

Discussion

Pierre Fortin

By far the most surprising macroeconomic development of recent years in Canada has been the lack of response of inflation to very high unemployment. In the five-year period 1992-96, unemployment has averaged 10.4 per cent, which is far above the standard estimates of 7 to 8 per cent that were available in the late 1980s for the non-accelerating-inflation rate of unemployment (NAIRU). Nevertheless, inflation in the CPI excluding food and energy has meanwhile remained very stable around 1.4 per cent. Since the traditional Phillips curve with a 7 to 8 per cent NAIRU significantly overpredicts inflation in the 1990s, something must be wrong with it, and it must be repaired. If we do not understand the macroeconomic impact of market slack on inflation, we do not understand how monetary policy works.

This is what sets the stage for the contributions by Dupasquier and Ricketts and by Crawford and Harrison to this conference, as well as by Fillion and Léonard (1997) in a recent Bank of Canada working paper. Three conjectures are investigated:

1. *Non-linear inflation response*: inflation could respond less to market slack when unemployment is high than when it is low. This would generate a convex *short-run* Phillips curve (Dupasquier-Ricketts).
2. *Shifting expectations*: changes in the monetary regime could have affected the inflation expectations process. This would produce shifts in the Phillips curve constant and autoregressive coefficients on past inflation (Fillion-Léonard).
3. *Wage resistance*: workers and firms could show strong resistance to nominal wage cuts. This would make the *long-run* Phillips curve

convex in the range of inflation rates for which average wage growth is low (Crawford-Harrison).

Two other relevant hypotheses were assessed in previous Bank of Canada research by Rose (1988) and by Cozier and Wilkinson (1991). They add a fourth and fifth point of view to the three listed above, as follows:

4. *Structural change*: changing labour-force composition, unemployment insurance regulations, and other such factors could affect the correspondence between measured market tightness and true (but unobserved) inflationary pressure. This would produce shifts in the Phillips curve constant (Rose).
5. *Unemployment hysteresis*: the disinflationary effect of persistent actual unemployment could decline over time for various reasons. This would imply that inflation depends on the rate of change of unemployment as well as on its level (Cozier-Wilkinson).

I marvel at the extraordinary amount of good work that has come out of the Bank's Research Department on these issues over the years. Unfortunately, however, the research effort has been a bit piecemeal. Each study tends to focus on one hypothesis or two, leaving the others untested. It is therefore vulnerable to damaging biases of omission. There is no way to judge whether the studies are all right, or all wrong, or in what proportion right and wrong. The missing piece is an integrative study that would allow all hypotheses simultaneously to compete within an encompassing macro-economic framework. My general suggestion is that Dupasquier-Ricketts, Crawford-Harrison, and Fillion-Léonard talk to each other and work together!

Let me now turn to specific comments on the Dupasquier-Ricketts and Crawford-Harrison contributions to this conference.

Dupasquier and Ricketts

The paper by Dupasquier and Ricketts uses U.S. and Canadian macro data to test whether the slope of the short-term Phillips curve is time-varying. The model they estimate with quarterly data from 1964 to 1994 is based on the following two equations:

$$\pi = \delta\pi_{-1} + (1 - \delta)\pi^* + \beta y + \varepsilon, \text{ and} \quad (2)$$

$$\beta = \alpha + \rho\beta_{-1} + \gamma x + \mu, \quad (3)$$

where π is total CPI inflation, π^* is expected inflation, y is the output gap, x is the vector of influences on the short-term slope β , and ε and μ are zero-mean error terms.

The expected inflation measure π^* is extracted from survey data or generated by the ex post application of three-state Markov switching models (MSM) to the inflation process. The potential output estimate that is the basis for the measurement of the output gap y is calculated with filtering techniques (Hodrick-Prescott or “extended multivariate”) or derived from three-variable structural vector autoregressions (SVAR). The content of x follows the authors’ summary of the literature on non-linearities in the output-inflation relationship, which is by itself a very useful contribution. Included in x are: any positive output gap (based either on the Phillips capacity constraint model or the Stiglitz monopolistic competition model); the average level of inflation (low, medium, or high, based on the Ball-Mankiw-Romer costly adjustment model); and the conditional volatility of inflation (based on the Lucas signal extraction model).

