What Happened to the Phillips Curve in the 1990s in Canada?

*Paul Beaudry and Matthew Doyle**

Introduction

There are at least two broad categories of interpretation to consider when examining observations on a country's output-inflation relationship (a country's Phillips curve). On the one hand is the traditional interpretation, which emphasizes how such a relationship mainly reflects a country's wage- and price-setting institutions. In this case, the Phillips curve is viewed primarily as a structural object, in the sense that its slope is governed foremost by the institutional aspects of the wage-setting mechanism—and hence is an object that constrains monetary policy. On the other hand is the view that the Phillips curve is essentially a reduced-form relationship, which mainly reflects rather than constrains the behaviour of monetary authorities. We will argue that this second view helps explain recently observed changes in the Phillips curve.

We begin by reviewing the changing nature of the Phillips-curve relationship in Canada and the United States from 1961–99. We define it as the statistical relationship between the change in inflation and the deviation of output from trend and, based on this definition, show that in both Canada and the United States the slope of the Phillips curve has become much smaller over the last 20 years, with a sharp reduction observed in the 1990s. This observation raises two related issues: (i) what explains the decline in

* We wish to thank Jeffrey Fuhrer, Steven James, and Pierre Duguay for helpful comments.

slope, and (ii) what does this decline imply for the proper conduct of monetary policy. Our goal is to provide new insight on these issues by presenting an explanation of the observed flattening of the Phillips curve, based on the notion that, since the 1970s, central banks have continually increased their awareness and understanding of the real forces that determine aggregate output. Hence, we believe that the current observation of a nearly horizontal Phillips curve may best be interpreted as a sign of well-executed, neutral stance, monetary policy.

Our explanation of the flattening of the Phillips curve is presented in a simple model that recognizes the role of both price rigidities and real disturbances in explaining macroeconomic fluctuations. In effect, our model extends the monopolistically competitive model of Blanchard and Kiyotaki (1987) in a manner that allows for real disturbances (as in the real business cycle literature) and for imperfect information.[1] However, in contrast to much of the macroeconomic literature with imperfect information (for example, Lucas 1972 and Barro and Gordon 1983), the information asymmetry we emphasize is such that the central bank is imperfectly informed regarding real developments in the economy, and hence is continually trying to infer the state of the economy while simultaneously affecting it.[2] We believe that this type of informational limitation is prevalent in all central banks and is important for understanding both the conduct of monetary policy and the co-movement between output and inflation.

In this simple model, we derive the properties of the output-inflation relationship under the assumption that monetary policy is conducted optimally, subject to the central bank's limited information. We show how a statistical Phillips curve can arise in this environment, with the causality running from real developments to nominal outcomes. Moreover, we show how the central bank will use observations of output and inflation to readjust the path of its monetary instruments.

We derive two main results from the model. Our first result is to show how, as the central bank becomes more aware of and sensitive to real developments in the economy, the slope of the Phillips curve will tend to approach zero. The intuition for this result is rather straightforward. The objectives of monetary policy should be to simultaneously support a well-functioning economy and to maintain price stability. However, in the

1. In the terminology of Goodfriend and King (1997), our model is a small-scale "new neoclassical synthesis" model.

2. In this respect, our model captures some of the elements present in Caplin and Leahy (1996) regarding the interaction between the central bank and private agents when the central bank is uninformed about the state of the economy.

absence of complete information on the state of the economy, the central bank cannot achieve this perfectly. The interaction between private agents and the central bank, both of whom are trying to learn from the other, gives rise to a Phillips-curve relationship. As the central bank learns to perform its information-gathering role more adequately, the positively sloped Phillips curve gradually disappears. We will argue that this mechanism helps to explain the observed flattening of the Phillips curve over the last 20 years, as the central banks in the United States and Canada first became aware of the importance of real shocks in the 1970s, and then learned to identify and react to them more appropriately throughout the 1980s and 1990s.

The second result we wish to highlight is that a flattening of the Phillips curve does not mean that the short-run output-inflation trade-off faced by the central bank has changed. In effect, we show why the Phillips curve and the output-inflation trade-off should be considered as two distinct objects, and why a flattening of the statistical Phillips curve can arise without a change in the relevant output-inflation trade-off faced by the central bank.

The remaining sections of the paper are structured as follows. In section 1, we document the changing nature of the Phillips curve for the United States and Canada from 1961 to 1999. In section 2, we present our model of the Phillips curve. In particular, we derive the properties of the output-inflation relationship under the assumption that monetary authorities are imperfectly informed about the state of the economy, but, nevertheless, try to conduct monetary policy optimally. We go on to compare the plausibility of our explanation of the flattening of the Phillips curve with one based on nominal wage rigidities. In the last section, we offer concluding comments.

1 Overview of the Output-Inflation Relationship in Canada and the United States

In this section, we review the evidence related to the existence of a positively sloped Phillips curve for both the United States and Canada from 1961 to 1999. We present evidence to suggest that the Phillips-curve relationship is robust to various specifications and roughly similar in Canada and the United States.[3] We also present evidence suggesting that the relationship between inflation and output has changed in recent decades. In particular, we show that the Phillips curve has flattened over the past

3. Fillion and Léonard (1997) present linear Phillips-curve estimates for Canada that resemble our own.

20 years. We find that the reduction in slope, which has occurred in both the United States and Canada, is quite substantive.

