicate significant and substantial nominal and real wage rigidity in the contract data.

Our approach is distinctively different from other recent studies that also test for the presence of both types of rigidity, including, among others, those of Bauer et. al. (2003) for Germany and Barwell and Schweitzer (2004) for the UK. Both studies find evidence for the presence of both types of rigidity, with real rigidity being more pervasive. Their approach builds upon the maximum-likelihood methodology originally proposed by Altonji and Devereux (2000) for the testing of DNWR alone, and requires parametric assumptions about the family of the rigidity-free nominal-wage-growth distribution. In contrast to them, we make no such assumptions, nor do we impose a symmetric structure, as it has often been done in studies that examine the presence of DNWR. For the identification of the shape of the rigidity-free distribution and the size of the distortions due to the presence of rigidity we exploit the fact that we have several yearly samples from nominal-wage-growth distributions whose shape is affected by Downward Wage Rigidity (DWR) differently from year to year.

The rest of the paper is organised as follows: In Section 2, we consider the effect of the presence of each type of rigidity on the wage-growth distribution and in Section 3 we present more details on the data and sources. The empirical specification and estimation issues are presented in Sections 4 and 5 respectively. The results obtained are described in Section 6, and concluding observations appear in Section 7.

2 Downward Wage Rigidity and Wage Growth Distributions

We take DNWR to describe that feature of the wage adjustment process where agents, individual employees or unions, are reluctant to accept a nominal-wage cut (negative wage growth) and instead would settle for a nominal-wage freeze (zero growth). Justifications for nominal rigidity range from the comparability and fairness arguments documented in Bewley (1999) to the theoretical papers by MacLeod and Malcomson (1993), Malcomson (1997), Holden (1994) and Holden (2004) which build on the notion that nominal wages can be changed only by mutual consent.

At the population level, this reluctance would mean fewer cuts in nominal wages and more nominal-wage freezes *relative to* the case of no rigidity. In terms of the distribution of nominal-wage-growth rates, this translates into a shift of probability mass from negative values of the support of the distribution towards the point zero. Therefore the rigidity-contaminated nominal-wage-growth distribution would show a deficit of probability mass for negative values of the support, and a surplus at point zero, relative to the rigidity-free (or notional) distribution. At the same time, the two distributions would be identical beyond the point zero.

We can see these effects in the two diagrams of the top row of Figure 1, where we have simulated the rigidity-free and the rigidity-contaminated nominal-wage-growth distributions for two particular types of DNWR mechanisms. In these two diagrams, as well as the rest in Figure 1, both distributions are represented by their probability histograms. The horizontal axis measures the nominal-wage growth, in percentage rates, and the vertical axis the probability mass that falls in the bins of the histogram. The rigidity-free distribution is always represented by the light-shaded bars, and the rigidity-contaminated by the dark-shaded bars. In these diagrams, the distortion in the nominal-wage-growth distribution is manifested by the difference between the height of the corresponding bars of the rigidity-contaminated and notional histograms. The

distortion could take the form of a (probability mass) deficit when the bars of the rigidity-contaminated histogram are shorter than the corresponding bars of notional histograms, and the form of a surplus when the opposite is true. An overall shift of probability mass to the right is detected when there is a collection of bins with a surplus that lies to the right of a collection of bins with a deficit.³

In the leftmost diagram, we consider the case of *absolute* DNWR, where all agents facing a nominal-wage cut succeed in settling for a nominal-wage freeze. Therefore, at the population level, there should be no nominal-wage cuts, but instead an ‘excessive’ amount of nominal-wage freezes. In the diagram we see that, although in the absence of rigidity there would be a number of wage cuts, indicated by the positive height of the bars of the notional probability histogram for the bins with negative values, the height of the corresponding bars of the rigidity-contaminated histogram is zero. At the same time, for the latter histogram, the missing mass is concentrated in the bin that contains the point zero, while the two histograms coincide beyond that. The case of absolute DNWR could be considered as the extreme scenario of *partial* DNWR. At the population level, the case of partial DNWR (rightmost diagram, row 1, Figure 1) would mean that there is a positive number of nominal-wage cuts, but also that there is an ‘excessive’ amount of nominal-wage freezes, although to a lesser extent than the case of absolute DNWR, other things being equal.⁴

DRWR can be defined in a similar way to DNWR. In particular, it is taken to describe the

³In order to make the comparison of the effects of the various types of wage rigidity easier, we have drawn the same rigidity-free distribution in all diagrams. We have also deliberately made the bins of the rigidity-contaminated histogram narrower in order to be easier to distinguish the two histograms.

⁴In all simulations presented in this paper where we look into cases of partial downward wage rigidity, we have assumed the ‘proportional’ type, where each agent facing a wage cut, either in nominal or real terms, depending on the type of rigidity we examine, faces the same probability of settling for the respective type of wage freeze. This assumption does not influence our conclusions.

situation where agents are reluctant to accept real-wage cuts but instead would settle for a real-wage freeze. In practice, this attitude takes the form of reluctance towards accepting reductions in the *anticipated* real wage since, at the time of bargaining, future inflation is typically unknown. The anticipated real-wage level is based on their belief, at the time of bargaining, about the future level of inflation and might be determined as described in theoretical constructs such as efficiency wages, efficient bargains, and implicit contracts. In this paper, we do not concern ourselves with how this anticipated real wage might be determined.

As in the case of DNWR, the presence of DRWR would distort the shape of the nominal-wage-growth distribution. To see how this could happen, we first note that DRWR could also be described as the situation where agents are reluctant to accept nominal-wage-growth rates that are below their anticipated rate of inflation for the period the wage is bargained for, and instead would settle for growth rates that are equal to that. At the population level, this would mean that agents who face nominal-wage growth at a rate below anticipated inflation would settle for a nominal-wage increase equal to the anticipated rate of inflation. Consequently, the presence of DRWR would shift probability mass to the right, from smaller values of nominal-wage growth towards the values of anticipated inflation in the population. The exact form of the shift of mass to the right towards the values of anticipated inflation depends on the nature of the rigidity mechanism and the joint distribution of the notional (nominal) wage growth and anticipated inflation among all agents.

Nevertheless, without any distributional assumptions, it is possible to distinguish three regions in the nominal-wage-growth distribution for which we can make qualitative predictions about the nature of the distortions introduced. Firstly, the interval of values that lies to the left of the support of the distribution of anticipated inflation, if one exists, could only loose mass to the right since all agents whose nominal-wage growth falls in this region face the prospect of a real-wage cut. Therefore, in this region, the rigidity-contaminated distribution can only

exhibit a deficit. Secondly, the interval of values that lies to the right of the support of the distribution of anticipated inflation, if one exists, would not be distorted, since all agents whose nominal-wage growth falls in this region face the prospect of a real-wage increase. Finally, the interval of values that corresponds to the support of the distribution of anticipated inflation, will attract mass from its left, and therefore for this interval the rigidity-contaminated distribution will exhibit a surplus, *in total*. However, it is possible that, in some parts of this interval, the rigidity contaminated distribution will exhibit a deficit. In terms of the probability histogram, we can understand how this could happen by noting that a particular bin that contains values of anticipated inflation can attract mass from bins to its left but at the same time loose mass to bins to its right that also contain values of anticipated inflation. The net effect cannot be clear without knowledge of how the notional-wage growth and anticipated inflation are jointly distributed. The only exception is the rightmost bin in this region, for which we know that it cannot exhibit a deficit since all other bins that contain values of anticipated inflation lie to its left. Despite this uncertainty, we could assume that it would be more likely that bins that lie further to the left in this interval will show a deficit and bins further to the right will show a surplus.

In order to see how in practice DRWR could distort the shape of the nominal-wage-growth distribution, we consider several examples of the presence of DRWR that differ from each other with respect to the characteristics of the distribution of anticipated inflation and the extent of the rigidity.⁵ First, we consider the case of *firm and uniform beliefs*, where all agents anticipate the same value of future inflation (marked by the vertical broken line in the diagrams of row 2, Figure 1). In the case of *absolute* DRWR, the presence of DRWR will shift probability mass towards the value of anticipated inflation from its left, in a similar way that the presence of

⁵In all simulations we assume that nominal-wage growth and inflation beliefs are independent. This does not affect our conclusions.

