Downward Nominal-Wage Rigidity: A Critical Assessment and Some New Evidence for Canada

Jean Farès and Thomas Lemieux

Introduction

Labour market observers have long suspected that, for a variety of reasons, employers are unwilling to reduce the nominal wages paid to their workers even when employers experience severe financial difficulties. (See Bewley 1999 for recent evidence.) Starting with Keynes' *General Theory,* this presumed downward nominal-wage rigidity (DNWR) has played a prominent role in many models of the labour market and the macroeconomy. One of Keynes' conjectures was that in a period of deflation, such as the Great Depression of the 1930s, DNWR resulted in higher real wages, which made the Depression longer and deeper.

There has been renewed interest in DNWR over the last decade for several reasons. From a research perspective, the availability of rich, longitudinal sets of microdata has enabled researchers to formally test for the existence of DNWR. From an economic-policy perspective, DNWR has become potentially more relevant to the conduct of economic policy as a number of countries have experienced very low inflation rates in the 1990s. One argument that is closely related to Keynes' conjecture is that when inflation is very low, DNWR may prevent real wages from falling by as much as they should when the economy experiences negative shocks. For instance, Fortin (1996) uses this argument to explain why the recession of the 1990s was much longer and deeper in Canada, where consumer price index (CPI) inflation averaged 1.4 per cent from 1992 to 1997, than in the United States, where CPI inflation averaged 2.9 per cent during the same period.

The objective of this paper is twofold. Our first goal is to critically review existing literature on the extent and consequences of DNWR. From this review, we conclude that recent studies, mostly based on U.S. longitudinal microdata, provide compelling evidence that DNWR is an important labour market phenomenon. The main finding from this literature is that there is a sharp concentration of nominal-wage changes at zero. The answer to the question of whether DNWR does in fact exist, is a decisive yes.

Much less clear from the literature, however, is whether DNWR has significant consequences for aggregate wage and employment (or unemployment) determination. The second goal of the paper, therefore, is to take a new look at the effect of DNWR on wage and employment determination in Canada during periods of low inflation.

One reason why little research has been conducted on this topic in Canada is that wage data here are limited relative to the United States. This lack explains why researchers, such as Fortin (1996) and Crawford and Harrison (1998) have used wage-settlement data from collective agreements to examine the extent and consequences of DNWR in Canada. Unfortunately, these data from large firms in the unionized sector may not be representative of the entire Canadian labour market.

To overcome these data shortcomings, we first develop a new wage series based on individual data files from Statistics Canada's Survey of Consumer Finance (SCF) for the period 1981–97. This new series has several important advantages over what was previously available. First, it is based on a representative survey that can also be used to compute separate wage series by province, industry, and so on. Second, it is possible to adjust wages for secular or business cycle changes in the composition of the workforce, since detailed information is available on human capital (e.g., age, education) and job characteristics (e.g., industry, occupation, seniority) in this survey. This is an important issue, since existing studies such as that of Solon, Barsky, and Parker (1994) suggest that changes in the composition of the workforce tend to understate the cyclicality of wages over the business cycle.

Controlling for changes in the composition of the workforce is particularly important in the context of the impact of DNWR, which is believed to apply only to workers who remain with the same employer. In a recession, aggregate wages may incorrectly look downward-rigid if workers who lose their jobs earn consistently less than those who keep them. This composition effect leads to an upward bias in aggregate wage changes, which could mask real-wage declines among workers who remain employed.

We then use this new wage series to analyze the relationship between real-wage changes and economic conditions. One key empirical implication of DNWR is that, in response to a given negative shock, the real wage should decline less in periods of lower than higher inflation, because DNWR is not likely to bind in the latter case. We test this implication by estimating “real-wage Phillips curves,” which link the unemployment rate to the change in real wages. If DNWR prevents real wages from adjusting (downwards) in periods of low inflation, the Phillips curve should be *flatter* in periods of lower inflation.

We use several empirical strategies to test whether the Phillips curve became flatter in the 1990s, when inflation dropped below 2 per cent. First, we analyze the aggregate time-series behaviour of real wages and find that it is partly consistent with this hypothesis. Until 1992 (when the inflation rate dropped “permanently” below 2 per cent), there was a negative and statistically significant relationship between the unemployment rate and changes in real wages. This relationship no longer holds since 1992, suggesting that real wages did not fall as much as they should have in the depths of the 1990s recession. One concern with these time-series results, however, is that other unmodelled factors, such as supply shocks or changes in the formation of expectations, may also have changed during this period. Furthermore, the relationship between real-wage changes and the unemployment rate is estimated imprecisely in the 1990s, because of small sample sizes.

Our second empirical strategy relies on variation in economic conditions across both time and Canadian provinces to identify potential changes in the relationship between unemployment rates and changes in real wages. Since different provinces are subject to different shocks at different times, it is possible, in principle, to identify the connection between (provincial) wage changes and (provincial) unemployment rates, while controlling for nation-wide factors using unrestricted year effects. Consistent with our expec-tations, we find that provinces that experience an increase in relative unemployment rates tend to experience a decline in relative wage growth. However, we do not find that this relationship has changed over time. In other words, these “provincial Phillips curves” did not become flatter in the years of very low inflation.

Finally, we use the richness of the SCF data to better understand the cyclical behaviour of real wages in Canada from 1981 to 1997. We find that during the recessions of 1981–83 and 1990–92, the real wages of older and more senior workers remained relatively constant. Most of the decline in real wages was concentrated among young workers and those having just started a new job. Irrespective of the inflation rate, new entrants seem to bear a disproportional share of the adjustments in real wages over the business cycle. This may explain why DNWR, which most likely binds for older and more senior workers, seems to have only a modest impact on aggregate wages and employment.

The paper is set out as follows. In section 1, we present a critical assessment of the existing literature and highlight the major knowledge gaps on the effect of DNWR on wages and employment. In section 2, we describe the SCF data and explain how we construct the wage series. In section 3, we estimate real-wage Phillips curves, using both aggregate data for Canada as a whole, and disaggregate provincial data. We also attempt to reconcile different pieces of evidence by analyzing the evolution of real wages by job seniority. We offer our conclusions in the final section.

1 Literature Review

This section will review some of the recent studies that document asymmetries in the wage-change distribution, based on micro-level data. While DNWR could clearly be a source of asymmetry in the wage-change distribution, other factors, such as menu costs, may also explain the observed asymmetries. We discuss the evidence related to the two hypotheses; we then argue that, from the monetary policy perspective, it is more interesting to examine the impact of DNWR on aggregate wages and, consequently, on employment. We briefly summarize current literature that examines this question.

1.1 Asymmetric wage-change distribution

The empirical literature using data at an individual level is expanding very quickly. We will restrict our attention to a few representative papers based on U.S. data, and more recent studies using U.K. household data and Canadian data. This literature typically considers the distribution of nominal-wage growth in an average year (in low-inflation years, for the most part) and highlights the following visual observations:

- There are relatively few wage cuts.
- There is a mass point in the wage-change distribution at zero.

1.1.1 How frequent are wage cuts?

McLaughlin (1994) documents that nominal-wage cuts were not rare in the United States between 1976 and 1986. Using survey data from the Panel Study of Income Dynamics (PSID), he finds that 17 per cent of workers with the same employers suffered nominal cuts. Subsequent studies using PSID confirmed these results. In particular, Card and Hyslop (1997) show that in a typical year in the 1980s, 15 to 20 per cent of non-job changers had measured nominal-wage declines, while Lebow et al. (1995) find a similar proportion of 18 per cent, on average, between 1971 and 1988.

Stylized facts from other data sources tend to show similar patterns. Using data from the British Household Panel Study (BHPS), Smith (2000) finds that, on average, 23 per cent of workers suffered nominal-wage cuts in their weekly pay over a one-year span in the 1992–96 period. In Canada, however, the evidence is less conclusive. The Labour Market Activity Survey (LMAS, 1988–90) and the Survey of Labour and Income Dynamics (SLID, 1993) results are similar to the PSID, with the SLID showing a surprisingly large number of wage cuts in 1993. On the other hand, the distribution of wage changes in the wage settlements from the unionized sector's collective bargaining agreements shows virtually no mass below zero wage change.

Akerlof et al. (1996) argue that the variation in the reported wages in the PSID is an artifact of measurement errors. Although no careful treatment of the measurement error has been conducted on the Canadian data, McLaughlin (1994) and Smith (2000)[1] found that about 5 percentage points of the fraction of wage cuts could be attributed to measurement error, decreasing the frequency of pay cuts to still significant levels of 12 per cent in the PSID and 18 per cent in the BHPS.

1.1.2 The spike at zero wage change

In all of these studies, the distribution of nominal-wage growth exhibits a large mass point at zero. In the PSID sample, Card and Hyslop (1997) report that the fraction of workers on the same job who experience a one-year wage change of zero is 8.3 per cent in the 1970s and 16 per cent in the 1980s. In the United Kingdom, Smith (2000) shows that this fraction is equal to 9 per cent between 1992 and 1996. Crawford and Harrison (1998) report that the fraction of wage freezes is 19.4 per cent in the unionized private sector in Canada between 1992 and 1996.

Some institutional factors, unrelated to any underlying rigidities, could, however, exaggerate the size of the mass point at zero. Long-term contracting or rounding could also explain part of the excess mass at zero wage change.

To control for the effect of long-term contracts, one can calculate the fraction of workers who received zero wage change over varying horizons. Card and Hyslop (1997) show that the mass point at zero in the two-year

1. It is very interesting to note that the BHPS gives interviewees a chance to check their pay slip when reporting their pay, thus substantially reducing the possibility of measurement errors.

wage-change distribution is reduced to 2.6 per cent in the 1970s and 8.1 per cent in the 1980s. Over three years, these fractions drop to 1.2 per cent and 4.7 per cent in the 1970s and 1980s.[2] In the United Kingdom, between 1992 and 1996, Smith (2000) shows that the mass at zero drops to 4 per cent for wage growth defined over two years, and to 2.5 per cent over three years. In Canada, Crawford and Harrison (1998) report a similar drop in the spike at zero when changing the wage-cut definition. The fraction of wage freezes in the unionized private sector between 1992 and 1996 drops to 12.9 per cent in the wage-change distribution over the life of the contract.[3]

After controlling for rounding problems and measurement errors, Lebow et al. (1995) calculate that almost 40 per cent of the spike at zero in the one-year wage-change distribution is due to rounding, while Smith argues that eliminating measurement error could cut the spike by half. This evidence, however, still indicates a substantial fraction of zero wage changes.

1.2 The source of asymmetries

Since the underlying "true" distribution of wage (or productivity) growth is unobservable, it is difficult to identify the source of distortions to the observed distribution. Two hypotheses, DNWR and menu costs, are usually considered. While both types of rigidities lead to a thinning in the left tail of the distribution and a piling up at zero wage change, menu costs also prevent small, positive wage changes from occurring.

If DNWR is only binding to the left of the median wage change in the wage-change distribution, and assuming symmetry around the median, then the difference between the two tails of the distribution is important in identifying the source of the rigidity. Alternatively, time variation may help disentangle the effects of DNWR from other sorts of institutional factors that might generate asymmetry in the observed wage-change distribution. For example, if the spike at zero is due to a downward constraint on wages, then, assuming that the shape of the underlying distribution does not vary over time, this constraint should be more binding in low-inflation years, and less binding in high-inflation years.

Card and Hyslop (1997) use the assumption of symmetry to construct counterfactual distribution of wage growth in the absence of rigidities. Their estimate of the fraction of people affected by DNWR, adjusted for the effect of menu costs, is around 10–12 per cent in the mid-1980s. Their estimates

2. Lebow et al. (1995) perform the same calculation and get slightly smaller numbers.
3. This fraction is higher in the public sector settlement, where wage freezes are between 56 per cent and 45 per cent, using different wage-change definitions.

also imply that DNWR may have increased by about 1 per cent the average wage growth for hourly-rated non-job changers, with a reduced effect in the later years of the sample. They conclude that DNWR exerts a small but measurable effect on average wage growth, with a greater effect in low-inflation years.

Lebow et al. (1995) use the difference between the cumulative frequency of the wage-change distribution above twice the median and the cumulative frequency of the distribution below zero as an alternative measure of asymmetry. They find that the frequency of wage changes below zero is nearly 4 percentage points lower than expected on the basis of their assumptions. The correlation between this measure of asymmetry and inflation constitutes a better test of the DNWR hypothesis. They find that this correlation is negative and significant only for job stayers paid by the hour.

This evidence could overstate the effect of DNWR if the underlying assumption of a symmetric distribution of wage changes was not satisfied. In fact, McLaughlin (1999) shows that the skewness of wage changes is not limited to the censoring of would-be wage cuts and small wage changes. There is even evidence of skewness close to the median. These results challenge the estimates of Lebow et al. and Card and Hyslop.

Intertemporal variation of the wage-change distribution provides another way to identify thinning of the distribution below zero. Under the assumption that the shape of the underlying distribution does not change over time, Kahn (1997) estimates that, in the PSID sample years of 1970–88, DNWR prevented 9.4 per cent of wage earners from receiving nominal-wage cuts.[4] However, if the sample in low wage-growth years has lower variance of wage changes, then the tails of the distribution would be thinner even if would-be wage cuts were not censored at zero.[5] To address this issue, McLaughlin (1999) uses a difference-in-difference estimator. His results still confirm those of Kahn, pointing to a thinning of tails below nominal zero of one-third to one-half of would-be cuts.

In summary, both DNWR and the menu-costs hypotheses are supported in the data analysis. DNWR clearly acts as a constraint on nominal-wage changes at the micro level. Section 1.3 discusses the evidence on how these two hypotheses are reflected in aggregate wages and employment.