1.1 Basic estimation and results

In its simplest form, the Phillips curve can be expressed as a relationship between inflation, inflation lags, and the deviation of output from its trend level (referred to as the output gap). In the absence of clear, theoretical guidance on the appropriate measure of prices, the Phillips-curve literature uses various measures, from broad ones like the GDP deflator, to measures that attempt to capture the notion of core inflation. In our baseline estimations, we use the percentage change in the GDP deflator as our measure of inflation.[4]

Measuring the output gap raises further issues. The literature arrives at output gap series by employing a variety of techniques, including H-P filters, structural vector autoregressions (VARs), structural macroeconomic models, and simple time trends, to infer the trend level of output. We explored several alternatives and chose as our baseline measure the output gap series created by applying an H-P filter to the natural logarithm of real GDP.[5] Since we recognize that the Phillips curve can be expressed as a relationship between inflation and unemployment, we also explored the nature of the inflation-unemployment Phillips curve to provide a check on our results.[6]

As a starting point, we estimate the following very simple Phillips curve:

$$\Delta\pi_t = \alpha + \beta * (GAP_t + \varepsilon_t) .$$

In Figures 1 and 2 we plot this relationship, along with the associated regression line, for Canada and the United States from 1961 to 1999. The slope of the estimated Phillips curve for the United States is 0.256, suggesting that a positive output gap of 1 per cent is associated with an increase in inflation of around one quarter of 1 per cent on average. The Canadian estimate of 0.214 is similar to that of the United States. In both countries, we reject the hypothesis that the slope of the Phillips curve is zero

4. We also used the consumer price index (CPI) as an alternative measure of prices to check the robustness of our results.
5. We verified that the results are robust to various values of λ, the smoothing parameter of the H-P filter. The results presented in the paper set λ to 1600, which, with annual data, implies that we are unlikely to be over-smoothing.
6. We obtained our data for Canada from Cansim, and our U.S. data from Basic Economics (formerly Citibase).

Figure 1
Phillips curve, Canada: 1961–99

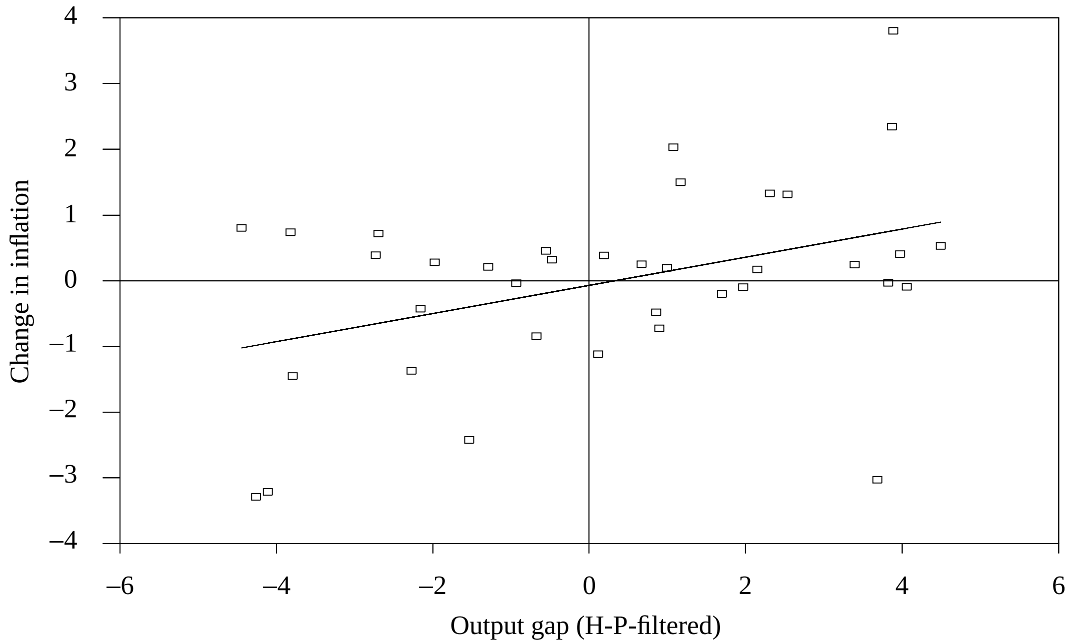

Figure 2
Phillips curve, United States: 1961–99

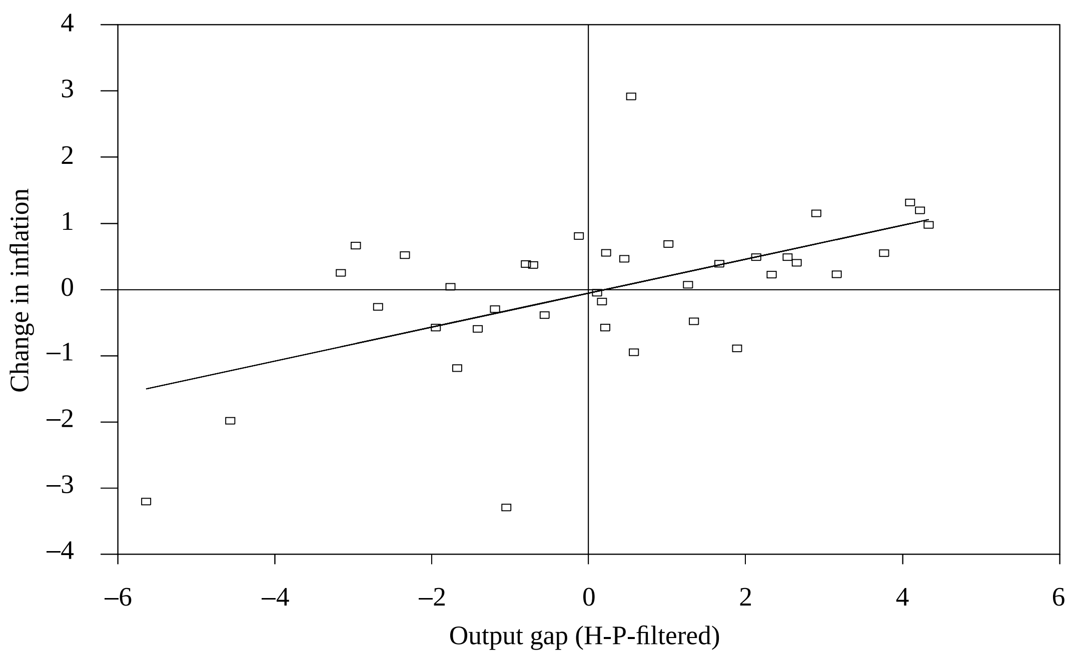

Table 1
Basic Phillips-curve estimates for Canada: 1961–99

	$\Delta\pi_t$	$\Delta\pi_t$	$\Delta\pi_t$
Constant (Std. error)	**–0.0360 (0.2259)**	**–0.0690 (0.2231)**	**0.0021 (0.2219)**
H-P-GAP$_t$	0.2141 (0.0828)	—	—
H-P-GAP$_{t-1}$	—	0.1890 (0.0814)	—
T-GAP$_{t-1}$	—	—	0.2599 (0.0998)