DNWR will shift probability mass from negative values towards the point of zero nominal-wage increase (leftmost diagram, row 2, Figure 1). In the case of *partial* DRWR, there remain a number of anticipated real-wage cuts but there is also an ‘excessive’ number of wage increases just equal to the anticipated rate of inflation (rightmost diagram, row 2, Figure 1).

More realistically, agents disagree on their beliefs of the future level of inflation, and thus there is a distribution of values of anticipated inflation among the population members.⁶ First, we consider as a benchmark the implausible scenario where the marginal distributions of notional-wage growth and anticipated inflation coincide on a point-by-point basis; that is, each and every agent faces a notional-wage growth exactly equal to the agent’s level of anticipated inflation. In this case, all agents effectively experience an anticipated real-wage freeze, and, therefore the actual distribution will coincide with the notional. If we relax the assumption that all agents face the prospect of an anticipated real-wage freeze and, instead, assume that some are faced with an anticipated real-wage reduction (and at the same time keep the assumption that the two marginal distributions coincide), then the presence of DRWR could shift some probability mass to the right (leftmost diagram, row 3, Figure 1).⁷

More likely, the two marginal distributions will not coincide, the support of the anticipated-inflation distribution can be assumed to lie within the support of the nominal-wage change distribution, and there will be agents in the population that face the prospect of an anticipated real-wage decrease. Examples of this case are depicted in the rightmost diagram of the third row, and the leftmost diagram of the fourth row of Figure 1. In the first diagram we consider the case where all the agents who face the prospect of a real-wage cut manage to avoid it (diagram with caption Absolute DRWR(b)), while in the second diagram only some agents manage to

⁶Firm and uniform beliefs could be seen as a special case of this.

⁷In this diagram, as well as the remaining diagrams in Figure 1, the probability density function of the non-degenerate anticipated-inflation distribution is depicted by the solid thick line.

avoid it (diagram with caption Partial DRWR(b)). The nature of the distortions in the two cases is qualitatively the same but they differ in terms of the size. Also note the nature of the distortion in the three regions of the rigidity-contaminated distribution relative to their position to the support of anticipated inflation (i.e. the absence of mass to the left of the minimum of the support of anticipated-inflation distribution, the absence of any effects to the right of its maximum, and the diverse patterns within it).

It is interesting to see what the presence of DRWR means for the distribution of *actual* real-wage growth. If we accept that typically the distribution of anticipated inflation extends below and above the realised inflation value, then the presence of DRWR is consistent with observing real-wage cuts (relative to the realised value of inflation), even in the case of absolute DRWR. Therefore, the occurrence of real-wage cuts does not, in general, suggest that DRWR does not exist; real-wage cuts would, however, rule out the case of absolute DRWR and perfect foresight.

Finally, when some collective agreements are affected by DNWR and others by DRWR, then both types of distortions will be present in the shape of the actual-wage-growth distribution. This case is depicted in the rightmost graph of the bottom row of Figure 1, where there is both a spike at the bin containing the point zero and deficit in probability mass for bins to the left as well as to the right of that bin. Note that the two types of distortions have similar effect at the bins below zero, i.e. they reduce the probability mass concentrated there. On the other hand, they have opposite effects at the bin containing zero, since the presence of DRWR shifts mass from that bin to other bins to its right (negative effect), while the presence of DNWR shifts mass to that bin from bins to its left (positive effect). The nature of the combined effect will depend on the proportion of agreements affected by each type of rigidity, as well as the intensity of each type. Moreover, there is probability surplus for the bins that lie towards the right tail of the distribution of anticipated inflation and no effect to the bins that lie beyond

the maximum level of anticipated inflation.

3 Data

The contract data used in this paper are compiled by HRDC, the federal ministry responsible for monitoring agreements between firms and unions. The database⁸ contains information on provisions for 10,945 wage contracts signed in the Canadian unionised sector and involves settlement dates as early as 1976 and as late as 1999. The agreements cover bargaining units involving 200 to nearly 80,000 employees, in both the private and the public sector, and their duration ranges from a few months to several years. Because reporting requirements apply, this information is thought to be very accurate. The data set that is used for the empirical analysis contains one observation for each contract, and the recorded information is the growth rate of the total nominal-wage adjustment ($WNC + COLA$)⁹ over the whole of the life of the contract, calculated in *annual* terms. The observation for each contract is allocated to the year the contract became effective.

Table 1 shows the number of contracts and the sample mean, median and standard deviation of the nominal-wage growth rate for each year in the observation period.¹⁰ Also, the corresponding annual rate of Consumer Price Index inflation (CPI) and an estimate of anticipated inflation (\widehat{P}^e) for that year.¹¹ From the CPI figures in column 6 one can see that

⁸See Christofides and Stengos (2003) for a detailed description.

⁹Our analysis deals with the total wage adjustment, which is composed of the non-contingent wage adjustment (WNC) and $COLA$. It should be noted, however, that, because the incidence and intensity of $COLA$ clauses is limited throughout the observation period, the results we obtain are similar to those that are based on the analysis of the non-contingent wage adjustment data.

¹⁰Because of the smaller number of contracts, the first two and the last three years in the sample are considered together in everything that follows.

¹¹The proxy for anticipated inflation is the one-year-ahead forecast from an AR(6) regression model with a GARCH(1,1) error process.

the observation period can be divided into three consecutive periods relative to the level of inflation: 1977-1983 could be considered as a high inflation period, with average inflation at 9.58%, 1984-1992 a medium inflation period, with average inflation at 4.67%, and 1993-1997 a low inflation period, with average inflation at 1.46%. The comparison of the mean (or median) wage-growth figures in columns 3 (or 4) with the *CPI* figures reveals that there exists a positive relationship between the level of realised inflation and the *location* of the wage-growth distribution across years. Also, the comparison of the standard deviation figures in column 5 with the *CPI* figures reveals a positive relationship between the level of realised inflation and the *spread* of the wage-growth distribution. We note that DWR leads to a compression of the wage-change distribution. Since DNWR is less likely to hold at higher rates of inflation, this mechanism could explain the positive relationship between the level of realised inflation and the spread of the wage-growth distribution. This positive relation may also arise if expectations about inflation are more diverse during high inflation periods.

Table 2 shows the incidence of nominal-wage adjustments relative to the value of zero and the realised level of inflation, by year. Only 102 (or 0.9%) of the contracts in the entire observation period show nominal-wage cuts, while a substantial number (1142 or 10.4%) show a wage freeze; jointly both figures could be considered as strong evidence in favour of the presence of DNWR. The wage freezes are particularly pronounced during the low inflation years; for each of the years 1993-1996 the proportion of contracts with a wage freeze was above 35%, peaking at 51.0% in 1993. On the other hand, 6045 (or 55.2%) of the contracts exhibit negative real-wage growth, while 4801 of them had at the same time positive nominal-wage growth. As expected, the number of contracts that had exactly zero real-wage growth is negligible, just 1 in this case, and the remaining 4899 (or 44.8%) contracts showed both nominal and real-wage increase.