---

4. In contrast, salary earners do not receive pay cuts less frequently than expected.

5. The changes in the shape of the wage-growth distribution are well-documented in Crawford (2000).

1.3 Aggregate effects of DNWR

Few papers address the macroeconomic implications of nominal-wage rigidity on aggregate wages and employment (or unemployment). Applying a hazard model to data for union wage settlements in Canada, Crawford (2000) estimates that the net effect of rigidity on average wage growth between 1992–97 is less than 0.2 per cent for the unionized private sector. These estimates are significantly lower than those reported in Simpson, Cameron, and Hum (1998) for the same data. Using a Tobit model for wage growth, Simpson et al. estimate that DNWR raised the average wage growth by 0.67 per cent between 1993 and 1995. On the other hand, Farès and Hogan (2000) conclude that, consistent with menu costs, nominal rigidities have a symmetric effect on wage changes above and below zero. Overall, they conclude that nominal rigidities result in lower than expected wage changes.[6]

Simpson et al. also provide some estimates on the effect of pay-cut resistance on employment growth and the unemployment rate. They use ordinary least squares (OLS) estimation of employment growth on pay freeze incidences and output growth, in different periods of high and low inflation. Their results indicate that, between 1993 and 1995, DNWR reduced mean employment growth across industries by more than half. However, the wage-freeze variable in this regression might be capturing some adverse shocks, particularly since the output growth estimated effect between 1993 and 1995 is significantly lower than in previous periods. Farès and Hogan and Faruqui (2000) show that, once adjusted for this endogeneity problem, the effect of wage freezes on employment growth becomes not statistically significant. Using a Tobit specification, Simpson et al. calculate that the unemployment cost of pay-cut resistance exceeds 2 per cent throughout the 1990s. One underlying assumption of these estimates is that the variance of the wage growth is time-invariant. As discussed, this assumption could exaggerate the effect of DNWR, given the noticeable compression in the wage-change distribution in the 1990s, a period of low inflation.

Card and Hyslop use average wage and unemployment data on a state level from 1976 to 1991 to estimate the effect of DNWR on unemployment. They use wage data constructed from the annual March Current Population Survey (CPS) that they adjust to reflect the varying composition of the workforce in each state in different years. They estimate the cross-state Phillips curve and find little evidence that the wage-adjustment rate across

6. Crawford (2000) discussed these results and suggested that a different treatment of inflation expectations could reconcile the results of Simpson et al. (1998) and Farès and Hogan (2000).

local markets is faster in a higher-inflation environment. Taken in combination with their micro-level findings, they argue that nominal rigidities have a small impact on the aggregate economy.

Overall, the micro-level evidence based on the distribution of individual wage changes reveals that, although nominal-wage cuts are not rare, there is a substantial spike at zero in the distribution of nominal-wage changes. Furthermore, there is evidence that the magnitude of the spike is correlated with inflation. It is much less clear from the literature, however, that DNWR has significant consequences for aggregate wage and employment (or unemployment) determination. We will attempt to fill some of these knowledge gaps by taking a new look at the effect of DNWR on wage and employment determination in periods of low inflation in Canada.

2 Wage Data

2.1 Survey of consumer finances

We assembled 16 annual microdata files from Statistics Canada's SCF to construct a consistent wage series over the years 1981 to 1997. The SCF provides large samples of around 40,000 workers for each of these years, with the exception of 1983, when the survey was not conducted.[7] For all available years, the SCF was conducted in April as a supplement to the Labour Force Survey (LFS), and asked a battery of questions about income in the previous year, in addition to the usual LFS questions that pertain to the reference week.[8]

The SCF contains information on annual income, as well as personal and labour-related characteristics of individuals aged 15 years and over. In particular, information is available on wages and salaries and income from self-employment in the previous year, labour force status, number of weeks worked in previous year, full-time/part-time status last year, number of hours in the reference week, occupation and industry, years of experience and seniority, and educational attainment.[9] Other demographic characteristics,

---

7. Public-use samples are also available for heads of households and spouses every other year during the 1970s. Data for all workers are only available starting in 1981. The survey was discontinued after 1997.
8. The reference week is the week immediately preceding the two-week period when the SCF is conducted.
9. One major concern using these data arises from changes in the way educational achievement is classified starting with the 1989 income file. Fortunately, the highest (university degree) and lowest (grade 8 or less) education categories appear to be quite comparable (in terms of sample proportions and average wages) under the two definitions. We use this feature later to ensure that our adjusted wage measures are comparable over time.

such as age, gender, marital status, language spoken, immigration status, and geographic location, are also available.

The wage measure we use is average weekly earnings, expressed in 1991 dollars.[10] For each individual in a given sample year, average weekly earnings are calculated as the ratio of annual wages and salaries, excluding income from self-employment and rental property, to the total weeks worked in that year. We only compute this wage measure for paid workers who report zero net income from self-employment to obtain a cleaner measure of wages for employed workers, since theories of DNWR are not relevant for self-employed workers. We also restrict the sample to workers aged 20 to 65.

Table 1 presents the distribution of workers across provinces, industries, and sectors. About 65 per cent of the (weighted) observations are concentrated in Quebec and Ontario, while more than half of the individuals work in the manufacturing, trade, and service industries. About 19 per cent of the sample is in the public sector.

The distribution of individual characteristics is presented in Table 2. In addition to standard demographic characteristics, the table provides information on full-time status and on the distribution of job tenure. Since job tenure is measured at the time of the survey in April, some workers (15.31 per cent of the sample in the "lost their job" category) report earnings in the previous year, despite the fact that they no longer work at the time of the survey. Approximately 15 per cent of workers have one year or less of tenure at the time of the survey, which indicates a fair amount of labour-market turnover.

Table 3 shows the provincial means of log-average weekly earnings for each year. Total average wages vary substantially across provinces (see last row in table), with a maximum gap of 26 per cent between Prince Edward Island and British Columbia. By contrast, real wages show relatively little variation over time. In fact, as shown in the last column, wages are very stable around their sample average, with the largest difference of 7 per cent (drop) between the first and the last years of the sample.

10. Earnings are defined as the sum of wages and salaries from all types of civilian employment. Included are gross cash wages and salaries received in the reference year from all jobs, before deductions for pension funds, hospital insurance, income taxes, CSBs, etc. Tips and net commissions are also included; taxable allowances and benefits provided by employers are not.

Table 1
Provincial and industrial distribution of the workforce, 1981–97

Sample composition (percentage)			
Province		**Industry**	
Newfoundland	1.91	Agriculture	1.31
Prince Edward Island	0.46	Other primary	2.39
Nova Scotia	3.11	Manufacturing (non-durables)	9.00
New Brunswick	2.61	Manufacturing (durables)	8.67
Quebec	25.37	Construction	5.74
Ontario	38.44	Transportation, communication	8.02
Manitoba	3.62	Wholesale trade	4.71
Saskatchewan	2.86	Retail trade	11.77
Alberta	9.44	Finance, insurance, real estate	5.95
British Columbia	12.19	Community services	19.16
		Personal services	7.94
Sector		Business and miscellaneous	7.88
Private	81.35	Public administration	7.47
Public	18.65		

Note: The estimated frequency distributions are all weighted.

Source: Statistics Canada, Survey of Consumer Finances. Cross-sectional files from 1981 to 1997. No data are available for 1983. Sample size is 623,875.

2.2 Adjusted vs. unadjusted wages

Two potential drawbacks arise when using average weekly earnings from the SCF as a measure of the wage rate over the business cycle. First, average weekly earnings may vary because of changes in the underlying (hourly) wage rate or because of changes in hours worked per week. Unfortunately, an hourly wage rate cannot be computed directly, since the SCF does not provide direct information on the number of hours worked per week in the previous year. Fortunately, several indirect measures of hours worked per year can be used to control for variation in hours. As mentioned earlier, the SCF collects information on hours worked during the reference week and on whether the worker worked full-time during the previous year.

We have also computed direct measures of actual hours worked per week by detailed category of worker, using the monthly microdata files from LFS, from 1981 to 1997. Matching these hours measures to workers in the SCF provides an additional proxy for weekly hours of work in the previous year. Our strategy, explained in detail below, is to use regression methods to "adjust" average weekly wages for changes in weekly hours of work, as proxied by these different measures.

The second potential drawback is that changes in the composition of the workforce may understate the cyclicality of real wages, since the skill

Table 2
Distribution of worker characteristics, 1981–97

Sample composition (percentage)			
Age group		**Job tenure**	
20–30	32.11	Less than 7 months	8.89
31–40	29.68	7 to 12 months	7.73
41–50	22.91	1 to 5 years	26.28
51–65	15.30	6 to 10 years	16.25
		11 to 20 years	16.92
Education		Over 20 years	8.62
No schooling or grade 8 or lower	6.64	Lost their job	15.31
Grade 9–10	9.27		
Grade 11–13 (did not graduate)	10.73	**Status**	
Grade 11–13 (graduate)	18.59	Full-time	83.81
Some post-secondary (no diploma)	10.59	Part-time	16.19
Post-secondary (diploma or certificate)	28.08		
University degree	16.11	**Gender**	
		Male	53.25
Mother tongue		Female	46.75
English	59.85		
French	21.19	**Marital status**	
Other	18.96	Single	24.47
		Married	67.61
		Other	7.92

Note: The estimated frequency distributions are all weighted.

Source: Statistics Canada, Survey of Consumer Finances. Cross-sectional files from 1981 to 1997. No data are available for 1983. Sample size is 623,875.

level of the workforce tends to decrease during expansions and increase during recessions, as younger and less educated workers are the first to lose their jobs in periods of economic downturn (Bils 1985; Solon, Barsky, and Parker 1994). As in the case of hours, we control for changes in the composition of the workforce by computing alternative “regression-adjusted” measures of the wage rate. More specifically, we use OLS to estimate the following wage equation:

$$w_{it} = \beta X_{it} + \sum_{t=1}^{16} \delta_t Year_t + \varepsilon_{it}, \tag{1}$$

where w_{it} is log real average weekly earnings of individual i in year t (earnings are deflated by total annual CPI); X_{it} includes various observable characteristics such as age, education, sex, marital status, language spoken, tenure, industry dummies, province dummies, full-time dummy, and actual hours of work (in the survey week or for similar workers in the LFS); $Year_t$ is a dummy variable for each year in the sample. The estimated coefficients

Table 3
Log-average real weekly earnings by province, 1981–97

	Provinces										
Year	**Nfld**	**PEI**	**NS**	**NB**	**QC**	**ON**	**MB**	**SK**	**AB**	**BC**	**Total**
1981	1.49	1.30	1.40	1.42	1.56	1.57	1.45	1.51	1.66	1.68	1.50
1982	1.46	1.21	1.38	1.42	1.53	1.50	1.43	1.45	1.65	1.64	1.47
1984	1.38	1.27	1.36	1.43	1.51	1.50	1.47	1.45	1.58	1.53	1.45
1985	1.39	1.31	1.38	1.41	1.50	1.52	1.45	1.43	1.55	1.54	1.45
1986	1.37	1.31	1.36	1.39	1.50	1.56	1.40	1.41	1.57	1.53	1.44
1987	1.39	1.32	1.40	1.35	1.51	1.57	1.40	1.38	1.52	1.52	1.43
1988	1.41	1.28	1.40	1.39	1.48	1.60	1.41	1.42	1.55	1.57	1.45
1989	1.45	1.34	1.42	1.41	1.52	1.58	1.42	1.40	1.53	1.55	1.46
1990	1.39	1.33	1.44	1.40	1.54	1.58	1.40	1.33	1.53	1.58	1.45
1991	1.37	1.31	1.36	1.39	1.50	1.56	1.36	1.33	1.50	1.59	1.43
1992	1.40	1.33	1.34	1.42	1.49	1.59	1.42	1.37	1.51	1.56	1.44
1993	1.37	1.34	1.40	1.40	1.48	1.58	1.40	1.37	1.50	1.52	1.43
1994	1.38	1.35	1.37	1.40	1.50	1.58	1.41	1.39	1.51	1.57	1.44
1995	1.42	1.33	1.35	1.37	1.49	1.54	1.40	1.40	1.47	1.62	1.44
1996	1.41	1.38	1.34	1.39	1.53	1.58	1.42	1.42	1.52	1.58	1.46
1997	1.30	1.31	1.33	1.37	1.51	1.59	1.43	1.38	1.51	1.60	1.43
Total	1.40	1.31	1.38	1.40	1.51	1.56	1.42	1.40	1.54	1.57	1.45

Notes: Average weekly earnings are calculated by dividing reported wages and salaries (in hundreds of dollars) by the number of weeks worked. Individual weights are used to calculate yearly averages. Total consumer price index (CPI = 100 in 1991) was used to deflate nominal wages.

Source: Statistics Canada, Survey of Consumer Finances. Cross-sectional files from 1981 to 1997. No data are available for 1983. Sample size is 623,875.

of the year dummies $\hat{\delta}_t, t = 1,...,16$, can then be interpreted as the regression-adjusted measures of the wage rate, i.e., the predicted yearly wage rate of an individual with a fixed set of characteristics.

Figure 1 illustrates the difference between the adjusted and unadjusted wage series in Canada. Except for the sharp drop during the 1981–83 recession, the unadjusted wage shows very little variation throughout the sample horizon. In particular, from 1988 to 1997, this series looks almost flat. By contrast, movements in various "adjusted" measures of the real wage follow a much more cyclical pattern, with a sharp increase in wages in the late 1980s, and a sharp decrease in the early 1990s. The figure shows three different adjusted measures of the real wages (all series are normalized to zero in 1997 for the sake of comparison). The top line on the graph is the wage adjusted only for changes in human capital (identified by HC in figures) and other socio-economic characteristics, while the two other wage series are based on models that also control for changes in hours, using the hours proxies available in the SCF and the LFS.

