

Working Paper 93-7/Document de travail 93-7

**Is the Canadian Phillips Curve Non-linear?**

by  
Douglas Laxton, David Rose and Robert Tetlow

Bank of Canada



Banque du Canada



July 1993

## **Is the Canadian Phillips Curve Non-linear?**

by

Douglas Laxton, David Rose and Robert Tetlow

Research Department  
Bank of Canada  
Ottawa, Ontario  
Canada  
K1A 0G9

Telephone: (613) 782-8550, 782-8728, 782-8670

Fax: (613) 782-7163

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

### **Acknowledgments**

Irene Chan, Jean Xie and Chris Lavigne provided excellent assistance in this research. We thank for their comments Pierre Duguay, David Longworth, Steve Poloz and participants of a seminar at the Bank of Canada and a seminar hosted by the Board of Governors of the Federal Reserve System. Helen Meubus provided patient editorial advice that improved the exposition. Responsibility for errors remains ours.

ISSN 1192-5434

ISBN 0-662-20808-0

## Contents

Abstract .....	v
Résumé .....	vi
1 Introduction and summary .....	1
2 The model and the data .....	3
2.1 The model .....	3
2.2 The data .....	4
2.2.1 Expectations .....	5
2.2.2 Potential output and the output gap .....	7
3 Econometric results .....	11
4 Implications of the estimates for monetary policy .....	16
5 Conclusions .....	21
Appendix: Data used in the regressions .....	23
References .....	25



## Abstract

In this paper, we provide econometric evidence that suggests that the Canadian Phillips curve is non-linear. The non-linearity has two important characteristics. First, the non-linearity is such that excess supply has a smaller effect on inflation than does excess demand. Second, there are significantly longer lags from excess supply gaps to inflation than from excess demand gaps to inflation.

There are two noteworthy points concerning the data we use in our estimation. First, we have assembled a proxy for inflation expectations from the survey of inflation forecasts published by the Conference Board of Canada. Second, we use a measure of the output gap that we have constructed using a multivariate filtering technique that exploits information about inflation in the identification of potential output and the output gap. Our empirical results indicate that both innovations help in identifying the structure and parameters of the Phillips curve.

We also simulate a simple macroeconomic model containing one of the estimated Phillips curves to investigate the implications of our results, focussing on the costs of reducing inflation. Our purpose is not to give precise estimates of the costs of disinflation, but rather to caution the reader that such calculations are subject to considerable uncertainty and need to be addressed as general macroeconomic issues and not as properties of the Phillips curve alone. Our results indicate that it is relatively hard to get inflation down once it has become entrenched in expectations, and that inflation can escalate rapidly if excess demand conditions are allowed to persist. The results also indicate that the output that is foregone in the process of disinflation is very sensitive to the manner in which expectations adapt during the transition.

## Résumé

Les résultats empiriques présentés par les auteurs indiquent que la courbe de Phillips n'est pas linéaire au Canada. Cette non-linéarité revêt deux caractéristiques importantes. Premièrement, un excédent de l'offre a une incidence plus faible sur l'inflation qu'un excédent de la demande. Deuxièmement, le premier se répercute sur l'inflation avec des retards bien plus longs que le second.

Deux points méritent d'être signalés concernant les données retenues pour l'estimation. D'une part, les auteurs utilisent les données de l'enquête du Conference Board du Canada relatives aux taux d'inflation prévus pour élaborer la mesure qui leur sert à représenter les attentes au sujet de l'inflation. D'autre part, ils construisent une mesure du déséquilibre de la production à l'aide d'un filtre à plusieurs variables qui met à contribution les renseignements relatifs à l'inflation pour évaluer la production potentielle et le déséquilibre de la production. Les résultats empiriques obtenus indiquent que l'emploi de ces deux mesures aide à déterminer la structure et les paramètres de la courbe de Phillips.

Les auteurs effectuent aussi quelques simulations à l'aide d'un modèle macroéconomique simple formalisant l'une des courbes de Phillips estimées en vue d'étudier la portée des résultats, en particulier en ce qui concerne les coûts d'une réduction de l'inflation. Leur objectif n'est pas d'estimer avec précision le coût de la désinflation, mais bien de prévenir le lecteur qu'une telle estimation est entachée d'une grande incertitude et qu'elle devrait être envisagée dans un cadre d'analyse macroéconomique général et non pas uniquement dans le cadre d'analyse restreint d'une courbe de Phillips. Selon les résultats obtenus par les auteurs, il est relativement ardu de faire baisser l'inflation une fois qu'elle s'est enracinée dans l'esprit des gens, et elle peut rapidement s'accélérer si un excédent persistant de la demande est toléré. Les résultats révèlent également que le pourcentage de la production auquel il faut renoncer pour abaisser le taux d'inflation dépend beaucoup de la manière dont les attentes s'adaptent au cours de la période de transition.



## 1 Introduction and summary

In this paper, we provide some econometric evidence that suggests that the Canadian Phillips curve is non-linear. The form of non-linearity that we find has two important characteristics. First, there is asymmetry in the functional form such that excess supply has a smaller effect on inflation than does excess demand. Second, there are significantly longer lags from excess supply gaps to inflation than from excess demand gaps to inflation.

Relative to other recent research, there are two notable features in our work. First, instead of relying on distributed lags to represent inflation expectations, we use survey data on inflation forecasts as a proxy for expectations. Because the survey data are not available on a quarterly basis, we estimate the Phillips curve using annual data. Second, we use measures of potential output derived from a multivariate filter that we have developed for this purpose (Laxton and Tetlow 1992). Our results show that the proxy for inflation expectations improves markedly the properties of the estimated model, relative to an alternative model limited to using lags of inflation to represent expectations. The results also suggest that there is a gain in estimation efficiency from using the output gaps from the multivariate filter, consistent with Monte Carlo evidence reported in Laxton, Rose and Tetlow (1993a).