The idea of simultaneously testing these competing theories of the non-linear output-inflation relationship within a common general framework is the main innovation of the paper—a most welcome one. The relationship between π and y will be non-linear if and only if any of the parameters ρ , γ , or $\text{var}(\mu)$ are statistically different from zero, which will mean that β is time-varying. The estimation method is maximum likelihood, and the unobserved slope β is generated with a Kalman filter and appropriate initial conditions.

Spanning various combinations of expected inflation and output gap measures, four models are estimated for the United States and two for Canada. The test results for the determinants of the slope β are reported in column 4 of Tables 2 to 7 of the paper. The main characteristic of these results is that they are inconclusive. They are particularly sensitive to how the output gap is measured.

The two U.S. models based on the Hodrick-Prescott filter for potential output and the Canadian model based on the SVAR measure of potential output support none of the four theories represented by components of x . The two U.S. models based on the SVAR measure of potential output and the Canadian model based on the Bank of Canada “extended multivariate” filter for potential output are all consistent with the old capacity constraint hypothesis. The Lucas signal extraction hypothesis is also a possibility in the latter Canadian model. The fragility of the results is disappointing. This calls into question the robustness of previous non-linearity results obtained at the Bank by Laxton, Rose, and Tetlow (1993) and by Fillion and Léonard (1997).

I see two important problems. First, several right-hand variables—such as expected inflation π^* , the output gap y , and components of the vector of non-linear influences x —are constructed artificially on the basis of many restrictive assumptions. The likelihood of measurement errors seems

high. Second, potentially important additional causal factors—such as hysteresis, supply shocks, and downward wage-price rigidity—are simply ignored. Measurement errors and omitted variables produce biased estimates and invalid tests.

I would suggest more transparency, more flexibility, and more comprehensiveness in model specification. For expected inflation, direct estimation of flexible distributed lags on past inflation, possibly with varying coefficients across monetary regimes (as in Fillion-Léonard), would be simpler and less restrictive. The monetary regimes can be identified by the Markov switching model applied to the inflation process, which is the one bit of positive value added I would attribute to this technique.

For potential output, I would recommend against the use of filters—however “extended multivariate” they are—and SVAR techniques. These methods are either purely mechanical, or based on combinations of mechanical filters and restrictive economic assumptions. None of them is rooted in any particular economic theory of wage-price disequilibrium, which is the context for the use of the output gap here. This point takes on more importance from the fact that the Dupasquier-Ricketts results are so sensitive to the particular choice of output gap measure.

One solution is to use a stationary variable such as the unemployment rate on the right-hand side of the Phillips equation and allow the regression constant to vary on account of structural factors leading to changes in equilibrium unemployment. There are many ways of implementing this, one being to add a transition equation for equilibrium unemployment with the lagged value of the unemployment rate and the various structural influences as right-hand variables.

In the vector x of influences on the slope β , I would include any positive unemployment rate gap (or some other function of the unemployment rate), a continuous moving average of recent inflation experience, and (even if the Lucas hypothesis has been falsified many times in the past) an ex ante measure of the recent conditional volatility of inflation. Markov switching techniques are not recommended here because they are based on ex post processing of inflation data for the full sample.

It would also be easy to repair the major omissions. A long distributed lag of changes in the unemployment rate is the appropriate way of testing for the presence of hysteresis. Supply variables such as changes in the effective indirect tax rate, in the relative prices of food, energy, and imports, and in the real effective exchange rate should not be forgotten. Concerning downward wage-price rigidity, the non-linear estimation method developed by Akerlof, Dickens, and Perry (1996) could be used. An alternative, simpler approach would be to allow the Phillips curve constant

and the coefficients on past inflation and unemployment to interact with an indicator variable for periods where wage growth is relatively slow.