Figure 1 also shows that using proxies for hours from the SCF or the LFS yields very similar adjusted wage series. The adjusted wage series for which hours are not controlled exhibits more of a downward trend, but its cyclical behaviour is similar to that of the two other adjusted wage series. In the remainder of the paper, we will use the wage series adjusted for human capital, other socio-economic characteristics, and hours as measured in the LFS. Note that the results obtained using the different adjustment schemes are all qualitatively similar.

We use a similar procedure to construct adjusted measures of real wages at the provincial level. More specifically, we estimate a model with a full set of province-year interactions:

$$w_{ijt} = \beta X_{ijt} + \sum_{j=1}^{10} \sum_{t=1}^{16} \delta_{jt} Prov_j * Year_t + \varepsilon_{ijt}, \tag{2}$$

where $Prov_j$, for $j = 1,...,10$, is a set dummy variable for provinces. The estimated province-year effects ($\hat{\delta}_{jt}$) can be interpreted as regression-adjusted measures of the wage rate in a province j in year t (i.e., the wage in different provinces and different years for an individual with a specified set of characteristics).

2.3 Comparison with U.S. wage series

As an additional check on the quality of our wage series, we compare our results to those obtained using similar data for the United States. In March of every year, the U.S. Bureau of Census conducts an income supplement to

Figure 1
Adjusted vs. unadjusted wages in Canada

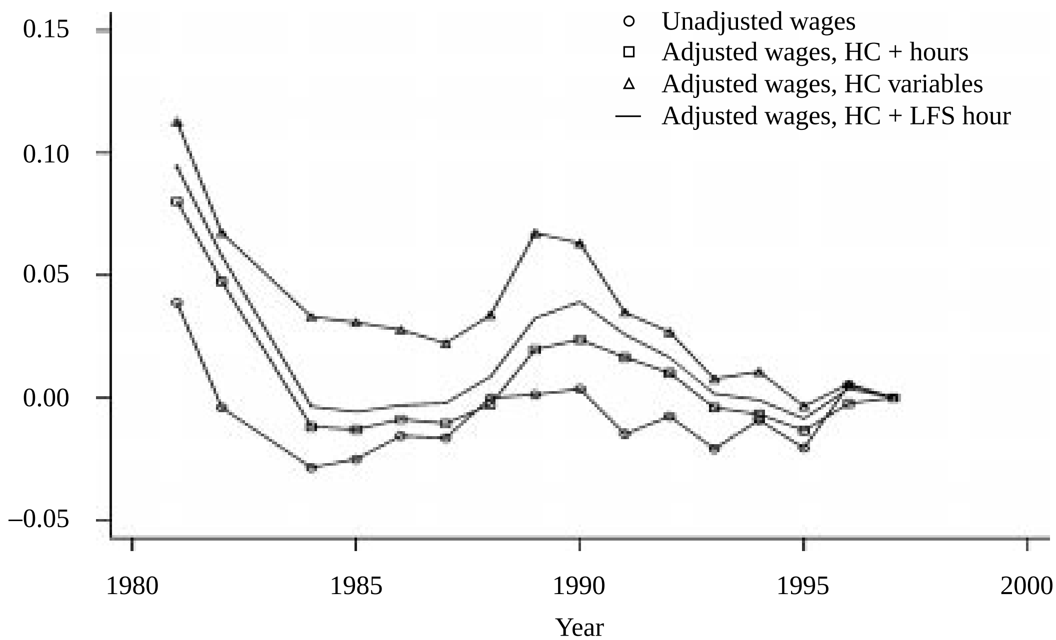

Notes: All wage indexes are normalized to zero in 1997.
HC = human capital.

the Current Population Survey (CPS), which is very similar to the SCF. Since 1976, the March CPS asks respondents about their usual weekly hours of work in the previous year. It is thus possible to compute a direct measure of hourly wage rates in the United States, by dividing annual wage and salary earnings by total hours of work (product of weeks worked and hours per week), and to compare this direct measure to the regression-adjusted methodology we use for Canada.

Figure 2 shows the unadjusted U.S. series for weekly and hourly wages, as well as the corresponding series adjusted for changes in individual characteristics and hours (in the case of weekly wages).[11] All wage series are procyclical although the timing of peaks and troughs in wages tends to slightly precede the peaks and troughs in overall economic activity. Interestingly, the adjusted wage series for hourly wages and weekly wages (see top of figure) are very close to each other, suggesting that weekly wages adjusted for the kind of hours measures available in the SCF are a very good proxy for the series based on actual hourly wage rates. Extrapolating from these U.S. results for Canada suggests that the time-series pattern of the

11. We perform the hours adjustment for the U.S. weekly wage series using the same variables as available in the SCF, namely full-time status in the previous year and hours worked in the reference week.

Figure 2
Adjusted vs. unadjusted earnings in the United States

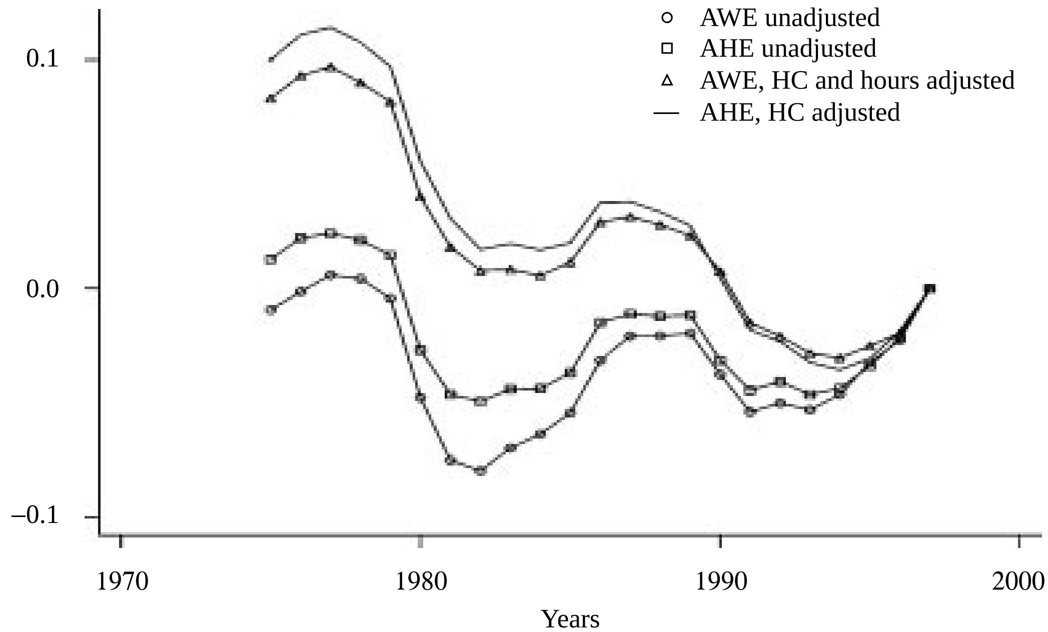

Notes: All wage indexes are normalized to zero in 1997.
AWE = average weekly earnings; AHE = average hourly earnings; HC = human capital.

Canadian wage series based on adjusted weekly wages mostly reflects true movements in hourly wages, as opposed to changes in weekly hours of work.

It is also interesting to explicitly compare the Canadian and U.S. wage series. Figure 3 plots the adjusted real weekly wage series (adjusted for individual characteristics and hours of work) for Canada and the United States. The two series are deflated by their own-country CPI. In both countries, wages drop in the late 1970s and early 1980s, increase during the recovery of the 1980s, and drop again in the early 1990s. Wage changes in the United States tend to precede those in Canada by a few years. For example, real wages drop dramatically between 1979 and 1982 in the United States, while this decline only occurs between 1981 and 1984 in Canada. In the 1980s, U.S. wages peak between 1986 and 1989, while in Canada the peak is reached only in 1989–91. Finally, U.S. real wages fall sharply between 1989 and 1991, while they start declining (at a slower pace) in Canada only after 1990.

One question raised by Figure 3 is whether the very low rates of inflation experienced by Canada in the 1990s prevented real wages from adjusting as quickly as they should have because of DNWR. Table 4 shows that starting in 1991–92, the inflation rate (CPI all items) dropped below

Figure 3
U.S. and Canadian adjusted wages

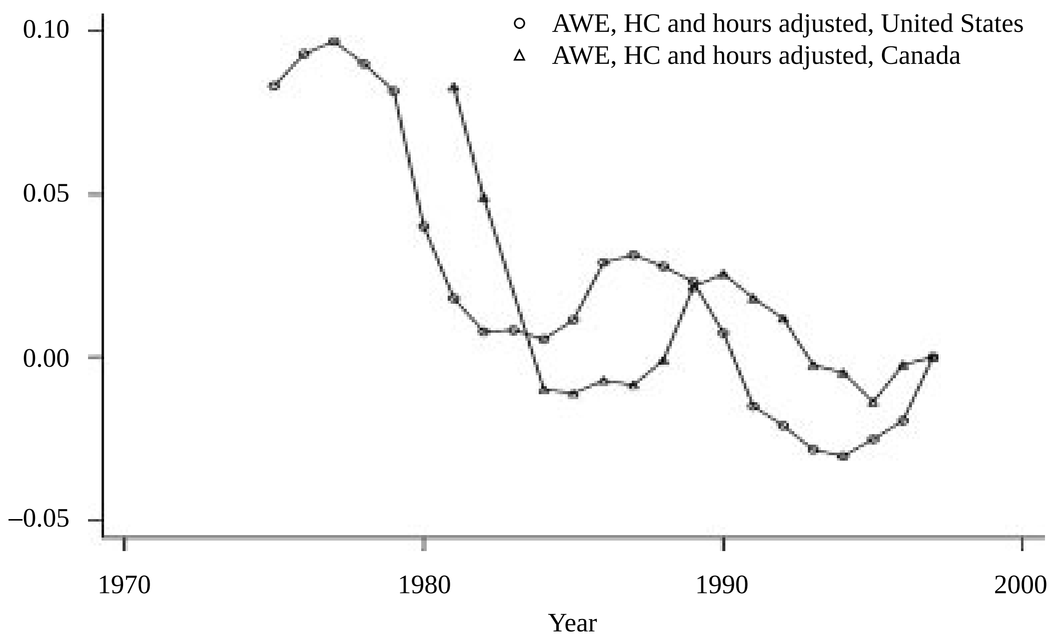

Notes: All wage indexes are normalized to zero in 1997.
AWE = average weekly earnings; HC = human capital.

2 per cent a year in Canada, while it remained around 3 per cent in the United States. By contrast, inflation rates in the two countries were roughly comparable during the 1980s. Therefore, if low inflation prevented real wages from declining quickly enough in Canada relative to the United States, this phenomenon should have occurred only after 1991. Figure 3 indicates, however, that real wages fell at least as quickly in Canada as in the United States after 1991. The big difference between Canada and the United States is that real wages remained constant between 1989 and 1991 in Canada, while they declined sharply in the United States during the same period. Since inflation rates in the two countries were comparable during this period, it is unlikely that DNWR can explain the relative evolution of real wages in the two countries after 1989.

A more direct way of assessing the role of DNWR in wage determination might be to look separately at the evolution of nominal wages and the price level (the two elements used to compute real wages). Figures 4 and 5 plot these two series for Canada and the United States. The figures show a much sharper break in the trends in these two series after 1991 in Canada than in the United States. In fact, there is almost no nominal-wage growth in Canada between 1991 and 1994, which is quite remarkable when compared to other time periods or to the United States. Taken at face value, this suggests that DNWR was quite "binding" in Canada in the early 1990s.

Table 4
The aggregate data

	Canada		United States	
Year	Δp_t	**UR**	Δp_t	**UR**
1981	11.70	7.58	9.48	7.60
1982	10.26	10.97	6.30	9.70
1984	4.22	11.31	4.22	7.50
1985	3.89	10.68	3.49	7.20
1986	4.09	9.66	1.84	7.00
1987	4.25	8.83	3.58	6.20
1988	3.97	7.77	4.05	5.50
1989	4.88	7.56	4.70	5.30

	Canada		United States	
Year	Δp_t	**UR**	Δp_t	**UR**
1990	4.65	8.13	5.26	5.50
1991	5.47	10.33	4.12	6.70
1992	1.48	11.15	2.96	7.40
1993	1.83	11.36	2.94	6.80
1994	0.17	10.38	2.52	6.10
1995	2.14	9.44	2.79	5.60
1996	1.56	9.65	2.91	5.40
1997	1.61	9.12	2.26	4.90

Notes: Price changes are calculated as log differences. Annual changes in total CPI is our inflation measure.
UR = unemployment rate.
Sources: CANSIM for Canada, Bureau of Labor Statistics for the United States.

In summary, the evidence on the role of DNWR in the evolution of real wages in Canada relative to its role in the United States is mixed. While the evolution of nominal wages between 1991 and 1994 suggests that DNWR was quite important, the fact that real wages fell as rapidly in Canada as in the United States during the same period suggests that DNWR did not prevent real wages from adjusting “fast enough.” In light of these ambiguities, we now turn to a more detailed analysis of how DNWR may affect the relationship between real-wage changes and economic conditions (unemployment rate).

3 Estimating Real-Wage Phillips Curves

As mentioned earlier, a key empirical implication of DNWR is that, in response to a given negative shock, the real wage should decline less in periods of lower inflation. We test this implication by estimating “real-wage Phillips curves” that link the unemployment rate to the change in real wages. If DNWR prevents real wages from adjusting (downwards) in periods of low inflation, the Phillips curve should be *flatter* in periods of lower inflation. These models are in the spirit of the traditional Phillips-curve approach, since *changes* in real wages, as opposed to their level, are expressed as a function of the unemployment rate.[12]

12. Blanchflower and Oswald (1994) suggest estimating a “wage curve” (wage level as a function of the unemployment rate) instead of a Phillips curve, while Card (1995) and Blanchard and Katz (1997) suggest otherwise.