To elaborate the implications of our estimation results, we simulate a small macro model containing one of our estimated, non-linear Phillips curves, using a range of assumptions about the degree of inertia in expectations. The results are shown to be quite sensitive to the specific assumptions made about the degree of rigidity in expectations. For example, over the range considered, these simulations indicate that the cumulative output loss necessary to reduce inflation by 1 percentage point, beginning from a steady state, may vary from close to 0 to just over 5 per cent.<sup>1</sup> However, we argue, based on the empirical properties of our proxy for inflation expectations, that the range of interest is 1.9 to 3.5 per cent. On the other hand, based on the latter assumptions, a cumulative excess demand gap of just 0.8 to 1.2 per cent will cause inflation to rise by one percentage point. These

---

1. This encompasses many of the estimates in the literature, for example Cozier and Wilkinson's (1990) 2 per cent and Howitt's (1990) 4.7 per cent. The model does not exhibit hysteresis, which some researchers, for example Fortin (1991), argue exists in Canadian labour markets. In models with hysteresis, reducing inflation results in permanent losses of output (unbounded cumulative output losses).

results indicate that it is relatively hard to get inflation down once it has become entrenched in expectations, and that inflation can escalate rapidly if excess demand conditions are allowed to persist.

We do not investigate other aspects of the sensitivity of the results to, for example, the calibration of the model or the assumptions concerning the way the policy is implemented. Our purpose is not to give precise estimates of the costs of disinflation, but rather to caution the reader that such calculations are subject to considerable uncertainty and need to be addressed as general macroeconomic issues and not as properties of the Phillips curve alone.<sup>2</sup>

The remainder of this paper is organized as follows. The second section discusses the model we estimate and the data. The third section describes the estimation results. In section 4, we consider the implications of our estimation results, focussing primarily on the implied costs of reducing the rate of inflation. In this discussion, we report the simulations described above. The final section provides some concluding comments.

---

2. It is important to note that our simple model does not include links between the rate of inflation and the level of productivity. The measures of cumulative gaps are relative to potential output. If potential output is lowered as inflation rises, as indicated by some recent research (for example, Novin 1991, Cozier and Selody 1992), then one cannot draw overall welfare conclusions from the type of partial experiments we report here.

## 2 The model and the data

### 2.1 The model

The equations estimated for this paper are all relatively standard “price” Phillips curves. There are two variants. The basic model is

$$\Pi_t = \Pi_{t+1}^e + \alpha \cdot YGAP_t. \quad (1)$$

Inflation is a function of expected inflation and the output gap. Expectations enter as in Rotemberg (1982), Cozier (1989) and Lebow, Roberts and Stockton (1992). In these models, current inflation is influenced by where inflation is expected to go in the future. Reaction to shocks is not instantaneous because of costs of adjustment. However, all such intrinsic dynamics are captured by the output gap, and lagged inflation effects are interpreted as part of expectations formation.

The second variant adds another element to intrinsic dynamics, similar to the discussion in Buiter and Miller (1985):

$$\Pi_t = \lambda \cdot \Pi_{t+1}^e + (1 - \lambda) \cdot \Pi_{t-1} + \alpha \cdot YGAP_t. \quad (2)$$

The basic model is extended to include an autoregressive component to inflation dynamics. In the Buiter-Miller model this is not part of expectations. It could be the result, for example, of asynchronous price adjustment in two markets. It could arise from contracts or any other rigidity in the adjustment process. The weight on the backward-looking component then depends on how important these things are—how much asynchronization there is in price adjustment, for example. Note that when such “intrinsic” dynamics are part of the model, there will be inertia in the adjustment process and it will be necessary to suffer an output loss in order to reduce inflation, even if expectations are fully consistent with the predictions of the model. Expectational rigidities provide an additional source of stickiness in the price adjustment process in this model. In the Buiter-Miller discussion, expectations themselves have backward- and forward-looking components.<sup>3</sup>

---

3. We prefer the Buiter-Miller interpretation. However, for readers who prefer to think of expectations as the only source of inertia in the process, an alternative interpretation can be given to equation (2), wherein there are components of expectations and  $\Pi^e$  represents, for example, a model-consistent forecast. One can then think of the overall representation as a weighting scheme reflecting the proportion of agents who form expectations in this way.

Although we show a contemporaneous output gap in equations (1) and (2), we experiment with various specifications of the timing of the effects of output gaps on inflation. In particular, many of the estimations include a lagged gap. The estimated equations also always include a freely estimated constant. The constant has no place in the model, from a theoretical perspective, but it allows us to control for level measurement errors, particularly with respect to the output gaps, and for the mean effect of omitted variables.

The tests for non-linearity involve adding to the above models either the positive gaps, such that there is a kink at zero excess demand (but the model remains locally linear elsewhere), or the squared positive gaps, such that the function has additional quadratic curvature in the region of excess demand. The kinked version involves adding to the basic model a variable called POSGAP, defined as

$$\begin{aligned} \text{POSGAP} &= \text{YGAP}, \text{ for } \text{YGAP} > 0 \\ &= 0, \text{ otherwise.} \end{aligned} \tag{3}$$

The quadratic version involves adding a variable called SQPOSGAP, defined as

$$\text{SQPOSGAP} = (\text{POSGAP})^2. \tag{4}$$

## 2.2 The data

The estimation of a Phillips curve of the kind described above is made doubly difficult by the absence of direct observations on potential output (and, hence, the output gap) or on inflation expectations. We have shown elsewhere that difficulties in identifying potential output can create bias in standard OLS estimates and severe problems of statistical inference.<sup>4</sup> It is likely that errors in representing expectations would create similar empirical complications. Our work here is subject to these problems, but we have made an effort to minimize them. We use a direct proxy for expectations of inflation and measures of the output gap coming from our work on a multivariate filter, which attempts to exploit information on inflation, as well as output itself, in identifying potential output.<sup>5</sup>