Crawford and Harrison

The paper by Crawford and Harrison can best be viewed as systematic criticism of three propositions:

1. There is strong resistance to nominal base-wage cuts in major union wage settlements (with at least 500 employees) when inflation is very low.
2. Flexible compensation schemes (such as bonuses and profit-sharing plans) do not really allow firms to avoid the downward nominal wage rigidity.
3. Downward nominal wage rigidity is also characteristic of non-unionized and smaller firms.

Sections 1 and 2 of the paper deal with the first proposition. Section 3 examines the other two statements.

Section 1: How strong is the resistance to wage cuts in unionized firms?

Section 1 disputes the practical significance of the fact that, during the low-inflation period 1992-96, 49 per cent of the first-year base-wage changes in large union contracts were freezes and only 4 per cent were rollbacks, which produced a very pronounced spike at zero for the wage-change distribution. Crawford and Harrison make the following two arguments, which have the combined effect of reducing the 49 per cent incidence of wage freezes to a much more modest 13 per cent:

1. The public sector is over-represented in the major wage settlements data (60 per cent). Moreover, the incidence of first-year freezes in that sector (also 60 per cent) is much higher than in the private sector (at 33 per cent). It follows that, if the data are re-weighted according to the relative shares of the private and public sectors in the total economy, the estimated incidence of first-year wage freezes over the 1992-96 period is not 49 per cent, but 37 per cent.
2. Furthermore, first-year wage freezes are not indicative of *persistent* downward wage rigidity. For this reason, the incidence of wage freezes over the lifetime of the contract is a better indicator of rigidity. But cases of persistent rigidity are infrequent in the private sector, where only 13 per cent of contracts imposed lifetime wage freezes during the 1992-96 period.

With respect to the breakdown between the public and the private sector, if the degree of downward nominal wage rigidity were different in

the two sectors, distinguishing between the two sectors would be appropriate. This, however, does not appear to be the case. The problem with looking at the incidence of freezes in the private and the public sectors is that, without an analytical basis, their practical significance cannot be assessed properly. The simplest way to make sense of the reported incidence of wage freezes is to observe that the most important determinant of the percentage of union contracts with wage freezes is the median wage change. As a check, the following two AR(1) regressions use aggregate quarterly data on wage settlements from 1984Q4 to 1996Q3 to establish the statistical link between the percentage of first-year wage freezes (PWF) and the median wage change (DWMED) in private and public sector agreements. The results (with standard errors in parentheses) are

$$\text{PWF} = 43.6 - 6.9 * \text{DWMED}, R^2 = 0.81$$

(3.9) (1.0)

for the private sector, and

$$\text{PWF} = 52.0 - 7.4 * \text{DWMED}, R^2 = 0.87$$

(10.5) (2.1)

for the public sector.

The important fact to observe here is that the differences between the two sets of coefficients are statistically negligible (p -value = 0.69). This indicates that the resistance to nominal wage cuts is just as strong in the private sector as in the public sector. With inflation averaging 1.5 per cent in the last five years, the median wage change has fluctuated between 0 and 3 per cent, giving a “normal” range of variation of 25 to 50 per cent for the percentage of first-year wage freezes in each sector. This is the important structural result. That the re-weighted percentage of freezes for the entire 1992-96 period (37 per cent) falls exactly on the mid-point of this interval is purely a chance occurrence, due to labour-market conditions that happened to be extreme in the public sector and much less so in the private sector.

The second argument is that lifetime wage freezes are the best indicator of rigidity because they capture persistence. This is a misleading argument. We are looking for an estimate of the percentage of new and ongoing contracts imposing wage freezes in any given year. This corresponds neither to the “first-year” nor to the “lifetime” definition of a wage freeze, but to the “year-over-year” definition given by Crawford and Harrison. For example, over the three-year period 1992-94 (which is where my micro data end), private sector contracts contained 36 per cent first-year freezes, 19 per cent lifetime freezes, and 26 per cent year-over-year freezes.