4 Empirical Specification

The problem of testing for the presence of a particular type of rigidity using micro data could be stated as one where, having several yearly samples of observations on nominal-wage growth

$$\{\dot{w}_{ti}\}_{t=1,\dots,T}^{i=1,\dots,n_t} \quad (1)$$

where \dot{w}_{ti} represents the nominal-wage growth agreed by the i 'th bargaining unit in year t , T is the number of yearly samples and n_t the number of observations in sample t , one would want to test whether these were generated from rigidity-free or rigidity-contaminated yearly distributions. Formally, the hypotheses to be tested could be stated as follows

$$\begin{aligned} H_0 & : F_t(\dot{w}) = F_t^N(\dot{w}) \\ H_1 & : F_t(\dot{w}) = G^R(F_t^N(\dot{w})) \end{aligned} \quad (2)$$

where $F_t(\dot{w})$ is the cdf of the actual-wage-growth distribution, $F_t^N(\dot{w})$ the cdf of the notional-wage-growth distribution, and $G^R(F_t^N(\dot{w}))$ the cdf of the rigidity-contaminated wage-growth distribution, in year t . The functional $G^R(\cdot)$ is used generically to represent the distortions introduced by the presence of rigidity, which can be either DNWR ($R = n$), or DRWR ($R = r$), or both ($R = nr$).¹²

Exploiting the distinct nature of the distortions in the shape caused by the presence of each type of rigidity, a test for the presence of rigidity of type R could be based on the comparison of the shape of the estimated actual-wage-growth distribution with the shape of the notional distribution (the counterfactual): if there were statistically significant differences in their shape of similar nature to those one would expect to find if rigidity of type R were present, this could be considered as evidence in favour of the presence of rigidity of type R . Formally, this would require one to have information on both $F^N(\cdot)$, that describes the counterfactual distribution,

¹²In this setup, we ignore the presence of measurement error in the wage-growth data. This is a realistic assumption when we work with the Canadian contract data which are collected by the regulating agency HRDC.

and $G^R(\cdot)$, that characterises the differences due to the presence of rigidity. Obtaining information on the nature of $G^R(\cdot)$ is relatively straightforward, as we have already done informally in Section 2. This is not the case for $F^N(\cdot)$, since typically we do not observe the notional-wage growth and thus we cannot make inference about the shape of its distribution *directly*.

The way we proceed here is to use the available actual-wage-growth data to infer information about it *indirectly*, estimating jointly the notional distribution and the distortions due to DWR. The basic idea is to test the hypotheses about the shape of the actual-wage-growth distribution in terms of the heights of the bars of the corresponding probability histogram. Its implementation can be organised in two stages:

Stage 1: Formulation of hypotheses in terms of the parameters of the probability histograms The aim in this stage is to transform the original problem of testing hypotheses about the cdf of the distribution of the actual-wage-growth data from each year in the observation period, as described in (2), to one where we test equivalent hypotheses about the corresponding probability histogram.

First we define the probability histograms. Let $P_{jt} \equiv F_t(h_{j+1,t}) - F_t(h_{j,t})$ be the height of the bar of the probability histogram of the actual-wage-growth distribution in year t that corresponds to the j 'th bin, denoted by $\mathcal{B}_{jt} \equiv [h_{j,t}, h_{j+1,t}]$, where the bin index $j \in \{-J, \dots, 0, \dots, J\}$ indicates the position of the bins in the probability histogram.¹³ Given that our analysis aims to examine the shape of the distributions but not their location, j is defined to indicate the position of the bins relative to each other rather than relative to values on the real line. In particular, the bin indexed by $j = 0$ contains the median of the actual-wage-growth distribution, bins indexed by a negative j lie $|j|$ positions to the left of the median bin, and bins indexed by a positive j lie j positions to its right. We refer to the probability histograms defined in this

¹³Then, the collection of the $2J + 1$ bars defines the probability histogram for that year.

way as ‘standardised’.

Having defined the probability histograms in this particular way, in the next step we parameterise P_{jt} under the two hypotheses, i.e.

$$P_{jt} = \begin{cases} p^N(z_{jt}^N; b_j^N) & , \text{ if } H_0 \text{ is true} \\ p^R(z_{jt}^R; b_j^R) & , \text{ if } H_1 \text{ is true} \end{cases} \quad (3)$$

where $p^N(\cdot)$ is the function of a vector of observables z_{jt}^N that gives the height of the j 'th bar of the probability histogram of the notional distribution in year t , $p^R(\cdot)$ the function of observables z_{jt}^R that gives the height of the corresponding bar of the probability histogram of the rigidity-contaminated distribution in the same year, and b_j^N and b_j^R the corresponding vectors of parameters. Typically both z_{jt}^N and z_{jt}^R will contain dummy variables that indicate the relative position of bin j in the probability histogram,¹⁴ and additional variables that capture characteristics of the year t , while z_{jt}^R will additionally contain variables that indicate the position of bin j relative to the position of the bins containing the values taken by the rigidity bounds in the population.¹⁵

Given that the shape of the actual-wage-growth distributions is reflected by the height of the bars of the corresponding probability histograms, then, in principle, we could formulate hypotheses about it in terms of the values of the parameters of the functions that describe these heights. Suppose that there is a set of restrictions on the vector of parameters b_j^R , namely $H(b_j^R) = 0$, such that the two functions $p^N(\cdot)$ and $p^R(\cdot)$ coincide.¹⁶ Then the proposed

¹⁴Therefore, these variables will be functions of j .

¹⁵Therefore, these variables will be functions of both j and the corresponding indices of the bins that contain the point zero, i.e. the rigidity bound for DNWR, and the anticipated inflation values, i.e. the rigidity bounds for DRWR.

¹⁶It is natural to think of $G^R(F_t^N(\cdot))$ as the unrestricted case of $F_t(\cdot)$ since $G^R(F_t^N(\cdot)) = F_t^N(\cdot)$ in the special case that $G^R(\cdot)$ is the ‘identity’ functional. Consequently, we think of $p^R(z_{jt}^R; b_j^R)$ as the unrestricted model of P_{jt} .

strategy to test for the presence of rigidity of type R is, firstly, to estimate the parameter vector b_j^R , and then to test the hypotheses

$$\begin{aligned} H_0 : H(b_j^R) &= 0 \\ H_1 : H(b_j^R) &\neq 0 \end{aligned} \tag{4}$$

Stage 2: Estimation of the probability histogram parameters and hypothesis testing

In this stage we estimate b_j^R and test the hypotheses stated in (4). The estimation is done in two steps, and exploits the fact that we have multiple samples on actual-wage-growth.

In Step 1, using the actual-wage-growth data from each year in the observation period, we produce estimates of the heights of the bars of the corresponding probability histograms. Let \hat{P}_{jt} be the estimator of P_{jt} and \hat{p}_{jt} the corresponding estimate. Then this exercise produces a collection of estimates $\{\hat{p}_{jt}\}_{\substack{t=1,\dots,T \\ j=-J,\dots,J}}$, which includes an estimate for the height of each bar in the histogram for each year in our sample.

In Step 2, for each j , using the set of T estimates of the height of bar j from all years, i.e. $\{\hat{p}_{jt}\}_{t=1,\dots,T}$, as the set of ‘observations’ on \hat{P}_{jt} ,¹⁷ we estimate the regression of \hat{P}_{jt} on the vector of observables z_{jt}^R . When the estimator \hat{P}_{jt} is unbiased, the regression function will coincide with $p^R(z_{jt}^R; b_j^R)$, and the regression equation will look like this

$$\hat{P}_{jt} = E(\hat{P}_{jt} | z_{jt}^R) + \varepsilon_{jt} = p^R(z_{jt}^R; b_j^R) + \varepsilon_{jt} \tag{5}$$

Therefore, the estimation of this equation would give estimates of the parameter vector b_j^R and its variance-covariance matrix, enabling us to test the restrictions stated in (4). In practice the regression equations corresponding to all bar heights are estimated jointly since this is typically more efficient.¹⁸

¹⁷Now $t = 1, \dots, T$ becomes the observation index.

¹⁸In such a case, the system would consist of $2J + 1$ equations. The dependent variable corresponding to the

For the identification of the parameters of the model it is required that each type of rigidity distorts different parts of the wage-growth distribution at least for some of the years in the sample. In this way, there will be sufficient variation in the dummy variables that indicate the bins that are affected by the distortions, so that these will not be collinear with the dummy variables that indicate the position of the bins in the notional probability histogram.