---

4. See Laxton, Shoom and Tetlow (1992), and Laxton, Rose and Tetlow (1993a).

5. See Laxton and Tetlow (1992).

### 2.2.1 Expectations

It seems generally accepted that we would be better able to test models of inflation dynamics if we had actual measures of the inflation expectations. Recent work at the Board of Governors of the Federal Reserve (Lebow, Roberts and Stockton 1992) has reported some gains from using the Livingston survey data as an expectations proxy. The only comparable measure that exists for Canada is a collection of forecasts of inflation published by the Conference Board of Canada. Unfortunately, it is not possible to derive a consistent quarterly time series. However, it is possible to get an annual measure. We have assembled one-year-ahead forecasts of the rate of inflation (for the GDP deflator) from the autumn publications of the Conference Board's *Canadian Business Review*. The measure is the average of the forecasts reported. These forecasts were based on information available in the early summer of the previous year. That is, the forecasts of inflation in 1989, say, were based on information that was available in early summer in 1988, including, at best, the first-quarter national accounts for 1988. The forecasts are available starting in 1975; we have constructed the series through to 1991 for this paper.

The series is shown in Figure 1 as PDOTE.<sup>6</sup> The series is shown lagged once, relative to the form used in the regressions, so that we can show it as a forecast of PDOT, the actual outcome. The series around the zero line shows the implied forecast errors. Evidently, forecasters have made some large and systematic errors in predicting inflation. Nevertheless, the series seems to compare favourably as a forecast to the results from a standard static model, where inflation is expected to be next year what it was last year. A similar picture for the static model is shown in Figure 2. In Figure 3, we compare the forecast errors of the two measures of expectations. Note, in particular, that the forecasters do systematically better than the static model when inflation is changing rapidly. They still make large errors, but there seems to be adaptation to information over and above what the static model would provide.<sup>7</sup>

The predictions of inflation from the Conference Board's survey of forecasters may not be a precise representation of the expectations held by the public at large, or rather, by those agents involved in the process determining the actual outcome


---

6. The data are also listed in the Appendix.

7. Because of the lags involved in an annual comparison, this argument is not as convincing as it would be if higher frequency information were used in the static model. Unfortunately, it is not possible to construct a consistent historical series on a quarterly basis.


**Figure 1**  
**The forecasters' expectations and errors**

(all variables measured as percentages)



for inflation. We have no direct way of telling. Nevertheless, from an econometric perspective, it may still be preferable to use these data as a proxy or instrument variable in estimation. The econometric and inference difficulties of dealing with specifications with lagged dependent variables give reason to prefer a direct proxy measure. For our purposes here, the reader does not have to believe that the forecasts are "good" measures, in some absolute sense, of expectations. To find the exercise interesting, however, the reader will have to accept that the proxy appears to have reasonable properties as an instrument variable.


**Figure 2**  
**Static expectations and errors**  
 (all variables measured as percentages)



### 2.2.2 Potential output and the output gap

Our main output gap measures are generated using the multivariate filtering technique described in Laxton and Tetlow (1992). The basic idea of the multivariate filter is that when an unobservable variable is thought to influence two or more observable variables, information from all the observable variables may be useful in identifying the unobservable value. In the case of potential output, it has been common practice to fit trend lines through the output data, or to use a production

**Figure 3**  
**Comparison of expectations errors**  
 (all variables measured as percentages)



function with a univariate detrending technique applied to its components, such as total factor productivity. In this case, the idea of the multivariate filter is to recognize that if one thinks that the state of excess demand has an influence on the course of inflation, then one may be able to use information about what is happening to inflation to help infer the level of potential output. In this particular application of these ideas, we use a quarterly, linear Phillips curve, based on the results reported by Cozier and Wilkinson (1990), along with an Okun's Law relationship and a production function, in a system similar to that described by Ford



and Rose (1989), to condition the estimates of potential output.<sup>8</sup> In the version used here, we estimate the NAIRU and use the results to infer the output gap via the Okun's Law relationship.<sup>9</sup> The annual values used for this paper are the simple averages of our quarterly estimates.

As a check on the sensitivity of the results to the methodology chosen to gauge the output gap, we have also estimated the model using a gap measure derived from an application of a univariate technique, the Hodrick-Prescott (H-P) (1980) filter, applied to output directly to represent potential output. We use two variants of this filter. One is an unconstrained application of the H-P methodology. However, the H-P filter, like any two-sided filter, has difficulty dealing with data at the end of a sample, especially when the end of the sample is an extreme point in the business cycle. The same problem affects the multivariate filter, albeit to a lesser extent, and our standard application of the multivariate technique allows judgment to be applied to the estimates. To assure comparability, we also prepared a version of the H-P gap constrained to be the same as the multivariate gap in 1992—the period following the sample we use for estimation.


The three output gap measures are shown in Figure 4 and listed in the Appendix. In the Appendix, the gap from the multivariate filter is called YGAP, while the raw and adjusted H-P measures are called HPGAP and HPGAPA, respectively. The multivariate filter shows greater excess demand in the early 1970s and greater excess supply in the recessions of 1981-82 and 1990-91. It shows smaller excess demand in the mid-to-late 1980s.

---

8. There is an inconsistency here that requires comment. We have used a linear Phillips curve to derive measures of potential output that are then used to identify a non-linear Phillips curve. A second round could be implemented to refine the estimates, but the free constant in the estimated Phillips curve does allow for simple level errors in the measure of potential output. Iteration would produce a proximate level shift in the estimates of potential output, so we do not think that the basic results in this paper with respect to the non-linearity would change if we were to recalculate the gap measures.

9. This research is based on work completed at an early stage in the development of a new macroeconomic model for projection and policy analysis at the Bank of Canada. We intend to update this work on potential output as part of the modelling of the supply side of the economy. In particular, more information will be used to condition the measures from the multivariate filter. The current measures of the gap should be treated as preliminary. Nevertheless, we see no reason to think that the new measures would change the character of the results in this study.