Table 2
Basic Phillips-curve estimates for the United States: 1961–99

	$\Delta\pi_t$	$\Delta\pi_t$	$\Delta\pi_t$
Constant (Std. error)	**–0.0014 (0.1567)**	**–0.0531 (0.1567)**	**0.0133 (0.1576)**
H-P-GAP$_t$	0.2560 (0.0650)	—	—
H-P-GAP$_{t-1}$	—	0.2404 (0.0619)	—
T-GAP$_{t-1}$	—	—	0.2719 (0.0714)

at conventional levels. To allow for the possibility that inflation responds to real developments with some delay, in the second column of Tables 1 and 2, we allow for lagged values of the output gap to enter as the right-hand side variable. This specification will be particularly relevant when discussing our theoretical model. As can be seen in the tables, our estimated Phillips-curve relationship is not strongly affected by the choice of the lag of the output gap rather than its contemporaneous value as a regressor.

To illustrate the robustness of these results, we consider various alternative Phillips-curve specifications. As mentioned above, one specification issue concerns our measure of the output gap. Since we derive our output-gap series by decomposing the output level into trend and gap components using an H-P filter, we wish to repeat our analysis using alternative detrending methods. In the third column of Tables 1 and 2, we report the results of estimating our simple Phillips-curve equation, using a cubic time trend to create the output-gap measure. Our point estimates of the slope differ depending on the choice of gap measure, as illustrated in the tables, but the differences are not very large.

We also wish to check the robustness of these results when we allow for a freer specification of the inflation process and when we control for supply-side factors. In short, we estimated several variants of the following equation:

$$\pi_t = const + a(L)\pi_{t-1} + b(L)(GAP_t) + cX_t + \varepsilon_t,$$

where π_t is inflation in period t, GAP_t is a measure of the output gap, and X_t is a vector of supply-side variables. We present a set of such results in Tables 3 and 4.

As can be seen in Tables 3 and 4, allowing for lags of the change in inflation as regressors can have noticeable effects on our estimated slope coefficients. Comparing the first column of Table 4 to our base results from Table 2 shows that the respecification has the effect of increasing the coefficient on the output gap for the United States. The same respecification, however, has almost no effect on the Canadian estimate. While the addition of lags of the change in inflation as regressors affects our slope estimates, in no case does the respecification overturn our initial results that indicate a positive and statistically significant co-movement between output and changes in inflation over the period 1961–99.

Respecifying the problem in terms of inflation rather than the change in inflation, as shown in the third and fourth columns of Tables 3 and 4, allows for a freer specification of the inflation process. We find that this specification of the inflation process also affects our slope estimates. In general, this results in a higher estimated coefficient on the output-gap term, where the estimate tends to increase with the number of inflation lags included.

Finally, the inclusion of supply-side variables appears to have moderate effects on our slope estimates. Columns two and four report results where inflation in relative energy prices is included as a regressor.[7] We find that the inclusion of energy prices has a small to moderate effect on our estimates of the coefficient on the gap variable, and that this effect differs in size and sign depending on the specification and country.

To summarize, we find that the data since 1960 strongly support the existence of a positively sloped Phillips curve in both Canada and the United States, and that this observation is robust to alternative specifications. Our estimates of the slope of the Phillips curve vary mostly between 0.2 and 0.3.

---

7. We define inflation in relative energy prices as the percentage change in the ratio of the CPI for energy to the total CPI.

Table 3
Basic Phillips-curve estimates for Canada: 1961–99

	$\Delta\pi_t$	$\Delta\pi_t$	π_t	π_t
Constant	**–0.0787**	**–0.0184**	**1.2991**	**1.7040**
	(0.2310)	**(0.2319)**	**(0.4042)**	**(0.5207)**
H-P-GAP$_{t-1}$	0.1935	0.2301	0.3646	0.3799
	(0.0922)	(0.0947)	(0.0889)	(0.0909)
$\Delta\pi_{t-1}$	0.0975	0.0743	—	—
	(0.1725)	(0.1718)		
π_{t-1}	—	—	0.8332	0.7391
			(0.1596)	(0.1739)
π_{t-2}	—	—	–0.1218	–0.1280
			(0.1445)	(0.1471)
π_t ENERGY	—	–0.0142	—	0.0665
		(0.0615)		(0.0570)
π_{t-1} ENERGY	—	–0.0886	—	0.0427
		(0.0590)		(0.0623)

Table 4
Basic Phillips-curve estimates for the United States: 1961–99

	$\Delta\pi_t$	$\Delta\pi_t$	π_t	π_t
Constant	**–0.0559**	**–0.0299**	**0.4776**	**1.2633**
	(0.1556)	**(0.1474)**	**(0.3095)**	**(0.2969)**
H-P-GAP$_{t-1}$	0.3161	0.2724	0.3054	0.2734
	(0.0771)	(0.0779)	(0.0742)	(0.0604)
$\Delta\pi_{t-1}$	–0.1449	–0.1726	—	—
	(0.1637)	(0.1792)		
π_{t-1}	—	—	0.8042	0.4578
			(0.1593)	(0.1595)
π_{t-2}	—	—	0.0694	0.2429
			(0.1618)	(0.1397)
π_t ENERGY	—	0.0719	—	0.0980
		(0.0300)		(0.0239)
π_{t-1} ENERGY	—	–0.0072	—	0.0723
		(0.0328)		(0.0305)

In all cases, the estimated slope is positive and significantly different from zero at conventional levels.