Our chosen parameterisation for the heights of the bars of the probability histograms under the null hypothesis (i.e. for the notional¹⁹ distribution), is the following

$$\begin{aligned} p^N(z_{jt}^N; b_j^N) &= \beta_{1|j|} + \beta_{2|j|} \times up_{jt} + (\beta_{3|j|} + \beta_{4|j|} \times up_{jt}) \times m_t, & j \neq 0 \\ &= \beta_{10} + \beta_{30} \times m_t, & j = 0 \end{aligned} \quad (6)$$

where m_t denotes the median of the actual-wage-growth data in year t , up_{jt} is a dummy variable that is equal to 1 if bin \mathcal{B}_{jt} lies to the right of the bin containing the median ($j > 0$), and the β 's are coefficients to be estimated. With this parameterisation the $2J + 1$ probability bars in each histogram can have different height from each other, therefore the notional distribution is not restricted to have any particular shape, and, in particular, to be symmetric. Furthermore, by making the bar height to be a linear function of the location of the actual-wage-growth distribution, and therefore of the location of the notional distribution itself, we allow for the shape of the notional distribution to vary with its location. For example, suppose that the notional distribution is symmetric around the bin containing m_t and, further, that its spread increases as its centre moves to higher values.²⁰ Then $\beta_{2|j|}$ and $\beta_{4|j|}$ will be equal to zero due to

equation for a particular observation would be \hat{P}_{jt} , where j is the equation index, and t the within equation observation index. To estimate the system we would have in total $(2J + 1) \times T$ observations, with T observations on each equation. More details of this case are discussed below.

¹⁹Given the parameterisation of the probability histograms under the alternative, which we discuss later, here we take the notional distribution to be the the nominal-wage-growth distribution free of any kind of distortions, either due to DWR or menu costs.

²⁰This would imply a positive relationship between the spread and location of the histograms of the actual-wage-growth data irrespective of whether any type of rigidity is present or not.

the symmetry assumption, $\beta_{1|j|}$ will be non-negative, and $\beta_{3|j|}$ will be negative for the bins in the middle of the distribution, i.e. for small $|j|$, and positive for the bins that lie to the tails of the distribution, i.e. for large $|j|$. Alternatively, if we allow $\beta_{4|j|}$ to be non-zero for some values of j , then the skewness of the notional distribution will also vary with the location.²¹

In order to test for the presence of both types of rigidity, the parameterisation of the probability histogram under the alternative hypothesis should reflect the distortions due to the presence of both. We assume that

$$p^R(z_{jt}^R; b_j^R) = p^N(z_{jt}^N; b_j^N) + D^u(z_{jt}^u; \mu) + D^n(z_{jt}^n; \gamma) + D^r(z_{jt}^r; \delta) \quad , \text{ for } R = nr \quad (7)$$

where $D^n(z_{jt}^n; \gamma)$ is defined to be the difference between the height of the j 'th bar of the rigidity-contaminated probability histogram and the height of the corresponding bar of the notional probability histogram in year t that is due to the presence of DNWR, and $D^r(z_{jt}^r; \delta)$ the corresponding difference that is due to the presence of DRWR. We also allow for distortions due to the presence of menu costs, captured by the term $D^u(z_{jt}^u; \mu)$.

For the effect of DNWR we write

$$D^n(z_{jt}^n; \gamma) = (\gamma_1 + \gamma_2 \times m_t) \times d0_{jt} + (\gamma_3 + \gamma_4 \times m_t) \times dn_{jt} + \gamma_5 \times dz1_{jt} \quad (8)$$

where $d0_{jt}$ is a dummy variable that is equal to 1 if bin \mathcal{B}_{jt} contains the point zero, dn_{jt} a

²¹The assumption in the original Kahn (1997) methodology that the shape of the notional distribution is the same across years, has often been cited as one of the main drawbacks of this methodology as in most actual-wage-growth data sets there appears to exist a variation in the spread of the distribution across years characterised by different levels of inflation. This point is raised by Nickell and Quintini (2003) who go on to propose a flexible way of studying DNWR. They exploit the fact that if DNWR is important, the distribution of real-wage change across individual agents should, other things equal, be influenced by inflation and study the proportion of real-wage changes which is below -x%, where x ranges from 2 to 9. This aspect of their analysis focuses on the overall shape of the real-wage-change distribution, rather than the data of the histogram heights used here. Both approaches are very flexible.

dummy variable that is equal to 1 if bin \mathcal{B}_{jt} is to the left of the bin containing the point zero, and $dz1_{jt}$ a dummy variable that is equal to 1 if bin \mathcal{B}_{jt} is the first bin to the right of the bin that contains the point zero. With the inclusion of the first term we can capture the distortion that applies to the bin that contains the point of zero nominal-wage-growth, and, with the second term, the distortion that applies to each one of the bins that contain negative values of wage growth.²² In particular, γ_1 accounts for the distortion associated with the bin that contains the point of zero nominal-wage-growth and γ_3 the distortion associated with the bins that lie to the left of this bin in the special case where the centre of the notional distribution, which we proxy by m_t , is located at the point zero (i.e. $m_t = 0$). In that case, and, in the presence of DNWR, we would expect γ_1 to be positive, signifying the concentration of probability mass surplus in the zero nominal-wage-growth bin, and γ_3 negative, signifying the loss of probability mass from the bins that contain negative values of notional-wage growth. When the centre of the notional distribution is located further to the right ($m_t > 0$), a smaller part of the left tail of the notional distribution lies below zero, i.e. the proportion of notional-wage cuts falls, and, therefore the proportion of notional-wage changes that become wage freezes due to DNWR is expected to fall. In that case, γ_2 must be negative, signifying the reduction in the probability mass surplus in the zero-nominal-wage-growth bin, while γ_4 could be either positive or negative or zero, as the amount of mass deficit from each bin containing negative values could change in any direction relative to its level at $m_t = 0$. The inclusion of the last term enables us to test the hypothesis that, apart from shifting mass to the point of zero nominal-wage-growth, the presence of DNWR could also induce a shift of mass beyond the point zero, towards small

²²Our approach subtracts a constant amount from the height of each bin below zero. This, as a percentage of mass, becomes larger the further to the left we go implying that large notional cuts are less likely than would be in the case under a proportional specification. Another advantage of this specification is that it is linear in parameters, and therefore easier to estimate, compared to a ‘proportional’ type of rigidity.

positive values (in that case, $\gamma_5 > 0$).²³

The distortion in the height of the probability bar of bin \mathcal{B}_{jt} due to DRWR is assumed to be given by

$$D^r(z_{jt}^r; \delta) = \begin{cases} \delta_{1k} + \delta_{2k} \times J_t^P, & k = j - J_t^P, \text{ if } k_{\min} \leq k \leq k_{\max} \\ 0 & \text{, otherwise} \end{cases} \quad (9)$$

where J_t^P is the value of the index of the bin in year t that contains the centre of the anticipated-inflation distribution in that year²⁴ and, thus, k the distance between bin \mathcal{B}_{jt} and that bin.²⁵

It is more convenient to write (9) more compactly as follows

$$D^r(z_{jt}^r; \delta) = \sum_{\nu=k_{\min}}^{k_{\max}} (\delta_{1\nu} + \delta_{2\nu} \times J_t^P) \times dp_{\nu,jt} \quad (10)$$

where $dp_{\nu,jt}$ are dummy variables indicating whether bin \mathcal{B}_{jt} is located k positions from the bin that contains the centre of the anticipated-inflation distribution in year t ,

$$dp_{\nu,jt} = \begin{cases} 1 & \text{if } \nu = k (= j - J_t^P) \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

With this specification we allow for the size of the distortions to differ according to the location of the bin in the support of the anticipated-inflation distribution (through the indexing by k), and its location in the support of the notional-wage-growth distribution (through the dependence on J_t^P). In the presence of DRWR, the δ_{1k} 's, which account for the distortion when the centre

²³Holden (1989,2004) and Cramton and Tracy (1992) describe mechanisms of DNWR where unions can not only resist nominal-wage cuts but also induce small positive nominal-wage changes by threatening to work less efficiently during bargaining (holdout). Cramton and Tracy (1992) find empirical support for this model in US wage contract data. Holden (1989,1998) do the same for wage setting in the Nordic countries.

²⁴In the empirical application we have proxied this value either with the realised inflation in year t , measured by the CPI in year t , or a GARCH estimate of anticipated inflation in year t . See Table 1 for their values.