Undesired wage freezes raise real wages and unemployment every time they occur, whatever their eventual duration. The basic view is that of an annual stochastic process in which the group of firms that are affected by freezes this year will be replaced by another, partly overlapping, group next year. The rigidity does not have to last several years in the same firm, it only has to threaten all firms every year. The macroeconomic significance of this Markov process of wage constraints depends on turnover and duration. To focus on duration and forget about turnover is to dismiss a large part of the problem without cause. The sense in which “persistence” might be a concern is if the stochastic process itself, involving a large turnover of short-duration freezes, were to persist if the very-low-inflation environment were maintained for a decade or more. But that is a different question.

An additional consideration is that often the only means by which firms can have workers swallow a freeze in year one of a contract is to offer a pay raise in years two or three or both. This implies that the year-over-year definition of the wage freeze will likely underestimate the true extent of the constraint imposed by downward nominal rigidity on *real* wages in the union sector. Assuming one-third of contracts with first-year freezes followed by second- or third-year increases were in this situation over the 1992-94 period, the percentage of all new and ongoing contracts that constrained real wages in any given year was in fact 30 per cent, not 26 per cent as reported above. In the non-union sector, the three measures almost coincide because most wage contracts in that sector last only one year.

To summarize, Crawford and Harrison could improve their Section 1 by conditioning the percentage of wage freezes on the median wage change. On that basis, there is no evidence that wage resistance is more prevalent in public sector contracts than in private sector contracts. With very low inflation, and wage growth falling in the 0 to 3 per cent interval, it is reasonable to estimate that the percentage of first-year wage freezes would fluctuate between 25 and 50 per cent, while the percentage of existing contracts under a real-wage constraint would range between 20 and 40 per cent. Crawford and Harrison’s focus on lifetime wage freezes is inappropriate. It leads to substantial underestimation of the extent of downward wage rigidity in the union sector.

Section 2: How important is the excess density at zero wage change?

Section 2 of the paper applies a proportional hazard model to private sector contract data from 1965 to 1996. The idea is to estimate the wage change distribution conditional on the CPI inflation rate. This generates a counterfactual experiment in which the increase in the density at zero as inflation declines should give an estimate of the rising importance of wage

rigidity. Figures 15 and 17 of the paper report that, when inflation falls from 6 per cent to 2 per cent, the proportion of wage changes that are stuck at zero in the private sector as a whole increases by 20 percentage points (from 13 to 33 per cent) in the case of first-year wage adjustments, and by 13 percentage points (from 4 to 17 per cent) in the case of lifetime wage changes. The two estimates for the 6 per cent inflation level (13 per cent first-year and 4 per cent lifetime) are surprisingly high.

The methodological innovation is interesting. Unfortunately, the experiment is vitiated by the choice of price inflation as the only covariate in the hazard model. This is a bad choice. The wage-change distribution is clearly influenced by price inflation, but also by other factors such as productivity growth, and supply and demand conditions. Omitting these factors creates a hazardous specification bias.

The inappropriate choice of price inflation as the unique hazard covariate is the most plausible source of the very high estimates obtained for the percentages of first-year and lifetime wage freezes at 6 per cent inflation. Because the data can make comparisons of densities only on the basis of differences in inflation levels, they must establish a spurious link between periods of high incidence of wage freezes, such as the mid-1980s and early 1990s, and the fact that the inflation rate then was in the range of 4 to 6 per cent. The authors do note, and seem surprised by, the sharp increases in the incidence of wage freezes during those periods.

One way to remove the specification bias, the spurious results, and the surprise would be, as I suggest above, to condition the wage change distribution directly on the median wage change. The latter statistic is a good summary measure of the joint effect of inflation, productivity growth, and circumstances of supply and demand on the aggregate labour market. Just as one would expect, the median wage growth in private sector contracts during the 1984-87 period was 3 per cent or less most of the time (using the first-year definition). Wage freezes were then quite frequent. The explanation does not require appealing to some obscure effect of inflation uncertainty. Conversely, periods of median wage growth in excess of 5 per cent, such as 1977-82 and 1988-90, are nearly exempt of first-year or lifetime wage freezes, just as predicted by the simple regressions presented above.