**Figure 4**  
**The output gap measures**  
(all variables measured as percentages)



### 3 Econometric results

Some econometric results are assembled in Tables 1 and 2. The first eight regressions use the output gap measures derived from the multivariate filter; the last two use the measures from the H-P filter. We have reported only the results with a unit-sum restriction imposed on the coefficients of PDOTE and the lag(s) of PDOT. In no case where PDOTE is included is this restriction close to being rejected based on the  $F$  test (reported at the bottom of the tables).<sup>10</sup> Moreover, for these cases the parameter estimates under the restriction are always very close to the unrestricted estimates.<sup>11</sup>

The PDOTE variable makes an important contribution to the explanatory power of the equation. Formulations with only autoregressive expectations do not fit as well. Compare, for example, regressions (5) and (8). The residuals are also cleaner when PDOTE is included in the regression. This pattern carries over to all the reported results, including regressions (9) and (10), where the H-P filter measures of the output gap are used. Note that the PDOTE variable is consistently given greater weight than the lagged values of PDOT in models with both components. Nevertheless, it is clear that one would want to retain both influences in explaining PDOT.<sup>12</sup>

Consider regressions (1) and (2) in Table 1. Note that the lagged gap seems to be preferred in regression (1). In regression (2), where we add the squared positive gaps (SQPOSGAP) as a separate variable to test for non-linearity in the effect of output gaps on inflation, the regression strongly prefers the SQPOSGAP variable to the contemporaneous YGAP variable. In fact, the sign of the estimated YGAP coefficient is "wrong." For regression (4), we drop the contemporaneous YGAP variable. In regression (4), the estimated coefficient on SQPOSGAP is statistically

---

10. In some cases where PDOTE is not included, the  $F$  test rejects the restriction on the lags of PDOT. To pursue this issue substantively, it would be necessary to report other results, such as Dickey-Fuller tests. We have not done so, since this question is not our concern here. All results we reference in the text have the restriction imposed.

11. The same cannot be said for the models with only lagged dependent variables and no PDOTE. There is no economic necessity to impose the restriction in these regressions. Expectations do not have to be formulated as if there were a unit root for the accelerationist property to be respected. However, for comparability with other work, we impose the restriction.

12. Whether the lag is useful because the economic forecasters are more forward-looking than average market agents or whether it reflects other intrinsic influences on inflation dynamics, such as contract lags, is an interesting issue. In our application of these results in the next section, we interpret the lag as an intrinsic-dynamic effect and not part of expectations.

**Table 1**  
**Regression results**

Dependent variable: PDOT (1975-91)			Model		
Regressors	(1)	(2)	(3)	(4)	(5)
Constant	-0.230 (0.261)	-0.622 (0.300)	-0.781 (0.357)	-0.614 (0.272)	-0.628 (0.287)
PDOTE	0.781 (0.150)	0.691 (0.140)	0.661 (0.140)	0.686 (0.124)	0.666 (0.149)
PDOT(-1)	0.219 (0.150)	0.309 (0.140)	0.339 (0.140)	0.314 (0.124)	0.293 (0.149)
PDOT(-2)					0.041 (0.152)
YGAP	0.074 (0.151)	-0.012 (0.133)			
YGAP(-1)	0.217 (0.152)	0.194 (0.135)	0.168 (0.134)	0.191 (0.125)	0.236 (0.198)
POSGAP			1.023 (0.538)		
SQPOSGAP		0.606 (0.293)		0.599 (0.269)	0.608 (0.281)
RBARSQ	0.89	0.91	0.91	0.92	0.91
p-value, <i>F</i> test (unit-sum restriction)	0.35	0.99	0.90	0.97	0.95
p-value, <i>Q</i> test	0.81	0.72	0.66	0.88	0.57
Durbin-Watson statistic	2.10	2.25	2.23	2.25	2.20

Note: The figures in parentheses are the estimated standard errors for the parameters.

**Table 2**  
**Regression results**

Dependent variable: PDOT (1975-91)		Model			
Regressors	(6)	(7)	(8)	(9)	(10)
Constant	-0.357 (0.425)	-1.652 (0.431)	-1.213 (0.402)	-0.770 (0.332)	-0.577 (0.302)
PDOTE				0.745 (0.127)	0.791 (0.128)
PDOT(-1)	0.626 (0.241)	0.629 (0.190)	0.617 (0.206)	0.255 (0.127)	0.209 (0.128)
PDOT(-2)	0.374 (0.241)	0.371 (0.190)	0.383 (0.206)		
YGAP	0.426 (0.213)				
YGAP(-1)	0.608 (0.362)	0.544 (0.282)	0.641 (0.298)	0.106 (0.163)	0.131 (0.170)
POSGAP		2.330 (0.633)		0.672 (0.357)	
SQPOSGAP			1.207 (0.388)		0.251 (0.176)
RBARSQ	0.70	0.81	0.78	0.90	0.89
p-value, F test (unit-sum restriction)	0.06	0.02	0.08	0.30	0.25
p-value, Q test	0.60	0.44	0.57	0.59	0.77
Durbin-Watson statistic	1.27	1.56	1.29	1.92	1.91

Note: The figures in parentheses are the estimated standard errors for the parameters. For models 9 and 10, the gap variables are the H-P measures, with end-of-sample judgment applied.

significantly different from zero at the 95 per cent confidence level based on the standard classical  $t$  test. In regression (2), the coefficient is not quite significant based on this test, but it is significant at the 93 per cent confidence level. Given the evidence in Laxton, Rose and Tetlow (1993a) that these tests are biased against identifying such a non-linearity, we interpret these results as providing strong support for an important asymmetry in the Canadian Phillips curve. Note that this asymmetry has two dimensions. The results indicate that a given level of excess demand creates more powerful pressures pushing inflation higher than the equivalent level of excess supply creates in the opposite direction. The results also indicate that excess demand acts *faster* to push inflation higher than excess supply acts to push it lower. The timing dimension is important because it implies that a monetary authority will have to act relatively quickly in the face of demand pressures to avoid an escalation of inflation.