²⁵The index k is assumed to take values from the set $\{k_{\min}, \dots, 0, \dots, k_{\max}\}$. The bin for which $k = 0$ contains the centre of the anticipated-inflation distribution, bins with positive values of k are located to the right of this bin, and bins with negative values to its left. The values taken by k_{\min} and k_{\max} are determined empirically.

of the anticipated-inflation distribution is located in the same bin as the median of the actual-wage-growth distribution ($J_t^P = 0$), are expected to be positive for the largest (and positive) values of k and negative for the smallest (and negative) values of k , signifying the shift of probability mass towards the right end of the support of the anticipated-inflation distribution - see discussion on p. 8. When J_t^P takes different values, the values of the δ_{2k} 's must be such that the distortions ($\delta_{1k} + \delta_{2k} \times J_t^P$) are qualitatively similar to the case where $J_t^P = 0$, however no specific statements can be made about their sign or size unless specific assumptions are made about the nature of the joint distribution of the notional-wage growth and anticipated inflation, and the rigidity mechanism.

Finally the effect of menu costs is parameterised as follows

$$D^u(z_{jt}^u; \mu) = \mu \times dnp1_{jt}$$

where $dnp1_{jt}$ is a dummy variable that is equal to 1 if bin \mathcal{B}_{jt} is either one position to the left or to the right of the bin that contains the point zero. Therefore we allow for a symmetric loss of mass ($\mu < 0$) around and close to zero.²⁶

5 Estimation

To produce the estimates of the heights of the bars of the probability histograms (\hat{p}_{jt}), we use the proportion of observations in the sample for year t that fall in bin j as the estimator of

²⁶Christofides and Leung (2003) found only weak evidence for the presence of menu costs for this data. However, accounting for their presence in our specification ensures that the magnitude of parameter γ_5 , which measures the distortion in the first bin to the right of the bin that contains the point zero, attributed to DNWR, is estimated correctly. Given the model setup, we are not able to test formally for the concentration of mass around the point zero within the bin that contains it. However data inspection shows that almost all mass accumulated in that bin is attributed to point zero.

P_{jt} .²⁷ This estimator, denoted by \hat{P}_{jt} , could be defined as

$$\hat{P}_{jt} \equiv \sum_{i=1}^{n_t} \frac{d_{jti}}{n_t} \quad (12)$$

where d_{jti} is a dummy variable that takes the value of 1 if \dot{w}_{ti} falls in bin \mathcal{B}_{jt} and 0 otherwise, and n_t is the number of observations in year t . Since $\Pr(d_{jti} = 1) = \Pr(\dot{w}_{ti} \in \mathcal{B}_{jt}) = P_{jt}$, then d_{jti} is a Bernoulli random variable with mean P_{jt} . Furthermore, assuming that $\dot{w}_{ti} \stackrel{iid}{\sim} F_t(\dot{w})$ within t , \hat{P}_{jt} becomes the sample mean of i.i.d. $Bernoulli(P_{jt})$ random variables and is thus an unbiased estimator of the true height P_{jt} , as well as consistent and asymptotically normal. We can further derive the exact algebraic expression for the covariance between any pair of estimators that correspond to bins from the same or different probability histograms. Treating the wage growth associated with different bargaining pairs as being independent also across years, this expression takes the form

$$Cov\left(\hat{P}_{jt}, \hat{P}_{\zeta\tau}\right) = \begin{cases} \frac{P_{jt}(1-P_{jt})}{n_t} & , t = \tau \text{ and } j = \zeta \\ -\frac{P_{jt}P_{\zeta t}}{n_t} & , t = \tau \text{ and } j \neq \zeta \\ \sum_{i \in \mathcal{I}_t \cap \mathcal{I}_\tau} \frac{\Pr(d_{jti}=d_{\zeta\tau i}=1) - P_{jt}P_{\zeta\tau}}{n_t n_\tau} & , t \neq \tau \end{cases} \quad (13)$$

where \mathcal{I}_t and \mathcal{I}_τ are the sets of indices denoting the bargaining pairs which appear in our sample to have a contract agreement in years t and τ respectively,²⁸ while $j, \zeta \in \{-J, \dots, J\}$

²⁷ We use the median of the actual-wage-growth data from year t , denoted by \hat{m}_t , as an estimate of m_t . Therefore the bin of the estimated probability histogram indexed by $j = 0$ is the one that contains \hat{m}_t . The bin width is set to be equal to 1% and the bin endpoints take values from the set $\{\dots, -1.5, -0.5, 0.5, 1.5, \dots\}$. Thus, the point zero is at the centre of the bin that contains it and small wage changes around zero also fall in the same bin. Furthermore, J is chosen to be equal to 8 so that each probability histogram consists of 17 bins that cover in total an interval of 17 percentage points. In this way we achieve coverage of more than 97% of the data points in each yearly sample.

²⁸With this notation we allow for the presence of agreements for the same bargaining unit in several of the available yearly samples.

and $t, \tau \in \{1, \dots, T\}$.²⁹

In Step 2 of the estimation stage we treat the $2J + 1$ equations as a system. After imposing the cross-equation parameter restrictions implied by the parameterisation of (7),³⁰ the equation for a typical observation for the stacked data can be written as follows

$$\hat{P}_{jt} = \sum_{q=-J}^J p^N(z_{qt}^N; b_q^N) \times d_{qt}^* + D^u(z_{jt}^u; \mu) + D^n(z_{jt}^n; \gamma) + D^r(z_{jt}^r; \delta) + \varepsilon_{jt} \quad (14)$$

where d_{qt}^* is a dummy variable that is equal to 1 if $q = j$, and 0 otherwise. In matrix form the system can be written as

$$\hat{\mathbf{P}} = \mathbf{Z}\mathbf{b} + \varepsilon \quad (15)$$

where $\hat{\mathbf{P}} \equiv \left[\hat{\mathbf{P}}_{-J} \quad \hat{\mathbf{P}}_{-J+1} \quad \dots \quad \hat{\mathbf{P}}_0 \quad \dots \quad \hat{\mathbf{P}}_{J-1} \quad \hat{\mathbf{P}}_J \right]'$ is the vector of dependent variables for the entire system, and $\hat{\mathbf{P}}_j \equiv \left[\hat{P}_{j1} \quad \hat{P}_{j2} \quad \dots \quad \hat{P}_{jT} \right]$ the vector of dependent variables that corresponds to equation j .

The choice of optimal estimation method for the parameters of the system depends on the nature of the variance-covariance matrix of the vector $\hat{\mathbf{P}}$ of estimators, denoted by $Var(\hat{\mathbf{P}})$, whose typical element is $Cov(\hat{P}_{jt}, \hat{P}_{\zeta\tau})$. Clearly from (13) we see that $Var(\hat{\mathbf{P}})$ is not spherical, since the diagonal elements are not identical (first line of the result), and since there are also some non-zero off-diagonal elements (second and third lines). Therefore the Ordinary Least Squares (OLS) procedure, despite producing consistent estimates of \mathbf{b} , gives wrong standard

²⁹Putting some intuition in this result, first we note that the estimators are sample means of non-identically distributed Bernoulli variables for samples of different sizes and, therefore, should be expected to have different variances, as shown in the first line. Also, since the heights of the bars for each histogram must sum up to one, then the estimators of these heights should be expected to be negatively correlated, as they appear to be in the second line. Finally, to the extent that the wage settlements reached by the same bargaining unit at different points in time are correlated with each other, then the probability estimators in different years could also be correlated, as suggested by the result in the third line.