Figure 4
Nominal earnings and CPI in Canada

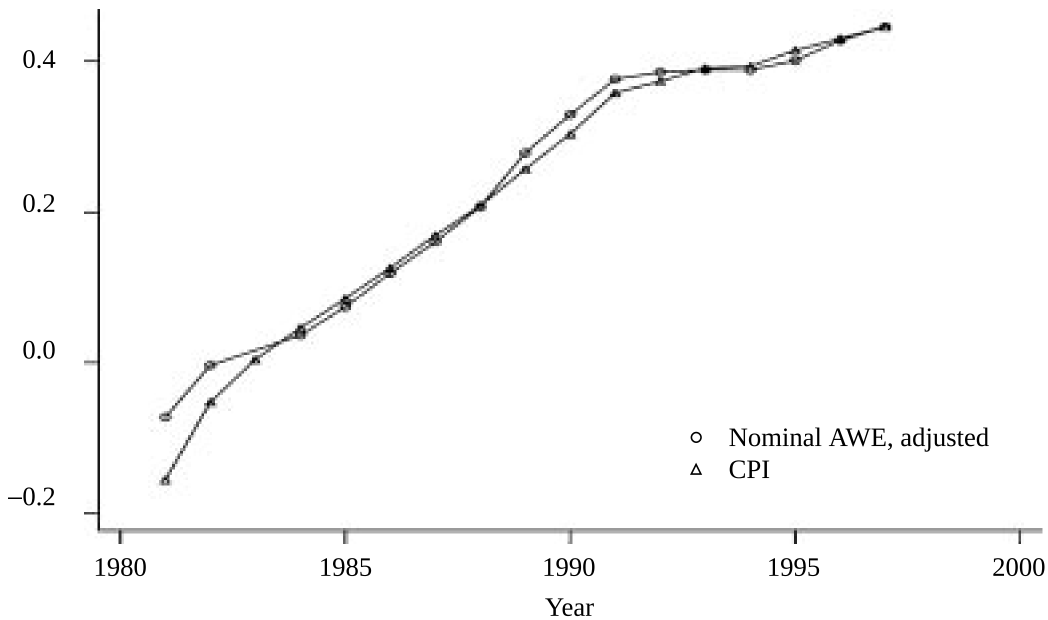

Note: AWE = average weekly earnings.

Figure 5
Nominal earnings and CPI in the United States

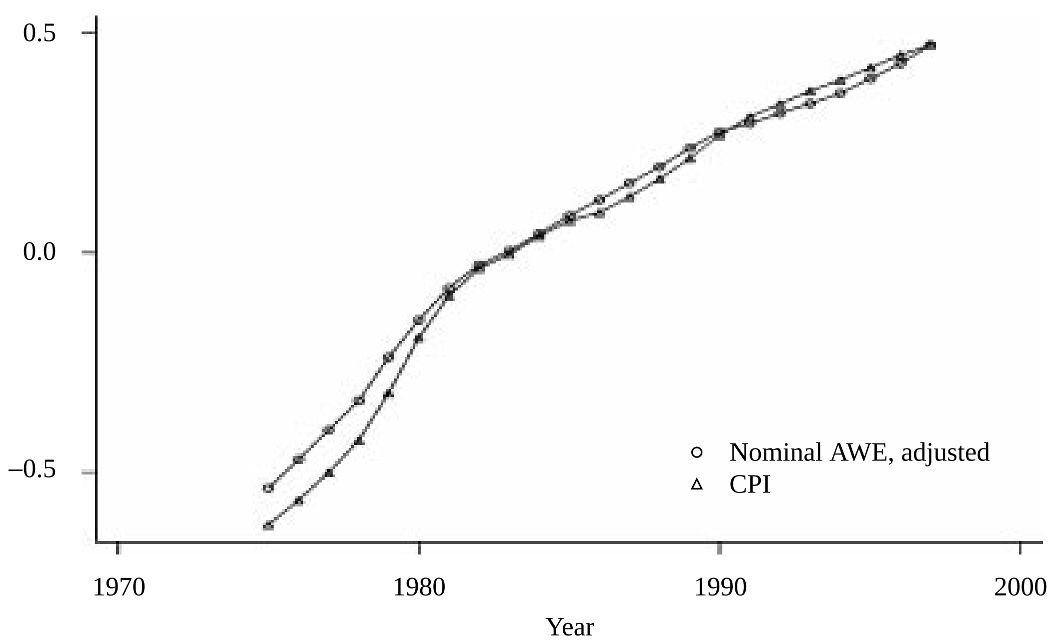

Note: AWE = average weekly earnings.

3.1 Aggregate Phillips curves

Figure 6 plots changes in (adjusted) real wages and the unemployment rate at the national level. Both series have been normalized, and the unemployment is plotted on an inverted scale to illustrate the co-movements between the two series. The figure indicates that the series track each other remarkably well. This close link is confirmed in Table 5, which reports OLS estimates of the Phillips curve. More specifically, column 1 reports estimates from a model in which the unemployment rate is the sole explanatory variable. The dependent variable used in all specifications is the change in adjusted (for individual characteristics and hours) real wages.[13] The estimated effect of the unemployment rate is negative and statistically significant. The estimated coefficient implies that real wages decline by 0.8 per cent each time the unemployment rate increases by 1 percentage point. The estimated effect is very similar when a linear time trend is also included in the model (column 2).

A closer look at Figure 6 suggests that the relationship between real-wage changes and the unemployment rate may have indeed changed after inflation dropped below 2 per cent a year in 1991. More specifically, changes in real wages stopped dropping and stabilized around −1 per cent a year after 1991, despite the fact that the unemployment rate kept rising between 1991 and 1993. Furthermore, real-wage declines in 1992 and 1993 were substantially smaller (around −1 per cent) than in the recession of 1981–83 (real-wage declines around −3 per cent), despite the fact that the unemployment rate was comparable (at around 11 per cent) in the two recessions.

This breakdown in the relationship between real-wage changes and the unemployment rate after 1991 is partly confirmed in the Phillips-curve estimates reported in column 3 of Table 5. The “low-inflation regime” is simply captured by a dummy variable equal to one in year 1992 and later, and to zero for earlier periods.[14] If the Phillips curve became flatter in this period, the interaction between this “low-inflation regime” dummy and the unemployment rate should be positive and statistically significant. The estimated interaction term reported in column 3 is positive, as expected, but is not significant at standard statistical levels.[15]

---

13. Since the SCF was not conducted for the (income) year 1983, we define the wage change for 1984 as the change between 1982 and 1984, divided by two.

14. This dummy captures most of the time-series variation in inflation, which hovered around 4 to 5 percentage points for most years until 1991, before declining permanently below 2 per cent.

15. The dummy for the low-inflation regime is also included by itself in the regression, since the intercept of the Phillips curve (real-wage change when the unemployment rate is zero) will likely be different during low- and high-inflation periods.

Figure 6
AWE growth and aggregate unemployment in Canada

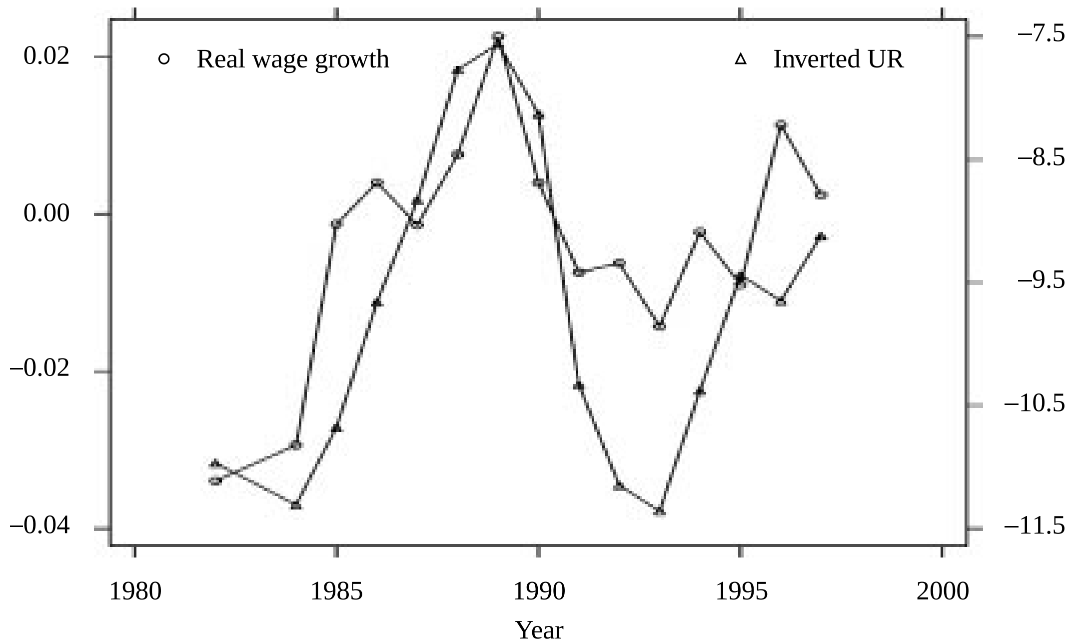

Note: UR = unemployment rate.

Quantitatively speaking, the estimated interaction term implies that the slope of the Phillips curve is about half as large during the post-1991 low-inflation period than earlier. However, no clear conclusion can be reached from the aggregate time-series analysis because of the imprecise results.

3.2 Provincial Phillips curves

The imprecision of the time-series results may not be surprising, since only six yearly observations are available in the "low-inflation regime" of the 1990s. Because different provinces experienced quite different economic conditions during the 1990s, this additional cross-provincial variation in unemployment rates (and potentially, real-wage changes) may help improve the precision of the parameters of interest.

One further concern with the aggregate time-series evidence is that other unmodelled economy-wide factors have also changed during this period. For example, inflation expectations may have changed after the Bank of Canada switched to a tighter (and low-inflation) monetary policy in the early 1990s. Supply shocks may have also shifted the Phillips curve during this period.

Table 5
Estimated aggregate Phillips curve
Sample years 1982–97

	Dependent variable: $\Delta\tilde{w}_t$ (change in adjusted wage)		
Control variables			
Constant	0.077	0.081	0.093
	(0.019)	(0.018)	(0.022)
u_t	–0.008	–0.008	–0.010
(Unem. rate)	(0.002)	(0.002)	(0.002)
Linear trend	—	0.0009	—
		(0.0005)	
*Y*1992	—	—	–0.037
			(0.049)
u_t^* *Y*1992	—	—	0.004
			(0.004)
$\bar{R}^2$	0.52	0.59	0.55

Notes: Standard errors are in parentheses. All regressions are weighted. Annual changes in log total CPI is the inflation measure. *Y*1992 is a dummy variable set to one if the year is greater than or equal to 1992.
For 1984, $\Delta\tilde{w}_{1984} = (\tilde{w}_{1984} - \tilde{w}_{1982})/2$.

Sources: Statistics Canada, Survey of Consumer Finances, for the wages. CANSIM for prices and aggregate unemployment.

A natural way to control for the economy-wide factors is to turn to cross-provincial analysis, which relies on variation in economic conditions across both time and provinces to identify potential changes in the slope of the (provincial) Phillips curve. Unrestricted year effects can be used to control for nation-wide factors, while provincial variations can identify the connection between provincial wage changes and unemployment rates.

More specifically, we estimate the following type of cross-provincial Phillips curve:

$$\Delta\tilde{w}_{jt} = a(j) + \gamma(t) + \beta_t U_{jt} + \varepsilon_{jt}, \qquad (3)$$

where $\tilde{w}_{jt}$ is the adjusted average real-wage index for province j at time t, with the first difference taken over time; $a(j)$, for $j = 1,...,10$, is a set of province dummies; $\gamma(t)$, for $t = 82,...,97$, is a set of year dummies; U_{jt} is the measured unemployment rate in province *j* at time *t*; ε_{jt} represents the residual error term.

In principle, a separate slope of the Phillips curve (β_t) could be estimated for each year. In practice, we estimate specifications similar to those for the aggregate time-series models in which the provincial unemployment rate is either interacted with the inflation rate or with a dummy variable for the "low-inflation regime" to test whether DNWR, combined with low inflation, has flattened the Phillips curve.

Before going to the regression models, it is useful to look at the main trends in real wages and unemployment rates across provinces. Figure 7 plots the unemployment rate and the change in real wages for the four largest provinces over the 1982–97 period. The lower panel shows that, as is well known, the recession of the early 1980s was more pronounced in the West (Alberta and British Columbia) than in central Canada (Quebec and Ontario). Interestingly, real wages also fell more precipitously in western Canada than in central Canada (upper panel). This illustrates a clear trade-off between the evolution in provincial unemployment rates and changes in real wages, i.e., a cross-provincial Phillips curve.

The regional patterns in the recession of the early 1990s are very different from those of the recession of the early 1980s. Quebec, and especially Ontario, experienced much steeper increases in unemployment than the western provinces. Unlike the 1980s, however, there is no clear visual evidence that real wages fell more precipitously in Ontario than in the West, suggesting that DNWR, coupled with low inflation, may have prevented real wages from adjusting as much as they should have in Ontario.[16]

Table 6 shows the OLS estimates of equation (3), using a variety of specifications. In all models, we include an unrestricted set of province dummies to absorb permanent differences in wage changes and unemployment rates across provinces. In columns 1 to 4, the slope of the Phillips curve is assumed constant over time. The model in column 1 includes no control for year effects, while column 2 includes a linear trend, and column 3 includes a set of unrestricted year effects. The model reported in column 4 includes different linear trends by province, in addition to the unrestricted set of year effects (at the national level). In all four cases, the unemployment rate has a negative effect on changes in real wages. The point estimates indicate that a 1 percentage point increase in the provincial unemployment rate reduces provincial real-wage growth by 0.3 to 0.6 per cent. The

16. Some could argue, however, that policies of the provincial government during this period may have also contributed to keeping real wages from falling more.