Regression (3) tries a different form of non-linearity—a kink at zero excess demand. The kinked function is locally linear, except at the point of zero excess demand (and hence it is easy to deal with analytically), but it has a very important global asymmetry. Although the coefficient of the POSGAP variable is not quite as well determined in regression (3) as is that of the SQPOSGAP variable in regressions (2) and (4), the results are nevertheless quite similar in all respects and support the case for a non-linearity. We can add that we found the same conclusions with respect to the *timing* of the YGAP variable in regressions using the POSGAP specification.<sup>13</sup>

Regression (5) illustrates that the basic finding persists if we add an additional lag of the dependent variable. When PDOTE is included, only one lag of PDOT is required. Regressions (6) to (8) show what happens if we omit the PDOTE variable and use only lags of the dependent variable. Note that a second lag of PDOT now seems to be required. As indicated above, we cannot get as good a fit without PDOTE, especially in the linear model in regression (6). The residual statistics give an indication that something is indeed missing from these regressions. The contemporaneous gap becomes more important as an explanatory variable, compared to the results in Table 1, in the linear model in regression (6). This is perhaps to be expected, since it is likely that PDOTE contains an influence of estimated

---

13. We also estimated a non-linear function with a parameter for the power coefficient on the gap. The results indicated that the best fits arise with *greater-than-quadratic* curvature. However, the estimated curvature is not well determined. This is not surprising, given the limited number of observations.

output gaps or the equivalent. However, regressions (7) and (8) show that either form of non-linearity still improves the results markedly and removes any case for retaining the contemporaneous YGAP variable in this model. Thus, our conclusion on the relative timing of response to excess demand and excess supply are confirmed in the regressions without PDOTE.

Regressions (9) and (10) replicate regressions (3) and (4) with the gap measure we obtain using the Hodrick-Prescott filter to measure potential output. The reported regressions use the version with judgment applied at the end of the sample, the variable listed as HPGAPA in the Appendix table. The results are not greatly different from those based on the multivariate filter. In particular, we find the same kind of results with respect to the role of PDOTE. However, the case for the non-linearity is not as clear. Although we have not included the results for the purely linear model with the H-P gaps, we can assure the reader that it is still true that the non-linear form fits the data better. But the estimated coefficients on the variables embodying the non-linearity are not statistically significant at the usual classical confidence levels. We repeat, however, that our Monte Carlo results show that if the true structure is non-linear with quadratic curvature, then this type of test will be unreliable—an econometrician relying on classical hypothesis testing will tend to reject non-linearity that is truly there in the process generating the data.<sup>14</sup> It is noteworthy that the Monte Carlo evidence also suggests that this problem is more severe if an H-P filter is used to measure potential output than if the multivariate filter is used for that purpose. Our econometric results exhibit the same patterns we saw in the Monte Carlo study. Thus, our use of the multivariate filter to measure potential output may be one reason why we find a significant asymmetry in the Phillips curve where others (for example, Cozier and Wilkinson 1990) do not.

We also computed these regressions using the raw H-P gaps (constructed without the end-of-sample judgment). The fits were slightly worse and the coefficients were slightly more poorly determined than the results reported in Table 2. However, the results were essentially the same in terms of the general conclusions. The differences in the results for the two filters (that is, multivariate filter versus H-P filter) do not seem to be sensitive to the treatment of the end-of-sample observations in this particular case.

---

14. See Laxton, Rose and Tetlow (1993a).

#### 4 Implications of the estimates for monetary policy

It has become common for researchers to consider the costs of reducing inflation that are implied by a Phillips curve. Often, this is presented in terms of the so-called "sacrifice" ratio, the cumulative output gap necessary to reduce inflation permanently by one percentage point. The notion that there is a fixed sacrifice ratio, independent of the path chosen by policy makers, is limited to the class of models known as integral-gap models. If there is a unit root in the inflation process, and if expectations are entirely backward-looking, with the unit-sum property reflected in the coefficients on lagged inflation terms (or, more generally, all terms reflecting the rate of inflation),<sup>15</sup> if the Phillips curve is linear and if potential output is independent of the cycle, then it follows that the pace of inflation reduction is immaterial to the output cost of reducing inflation. One can compute "the" sacrifice ratio. However, if any of these strong assumptions are relaxed, there is no such thing as a fixed sacrifice ratio. The output costs of reducing inflation depend, in general, on the speed of the path chosen by the monetary authority, on how expectations adapt to the change in policy, as well as on the functional form of the Phillips curve.<sup>16</sup>

Thus, in general, one must address questions concerning the output costs of reducing inflation in the context of a complete macro model. This is what we do in this section. We first provide some evidence and arguments as to how PDOTE should be modelled for the purposes of counterfactual policy simulations. We then describe a very simple macroeconomic model that is sufficient to enable us to create an interesting policy experiment. We stress that these results are offered for illustrative purposes. We investigate only one dimension of possible variation in the conclusions—the sensitivity of the results to the assumptions made about the degree of rigidity in expectations.

Since we have a historical proxy for inflation expectations, we need to address the issue of how expectations should be characterized as responding to a shock. We could, of course, resolve the model under the assumptions of the integral-gap

---

15. This is much stronger than assuming the original "accelerationist" restriction.

16. We remind the reader of the point made in footnote 2, that these calculations ignore any effect of the rate of inflation on the level of potential output or other real benefits from a lower rate of inflation. This issue must be addressed in an assessment of the welfare implications of a choice to move to a lower rate of inflation. The model we use here does not incorporate any benefits of lower inflation or costs of higher inflation. See Selody (1990) and Cozier and Selody (1992) for a review of the issues and some evidence on the benefits of lower inflation.



model, that is, we could assume that the expected rate of inflation is given by a distributed lag on past inflation, with a unit-sum restriction. However, we have already argued that this is not an acceptable view, given the empirical results. The proxy contains information not adequately captured by a purely autoregressive formulation. Moreover, even under these assumptions for expectations, our non-linear Phillips curve would not permit a globally constant sacrifice ratio.