³⁰Specifically, the parameter vectors that capture the effect of the rigidities, i.e. γ and δ , and the parameter that captures the effect of menu costs, i.e. the μ , are common to all equations.

error estimates. Therefore we opt for the Feasible Generalised Least Squares (FGLS) procedure, substituting the probabilities in the right-hand side of (13) with consistent estimates; in the case of the probabilities of the form P_{jt} with the estimates obtained in stage 2 (i.e. \hat{p}_{jt}), and for $\Pr(d_{jti} = d_{\zeta\tau i} = 1)$ with estimates produced in a similar way

$$\Pr(d_{jti} = \widehat{d_{\zeta\tau i}} = 1) = \sum_{i \in \mathcal{I}_t \cap \mathcal{I}_\tau} \frac{d_{jti} d_{\zeta\tau i}}{\#(\mathcal{I}_t \cap \mathcal{I}_\tau)} \xrightarrow{p} \Pr(d_{jti} = d_{\zeta\tau i} = 1) \quad (16)$$

where $\#(\mathcal{I}_t \cap \mathcal{I}_\tau)$ is the number of elements in the set $\mathcal{I}_t \cap \mathcal{I}_\tau$.³¹

6 Results

In Table 3 we present the estimation results when we apply the FGLS estimator (columns 2 and 3), and the OLS estimator with corrected (columns 4 and 5) and uncorrected standard errors (columns 6 and 7). To obtain these results we have used the GARCH approach, rather than the actual CPI growth, to estimate the mean anticipated inflation rate.³² The table is divided in three panels; the top panel includes the estimates associated with the notional distribution (β 's) and the distortion due to menu costs (μ), the middle panel those associated with the distortion due to the presence of DNWR (γ 's), and the bottom panel those associated with the distortion due to the presence of DRWR (δ 's). Furthermore, in Table 4, we present the results from testing joint hypotheses about the parameters of the model using the Wald and F statistics, which are based on the results from the FGLS estimation. Next we discuss the results from the FGLS estimation.³³

³¹In order to obtain the results described in the next section, we have assumed - a priori - that all these quantities are equal to zero. Given that our expectation is that the corresponding estimates would have been relatively small and close to zero, we do not think that this assumption has affected considerably these results.

³²This is more consistent with the conceptual approach of Section 2 and produces clearer results.

³³The OLS results with corrected standard errors are, at least qualitatively, similar; they show a shift of probability mass to the right towards the values of anticipated inflation, a spike at zero, and mass deficit below

Notional distribution & menu costs Overall, the majority of the estimates of the β 's in the top panel of Table 3 are statistically significant, suggesting that the shape of the notional distribution is not fixed but varies with its location.

In particular, the estimates of $\{\beta_{10}, \dots, \beta_{18}\}$ and $\{\beta_{21}, \dots, \beta_{28}\}$ suggest that, if the median of the actual-wage-growth distribution were equal to zero, then the corresponding notional distribution would be bell shaped, with the mode of the distribution being located close to the median of the actual distribution. Also, the negative sign of $\{\beta_{21}, \dots, \beta_{28}\}$ suggests that the distribution has less mass above the bin that contains the actual-wage-growth median.³⁴ As the median of the actual distribution increases, and therefore the location of the notional distribution also moves to the right, the estimates of $\{\beta_{30}, \dots, \beta_{38}\}$ and $\{\beta_{41}, \dots, \beta_{48}\}$ together suggest that there is a flattening of the shape of the distribution in the centre and that there is progressively more concentration of probability mass at the right tail of the notional distribution and less mass at the left tail. The overall effect is that, as the location of the notional distribution moves further to the right, the distribution becomes progressively skewed to the right and its spread increases.

The estimate of μ is statistically significant but suggests that the number of small positive and negative wage adjustments that do not take place, presumably due to menu costs, is relatively small. Specifically, the number of adjustments in the intervals $[-1.5, -0.5]$ and $[0.5, 1.5]$ that do not take place corresponds to around 1.5% of the total number of contracts in each year, confirming an earlier result by Christofides and Leung (2003) for modest menu costs.

zero. For the OLS results without correction of the standard errors we note that all of the parameters that measure the effect of DRWR, with the exception of one, are statistically insignificant.

³⁴If we believed that the median of the notional distribution were close enough to the median of the actual so that they were both located in the same bin, then we could interpret this result as one that suggests that the notional distribution is non-symmetric.

DNWR The estimates of the parameters measuring the effect due to the presence of DNWR in the middle panel are also significant and have signs that are consistent with the presence of this type of rigidity; That is, $\hat{\gamma}_1$ is positive, $\hat{\gamma}_2$ negative, and $\hat{\gamma}_3$ negative.

If the median of the actual-wage-change distribution were equal to zero, the bin containing the point of zero nominal-wage growth would attract an estimated excess probability mass of 10.43% ($= \hat{\gamma}_1$), while each bin that contained negative values of wage growth would show a probability adjustment of -1.74% ($= \hat{\gamma}_3$). For all other values of the actual median, the surplus concentrated at the bin of zero nominal-wage growth would change by -1.64% ($= \hat{\gamma}_2$) for each 1% increase in the actual median, while the deficit in the bins containing negative values would become smaller by 0.13% ($= \hat{\gamma}_4$). These results are similar to those reported by Christofides and Leung (2003) who apply a variant of the original Kahn (1997) approach to test for the presence of DNWR only. The estimate of γ_5 is positive, but small and statistically insignificant. Therefore we only find weak support for the hypothesis that agents can induce small positive nominal-wage changes, in the fashion described by Holden (1989,2004) and Cramton and Tracy (1992) .

DRWR The majority of estimates of the parameters in the bottom panel of Table 3 are statistically significant, suggesting that the shape of the wage-growth distribution is distorted in the region where we expect future inflation beliefs to lie. Furthermore, the pattern of the distortions suggests a shift of probability mass to the right towards these values, that is similar to what we would expect if DRWR were present.

In particular, when the centre of the anticipated-inflation distribution is located in the same bin as the median of the actual distribution, then this bin ($j = 0$ & $k = 0$), as well as the bin that lies immediately to its left ($k = -1$) and all the bins to its right ($k = 1, \dots, 5$) attract a surplus of probability mass that is statistically significant. At the same time, the remaining bins

that lie to the left of this group of bins ($k = -5, \dots, -2$) show a deficit of probability mass, that is also statistically significant in all cases except for the leftmost bin ($k = -5$). The bin indexed by $k = -1$ attracts the biggest surplus, equal to 8.71%, while the bin at the centre of both the actual and the anticipated-inflation distributions attracts the second biggest, equal to 7.78%. The surplus continues to diminish as we move further to the right, taking values between 5.44% to 0.76%. On the other hand, the deficit for the bins indexed by $k = -2, \dots, -4$, takes values between 0.25% to 1.61%.

This pattern remains qualitatively unchanged when the centre of the distribution of anticipated inflation is located at its sample values,³⁵ while quantitatively the variation is relatively very small. The signs of the estimates of $\{\delta_{-25}, \dots, \delta_{25}\}$ suggest that quantitatively the distortions become more pronounced at the centre of the anticipated-inflation distribution the further to the right the centre of this distribution is located relative to the centre of the actual-wage-growth distribution. In particular, the excess mass becomes progressively more concentrated in the two bins at its centre, indexed by $k = 0, -1$, while the deficit in the bins further to their left increases.

Joint tests The test results in Table 4 concerning joint hypotheses about subsets of the parameter set offer additional support to our main conclusions about the presence of the two types of rigidity and the shape of the notional distribution. In particular, we reject overwhelmingly the null hypotheses: (i) that the shape of the actual-wage-growth distribution coincides with the shape of the rigidity-free distribution (line 1), (ii) the absence of DRWR (line 2), and (iii) the absence of DNWR (line 3). We also reject the hypotheses: (iv) that the shape of the notional distribution remains the same as the centre of the distribution changes location (line 4), and (v) that the notional distribution is symmetric around the bin containing the median of the actual

³⁵These appear in the last column of Table 1

distribution (line 5).