With these corrections in mind, the results of the proportional hazard model essentially bring us back to the descriptive analysis of Section 1. When the median wage change is about 2 per cent, the incidence of first-year and lifetime wage freezes falls, respectively, in the mid-30s and the mid-teens. When the median wage change is 6 per cent, the incidence of freezes is zero. There is no reason to question my earlier summary that when median wage growth falls in the 0 to 3 per cent interval, the percentage of

existing contracts under a real-wage constraint is likely to fluctuate between 20 and 40 per cent.

Finally, Sections 1 and 2 try to dismiss “much of the density in the neighbourhood of zero” by appealing to the menu costs of changing the wage. Presumably, the argument is that firms will not bother to give a wage increase or decrease if the amount is too small to justify the fixed cost of changing the wage scale. This argument is hard to swallow in the case of unionized firms who are already incurring the gargantuan fixed cost of negotiating collective agreements, whose payrolls are nowadays entirely computerized, and who are concerned about the morale of employees. Reporting on the industry survey conducted by the Bank of Canada in 1996, Table 9 of the paper indicates that menu costs were mentioned as a source of wage freezes in 10 per cent of cases. This is very far from “much of the density in the neighbourhood of zero.”

Section 3: How significant are cash bonuses and how rigid are non-unionized and smaller firms?

Major wage settlements data cover only 20 per cent of total paid non-agricultural employment, and only 10 per cent of private non-agricultural employment. Their information content is also restricted to base wages. It is therefore perfectly legitimate to ask whether wages are more flexible in the non-union sector and at smaller firms, and whether flexible compensation schemes could allow firms to get around the downward inflexibility of straight pay. Section 3 of the paper tries to shed light on these two questions by looking at survey evidence from four micro data sets: Statistics Canada’s Survey of Labour and Income Dynamics (SLID), the annual compensation survey of Sobeco Ernst and Young (the Sobeco survey), the member survey of the Alliance of Manufacturers and Exporters (the Alliance survey), and the Bank of Canada industry survey.

Crawford and Harrison report that 10 per cent of respondents to the SLID indicated that their hourly wage was cut by at least 10 per cent in 1993. This attributes much more importance to wage cuts than implied by the major wage settlements data. However, it is well known that in surveys such as the Canadian SLID and the U.S. PSID (the Panel Survey of Income Dynamics) the size and frequency of wage cuts are greatly exaggerated by respondents. The responses cannot be corrected unless matched to other micro data sets—such as, indeed, the major wage settlements data bank. Matching PSID responses to specific union contracts, Shea (1996) recently found that, while 21.1 per cent of the matched PSID responses reported wage cuts, only 1.3 per cent of the corresponding contracts in fact showed them. It is not clear why we should believe SLID respondents to be more reliable.

That variable pay schemes (cash bonuses) allow more flexibility in total compensation than is found in base wages is a truism. The relevant question is: How much more flexibility? Crawford and Harrison use the Sobeco survey to produce evidence on this issue. The Sobeco questionnaire is sent to 10,000 organizations of which about 50 per cent have more than 250 employees. The response rate, usually 10 per cent, was only 4 per cent in 1996. Moreover, the micro data set analysed by the authors for that year is based on a subsample of 1 per cent—that is, 25 per cent of this 4 per cent. That the sample is not more representative is worrisome.

Tables 5 and 6 infer from this 1 per cent subsample that in 1996 there were, as expected, fewer freezes and more rollbacks for the change in total compensation than for the change in base salary. However, there is no indication of how large the average rollback was, so that we do not know *how much* additional flexibility was allowed by cash bonuses. This can be partly made up by looking at the data for the full 4 per cent sample reported in and around Table 4. They imply that no more than 40 per cent of employees received cash bonuses and that, when present, the average bonus did not exceed 6 per cent of the base wage. A 25 per cent cut in the average bonus would therefore generate less than a 0.6 per cent cut in total compensation in the aggregate (since $0.25 \times 0.06 \times 0.40 = 0.006$). It therefore seems that the amount of additional compensation flexibility allowed by cash bonuses, though not negligible for many organizations, is very modest in the aggregate.