1.2 The changing slope of the Phillips curve

Having reviewed the case for the existence of a positive co-movement between inflation and output in Canada and the United States from 1961 to 1999, we now turn our attention to whether or not the Phillips-curve relationship may have changed over time. As we will show, the slope of the Phillips curve in Canada and the United States has declined markedly from its peak in the late 1970s.

To examine the slope of the Phillips-curve relationship over time, we use a series of rolling regressions on a 15-year moving window of data. That is, for each year in our sample, starting in 1978, we estimate the Phillips curve for the most recent 15-year period. For example, the estimates for 1983 are derived from observations over the period from 1969 to 1983.

Figures 3 and 4 present results from running the change in inflation on the lag of the output gap. We use this as our baseline specification, since it is easily tied into the theoretical results presented in subsequent sections. The estimated slope of the Phillips curve peaks around 1982 in both countries. It then declines throughout the 1980s and 1990s. In the United States, the slope begins to fall around 1988 and declines smoothly through to the end of the sample. This decline does not occur until 1992 in Canada, leading to a much sharper decline in the late 1990s. By the end of the 1990s, the slope of the Phillips curve is not significantly different from zero in either Canada or the United States.

As described in the previous section, we performed a variety of robustness checks of our baseline specification. We find that the pattern of a flattening of the Phillips curve in the 1980s in both Canada and the United States is robust across different specifications. The results presented in Figures 3 and 4 also seem to suggest that the slope of the Phillips curve may have been quite low in the mid to late 1970s. This implication, however, is not robust to the choice of estimation framework.

Figures 5 and 6 present one example where we include as an additional regressor inflation in relative energy prices. As can be seen, the pattern of a declining slope in the 1980s and 1990s remains essentially unchanged. Note, however, that our estimate of the Phillips curve's slope in Canada at the end of the 1970s is almost triple under the new specification (relative to the estimate in Figure 3). As a result, we do not believe that the

Figure 3
Slope of Canadian Phillips curve over time
15-year rolling regression of change in inflation on lag of output gap

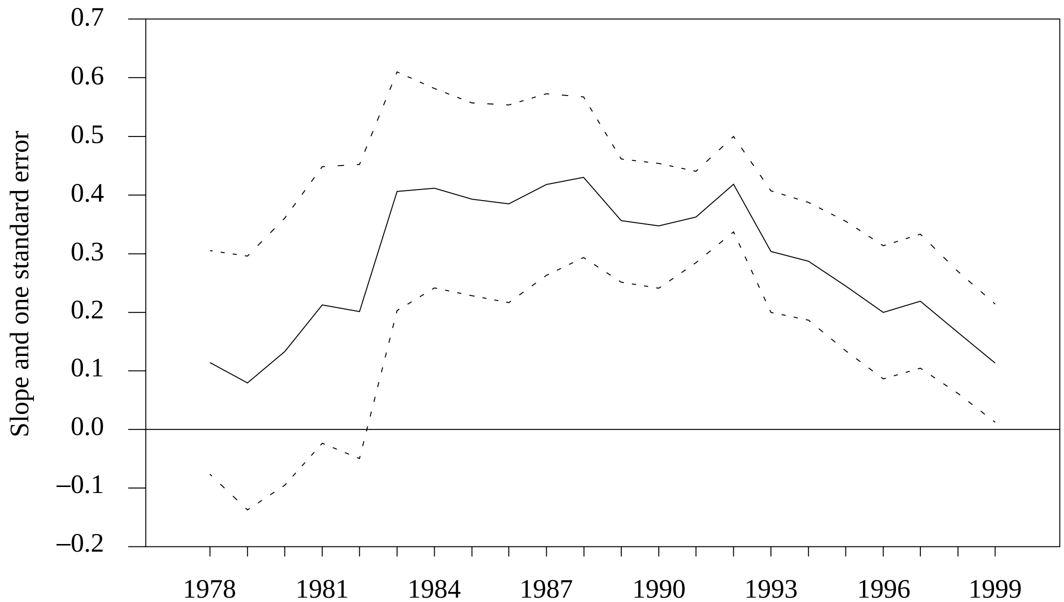

Figure 4
Slope of U.S. Phillips curve over time
15-year rolling regression of change in inflation on lag of output gap