Because of the large number of estimated parameters and the presence of interaction terms, the nature of the estimated distortions due to the presence of both types of rigidity may not be immediately clear. Therefore, in Figure 2, we draw the fitted probability histograms for the notional and actual-wage-growth distributions, for selected years of high, medium, and low inflation. The histograms in this figure are ‘standardised’, therefore the bin indexed by zero contains the median of the actual-wage-growth data from the relevant year. The light-shaded bars correspond to the notional distribution, i.e. $p^N(z_{jt}^N; \hat{b}_j^N)$, while the dark-shaded bars correspond to the actual distribution, i.e. $p^R(z_{jt}^R; \hat{b}_j^R)$. In the diagrams we also indicate (except for the top row when it is off the left point of the chart) which bins contain the point zero (“0”) and the estimated centre of the distribution of anticipated inflation (\widehat{P}_t^e). These diagrams clearly demonstrate a shift of probability mass to the right, towards the values of anticipated inflation in the population, that produces a similar type of distortions to those in the simulated diagrams that assume the presence of DRWR, in Figure 1. Furthermore, we can distinguish a spike at the bin containing the point zero, for example bin -5 in the graphs for years 1983 and 1989, bin -4 for 1984, and bin -2 for 1992, and deficit in the bins below zero, which is consistent with the presence of DNWR.³⁶

³⁶We have estimated several variations of the model in order to examine the sensitivity of our results to the particular specification of the model and found that these were robust, at least in a qualitative sense; in all cases we estimated a shift of mass to the right towards point zero and the interval we expect to contain the values of anticipated inflation, the presence of excess mass at zero, and a mass deficit in the part of the distribution that lies below zero.

7 Conclusion

In this paper, we study collective bargaining wage outcomes drawn from the Canadian unionised sector, over a long period of diverse inflation experience. Earlier studies involving this data found evidence for the presence of DNWR. The challenge was to specify mechanisms consistent with the notion of DRWR and to superimpose these mechanisms on the broad approach used to measure DNWR in the past. The results obtained suggest that DRWR is clearly present in the data and that it can be identified over and above substantial DNWR effects.

References

- [1] Altonji, J. G. and P. J. Devereux (2000). The extent and consequences of downward nominal wage rigidity. *Research in Labor Economics* 19, 383-431.
- [2] Barwell, R. and M. Schweitzer (2004). The incidence of nominal and real wage rigidity in Great Britain: 1978-1998. Mimeo, Bank of England.
- [3] Bauer, T.K., H. Bonin and U. Sunde (2003). Real and nominal wage rigidities and the rate of inflation: evidence from West German micro data. IZA DP 959.
- [4] Bewley, T. F. (1999). *Why Wages Do Not Fall During a Recession?* Harvard University Press.
- [5] Christofides, L.N. and M.T. Leung (2003). Nominal wage rigidity in contract data: a parametric approach. *Economica* 70, 619-638.
- [6] Christofides, L.N. and T. Stengos (2003). Wage rigidity in Canadian collective bargaining agreements. *Industrial and Labor Relations Review* 56(3), 429-448.
- [7] Cramton, P. and J. S. Tracy (1992). Strikes and holdouts in wage bargaining: theory and data. *American Economic Review* 82, 100-121.
- [8] Dickens, W. and E. Goshen (2004). The International Wage Flexibility Project (IWFP). (Proceedings of the Final Conference, European Central Bank, Frankfurt Am Main, Germany).
- [9] Friedman, M. (1968). The role of monetary policy. *American Economic Review* 58(1), 1-17.
- [10] Holden S. (1989). Wage drift and bargaining: evidence from Norway. *Economica* 56, 419-432.

- [11] Holden S. (1994). Wage bargaining and nominal rigidities. *European Economic Review*, 38, 1021-1039.
- [12] Holden S. (1998). Wage drift and the relevance of centralised wage setting. *Scandinavian Journal of Economics* 100, 711-731.
- [13] Holden S. (2004). The costs of price stability: Downward nominal wage rigidity in Europe. *Economica* 71, 183-208.
- [14] Kahn, S. (1997). Evidence of nominal wage stickiness from microdata. *American Economic Review* 87, 993-1008.
- [15] Macleod, W. B. and J. M. Malcomson (1993). Investment, holdup, and the form of market contracts. *American Economic Review*, 37, 343-354.
- [16] Malcomson, J.M. (1997). Contracts, hold-up, and labor markets. *Journal of Economic Literature*, 35 (4), 1916-1957.
- [17] Nickell S. J. and G. Quintini (2003). Nominal wage rigidity and the rate of inflation. *The Economic Journal* 113, 762-781.

Year	#	WNC+COLA			CPI	\widehat{P}^e
		<i>mean</i>	<i>median</i>	<i>s.d.</i>		
1977	226	8.69	8.20	2.99	7.55	7.22
1978	673	8.16	7.43	2.89	8.01	8.42
1979	569	10.64	10.11	3.18	8.95	8.45
1980	520	12.39	11.95	2.87	9.13	9.28
1981	450	13.64	13.10	3.20	10.16	11.66
1982	562	10.31	10.69	3.64	12.43	10.43
1983	643	4.89	5.00	2.63	10.80	6.05
1984	676	3.76	4.00	1.90	5.86	4.50
1985	519	3.78	4.04	2.14	4.30	3.81
1986	551	3.65	4.10	1.82	3.96	4.08
1987	557	3.90	3.83	1.77	4.18	4.37
1988	556	4.92	4.89	1.72	4.34	3.97
1989	493	5.68	5.22	1.82	4.05	4.83
1990	547	5.79	5.77	1.99	4.99	4.55
1991	530	3.89	4.19	2.19	4.76	5.91
1992	632	2.16	2.00	1.80	5.62	1.49
1993	516	0.75	0.00	1.35	1.49	2.00
1994	471	0.60	0.00	1.63	1.86	0.50
1995	460	0.86	0.68	1.16	0.16	2.24
1996	448	1.22	0.87	1.31	2.16	1.43
1997	346	1.87	1.87	1.35	1.62	1.95
Total	10945					

Table 1: Descriptive statistics.

Year	$\dot{w} < 0$		$\dot{w} = 0$		$0 < \dot{w} < CPI$		$\dot{w} = CPI$		$\dot{w} > CPI$	
	#	%	#	%	#	%	#	%	#	%
1977			2	0.9	86	38.1			138	61.1
1978					393	58.4			280	41.6
1979					198	34.8			371	65.2
1980					43	8.3			477	91.7
1981			1	0.2	38	8.4			411	91.3
1982	1	0.2	3	0.5	397	70.6			161	28.6
1983	4	0.6	26	4.0	597	92.8			16	2.5
1984	1	0.1	61	9.0	559	82.7			55	8.1
1985	1	0.2	26	5.0	286	55.1			206	39.7
1986	2	0.4	24	4.4	238	43.2	1	0.2	286	51.9
1987			17	3.1	307	55.1			233	41.8
1988			4	0.7	203	36.5			349	62.8
1989					60	12.2			433	87.8
1990			14	2.6	136	24.9			397	72.6
1991	2	0.4	57	10.8	243	45.8			228	43.0
1992	7	1.1	82	13.0	488	77.2			55	8.7
1993	18	3.5	263	51.0	116	22.5			119	23.1
1994	53	11.3	186	39.5	146	31.0			86	18.3
1995	9	2.0	162	35.2	2	0.4			287	62.4
1996	3	0.7	164	36.6	174	38.8			107	23.9
1997	1	0.3	50	14.5	91	26.3			204	59.0
Total	102	0.9	1142	10.4	4801	43.9	1	0.0	4899	44.8

Table 2: Wage-growth statistics.