Figure 7
Adjusted provincial wages and unemployment rates

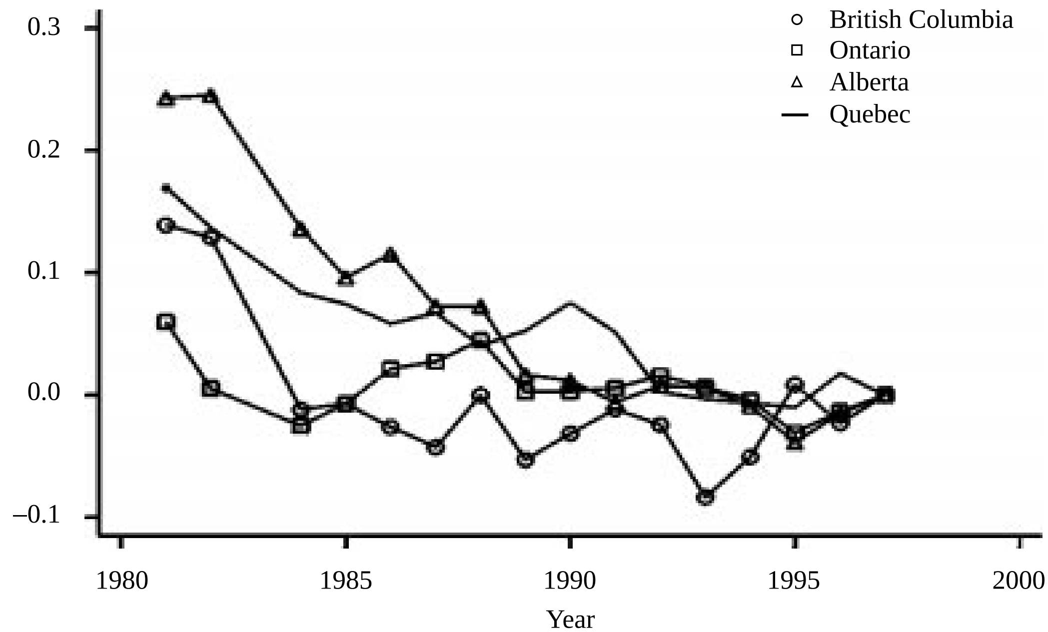

Provincial wages

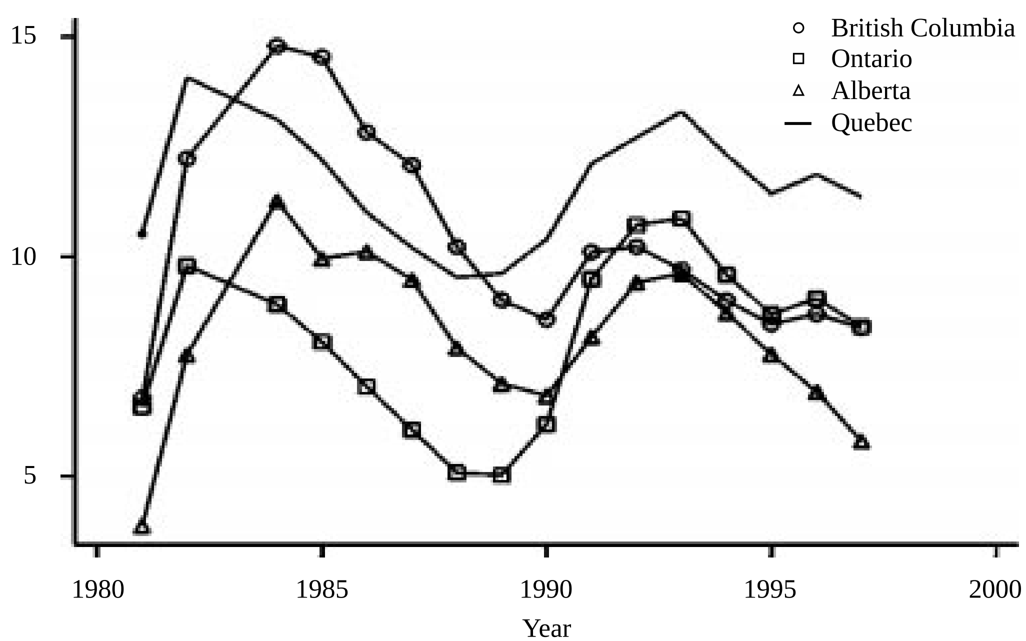

Provincial unemployment rate

Note: All wage indexes are normalized to zero in 1997.

Table 6
Estimated provincial Phillips curve
Sample years 1982–97

	Dependent variable: $\Delta\tilde{w}_{jt}$							
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Control variables								
Constant	0.066	0.070	0.042	0.043	0.067	0.072	0.033	0.051
	(0.013)	(0.013)	(0.019)	(0.023)	(0.014)	(0.014)	(0.022)	(0.020)
u_{jt}	–0.006	–0.006	–0.003	–0.003	–0.006	–0.005	–0.003	–0.005
(Unemployment rate)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Linear trend	—	0.0009	—	—	—	0.002	—	—
		(0.0004)				(0.000)		
Year effects	No	No	Yes	Yes	No	No	Yes	Yes
Province trends	No	No	No	Yes	No	No	No	No
*Y*1992	—	—	—	—	0.008	–0.010	0.017	—
					(0.017)	(0.019)	(0.019)	
u_{jt} * *Y*1992	—	—	—	—	–0.0003	0.000	–0.001	—
					(0.0017)	(0.001)	(0.001)	
$u_{jt}\Delta p_t$	—	—	—	—	—	—	—	0.042
								(0.036)
$\bar{R}^2$	0.13	0.15	0.19	0.14	0.13	0.15	0.19	0.19

Notes: Standard errors are in parentheses. All specifications include 10 province dummies. Regressions are weighted using province weights. Annual changes in log total CPI is the inflation measure. The number of observations is 150. Excluded year is 1997 and excluded province is British Columbia. For 1984, $\Delta p_{1984} = p_{1984} - p_{1983}$ and $\Delta w_{j,\,1984} = (w_{j,\,1984} - w_{j,\,1982})/2$.

Sources: Statistics Canada, Survey of Consumer Finances, for the wages. CANSIM for prices and provincial unemployment.

estimated effects are statistically significant for all specifications except the one in column 4.

Columns 5 to 7 report estimates for the same three specifications as in columns 1 to 3, when the provincial unemployment rate is interacted with the dummy variable for low inflation. As expected, the interaction term is estimated much more precisely using cross-provincial variation than when using only aggregate variation (see Table 6). The standard error is around 0.001, as opposed to 0.004 in Table 5. The point estimates of the interaction term are now small and not statistically significant for all of the reported models. The same conclusion is reached in column 8, where the actual inflation rate (as opposed to a dummy for low-inflation years) is interacted with the unemployment rate. All in all, the cross-provincial estimates do not support the view that the Phillips curve is flatter in years of very low inflation than in other years.

4 Reconciling the Pieces of Evidence: For Whom Does DNWR Bind?

We have touched on contradictory pieces of evidence regarding the importance of DNWR. On the one hand, we have shown that there was almost no nominal-wage growth in Canada during the 1991–94 period and that real wages did not fall as quickly in this period as in the 1980s recession. On the other hand, our estimates do not suggest that the slope of the Phillips curve decreased during years of very low inflation than during other years, as it should have if DNWR prevented real wages from adjusting enough in the face of negative unemployment rate shocks. Furthermore, real wages fell as quickly in Canada as in the United States, where the inflation rate was higher during the 1991–94 period.

One possible way of reconciling these apparently contradictory findings is to exploit the richness of the SCF data to better understand the dynamics of real-wage adjustment along the business cycle. As mentioned in the literature survey, DNWR theories are most relevant for more “stable” workers, who are most likely to stay with the same employer. By contrast, DNWR should not prevent employers from hiring new workers at lower nominal wages than they may have done in other circumstances. If the bulk of wage adjustments over the business cycle occur at the entry level, the presence of DNWR may not have much impact on (upward or downward) aggregate wage adjustments.

For example, Beaudry and DiNardo (1991) show that, consistent with implicit wage theory, *real* wages of workers who stay with the same employer are downward-rigid. Aggregate real wages only decline during recessions because of workers who start new jobs. During expansions, real wages may either increase because new workers obtain higher wages or because workers still with the same employer receive pay increases (to prevent other employers from "poaching" them).[17] Taking Beaudry and DiNardo's results at face value suggests that DNWR should have no effect on aggregate wages and employment. Of course, when inflation gets very close to zero, nominal rigidities are the same as real rigidities. They can appear to have an effect, to the extent that real rigidities also have an effect.

The SCF data allow us to examine these issues by looking at the evolution of real wages for different levels of job seniority. Figure 8 shows the adjusted wages between 1981 and 1997 for the different levels of seniority available in the SCF. The most noticeable feature of this figure is that real wages of more senior workers are much less cyclical than those of less senior workers. For example, the real wages of workers with 20 years or more of seniority hardly fall at all during the recession of the early 1980s. By contrast, real wages of workers with a year or less of seniority (workers on "new jobs") fell by almost 20 per cent during the same period.[18]

Real wages of workers with a year or less of seniority fell by much less in the recession of the 1990s than in the early 1980s. Since DNWR should not play an important role for these workers, this suggests that other factors were at play. For the most senior workers, real wages appear relatively rigid over the business cycle throughout the 1981–97 period. The years of very low inflation since 1991 are not different from other years in this regard.

The behaviour of real wages for the different groups may help explain why DNWR may not have much impact on aggregate wages and employment, despite the fact it is "binding" in some circumstances. As mentioned earlier, DNWR most likely matters for senior and stable workers who have long-term associations with their employers. For this group, however, Figure 8 suggests that real wages are quite rigid anyway (for other reasons, such as implicit contracts, for example). This means that DNWR matters most for workers whose real wages are relatively inflexible. By contrast, most of the real-wage adjustments over the business cycles are accounted for by workers on new jobs, whom DNWR should not affect to any great degree.

17. McDonald and Worswick (1999) find similar results for Canada (Beaudry and DiNardo 1991 use U.S. data).

18. Individuals in the "lost their job" category report earnings during the previous year despite the fact that they were no longer employed at the time of the survey. Their wages can be thought of as wages for workers who were about to lose their jobs.

Figure 8
Adjusted earnings for different job tenures

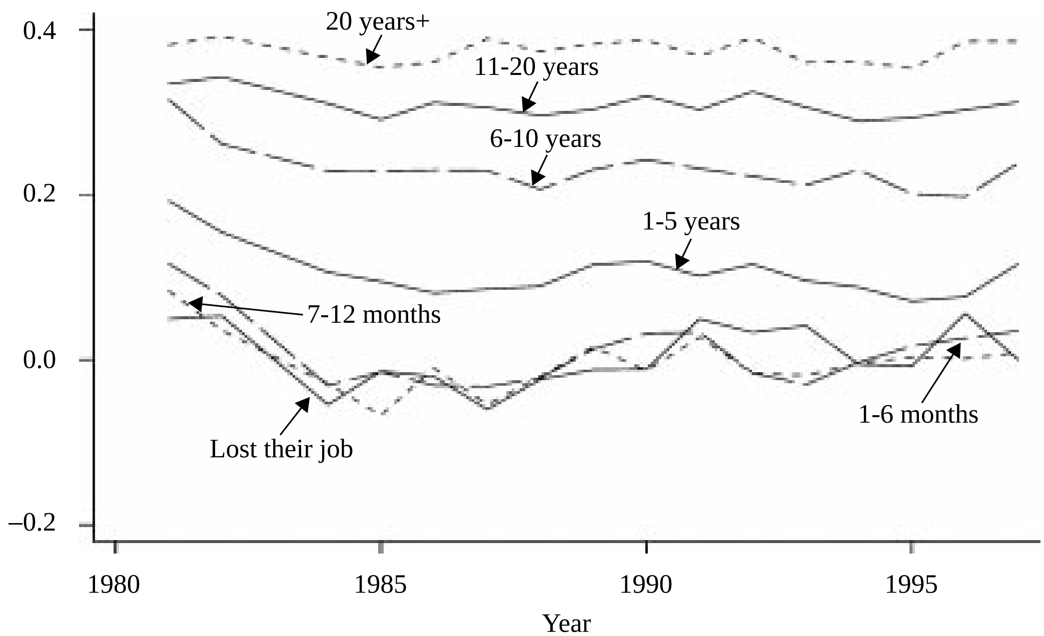

Conclusion

One main contribution of this paper is the development of a series of adjusted real wages for Canada from 1981 to 1997. This series is constructed using detailed data from the SCF that allow us to control (adjust) for composition effects over the business cycle. One first finding is that real wages are clearly procyclical in Canada, and that failure to adjust for changes in the composition of the workforce tends to understate the cyclicality of real wages.

We use these wage data to test whether DNWR tends to flatten the relationship between real wages and economic conditions as captured by the unemployment rate. While the aggregate results are indecisive because of small sample sizes, the results based on cross-provincial variation indicate that the slope of this real-wage Phillips curve has remained constant over time. These findings suggest that DNWR did not have a significant impact on wage and employment determination during the post-1991 period of very low inflation.