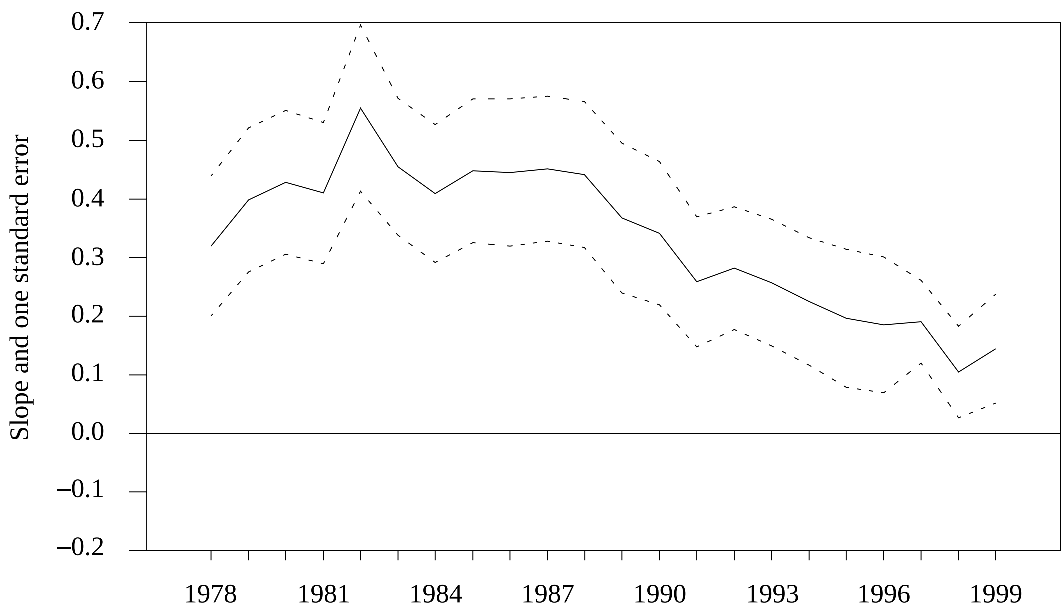

Figure 5
Slope of Canadian Phillips curve over time
Change in inflation on lag of output gap and inflation in relative energy prices

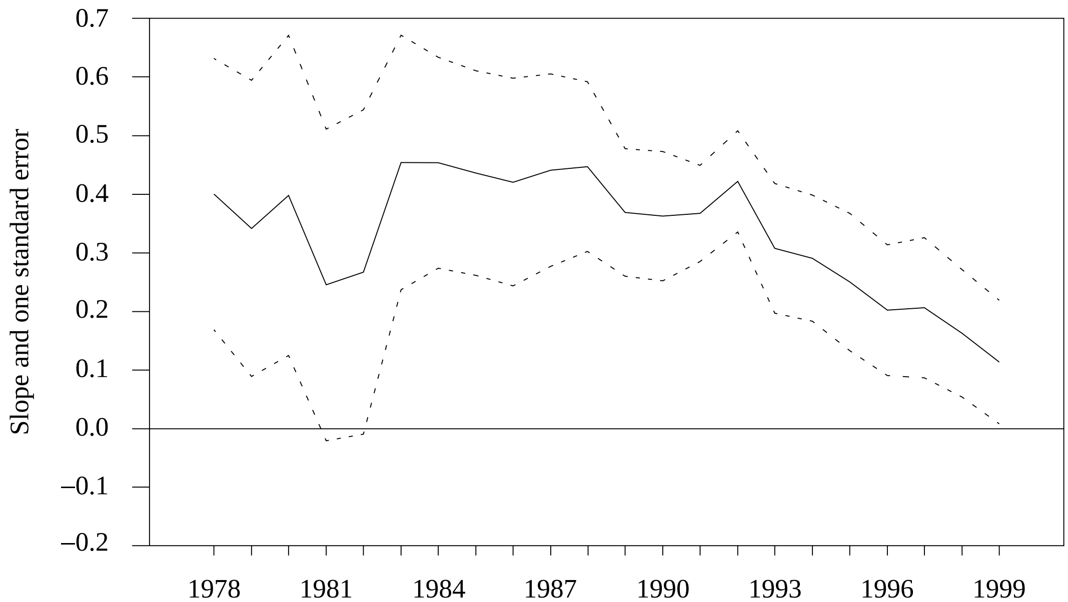

Figure 6
Slope of U.S. Phillips curve over time
Change in inflation on lag of output gap and inflation in relative energy prices

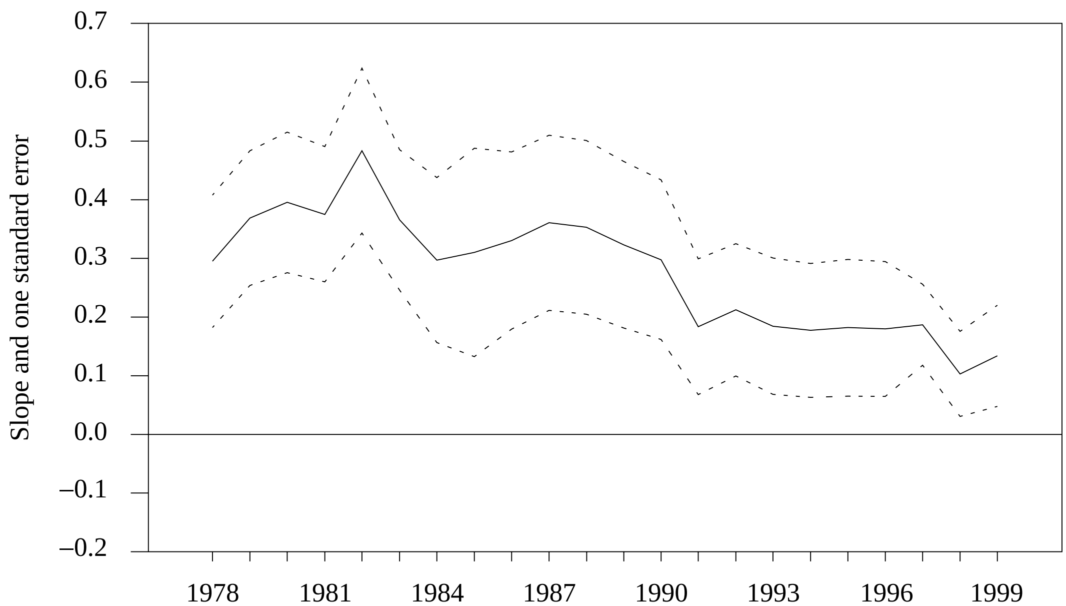

low slope coefficient observed in the 1970s (see Figure 3) is a robust feature of the data.[8]

Since we are attempting to examine changes in the Phillips-curve relationship over time, we also ran a series of weighted rolling regressions in which we imposed declining weights on more distant years. This procedure reduces the chance that one or two observations might unduly influence the profile of our estimates. We found that this approach yields similar results to those presented in Figures 3 to 6. That is, the slope of the Phillips curve appears to decline substantially over the 1980s and 1990s.