Parameter	FGLS		OLS-corrected		OLS	
	Estimate	(Std. Err.)	Estimate	(Std. Err.)	Estimate	(Std. Err.)
β_{10}	0.3112**	(0.0098)	0.2998**	(0.0128)	0.2998**	(0.0293)
β_{11}	0.0881**	(0.0057)	0.2047**	(0.0102)	0.2047**	(0.0282)
β_{12}	0.0418**	(0.0038)	0.1028**	(0.0077)	0.1028**	(0.0258)
β_{13}	0.0385**	(0.0032)	0.0875**	(0.0059)	0.0875**	(0.0233)
β_{14}	0.0250**	(0.0021)	0.0564**	(0.0046)	0.0564**	(0.0204)
β_{15}	0.0194**	(0.0021)	0.0520**	(0.0041)	0.0520**	(0.0195)
β_{16}	0.0155**	(0.0053)	0.0492**	(0.0040)	0.0492*	(0.0200)
β_{17}	0.0195**	(0.0028)	0.0492**	(0.0039)	0.0492*	(0.0204)
β_{18}	0.0183**	(0.0030)	0.0494**	(0.0039)	0.0494*	(0.0205)
β_{21}	0.1105**	(0.0101)	0.0012	(0.0124)	0.0012	(0.0228)
β_{22}	-0.0086	(0.0078)	-0.0273**	(0.0103)	-0.0273	(0.0278)
β_{23}	-0.0448**	(0.0058)	-0.0435**	(0.0084)	-0.0435	(0.0295)
β_{24}	-0.0369**	(0.0038)	-0.0382**	(0.0066)	-0.0382	(0.0277)
β_{25}	-0.0267**	(0.0031)	-0.0492**	(0.0050)	-0.0492 [†]	(0.0252)
β_{26}	-0.0187**	(0.0056)	-0.0541**	(0.0042)	-0.0541*	(0.0236)
β_{27}	-0.0174**	(0.0030)	-0.0513**	(0.0041)	-0.0513*	(0.0236)
β_{28}	-0.0185**	(0.0036)	-0.0504**	(0.0040)	-0.0504*	(0.0235)
β_{30}	-0.0204**	(0.0011)	-0.0140**	(0.0012)	-0.0140**	(0.0022)
β_{31}	0.0027**	(0.0008)	-0.0060**	(0.0011)	-0.0060**	(0.0023)
β_{32}	0.0064**	(0.0006)	0.0024**	(0.0009)	0.0024	(0.0024)
β_{33}	0.0018**	(0.0005)	-0.0011 [†]	(0.0007)	-0.0011	(0.0025)
β_{34}	0.0002	(0.0004)	-0.0012*	(0.0006)	-0.0012	(0.0026)
β_{35}	-0.0018**	(0.0003)	-0.0016**	(0.0005)	-0.0016	(0.0025)
β_{36}	-0.0009	(0.0006)	-0.0029**	(0.0004)	-0.0029	(0.0025)
β_{37}	-0.0013**	(0.0003)	-0.0033**	(0.0003)	-0.0033	(0.0023)
β_{38}	-0.0012**	(0.0004)	-0.0034**	(0.0003)	-0.0034	(0.0023)
β_{41}	-0.0113**	(0.0014)	0.0003	(0.0017)	0.0003	(0.0030)
β_{42}	-0.0012	(0.0011)	0.0019	(0.0014)	0.0019	(0.0033)
β_{43}	0.0024**	(0.0008)	0.0026**	(0.0010)	0.0026	(0.0033)
β_{44}	0.0037**	(0.0006)	0.0036**	(0.0008)	0.0036	(0.0033)
β_{45}	0.0044**	(0.0005)	0.0038**	(0.0007)	0.0038	(0.0033)
β_{46}	0.0027**	(0.0007)	0.0055**	(0.0005)	0.0055 [†]	(0.0031)
β_{47}	0.0017**	(0.0003)	0.0049**	(0.0004)	0.0049 [†]	(0.0029)
β_{48}	0.0020**	(0.0005)	0.0042**	(0.0004)	0.0042	(0.0029)
μ	-0.0077**	(0.0022)	-0.0196**	(0.0016)	-0.0196 [†]	(0.0115)
γ_1	0.1043**	(0.0053)	0.1735**	(0.0104)	0.1735**	(0.0176)
γ_2	-0.0164**	(0.0010)	-0.0309**	(0.0019)	-0.0309**	(0.0037)
γ_3	-0.0174**	(0.0016)	-0.0530**	(0.0038)	-0.0530**	(0.0154)
γ_4	0.0013**	(0.0003)	0.0057**	(0.0003)	0.0057**	(0.0023)
γ_5	0.0027	(0.0030)	0.0041	(0.0049)	0.0041	(0.0161)
δ_{-15}	-0.0025	(0.0016)	-0.0084**	(0.0022)	-0.0084	(0.0111)
δ_{-14}	-0.0099**	(0.0020)	-0.0236**	(0.0040)	-0.0236	(0.0151)
δ_{-13}	-0.0161**	(0.0030)	-0.0306**	(0.0059)	-0.0306	(0.0199)
δ_{-12}	-0.0118**	(0.0045)	-0.0323**	(0.0075)	-0.0323	(0.0235)
δ_{-11}	0.0871**	(0.0067)	0.0120	(0.0094)	0.0120	(0.0259)
δ_{10}	0.0778**	(0.0075)	0.0260**	(0.0100)	0.0260	(0.0265)
δ_{11}	0.0544**	(0.0067)	0.0111	(0.0092)	0.0111	(0.0254)
δ_{12}	0.0416**	(0.0058)	-0.0014	(0.0079)	-0.0014	(0.0228)
δ_{13}	0.0199**	(0.0044)	-0.0124*	(0.0062)	-0.0124	(0.0189)
δ_{14}	0.0114**	(0.0029)	-0.0094*	(0.0042)	-0.0094	(0.0142)
δ_{15}	0.0076**	(0.0016)	-0.0007	(0.0025)	-0.0007	(0.0106)
δ_{-25}	-0.0007	(0.0011)	-0.0039*	(0.0016)	-0.0039	(0.0079)
δ_{-24}	-0.0035**	(0.0012)	-0.0051*	(0.0021)	-0.0051	(0.0085)
δ_{-23}	-0.0094**	(0.0015)	-0.0094**	(0.0025)	-0.0094	(0.0087)
δ_{-22}	-0.0043*	(0.0018)	-0.0013	(0.0028)	-0.0013	(0.0085)
δ_{-21}	0.0302**	(0.0027)	0.0149**	(0.0036)	0.0149 [†]	(0.0082)
δ_{20}	0.0099**	(0.0032)	-0.0004	(0.0036)	-0.0004	(0.0080)
δ_{21}	-0.0016	(0.0028)	-0.0009	(0.0034)	-0.0009	(0.0081)
δ_{22}	-0.0062*	(0.0025)	-0.0009	(0.0033)	-0.0009	(0.0084)
δ_{23}	-0.0056**	(0.0021)	0.0013	(0.0031)	0.0013	(0.0086)
δ_{24}	-0.0037*	(0.0016)	0.0091**	(0.0026)	0.0091	(0.0084)
δ_{25}	-0.0018	(0.0013)	0.0018	(0.0022)	0.0018	(0.0077)
N	357		357		357	

Significance levels : † : 10% * : 5% ** : 1%

Table 3: Estimation results.

#	W	$Pr(\chi_q^2 > W)$	F	$Pr(F_{(q,n-k)} > F)$
1 ^a	1484.7550	0	54.99091	0
2 ^b	649.3472	0	29.51578	0
3 ^c	587.1228	0	117.42460	0
4 ^d	781.0962	0	45.94683	0
5 ^e	439.8713	0	27.49195	0

^a $H_0 : \gamma = 0 \cap \delta = 0, H_1 : \gamma \neq 0 \cup \delta \neq 0$ ($q = 27, n - k = 296$)

^b $H_0 : \delta = 0, H_1 : \delta \neq 0$ ($q = 22, n - k = 296$)

^c $H_0 : \gamma = 0, H_1 : \gamma \neq 0$ ($q = 5, n - k = 296$)

^d $H_0 : \beta_{3|j|} = \beta_{4|j|} = 0, H_1 : \beta_{3|j|} \neq 0 \cup \beta_{4|j|} \neq 0, \forall j$ ($q = 17, n - k = 296$)

^e $H_0 : \beta_{2|j|} = \beta_{4|j|} = 0, H_1 : \beta_{2|j|} \neq 0 \cup \beta_{4|j|} \neq 0, \forall j$ ($q = 16, n - k = 296$)

Table 4: Joint test results.




Figure 1: Shapes of simulated rigidity-free (notional) and rigidity-contaminated nominal-wage-growth distributions.




Figure 2: Notional Vs Actual nominal-wage-growth distributions (fitted values).