We attempt to reconcile this finding with the rest of the literature that clearly indicates the existence of DNWR by analyzing the evolution of real wages for different groups of workers. Our results suggest that DNWR binds most for more senior workers who would have relatively rigid real wages even in the absence of DNWR. By contrast, the bulk of real-wage adjustments over the business cycle is experienced by new entrants (young workers or workers on new jobs) for whom DNWR is least likely to bind. This may explain why DNWR has little effect on aggregate real-wage determination, despite the fact that it is a significant phenomenon for some groups, such as older and more senior workers.

References

Akerlof, G.A., W.T. Dickens, and G.L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 1: 1–59.

Beaudry, P. and J. DiNardo. 1991. "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data." *Journal of Political Economy* 99 (4): 665–88.

Bewley, T.F. 1999. *Why Wages Don't Fall During a Recession.* Cambridge: Harvard University Press.

Bils, M.J. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93 (4): 666–89.

Blanchard, O. and L.F. Katz. 1997. "What We Know and Do Not Know About the Natural Rate of Unemployment." *Journal of Economic Perspectives* 11 (1): 51–72.

Blanchflower, D.G. and A.J. Oswald. 1994. *The Wage Curve.* Cambridge: MIT Press.

Card, D. 1995. "The Wage Curve: A Review." *Journal of Economic Literature* 33 (2): 785–99.

Card, D. and D. Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labour Market'?" In *Reducing Inflation: Motivation and Strategy*, edited by C.D. Romer and D.H. Romer, 71–114. Chicago: University of Chicago Press.

Crawford, A. 2000. "How Rigid Are Nominal-Wage Rates?" Bank of Canada. Working Paper. Forthcoming.

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.

Farès, J. and S. Hogan. 2000. "The Employment Costs of Downward Nominal-Wage Rigidity." Bank of Canada Working Paper 2000–1.

Faruqui, U.A. 2000. "Employment Effects of Nominal-Wage Rigidity: An Examination Using Wage Settlements Data." Bank of Canada Working Paper 2000–14.

Fortin, P. 1996. "The Great Canadian Slump." *Canadian Journal of Economics* 29 (4): 761–87.

Kahn, S. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review* 87 (5): 993–1008.

Lebow, D.E., D.J. Stockton, and W.L. Wascher. 1995. "Inflation, Nominal Wage Rigidity, and the Efficiency of Labour Markets." U.S. Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series No. 95–45.

McDonald, J.T. and C. Worswick. 1999. "Wages, Implicit Contracts, and the Business Cycle: Evidence from Canadian Micro Data." *Journal of Political Economy* 107 (4): 884–92.

McLaughlin, K.J. 1994. "Rigid Wages?" *Journal of Monetary Economics* 34 (3): 383–414.

———. 1999. "Are Nominal Wage Changes Skewed Away from Wage Cuts?" *Federal Reserve Bank of St. Louis Review* 81 (3): 117–32.

Simpson, W., N. Cameron, and D. Hum. 1998. "Is Hypoinflation Good Policy?" *Canadian Public Policy* 24 (3): 291–308.

Smith, J.C. 2000. "Nominal Wage Rigidity in the United Kingdom." *The Economic Journal* 110 (462): C176–95.

Solon, G., R. Barsky, and J.A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?" *The Quarterly Journal of Economics* 109 (1): 1–25.

Discussion

Allan Crawford

The impact of downward nominal-wage rigidity on real wages and employment is one of the key issues in the debate over the appropriate level for the long-run inflation target. Since each wage series for Canada has some limitations for testing the effects of rigidity, it has proven useful to study this issue from a variety of perspectives, using different databases and statistical techniques. One branch of the literature has followed a microeconometric approach. Most of the Canadian studies in this area have used data for individual union contracts to estimate the effect of rigidity on wage growth and/or employment. Another branch of the literature has taken a more macro perspective by estimating a Phillips curve using some measure of aggregate wage growth. The focus of the latter group of studies is to test the prediction that downward nominal rigidity would cause the Phillips curve to become flatter at low rates of inflation.

The Farès-Lemieux paper adopts the macro perspective to look for a change in the slope of the Phillips curve in Canada. It proceeds in three stages. First, the authors construct a series for aggregate real-wage growth that incorporates adjustments for the effect of changes in the composition of employment. Second, they use the adjusted aggregate data to estimate a real-wage Phillips curve, and test for a change in its slope during the low-inflation years of the 1990s. Finally, to better understand the results from the Phillips curve, they carry out some informal analysis of real-wage movements for different categories of workers. I have some comments on each of these three themes and then compare their results with those from the micro studies of wage rigidity in Canada.

Data for Real-Wage Growth

The novel feature of the Farès-Lemieux paper is the use of an aggregate wage series constructed from the Survey of Consumer Finances (SCF), which collected information from approximately 40,000 individuals per year. Unfortunately, the SCF did not follow a given sample of individuals over time, so these data cannot be used to study wage rigidity at the micro level. Given this constraint, Farès and Lemieux combine the individual wage data to form a series for aggregate wage growth (either at the national or the provincial levels) and then use the aggregate series to estimate the Phillips curve. Since this wage variable is the focus of their analysis, some further comments on its construction are in order.

The aggregate wage series of Farès and Lemieux has several desirable attributes for a test of wage rigidity. First, since it is constructed from data for paid workers aged 20 to 65 years, the wage variable should be broadly representative of the overall labour market in terms of sectoral coverage and other dimensions such as union status. This characteristic makes it easier to draw general conclusions about the effect of rigidity on aggregate economic outcomes. Other potential candidates for an aggregate wage variable, such as total labour income per person-hour, would also have the advantage of broad sectoral coverage.

The real strength of the Farès-Lemieux series relative to other measures of aggregate wage growth is that it attempts to control for the effects of changes over time in the composition of the workforce. Farès and Lemieux are able to make these adjustments by using information on human capital and job characteristics from the individual data files of the SCF. As noted by the authors, these compositional shifts will probably cause the unadjusted measure of aggregate wage growth to overstate the rigidity of wage changes for workers of given characteristics. Consistent with their expectations, they show that the aggregate real wage in Canada is indeed more responsive to the cycle after adjusting for the effect of compositional shifts in employment. An important implication of this finding is that other studies in the literature that have used an aggregate wage variable without compositional adjustments may be biased in favour of finding rigidity.

Farès and Lemieux calculate the adjusted real wage using the consumer price index (CPI). Thus, all of their econometric work examines the relationship between changes in the consumer real wage and the unemployment rate. It is an empirical issue whether consumer prices or producer prices (or both) determine nominal-wage growth.[1] Whatever the

1. See Cozier (1991) for a comparison of the long-run movements in producer and consumer real wages in Canada.

case, it is ultimately the producer real wage (the nominal wage divided by a measure of producer prices such as the GDP price deflator) that is relevant for the employment decisions of firms and for analysing the employment effects of rigidity. This point suggests that a useful addition to the Farès-Lemieux paper would have been to extend their analysis to consider the case of producer real wages.

Is this distinction between consumer and producer real wages likely to matter empirically? CPI and producer price inflation tend to follow similar trends. Nevertheless, their inflation rates can diverge in the short run, and these differences can persist for long enough periods to give different trend movements for the levels of the two indexes. To illustrate this point, consider Figure 1, which shows the ratio of the GDP price deflator to the CPI. The vertical line indicates the beginning of the sample period used by the authors. The downward trend in this ratio over the sample period implies that consumer prices increased more rapidly than producer prices.

What are the implications of these movements in relative prices for the trend change in the real wage? Table 3 of the Farès-Lemieux paper shows a decline in the unadjusted real weekly earnings over the 1981–97 sample period, when the real wage is calculated using the CPI. Since the GDP deflator rose at a slower pace over those years, the real wage *increased* when calculated using producer prices.

Figure 1
Ratio of producer prices to the CPI (1961 = 1.0)

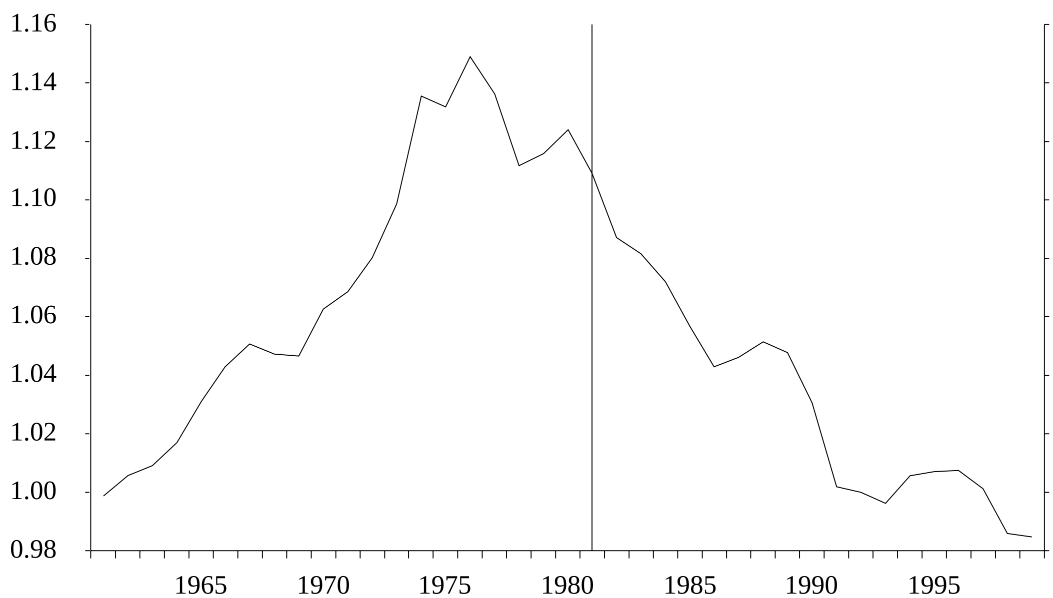

Potentially more important for analysing the effects of downward nominal rigidity at low inflation is how the choice of price variable affects the measured *cyclicality* of real wages. Using their series for consumer real wages, Farès and Lemieux note that the decline in real wages was smaller during the recession and recovery period in the early 1990s than during the comparable period in the early 1980s (when inflation was much higher). This pattern could be interpreted as informal evidence of a weaker relationship between the unemployment rate and real-wage growth at lower rates of inflation. However, the difference between the growth rates of real wages in the early 1980s versus the early 1990s is reduced when the real wage is calculated using producer prices.[2] Since employment decisions depend on producer real wages, basing the analysis on the consumer real wage could overstate the employment effects of rigidity at lower rates of inflation.

I have several additional comments on the Farès-Lemieux wage series and their informal analysis of these data. Farès and Lemieux restrict their sample to paid workers aged 20 to 65 on the grounds that the concept of downward rigidity is not relevant for self-employed workers (who account for approximately 15 per cent of the labour force). While this statement is valid, self-employed workers could have been included in the sample given the authors' goal of evaluating the impact of downward nominal-wage rigidity on aggregate wages and employment. By excluding the self-employed, they may be biasing the results in favour of finding rigidity. The net effect of excluding teenagers from their wage series is perhaps less clear-cut. Teenaged workers are more likely to be subject to minimum-wage floors, although evidence presented later in their paper suggests that real wages tend to be most flexible for new entrants to the labour force.

Finally, the authors report that there is almost no change in the adjusted nominal wage between 1991 and 1994, and conclude that this suggests, "at face value," that downward nominal-wage rigidity was significant in Canada in the early 1990s. It is not clear how a constant (aggregate) nominal wage is evidence of downward rigidity. In fact, it could just as easily be interpreted as reflecting a high degree of downward flexibility in nominal-wage rates. To give a hypothetical example, a constant aggregate wage would be consistent with a scenario in which half of all workers receive a wage cut and half receive a wage increase. As the authors acknowledge, a more formal analysis (such as estimation of Phillips curves)

---

2. Consumer price inflation was significantly greater than producer price inflation in the recession and recovery period in the early 1980s, whereas there was a smaller difference between the two measures of price change during the low-inflation years of the early 1990s.

must be used to assess the degree to which the observed outcome is affected by downward nominal rigidity.

Estimation of Phillips Curves

Downward nominal-wage rigidity implies that a negative shock would have less effect on real wages in a period of low inflation, because nominal-wage floors are more likely to bind under those conditions. The authors test this prediction by estimating a real-wage Phillips curve for Canada and examining whether it is flatter during the post-1991 period of low inflation. This test is incorporated in most of their equations by including a dummy variable that interacts with the unemployment rate for the low-inflation years in their sample (1992 to 1997). A limitation of this specification is that the dummy variable constrains rigidity to have the same effect on the slope of the Phillips curve for each year in the low-inflation period. A more flexible model would permit the slope of the Phillips curve to vary depending on the frequency of binding nominal-wage floors in each period.

Farès and Lemieux consider an alternative specification that does allow the effects of rigidity to vary systematically over time. In the section of their paper that examines cross-provincial Phillips curves, they estimate a model in which the provincial unemployment rate is interacted with the level of CPI inflation. This interaction term will have a negative sign (and the Phillips curve will become flatter at lower rates of inflation) if binding nominal-wage floors become more widespread at low inflation. Models of wage determination at the micro level help to identify the conditions under which this specification would capture the effects of rigidity. For example, in a Tobit model of wage rigidity, there is an inverse relationship between rigidity and inflation if the *ratio* between the mean and the standard deviation of the notional wage-change distribution decreases at lower rates of inflation.[3] The Farès-Lemieux specification (with rigidity modelled as a continuous function of inflation) can be interpreted as an attempt to approximate the relationship between this ratio and inflation.[4]

The authors estimate their model using either national or provincial data that have been adjusted to control for the effects of compositional shifts in employment. Both sets of results suggest that the slope of the Phillips

3. The "notional" distribution is the distribution that would be observed in the absence of downward nominal rigidity and menu-cost effects.