The point estimates presented in Figures 3 to 6 are not very precise, as can be seen from the size of the standard error bands. Since the imprecision of our estimates is a function of the size of our moving sample, we face a trade-off: we can increase the precision of our estimates only by including more distant years in our sample, in which case the composition of our sample tends to change much more slowly. Given this choice, we prefer to present estimates, which may be imprecise, but more fully capture any possible changes in the Phillips-curve relationship. We believe that the magnitude of the change in the point estimates is economically important enough to warrant interest, even if the statistical significance can be questioned.

As a further check on the robustness of our results, we pool our U.S. and Canadian data to increase the number of observations in each sample. The slope of the Phillips curve estimated on the full sample is 0.2239, which is in the same range as our previous estimates. As before, we find our estimate is robust to a variety of alternative specifications. Figure 7 reports the results of a series of rolling regressions, each on 15 years of pooled data and using the baseline specification.[9] We find the slope of the pooled U.S. and Canadian Phillips curve exhibits the same profile as in the individual samples. That is, the slope of the Phillips curve peaks in the early 1980s and declines thereafter. As in previous cases, the decline in slope from its peak to its 1999 level is substantial.

To illustrate the flatness of the Phillips curve since the mid-1980s, Figure 8 plots the relationship between the change in inflation and the output gap for the pooled sample from 1985 to 1999. The estimated slope for this sample is 0.1108, with a standard error of 0.0683. This is substantially lower than any of our full-sample estimates and is not significantly different from zero at conventional levels. As can be seen from the figure, one outlier drives

---

8. The outlying observation of 1975 (which was a year characterized by large movements in commodity prices) may explain why our estimates for the 1970s are sensitive to alternative specifications.

9. We regress the change in inflation on the lag of the output gap.

Figure 7
Slope of Phillips curve over time, pooled sample
Rolling regression of change in inflation on lag of output gap

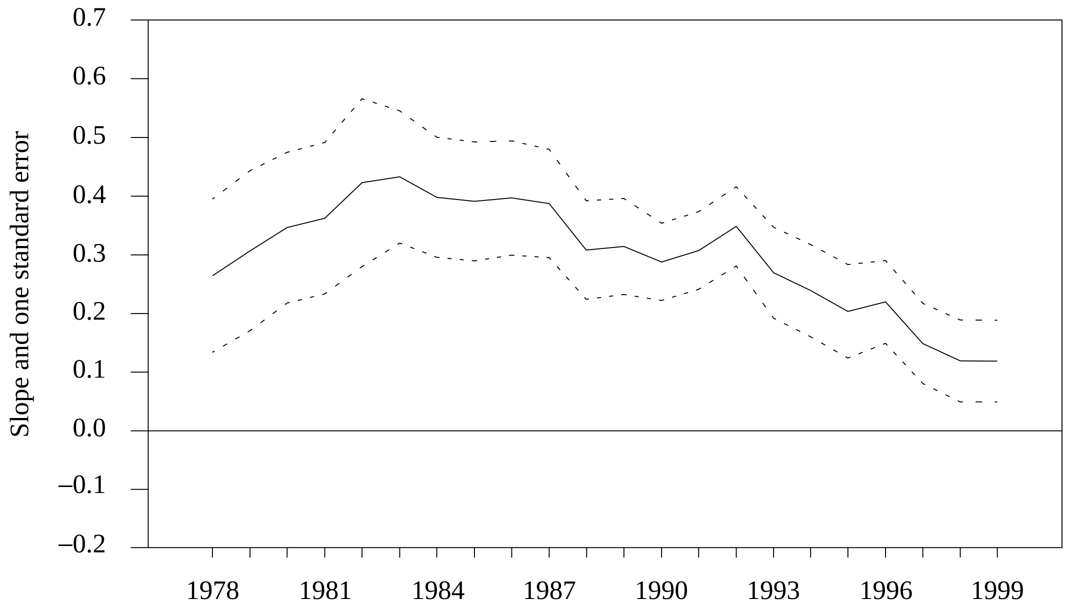

Figure 8
Change in inflation vs. gap 1985–99, pooled sample

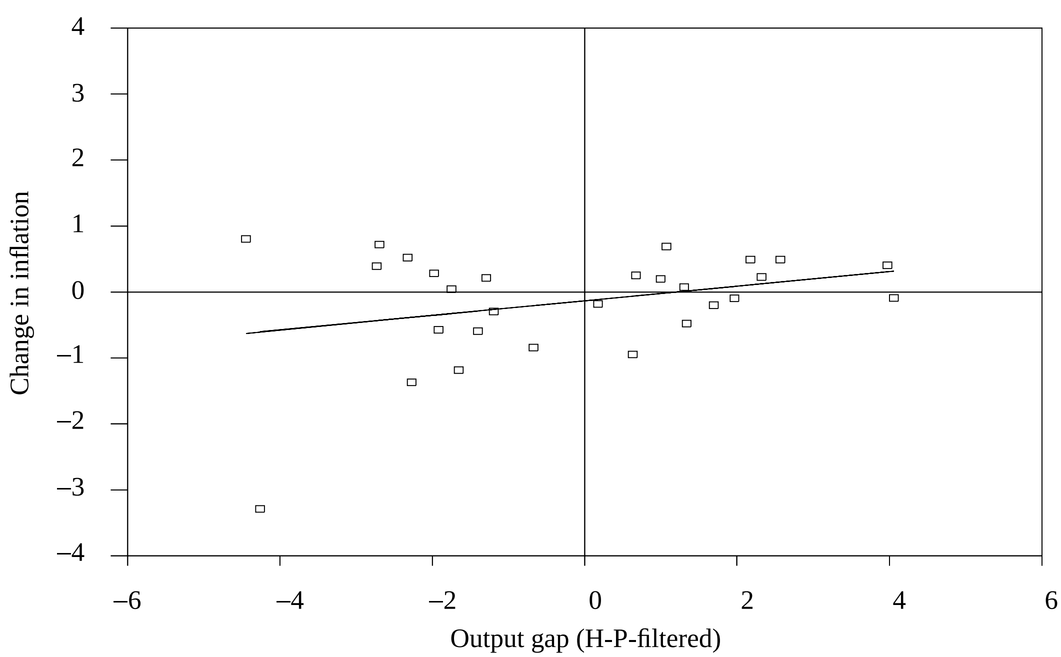

much of this slope: the Canadian observation for 1992. If we include a dummy variable to control for this observation, we find the slope of the Phillips curve for the United States and Canada since 1985 to be merely 0.0212.