For expectations, a key question is how much forward-looking behaviour is suggested by our historical proxy variable. If we estimate a simple regression of PDOTE on both lagged and actual future inflation, we get about equal weights on the backward- and forward-looking components. Using actual future inflation for the forward component makes the strong assumption of perfect foresight. We wanted to get closer to a model-consistent forecast for the forward component. To do this, we estimated a model for PDOT with two autoregressive lags and the lagged output gap. We used this model to generate one-year-ahead forecasts of PDOT. This variable, PDOTF, we then used as an instrument variable for the forward component of PDOTE. We regressed PDOTE on a lag of PDOT, for the backward-looking component, and the instrument, under the restriction that the coefficients sum to unity. This yielded a point estimate of 59 per cent weight on the forward component and 41 per cent weight on the lag:

$$\text{PDOTE} = 0.41 \text{ PDOT}(-1) + 0.59 \text{ PDOTF}; \quad \text{RBARSQ} = 0.75, \text{DW} = 0.72.$$

(4.3)                      (6.3)

The figures in brackets are *t* statistics; the 95 per cent confidence interval puts the weight on the forward-looking component in the interval 40 per cent to 80 per cent. However, since the equation used to generate PDOTF has a lag component, we may be getting too high a weight on the forward component of expectations for the purposes of counterfactual simulation. We would consider cases with a weight between 50 and 75 per cent on the backward component to be roughly what is indicated by our regression results for PDOTE.

To illustrate the importance of the assumptions about expectations, we have simulated the Phillips curve from regression (4),<sup>17</sup> embedded in a simple macroeconomic model, varying the weight on the backward component of expectations over the full range of 0.0 to 1.0. In these simulations, PDOTF becomes the one-

17. For this exercise we drop the estimated constant. The shock-control results would be identical if we were to retain it, but this would necessitate re-interpretation of the level of potential output in the simulation model.

year-ahead forecasts consistent with the model (and therefore includes the effects of the non-linearity). Note that the discussion here is about the weights for the components of expectations, but this is not the only source of autoregressive dynamics for inflation. Even in the extreme case of fully model-consistent expectations, we still have the intrinsic dynamics coming from the lagged PDOT term in the estimated Phillips curve.

The shocks are a reduction in the target inflation rate of 1 percentage point and an increase in the target inflation rate of 1 percentage point. Given a non-linearity in the Phillips curve, the sacrifice ratio will depend on the initial conditions of the experiment. Our simulations assume that the economy is initially in steady state with zero excess demand and a constant rate of inflation. We then compute the cumulative excess supply gap necessary to reduce inflation by 1 percentage point and label this the TOL (temporary-output-loss) or “sacrifice” ratio. We also compute the cumulative excess demand gap sufficient to raise inflation by 1 percentage point. We label this the TOG (temporary-output-gain) ratio; it will be lower in absolute value than the sacrifice ratio when there is a non-linearity of the type that we have identified in the Phillips curve.

To do these simulations, we need a monetary rule to give the desired pace of the adjustment and some link between the policy variable and excess demand.

The policy *effect* is provided by the following simple equation:

$$YGAP = 0.61 * YGAP(-1) - 0.98 * RGAP(-1), \quad (5)$$

where RGAP is the policy instrument. It is written as a deviation from some underlying level control, since we are interested in only shock-minus-control effects here. For our purposes here it does not matter what the instrument is, but the reader can think of it as reflecting policy influences on short-term interest rates. The calibration of this equation is based on estimation work we have done to calibrate equations of a similar nature for use with a more elaborate macro model.<sup>18</sup> There are two important features of this equation for the simulations we report here: policy affects output with a lag, and the output gap itself has autoregressive dynamics.

---

18. In these estimations, we have used the slope of the term structure, short-term rates relative to long-term rates, expressed as a “gap” by measuring this difference relative to a trend value.

The reaction function that determines the settings of the policy instrument is

$$RGAP = PGAP(1) + 0.5 * PGAP(2) + 0.25 * PGAP(3) + 0.5 * YGAP + 0.5 * RGAP(-1), \quad (6)$$

where PGAP is the difference between the actual inflation rate and the rate targeted by the monetary authority. Note that this control rule is forward-looking. The monetary authority acts based on what it expects to happen to inflation, relative to target, over the next three years. The current output gap is also used as a signal of inflationary pressure. Finally, the lagged RGAP term acts as a way of introducing an interest-rate smoothing property—it prevents short-term interest rates from moving too much in any one period. This rule describes policy reaction sufficient to bring inflation to the new target in three to five years, depending on how much inertia there is in inflation dynamics.<sup>19</sup>

The simulation results are shown in Table 3. We report results for five assumptions concerning the degree of forward-looking content in inflation expectations. The extreme value of 1.00 corresponds to the entirely backward-looking world, where agents give no weight to predictions about the likely future course of inflation and react in a static manner to its evolution, in this case treating inflation as if it were a random walk. The cumulative loss of output when the rate of inflation is reduced

Table 3 Simulation results					
	Weight applied to lagged inflation in expectations formation				
	1.00	0.75	0.50	0.25	0.00
Cumulative output loss from reducing inflation by 1 percentage point (TOL or "sacrifice" ratio)	5.2	3.5	1.9	0.7	0.2
Cumulative output gain from raising inflation by 1 percentage point (TOG ratio)	1.7	1.2	0.8	0.5	0.2
Average ratio	3.5	2.3	1.4	0.6	0.2
TOL ratio/TOG ratio	3.1	2.9	2.4	1.4	1.0

19. This policy rule is used and explained in greater detail in Laxton, Rose and Tetlow (1993b).

by 1 percentage point in this world, the TOL or "sacrifice" ratio is 5.2 per cent. The cumulative rise in output when the shock is to increase inflation by 1 percentage point is much smaller, 1.7 per cent, because of the faster and stronger response of inflation to excess demand. At the other extreme, where expectations are entirely forward-looking and fully consistent with the predictions of the model, there is not much change in output associated with the policy initiative. We do not consider either extreme case to be of practical interest.