4. Using a Tobit model, Crawford and Wright (2001) show that both the mean and the standard deviation of the notional distribution tend to decrease with inflation. The net effect is a decrease in their ratio.

curve did not become flatter in years of low inflation.[5] From these findings, they conclude that downward nominal-wage rigidity did not have a significant effect on aggregate wages and employment during the low-inflation years of the 1990s. As noted previously, the adjustment for compositional shifts is a welcome innovation, and studies that fail to make similar adjustments to aggregate measures of wage growth may be biased in favour of finding significant effects from rigidity. Thus, an interesting extension of the Farès-Lemieux paper would have been to report estimates of the Phillips curve obtained from the unadjusted wage data so that readers could assess the practical significance of this critique.

Real-Wage Movements for Individual Groups

Section 4 of the paper presents an informal analysis of movements in the real wages of different categories of workers in order to explain why there is no evidence of a flattening of the Phillips curve. A key finding from this disaggregated analysis is that the (consumer) real wages of older and more senior workers remained relatively constant at all inflation rates, which leads the authors to conclude that real wages are relatively rigid for this group for reasons other than downward nominal rigidity. Real wages of new entrants (young workers or those with a year or less of seniority) fell by less in the recession of the 1990s than in the early 1980s, but the authors argue that other factors should explain this pattern, since downward nominal rigidity should not be important for these types of workers.

The evidence from the disaggregated data provides useful insight for understanding the aggregate results. This analysis is limited, however, by the absence of a formal test or estimate of what real-wage movements would have occurred for each group in the absence of downward nominal rigidity. One direction for future research could be to estimate the wage equation for the different groups, although it may be difficult to obtain adequate measures of demand pressures at the disaggregated level.

5. The Farès-Lemieux approach is closely related to U.S. work by Card and Hyslop (1997) who estimated similar models using state-level wage data adjusted for compositional effects. They did not find a statistically significant relationship between the slope of the Phillips curve and inflation over the period 1976 to 1991. Thus, the methodology gives similar conclusions for Canada and the United States. The sample period of Farès and Lemieux should provide a better test of rigidity, because it includes years with lower inflation (CPI inflation in Canada averaged about 1.5 per cent from 1992 to 1997).

Comparison with Canadian Micro Studies

The introduction to my remarks noted that a variety of techniques and databases have been used to study the extent of downward nominal rigidity and its effect on employment. To conclude, I provide some comments on whether the aggregate results of Farès and Lemieux are corroborated by the findings of Canadian studies using microdata.

The authors' conclusion that downward rigidity had little effect on aggregate wages and employment stands in sharp contrast to the conclusions of Simpson, Cameron, and Hum (1998). Using a standard Tobit model and data for union wage settlements, they concluded that downward rigidity is quite widespread in Canada at low rates of inflation. Simpson et al. also estimated a reduced-form employment equation that suggests downward nominal rigidity (proxied by the percentage of contracts with wage freezes) reduced employment by a significant amount in the mid-1990s.

More recent micro studies suggest that the effects of rigidity on wages and employment are much closer to the conclusions from the aggregate analysis of Farès and Lemieux. In terms of the effects on wages, the models in these micro studies are structured to incorporate important stylized facts from the wage-change distribution. For example, a weakness of the standard Tobit model used by Simpson et al. is that it attributes all wage freezes to downward nominal rigidity. This assumption is questionable given the observation that the distribution of wage settlements contains few contracts with small wage increases or small wage decreases. This pattern suggests that some wage freezes are caused by symmetric menu-cost effects rather than asymmetric downward rigidity. Failure to consider menu-cost effects would lead to an overstatement of the impact of downward nominal rigidity on wage growth.

Another feature of the observed distribution of wage settlements is a decrease in variance in periods of lower inflation. It is sometimes suggested that this decrease reflects a thinning of the density in the left tail of the distribution owing to downward rigidity, rather than a change in the notional variance. However, the decrease in dispersion occurred on both sides of the distribution, which suggests that much of the downward trend in the observed variance can be attributed to a decrease in the notional variance at lower rates of inflation. If this interpretation is correct, a model that constrains the notional variance to be constant will tend to overstate the amount of the notional distribution below zero in periods of low inflation, and therefore overstate the effects of rigidity on wage growth. Accordingly, empirical models should test whether the notional variance is time-varying.

Crawford and Wright (2001) show that extending the Tobit model to include both menu-cost effects and a time-changing notional variance give significantly lower estimates of rigidity than those reported by Simpson et al. In these extended models, the estimated net effect of downward rigidity and menu-cost effects on wage growth in the 1990s is approximately 0.4 per cent for the average wage change in the first year of contracts and less than 0.1 per cent for the average annual change over the duration of contracts. Similarly, in hazard models reported in Crawford (2001), the estimated net effect of downward nominal rigidity and menu-cost effects on the average annual wage growth over the lifetime of contracts is within the 0.10 to 0.20 per cent range.

Recent studies also provide further evidence on the employment effects of rigidity. Faruqui (2000) extended the reduced-form employment equation of Simpson et al. in various ways to better control for the effects of demand shocks. In most of his specifications, the wage-freeze proxy for rigidity has no significant effect on employment growth. Another test for employment effects is to examine whether the long-run Phillips curve is non-vertical at low inflation. Using Tobit models and wage-settlements data, Crawford and Wright estimate that the long-run curve is close to vertical at inflation rates of 2 per cent or more if productivity growth is close to its average from recent decades.

On balance, the micro evidence for Canada suggests that any effect of downward nominal-wage rigidity on wages and employment was small during the low-inflation period in the 1990s. This conclusion is consistent with the results from the aggregate analysis in the Farès-Lemieux paper.

References

Card, D. and D. Hyslop. 1997. “Does Inflation ‘Grease the Wheels of the Labour Market’?” In *Reducing Inflation: Motivation and Strategy,* edited by C.D. Romer and D.H. Romer, 71–114. Chicago: University of Chicago Press.

Cozier, B. 1991. “Real Wages and Productivity in Canada: An Update.” *Bank of Canada Review* (July): 3–8.

Crawford, A. 2001. “How Rigid Are Nominal-Wage Rates?” Bank of Canada Working Paper. Forthcoming.

Crawford, A. and G. Wright. 2001. “Downward Nominal-Wage Rigidity: Micro Evidence from Tobit Models.” Bank of Canada Working Paper. Forthcoming.

Faruqui, U.A. 2000. “Employment Effects of Nominal-Wage Rigidity: An Examination Using Wage-Settlements Data.” Bank of Canada Working Paper 2000–14.

Simpson, W., N. Cameron, and D. Hum. 1998. “Is Hypoinflation Good Policy?” *Canadian Public Policy* 24 (3): 291–308.

Discussion

Wayne Simpson

What should the long-run target for monetary policy be? Absolute price stability or something else? These are questions that have occupied macroeconomists and central bankers for a long time. One issue in this debate has been how the labour market operates, particularly whether wages are flexible or not. In the 1990s, as inflation declined but Canadian unemployment remained high relative to the United States, the debate has centred on "downward nominal-wage rigidity" as a particularly strong form of wage inflexibility. The paper by Farès and Lemieux provides an overview of the debate in the Canadian context, offers some new evidence, and suggests new directions for research.

Suppose that, as the inflation rate approaches zero, adverse (below average) business conditions in some firms lead to negative real-wage offers, which amount to nominal pay cuts. This leads to questions that could affect the targets for monetary policy. First, will these workers successfully resist pay cuts? Second, will firms react by cutting employment and output? And third, what are the consequences for the conduct of monetary policy? My comments will examine each of these questions, with particular reference to the paper by Farès and Lemieux, and offer some concluding observations on the implications for further research in this area.

Is There Evidence of Downward Nominal-Wage Rigidity?

Farès and Lemieux provide a compact survey of recent empirical research in Canada and the United States and conclude that there is evidence of downward nominal-wage rigidity (DNWR). I agree with their assessment that "DNWR clearly acts as a constraint on nominal-wage changes at the micro level" (page 9). Less clear is *why* it is important and for which workers and firms.

In Canada, much of the research has responded to Fortin's (1996) comments in his address to the Canadian Economics Association, by more closely examining data on settlements in bargaining units with 500 or more workers. The advantage of these data is that they provide a clear measure of the base, or scale, increase to workers in a particular occupation or industry. The obvious disadvantage is that they represent wage settlements for only a collectively powerful minority of workers in the country. In that sense, this evidence was an important first test of DNWR: if there were no evidence of DNWR in the base settlements data, we shouldn't expect to find it elsewhere.

As Farès and Lemieux report, even when we look beyond the first year of multi-year contracts, which Fortin did not do, there is evidence of DNWR. I would add two recent unpublished studies to their evidence. I have used the model developed by Kahn (1997) to decompose annual settlements in the private sector into pure distributional effects, based on distance from the median settlement, and additional special effects associated with downward wage rigidity and menu costs (Simpson 1998). For the entire period, DNWR is estimated to apply to 10 per cent of all settlements, and about half of these may be attributed to menu costs associated with resistance to very small negative or positive wage changes. But the estimate for DNWR is not stable, as we might expect, rising to 15 per cent during 1993–97, when inflation is much lower. Menu costs can still account for no more than 5 per cent of the total, leaving an estimated 10 per cent of settlements to be explained by DNWR during the mid-nineties. Another recent paper, by Christofides and Stengos (2000), finds that wage settlements are less symmetric during 1992–96 than during earlier periods, using non-parametric tests of the symmetry of the wage distribution, consistent with the hypothesis of DNWR.

Does the evidence of DNWR extend beyond large bargaining units? Here the evidence is less clear, in large part because the data are more confusing. Farès and Lemieux note that a number of studies of household microdata sets, including the Panel Study of Income Dynamics in the United States and the Labour Market Activity Survey and Survey of Labour and Income Dynamics (SLID) in Canada, have found evidence of DNWR, despite the well-known bias towards over-reporting pay cuts in such data (Simpson et al. 1998). Although more direct questions on wage changes—such as, "Did you receive an adjustment to your basic hourly wage this year? If so, how much?"—would be more useful than reported earnings and hours, household microdata remain a promising general area for research on wage adjustment. Farès and Lemieux concentrate on the Survey of Consumer

Finance (SCF) cross-sectional files from 1981 to 1997, but future researchers will likely find SLID more valuable, particularly the forthcoming initial six-year panel (1993–98) and its successors. In addition to tests for DNWR among all workers, we would like to know whether it is primarily a phenomenon of the unionized sector, or whether it extends to non-unionized workers as well. Understanding who is affected by DNWR will help us test theories that would claim to explain it.

There seems to be a strong body of evidence to suggest that DNWR exists, at least for workers in large bargaining units. The approaches developed to measure DNWR will be useful as new data sets become available to test it further, and as macroeconomic conditions change.

How Do Firms React to DNWR?

If DNWR exists, does it affect employment and output decisions of firms? At a theoretical level it seems clear that DNWR must have some effect. If firms choose the level of employment once wages are determined, then DNWR implies a higher real wage and a lower level of employment and output than would be obtained in the absence of DNWR. An efficient bargain may lie off the labour demand curve, but Hum et al. (1999, Appendix) show that, when union preferences favour nominal-wage maintenance at any cost (as they must under DNWR), the contract must lie along the labour-demand curve, implying a lower level of employment than if nominal wages are negative.

But can this theoretical prediction be observed in aggregate data and, if so, how large is the employment effect of DNWR? Here the evidence is less clear and more controversial. Our paper (Simpson et al. 1998, Table 3) first looks across industries to determine whether a detectable partial correlation exists between the incidence of zero or negative settlements, which are indicative of DNWR, and employment growth. What we had in mind was a standard model of employment and real wages at the industry level, in which labour demand shifted as output changed, and labour supply was relatively stable. Then the reduced-form model for employment (and the corresponding model for wages) would depend on output. Industry-specific fixed effects could be eliminated by estimating the model in first differences, yielding a model of employment growth as a function of output growth. Our twist was to add a term, the incidence of pay freezes or cuts in the settlements data, to capture the incidence of DNWR. This term is significantly negatively correlated with employment growth, as we predicted. We argued then, and would argue now, that we are reluctant to place too much emphasis on this result because the measure of DNWR is a very poor one: it does not capture wage settlements in the industry beyond those in large bargaining units and, even for large bargaining units, it measures only the incidence, and not the extent or size, of the DNWR.

Bank of Canada papers by Farès and Hogan (2000) and Faruqui (2000) question the robustness of this result. They introduce the rate of wage change into the employment growth equation to account for labour-demand shocks, and find that the effects of DNWR disappear. We find this approach problematic, because wage change is endogenous and its introduction produces simultaneity bias that distorts the results. While we welcome attempts to test the fragility of our results, we are not convinced by this approach. Further tests and better data at the level of the individual firm are likely required. For example, Groshen and Schweitzer (1999) used a 40-year panel of wage changes reported by large employers in the U.S. midwest to find evidence for DNWR ("sand") at inflation levels below 5 per cent. Employer-based evidence of this sort would likely be needed to provide better tests of the impact of DNWR on employment and output.

Does DNWR Matter for Monetary Policy?

DNWR is important to policy-makers only if it can be demonstrated that it affects the trade-offs between inflation and other economic goals, such as unemployment. With apologies for stereotyping, I would characterize the official position among most macroeconomists and central bankers as being that inflation is costly and price stability is preferred in the long run. That is, the long-run Phillips-curve relationship between real-wage change and unemployment is vertical, and the short-run curve slopes downward fairly steeply at all rates of inflation. Thus, their interest is in models in which the aggregate impact of DNWR can be substantiated within this Phillips-curve framework.