This evidence leads us to believe that the Phillips-curve relationship has changed significantly in recent decades. In particular, we find that it has flattened substantially in both the United States and Canada. Furthermore, at least in Canada, this flattening occurred mainly in the 1990s.

2 Why Is There a Phillips Curve, and Why Might Its Slope Change over Time?

In this section, we explore the theoretical nature of the output-inflation relationship. Our goal is to illustrate the mechanism by which optimal monetary policy can give rise to a Phillips curve and how, in such a case, the slope of the Phillips curve relates to the fundamentals of the economy. In particular, we want to highlight the link between the slope of the Phillips curve and the degree to which monetary authorities are imperfectly informed about the state of the economy. We present this issue by building on a commonly used monopolistic competition macromodel (see, for example, Blanchard and Kiyotaki 1987), which we specify to allow for both real and nominal disturbances to affect output.

We consider an environment in which one final good, Y_t, is produced using a set of N intermediate goods, X_{it}, where $i = 1, \ldots, N$. The intermediate goods are produced by monopolistically competitive firms, which must pre-set prices at the beginning of each period, before the demand for intermediate goods is determined. The final good is produced by competitive firms, according to the constant returns to scale (CRS) production function given in equation (1).

$$Y_t = \left(\sum_{i=1}^{N} X_{it}^{\alpha} \right)^{\left(\frac{1}{\alpha}\right)} N^{\left(\frac{1-\alpha}{\alpha}\right)} \tag{1}$$

Each firm producing intermediate goods has access to a production technology given by equation (2).

$$X_{it} = A_t^{(1-\gamma)} L_{it}^{\gamma}, \tag{2}$$

where L_{it} is the quantity of labour employed in firm i, and A_t is the productivity index.

We assume that the productivity index, A_t, is common to all intermediate goods, and that the log of A_t follows the stationary stochastic process given by equation (3).[10]

$$a_t = \sum_{j=0}^{\infty} \psi_j \varepsilon_{t-j}, \qquad \psi_0 = 1, \qquad \sum_{i=1}^{\infty} \psi_i^2 < \infty, \tag{3}$$

where ε_t is assumed to be a normally distributed mean-zero random variable with variance σ_ε, and the ψ_i's are assumed to be positive. This last restriction is meant to capture the notion that deviations of technology from trend are positively autocorrelated.

To keep the presentation of the model as simple as possible, we do not explicitly include a trend in the process for A_t. Nonetheless, we think it is best to interpret the variables of the model as deviations from a trend induced by growth in A_t. Furthermore, our assumption of a common technology process across intermediate goods is clearly restrictive, but is justifiable on the grounds that we are interested only in aggregate fluctuations.

The representative household in this economy has preferences defined over consumption, labour supply, and real balances, as given by equation (4). We assume that the household's utility is linear in labour so as to generate a constant real wage. Hence, the model can alternatively be interpreted as a model with exogenously fixed real wages.

$$U\left(C_t, \frac{M_t}{P_t}, L_t\right) = C_t^{\theta} \frac{M_t}{P_t}^{1-\theta} - \Phi L_t \tag{4}$$

The household's budget constraint is given by equation (5), where P_t is the price of the final good, W_t is the nominal wage rate, $\overline{M}_t$ is money demanded, and M_t is the money balances distributed by the central bank at the beginning of each period.

$$P_t C_t + \overline{M}_t = W_t L_t + M_t \tag{5}$$

To solve for the private sector's equilibrium behaviour, we start by examining the household's decision problem. The representative household takes prices as given and chooses consumption, labour, and money balances to maximize utility. The first-order conditions associated with this maximization imply that money demanded satisfies equation (6), and that labour is supplied elastically at the real wage given by equation (7).

---

10. In all that follows, we use lower case letters to denote the logarithm of a variable.

$$\overline{M}_t = P_t C_t \frac{(1-\theta)}{\theta} \tag{6}$$

$$\frac{W_t}{P_t} = (1-\theta)^{1-\theta}\theta^{\theta}\phi \tag{7}$$

Producers of the final good also take prices as given and maximize profits by choosing the amount of intermediate inputs to use. This gives rise to a demand for intermediate goods given by equation (8), where P_{it} is the price of the *i*th intermediate good.

$$X_{it} = \left(\frac{P_{it}}{P_t}\right)^{\frac{-1}{1-\alpha}} \frac{Y_t}{N} \tag{8}$$

The problem facing a firm producing an intermediate good is more complicated, given that the prices of intermediate goods must be set before the realizations of either A_t or M_t. The firm's objective is therefore to set P_{it} to maximize expected profits conditional on the information set Ω_{t-1}, which contains all information dated $t-1$ or earlier, including realizations of past values of ε. Therefore, an intermediate-good producer's problem can be expressed as:

$$\max_{P_{it}} E[P_{it}X_{it} - W_t L_{it} / \Omega_{t-1}]$$

s.t. (2), (3), (7), (8).