The Alliance survey provides information on base-wage changes for hourly paid workers at a sample of 111 unionized and non-unionized firms in 1996. The subsample of large firms (with more than 500 employees) contains too few observations to be reliable. The nature and extent of the sampling and response biases are otherwise unknown.

The most useful summary statistic that can be retrieved from the sample to infer the extent of downward wage rigidity is the percentage of wage freezes among non-positive wage changes. Crawford and Harrison do not report it, but it can be calculated from the data in Tables 7 or 8. For Table 7, the occurrence of wage freezes among non-positive wage changes in the 1996 Alliance survey of 100 firms is as follows:

Number of employees	Union status of firm	
	Unionized	Non-unionized
< 100	73.6	32.8
100-500	77.5	45.4

The non-union sector in the sample seems less constrained by resistance to wage cuts than does the union sector, but allowing for sampling errors, smaller firms are not obviously less constrained than medium-sized firms, for a given union status.

The Bank of Canada industry survey of 1996 investigates the number of wage freezes and rollbacks ever imposed in the past by a sample of 62 organizations. The sample is said to be representative of the major industries and regions of Canada, but no further information is given on sampling design and response pattern. One can infer from Table 9 that wage freezes retrospectively represent 71.6 per cent of reported instances of non-positive wage changes among the 43 unionized firms, and 55.2 per cent among the 19 non-unionized firms. However, it is doubtful whether these two numbers are in fact statistically different, given the small number of non-unionized organizations in the sample.

Finally, determining whether a wage freeze was instituted because the previous wage was still appropriate or because the firm was somehow constrained not to decrease the wage seems an extremely fine distinction to ask respondents to make, particularly in a retrospective survey. These results should be treated with great caution.

Overall, Section 3 is a valiant attempt to determine whether wages are more flexible in the non-union sector and at smaller firms, and whether flexible compensation schemes could allow firms to get around the downward inflexibility of straight pay. To this end, the authors examine the evidence from four micro data sets. This is a difficult task because the significance of sampling, reporting, and response biases is often hard to judge.

I would conclude from this evidence that Crawford and Harrison demonstrate, as anyone would have expected, that variable pay schemes allow greater flexibility in overall compensation. But they are not convincing that the amount of additional flexibility is very significant in the aggregate economy. They are more successful in showing that, in their data sets, base wages are more flexible in the non-union sector than in the union sector. But they do not present serious evidence that wages are more flexible in smaller than in larger firms, for a given union status.

As the authors point out, other sources of compensation flexibility could be investigated: changes in non-wage benefits, better use of organizational slack, the flow of entry and exit of firms, the changing importance of self-employment, the process of hires and layoffs across firms and sectors, and so on. Studies of these phenomena would certainly be helpful. But I doubt that bickering on the micro evidence from now on will lead us anywhere. What we need is to test the wage floor hypothesis at the macro level, where the resistance to wage cuts and all the means of getting around it, obvious and subtle, will be taken into account automatically and simultaneously.

References

- Akerlof, G. A., W. T. Dickens, and G. L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* (1): 1-76.
- Cozier, B. V., and G. Wilkinson. 1991. *Some Evidence on Hysteresis and the Costs of Disinflation in Canada*. Technical Report No. 55. Ottawa: Bank of Canada.
- Fillion, J.-F., and A. Léonard. 1997. "La courbe de Phillips au Canada : un examen de quelques hypothèses." Working Paper No. 97-3. Bank of Canada, Ottawa.
- Laxton, D., D. E. Rose, and R. Tetlow. 1993. "Is the Canadian Phillips Curve Non-Linear?" Working Paper No. 93-7. Bank of Canada, Ottawa.
- Rose, D. E. 1988. *The NAIRU in Canada: Concepts, Determinants and Estimates*. Technical Report No. 50. Ottawa: Bank of Canada.
- Shea, J. 1996. "Comments on 'Does Inflation Grease the Wheels of the Labor Market?'" by D. Card and D. Hyslop. Prepared for a conference on monetary policy and low inflation, National Bureau of Economic Research, Cambridge, Mass.