The more interesting results are those of the intermediate cases. For the case of 75 per cent weight on the lag, we find a cumulative output loss of 3.5 per cent when inflation is reduced by 1 percentage point. The corresponding gain when inflation is increased by 1 percentage point is much less—1.2 per cent. For the case of 50-50 weighting, the individual numbers are 1.9 and 0.8 per cent, respectively, with an average of 1.4 per cent. The sacrifice ratio in this case is similar to that obtained by Cozier and Wilkinson (1990) using quarterly data and a linear model with entirely backward-looking expectations. While the evidence on the forward weight revealed by our PDOTE measure supports this result, we repeat our caution that there is reason to suspect that this puts too high a weight on the model-consistent component of expectations and therefore risks understating the output costs of disinflation. However, our main message remains that there is considerable uncertainty on this point. Further research into the determinants of PDOTE would be a promising next step in understanding better the process of disinflation.

Our estimation results indicate that inflation will escalate rapidly in the face of excess demand but will be more persistent in the face of excess supply. In itself, this will tend to make inflation difficult to reduce. However, we have argued that the costs of disinflation cannot be determined from the Phillips curve alone. We have illustrated with a simple macro model how sensitive the answer is to assumptions regarding the adaptation of expectations during the adjustment process. We find that a modest weighting on model consistency in expectations, a weighting defensible based on the properties of our data for inflation expectations, changes the results markedly from what one obtains with a static model of expectations.

Other aspects of the calibration and, in particular, the assumptions about how aggressively the monetary authority pursues the new inflation target, are likely to have an important effect on the calculations as well. We repeat our caution that not too much should be read into the precise numbers we report.

## 5 Conclusions

We have presented econometric evidence that supports the case for a non-linear response of inflation to excess demand.<sup>20</sup> The results indicate that inflation picks up faster and more strongly in the face of excess demand than it falls in the face of excess supply. Both points are important. The asymmetry in the functional form provides the stronger response to excess demand. The second point concerns the timing of response of inflation to the output gap. The results indicate a clear preference for a formulation with longer lags for the response to excess supply than to excess demand. The response to excess demand begins within the year—contemporaneously in the estimated model. The main response to excess supply comes with an important delay—a one-year lag in the estimated model.

In investigating the sensitivity of our conclusions to the approach to modelling expectations, we found that the results appear to be fairly robust. While the results obtained when we use the proxy measure of expectations are clearer, the case for non-linearity does not seem to depend critically on this choice. The result does seem to depend somewhat on the methodology chosen to measure the output gap, in that researchers using classical tests might reject non-linearity based on the results with the gaps derived using the Hodrick-Prescott filter. With our measures, based on the multivariate filter, the evidence for non-linearity is much clearer. Interpreted in the light of the Monte Carlo evidence reported in Laxton, Rose and Tetlow (1993a), which indicates that these tests are biased towards statistical rejection of non-linearity, our estimation results provide a coherent case in favour of a non-linear specification for the Canadian Phillips curve.

It would be interesting to check if the choice of an annual frequency is important to the result; this might explain why our results differ from those of Cozier and Wilkinson (1990), who rejected non-linearity, based on a different type of test, in a quarterly specification. Unfortunately, we cannot do this and retain our proxy for expectations. It might also be of some interest to introduce other common explanatory variables, such as oil prices.<sup>21</sup> However, given the small sample size, such

---

20. McCallum (1988) also finds some evidence in favour of a non-linear specification for a wage Phillips curve.

21. However, to the extent that such variables have an effect on potential output as well as on prices directly, even if these variables reduced the explanatory power of the non-linearity, it would be necessary to ask whether the effect on potential had been isolated.

additional specification searches would have to be quite limited, if the apparent gains from using the expectations proxy are to be retained.

We find the results from using the forecasters' expectations interesting. While we have not pursued the matter in detail here, it seems clear that publicly available forecasts contain more than autoregressive information, and, more importantly, that they respond more quickly to changing circumstances than permitted by fixed parameter models. We presented some evidence in support of this view from a regression designed to test for forward-looking content in our proxy measure of inflation expectations. To extend this inquiry and look more closely at what has influenced expectations historically, as revealed by this proxy, would be an interesting area for future research.

We have provided some illustrative simulations to indicate the implications of our estimated non-linear Phillips curve for the analysis of the effects on output of a decision to change the target inflation rate. We showed that the answer is very sensitive to the assumption made about the degree of rigidity in expectations formation. This point, though widely appreciated in theory, has not received the attention it deserves in much of the recent discussion of monetary policy in Canada, owing to the highly simplified framework in which the question of the costs of disinflation has been addressed.

**Appendix**  
**Data used in the regressions**

Year	PDOT	PDOTE	YGAP	HPGAP	HPGAPA
1973	8.84	NA	3.14	1.51	1.51
1974	14.42	NA	1.58	0.83	0.83
1975	9.90	9.10	-1.57	-1.30	-1.30
1976	8.70	7.30	-0.47	0.28	0.28
1977	6.25	6.90	-0.71	-0.19	-0.20
1978	6.03	7.10	0.88	0.72	0.71
1979	10.02	8.80	1.96	1.53	1.52
1980	10.61	9.90	0.53	0.57	0.55
1981	10.83	10.80	1.20	2.07	2.07
1982	8.69	8.70	-4.65	-3.35	-3.33
1983	4.99	5.30	-3.87	-2.98	-2.90
1984	3.13	5.20	-0.50	-0.31	-0.16
1985	2.58	4.40	0.82	0.45	0.70
1986	2.40	4.00	0.27	-0.22	0.10
1987	4.70	4.50	0.32	0.19	0.50
1988	4.65	4.30	0.85	2.05	2.10
1989	4.84	4.50	0.38	2.26	1.67
1990	3.26	4.80	-1.59	0.61	-1.16
1991	2.72	2.80	-5.44	-1.68	-5.24