Our approach was to model DNWR explicitly by treating pay freezes and pay cuts as censored data, using a Tobit model (Simpson et al. 1998). Our results suggested that, if the behaviour with respect to DNWR in the settlements data were representative of other wage behaviour in the economy, DNWR kept wages 2/3 of a percentage point higher than they would have been otherwise between 1993 and 1995. We therefore estimated that DNWR "cost" the economy a higher unemployment rate of 2 per cent to achieve the same inflation goals as would have been available in the absence of DNWR. These are certainly results that should be of interest to policy-makers if they are corroborated in further research.

Contrary to what Farès and Lemieux report (page 10), we did not assume that the variance of wage growth was constant over time. Rather, we allowed the variance to depend on time, and we reported the results of the Tobit model with and without heteroscedasticity (Simpson et al. 1998,

Table 5). Farès and Lemieux argue that there has been "noticeable compression in the wage-change distribution in the 1990s" (page 10), but it is not clear how this matters if, as our results suggest, this has been primarily as a result of the compression of wage settlements at zero due to DNWR itself.

The authors use the SCF cross-sectional data files to examine the evolution of Canadian wage rates from 1981 to 1997. They recognize that a disadvantage of the SCF is that, unlike its successor, SLID, it does not provide direct information on hours worked per week over the survey year. They therefore use both direct and indirect information on hours worked, along with information on other observable worker characteristics, to produce adjusted estimates of average weekly wages per year. While lack of direct information on hourly wage rates adds noise to the data, the direction of the bias is not clear, and the methodology does allow some potentially interesting new features of wage change to be explored.

One finding is that wages adjusted for worker characteristics demonstrate a cyclical pattern not apparent in the unadjusted series. This is somewhat surprising, however, because the unadjusted series reflects the movement of both hourly wages and hours, while the adjusted series is supposed to reflect only movements in hourly wages. If the adjusted series exhibits a cyclical pattern similar to that of the United States, as the authors' evidence suggests (Figure 3), does this mean that hours worked are less cyclical in Canada than in the United States? Or do other adjustments for worker quality differ between Canada and the United States? It would be useful to compare the unadjusted hours worked series for Canada and the United States and, if necessary, to decompose the differences in the movement of the average weekly earnings into movements in hours, changes in worker quality, and movements in hourly wages, to help readers understand the adjusted series.

The authors then estimate real-wage Phillips curves for Canada, following the approach taken by Card and Hyslop (1997) for the United States. That is, they regress the unemployment rate on the change in real wages to determine whether the relationship is flatter when inflation is low. The argument is: when inflation is low, either DNWR would lead to reductions in employment and higher unemployment, or a given level of slack in the labour market would have a smaller effect in reducing nominal and real wages. This seems like a more indirect method of testing for DNWR than our approach, which explained individual settlements in terms of prevailing economic conditions (monthly unemployment and inflation rates) and treated DNWR as a process of censoring. It is certainly less powerful econometrically, since it aggregates wage changes into an annual series with only 15 observations, compared with nearly 15,000 observations in our analysis. It is not surprising, then, that the authors' results (Table 5) are not statistically significant, although they are consistent with a flatter real-wage Phillips curve after 1991. Their test, however, is not the appropriate one. They focus solely on the interaction term between the unemployment rate and the time dummy (equal to 1 after 1991), but a flatter Phillips curve will imply both a change in the slope, captured by the interaction term, and the intercept, captured by the time dummy itself. These two coefficients should be tested jointly in what amounts to a Chow test for stability of the regression. Although the outcome of these results is not certain, I suspect that the sample is simply too small to detect a statistically significant shift in the curve after 1991. It is worth noting, however, that the results, while not significant, are consistent with the hypothesis of DNWR during this period.

To expand the sample size, Farès and Lemieux estimate provincial Phillips curves. I suspect that this is the only alternative available for this data set, but it raises some new concerns. For example, the authors use the national inflation rate when provincial inflation rates should be used. Since the premise for pooling is that there are provincial Phillips curves, provincial inflation rates will be negatively correlated with provincial unemployment rates which, the authors show, have varied considerably through the 1990s (Figure 7). The discrepancies between the national and provincial inflation rates, which would be part of the error term in their equation (3), would then be correlated with the unemployment rate in their regressions, resulting in biased estimates.

I also wonder whether it is reasonable to pool the provinces, i.e., are the provincial Phillips curves sufficiently similar in structure? If not, provincial dummy variables and year dummy variables may not adequately capture the differences, rendering pooling inappropriate. I think much more work needs to be done to determine why the results are different when the pooled data are used. Why, for example, is the Phillips curve so much flatter in general when the data are pooled? Although the coefficient on the unemployment rate was quite accurately measured in the aggregate data at -0.008 to -0.010 (with a t-value of 4 to 5), it has dropped sharply to as low as -0.003 when the provincial data are used.

Although there are good theoretical reasons to link the evidence of DNWR with employment and unemployment, the task of finding those links in the available data is difficult, and the results will naturally be more controversial. I do not see conclusive evidence that DNWR has changed the shape of the Phillips curve, although there is some evidence that it might have, including the aggregate results in Farès and Lemieux's paper. The provincial results appear to be unreliable at this stage.

In perhaps the most interesting part of the paper, they use the SCF microdata to disaggregate wage growth by job seniority. They show that the sensitivity of real wages to cyclical fluctuations, measured by movement of the unemployment rate, declines with increasing seniority, as we would expect. That is, wage offers to incoming and probationary workers are more sensitive to labour market conditions than wage offers to their more senior colleagues. At the same time, it is difficult to tell from the data whether DNWR matters for some workers regardless of seniority, such as unionized workers in large bargaining units. In other words, what other characteristics of workers might account for wage growth or lack thereof? This is clearly an important area for further research, for which microdata will be indispensible.

Where Do We Stand on DNWR?

I see the Farès-Lemieux paper as consistent with my assessment of the evidence to date. We generally observe DNWR in microdata on wage changes, but we find it much more difficult to find conclusive evidence that it affects the decisions of employers and, in the aggregate, the trade-offs between inflation, unemployment, and other economic goals that guide monetary policy formulation. Different responsibilities and prior beliefs lead to different reactions to this evidence, either to move on to something more interesting or to look further into the issue. My inclination is definitely towards the latter.

Regardless of the importance of DNWR, the issue has concentrated attention on important links between macroeconomic policy and labour market behaviour. As new data sets are available, such as the SLID six-year panel and the Workplace and Employee Survey from Statistics Canada, they will provide new opportunities to test for DNWR and to improve our understanding of wage determination. For those who think that the labour market is a fascinating and important part of economics and economic policy, this cannot be a bad development.

References

Card, D. and D. Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labor Market'?" In *Reducing Inflation: Motivation and Strategy*, edited by C.D. Romer and D.H. Romer, 71–114. Chicago: University of Chicago Press.

Christofides, L.N. and T. Stengos. 2000. "The Symmetry of the Wage-Change Distribution: Non-Parametric Tests Using Contract and Survey Data." University of Guelph. Manuscript.

Farès, J. and S. Hogan. 2000. "The Employment Costs of Downward Nominal-Wage Rigidity." Bank of Canada Working Paper 2000–1.

Faruqui, U.A. 2000. "Employment Effects of Nominal-Wage Rigidity: An Examination Using Wage-Settlements Data." Bank of Canada Working Paper No. 2000–14.

Fortin, P. 1996. “The Great Canadian Slump.” *Canadian Journal of Economics* 29 (4) 761–87.

Groshen, E.L. and M.E. Schweitzer. 1999. “Identifying Inflation’s Grease and Sand Effects in the Labor Market.” In *The Costs and Benefits of Price Stability*, edited by M. Feldstein, 273–308. NBER Conference Report. Chicago: University of Chicago Press.

Hum, D., W. Simpson, and N. Cameron. 1999. “Reducing Inflation in an Era of Low Inflationary Expectations.” In *Macroeconomic Causes of Unemployment: Diagnosis and Policy Recommendations*, edited by W. File and C. Kohler, 271–88. Belin, Duncker, & Humblot.

Kahn, S. 1997. “Evidence of Nominal Wage Stickiness from Microdata.” *American Economic Review* 87 (5): 993–1008.

Simpson, W. 1998. “A Test for Nominal and Real Wage Stickiness in Canadian Wage Settlements.” University of Manitoba (October). Photocopy.

Simpson, W., N. Cameron, and D. Hum. 1998. “Is Hypoinflation Good Policy?” *Canadian Public Policy* 24 (3): 291–308.

General Discussion*

Thomas Lemieux thanked both discussants for their constructive comments. He indicated that he and Jean Farès would follow up on many of the suggestions, including the use of provincial CPI, unadjusted data, and the producer price index (PPI). He noted that they had conducted specification checks but that they could do further work.

Lemieux also responded to Wayne Simpson's comment on how the adjustment of hours affected the finding of cyclicality of real wages. He pointed out that some confusion may have resulted from the fact that they adjust for other characteristics such as age and education at the same time as they adjust for hours. The adjustment for other characteristics makes real wages more cyclical, whereas the adjustment for hours actually makes real wages less cyclical, just as one would expect. In fact, in the raw data hours rise a bit at the end of the expansion—over the 1989–90 period—and fall over the 1990–91 period.

In response to Simpson, Farès pointed out that the standard errors on the slope estimates are smaller in the provincial estimates than in the aggregate estimates. He also mentioned that Allan Crawford's suggested use of the PPI as a deflator for real wages appealed to him. He expected that real wages calculated in this manner would decline less in the 1980 recession, making the graph of real wages in the 1980s much more appealing. He then turned to the subject of looking at different tenure profiles and indicated that these profiles could benefit from more analysis. In particular, the wages of young, junior workers are not bound by DNWR. For example, if firms in the 1990s were trying to adjust their wages by adjusting over these young

* Prepared by Marianne Johnson.

workers, there would have been a much more aggressive reaction in their real wages than is evident in the data. Real wages of young workers actually drop a lot less in the 1990s than in the 1980s, so these movements in real wages are probably not due to DNWR.

Michael Parkin suggested as a reference a 1976 Carnegie-Rochester conference volume article by Michael Sumner, Robert Ward, and himself, which examines nominal Phillips curves but addresses the price-index issue. He pointed out that if you assume that the equilibrium-wage rate is determined by supply and demand, and if the demand for labour depends on the real producer wage and the supply of labour depends on the real consumer wage, then the equilibrium nominal-wage rate depends on both the producer price index and the consumer price index, and the weights depend on the relative elasticities of supply and demand. If you then deflate the nominal-wage rate by one of these two indexes—and the choice is arbitrary—then the ratio of the two should appear as one of the explanatory variables. Therefore, the real-wage Phillips curve needs both of these indexes.

Tim Sargent highlighted the issue of measurement error in provincial Phillips-curve equations. He noted that the authors conclude from their regressions that β—the slope of the Phillips curve—has not fallen in the 1990s. The regressions control for constant differences in natural rates across provinces and changes in the national natural rate, but do not control for important differences in the behaviour of provincial natural rates in the 1990s in particular, because of the differing impact of EI and welfare reform across provinces. This will bias up the estimates of β in the 1990s, because it is as if there is measurement error in the measure of U, since β should be attached to $(U–U^*)$ where U is the natural rate. The authors cannot conclude that the slope of the Phillips curve has not fallen in the 1990s, because the estimate of β is biased up in the 1990s relative to β in the 1980s.

Paul Beaudry highlighted the importance of focusing on real wages. He noted that the spikes in nominal data, such as the settlements data, could either be a reflection of the absence of real changes or a reflection of even more wage changes. When inflation is stable at 1 to 2 per cent, perhaps there can be a nominal-wage change at zero without much difficulty, although this is actually a 1 1/2 per cent decrease in real wages. When there is inflation at 10 per cent, firms are usually forced to index wages because inflation is so extreme. Beaudry mentioned that he has looked at cross-country data and the evidence in a broad sense suggests that in high-inflation periods there appears to be less adjustment in real wages than in periods of low inflation.

Serge Coulombe suggested that the authors check for heteroscedasticity, which is often a problem with pooled cross-sectional data.

William Robson raised the issue of the composition of the workforce. He recommended separating out public sector wage changes, since federal and provincial wage freezes were important in the early 1990s. It may prove more enlightening to exclude the public sector.

Thomas Lemieux thanked the speakers for their comments and suggestions. He was interested, in particular, in following up on the PPI and CPI index suggestions. He responded to Tim Sargent's point, agreeing that natural rate, EI reform, and other issues such as changes in minimum wages are important. He mentioned that they did add year dummies as well as provincial trends to capture some of these effects. He admitted that the natural rate might not be moving the same way in all provinces and that this may warrant further investigation. To Beaudry he responded that the percentage of contracts with explicit indexation does rise in high-inflation periods, a stylized fact that is consistent with his point. He thanked Coulombe for his suggestion. They had done specification tests and could add the results to the paper. He also agreed that Robson's idea of a public-private sector breakdown might be interesting.

Crawford responded to Simpson's analysis of a 1978–97 graph that looked at distance from the median using wage-settlements data. Crawford himself had conducted a similar experiment for private sector settlements for the low-inflation years, from 1992 onward. When one looks at distance from the median and compares the right and left sides, there are two equally prominent distinguishing characteristics. One—the spike in the interval containing wage freezes—is undeniable. The second characteristic is the degree to which there are very few small wage increases. If you assume symmetry and look at the private sector in low-inflation years, there is nothing unusual, like a shortage of wage cuts, for example. He admitted that the symmetry assumption is strong.

Simpson mentioned that the advantage of the menu-cost approach used in a recent paper by Shulamit Kahn, where she did not impose symmetry, is that you can disentangle the piling up at zero and the depressions around zero associated with menu costs.

Simpson also raised the question of how to treat individuals in a microdata set, when the individuals have no reported wage in one period but do report a wage in the next, or vice versa. Their wages go from zero to a positive or from a positive to zero. These individuals are usually dropped, so there will be compositional changes in the data over time, which is not ideal.