## References

- Buiter, Willem and Marcus Miller. 1985. "Costs and Benefits of an Anti-inflationary Policy: Questions and Issues." In *Inflation and Unemployment: Theory, Experience and Policy*, edited by V. Argy and J. Nevile, 11-38. London: George Allen and Unwin.
- Chadha, Bankim, Paul R. Masson and Guy Meredith. 1992. "Models of Inflation and the Costs of Disinflation." *IMF Staff Papers* 39 (June): 395-431.
- Cozier, Barry V. 1989. "On the Aggregate Implications of Optimal Price Adjustment." Working Paper 89-4. Bank of Canada, Ottawa.
- Cozier, Barry, and Gordon Wilkinson. 1990. "How Large are the Costs of Disinflation in Canada?" Working Paper 90-6. Bank of Canada, Ottawa.
- Cozier, Barry, and Jack Selody. 1992. "Inflation and Macroeconomic Performance: Some Cross-Country Evidence." Working Paper 92-6. Bank of Canada, Ottawa.
- Dupasquier, Chantal and Nathalie Girouard. 1993. "La courbe de Phillips : nouvelles considérations empiriques." Forthcoming working paper. Bank of Canada, Ottawa.
- Ford, Robert and David E. Rose. 1989. "Estimates of the NAIRU Using an Extended Okun's Law." Working Paper 89-3. Bank of Canada, Ottawa.
- Fortin, Pierre. 1991. "The Phillips Curve, Macroeconomic Policy, and the Welfare of Canadians." Innis Lecture, 1991 Meetings of the Canadian Economics Association. *Canadian Journal of Economics* 24: 774-803.
- Hodrick, Robert and Edward C. Prescott. 1980. "Postwar U.S. Business Cycles: An Empirical Investigation." Working Paper No. 451. Carnegie-Mellon University.
- Howitt, Peter. 1990. "Zero Inflation as a Long-run Target for Monetary Policy." In *Zero Inflation: The Goal of Price Stability*, edited by Richard Lipsey, 67-108. Ottawa: C. D. Howe Institute.
- Laxton, Douglas and Robert Tetlow. 1992. *A Simple Multivariate Filter for the Measurement of Potential Output*. Technical Report No. 59. Ottawa: Bank of Canada.

- Laxton, Douglas, David Rose and Robert Tetlow. 1993a. "Problems in Identifying Non-linear Phillips Curves: Some Further Consequences of Mismeasuring Potential Output." Working Paper 93-6. Bank of Canada, Ottawa.
- Laxton, Douglas, David Rose and Robert Tetlow. 1993b. *Monetary Policy, Uncertainty and the Presumption of Linearity*. Forthcoming technical report. Bank of Canada, Ottawa.
- Laxton, Douglas, Kevin Shoom and Robert Tetlow. 1992. "Should the Change in the Gap Appear in the Phillips Curve? Some Consequences of Mismeasuring Potential Output." Working Paper 92-1. Bank of Canada, Ottawa.
- Lebow, D., J. Roberts and D. Stockton. 1992. "Economic Performance Under Price Stability." Working Paper No. 125. Division of Research and Statistics, Board of Governors of the Federal Reserve System, Washington.
- McCallum, John. 1988. "Les taux de chômage canadien et américain dans les années 1980 : un test de trois hypothèses." *L'Actualité Économique, Revue d'analyse économique* 64: 494-508.
- Novin, Farid. 1991. "The Productivity-Inflation Nexus Revisited: Canada, 1969-1988." Working Paper 91-1. Bank of Canada, Ottawa.
- Rotemberg, J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90: 1187-1211.
- Selody, Jack. 1990. *The Goal of Price Stability: A review of the issues*. Technical Report No. 54. Ottawa: Bank of Canada.

# Bank of Canada Working Papers

1992

- |      |   |                               |
|------|---|-------------------------------|
| 92-2 | Determinants of the Prime Rate: 1975-1989   | S. Hendry                     |
| 92-3 | Is Hysteresis a Characteristic of the Canadian Labour Market? A Tale of Two Studies | S. Poloz<br>and G. Wilkinson  |
| 92-4 | Les taux à terme administrés des banques  | J.-P. Caron                   |
| 92-5 | An Introduction to Multilateral Foreign Exchange Netting                            | W. Engert                     |
| 92-6 | Inflation and Macroeconomic Performance: Some Cross-Country Evidence                | B. Cozier<br>and J. Selody    |
| 92-7 | Unit Root Tests and the Burden of Proof   | R. Amano<br>and S. van Norden |

1993

- |      |  |  |
|------|--|--|
| 93-1 | The Implications of Nonstationarity for the Stock-Adjustment Model   | R. Amano<br>and T. S. Wirjanto         |
| 93-2 | Speculative Behaviour, Regime Switching and Stock Market Fundamentals  | S. van Norden<br>and H. Schaller       |
| 93-3 | Terms of Trade and Real Exchange Rates: The Canadian Evidence  | R. Amano<br>and S. van Norden          |
| 93-4 | State Space and ARMA Models: An Overview of the Equivalence  | P. D. Gilbert                          |
| 93-5 | Regime Switching as a Test for Exchange Rate Bubbles   | S. van Norden                          |
| 93-6 | Problems in Identifying Non-linear Phillips Curves: Some Further Consequences of Mismeasuring Potential Output | D. Laxton,<br>D. Rose and<br>R. Tetlow |
| 93-7 | Is the Canadian Phillips Curve Non-linear?   | D. Laxton,<br>D. Rose and<br>R. Tetlow |

Single copies of Bank of Canada papers may be obtained from

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