Using the market-clearing conditions for both the goods market and the money market, and imposing symmetry on the behaviour of intermediate goods producers, one can easily derive equations (9) and (10), which describe the behaviour of the aggregate price level and aggregate output.[11] In these two equations, constant terms have been dropped.

$$p_t = E[m_t / \Omega_{t-1}] - \sum_{i=1}^{\infty} \psi_i \varepsilon_{t-i} \tag{9}$$

11. To derive equation (9) from the intermediate-good firm's problem, it is easiest to first use equations (2) and (7) to eliminate W_t and L_{it} from the firm's objective function. Then, using the market-clearing conditions $C_t = Y_t$ and $M_t = \overline{M}_t$ in combination with equations (6) and (8), the demand facing the firm can be written simply as a function of P_{it}, P_t, and M_t. Finally, imposing that $P_t = P_{it}$ in the first-order condition associated with the firm's optimal choice of P_{it}, and taking logs leads to equation (9).

$$y_t = m_t - p_t = (m_t - E[m_t/\Omega_{t-1}]) + \sum_{i=1}^{\infty} \psi_i \varepsilon_{t-i} \quad (10)$$

Equations (9) and (10) represent the equilibrium behaviour of private agents, for arbitrary processes of money supplied. Note that both prices and output depend on real and monetary forces. In particular, the aggregate price level depends on real shocks and expected money, while aggregate output depends on real shocks and unexpected money. Also note that the ε_t terms in equations (9) and (10) can be interpreted very broadly as reflecting any real shocks that affect the potential gains from trade, as opposed to the narrow technology shock representation.

The model thus far is a typical pre-set prices macromodel and generates a structure common to models of this type. The novel aspect of our analysis concerns the nature of the interaction between the private sector and the central bank. We now introduce the objectives and constraints facing the central bank, and highlight our model's key elements.

We assume that the central bank's objective is to minimize deviations of output and prices from target levels y_t^* and p_t^*, as given by equation (11).

$$\sum_{i=0}^{\infty} \beta^i E[(y_t - y_t^*)^2 + \Phi * (p_t - p_t^*)^2 \big| \bar{\Omega}_t] \quad (11)$$

In (11), Φ is the weight the central banker places on deviations of inflation from its target, relative to output deviations.

With respect to the output target, we assume that it is the level of output that would arise in the competitive equilibrium in the absence of any price rigidities or informational imperfection, that is,

$$y_t^* = \sum_{i=0}^{\infty} \psi_i \varepsilon_{t-i} .$$

(Note that we have again dropped the constant term.)[12] Although this choice of output target may be controversial, we believe it is the most reasonable assumption for the model. We assume that the price-level target is driven by an exogenous inflation target, $\bar{\pi}_t$, such that $p_t^* = p_{t-1} + \bar{\pi}_t$. For our purposes, the process for the inflation target can be thought of as being either stochastic or deterministic; the key simplifying assumption is that it is

12. By assuming that the central bank's objective is to attain the competitive-equilibrium outcome, we are eliminating a standard channel that gives rise to time-consistency problems (and inflationary bias).

exogenous. To allow for the possibility that the inflation target is stochastic, we denote the agents' expectation of target inflation as of time $t-1$ by ${}_{t-1}\bar{\pi}_t$. By assuming that the inflation target follows a known exogenous process, we are obviously sidestepping important issues related to the signalling of inflation targets.

The key assumptions of our model relate to the timing of moves and the information available to the central bank and private agents when making decisions. The assumptions are chosen to capture the notion that, in the short run, because of sticky prices, the central bank has the important but difficult task of helping private agents achieve gains from trade by providing the right amount of liquidity to the system. In effect, we model the central bank as having both an informational disadvantage and a timing advantage relative to the private sector. The central bank's disadvantage is that it does not directly observe the ε_t terms, and therefore must infer their values from past developments in the economy. Its advantage is that it has information on the current state of the economy, which it can use during the period over which prices are pre-set.

In effect, we assume that the central bank receives a signal, s_t, from its research department each period. This signal is an unbiased indicator of real developments in the economy, as captured by equation (12), where μ_t is a normally distributed mean-zero random variable with variance σ_μ^2.

$$s_t = \varepsilon_t + \mu_t \tag{12}$$

We denote by τ^2 the noise-to-signal ratio $\frac{\sigma_\mu^2}{\sigma_\varepsilon^2}$.

The timing of moves is as follows. At the beginning of a period, firms producing intermediate goods set prices, and the central bank simultaneously decides on the money supply. However, since private agents and the bank are differentially informed, the information used to determine these elements is different. Private agents know all past developments in the economy, but do not know the realization of ε_t that will arise during the period. In contrast, the central bank has past information on only output and prices (not the supply shocks), but has the advantage of observing s_t. We will denote the information set of the central bank at the beginning of time, t, by $\bar{\Omega}_t = \{s_t, s_{t-1}, \ldots, p_{t-1}, \ldots, y_{t-1}, \ldots\}$, and the information set of the private agents as $\Omega_{t-1} = \{\varepsilon_{t-1}, \ldots, s_{t-1}, \ldots, p_{t-1}, \ldots, y_{t-1}, \ldots\}$.

Our justification for giving the central bank an informational advantage through s_t captures the notion that the central bank has a timing advantage over the private sector. Since the private sector has pre-set prices, the bank has more flexibility within a period to react to current shocks, but is, nevertheless, imperfectly informed regarding the right way to act.

The problem facing the central bank is to choose a monetary policy rule to minimize equation (11), subject to its informational restrictions and the optimizing behaviour of the private economy, given by equations (9) and (10). This problem is more intricate than standard optimal policy problems, since the information sets of the private agents and the central bank are neither identical nor subsets of each other. In fact, our set-up is similar to a simultaneous-move game in which both sides have private information. As discussed in Townsend (1983), this can give rise to infinite regress problems. In this case, however, we have kept the problem simple enough to allow for an explicit solution.

The policy rule that solves the central bank's problem is given by equation (15), with the implied equilibrium solution for inflation (π_t) and output given by equations (13) and (14), respectively.

$$\pi_t = {}_{t-1}\bar{\pi}_t + \psi_1(\Theta s_{t-1} - \varepsilon_{t-1}), \qquad \Theta = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\mu^2} \tag{13}$$

$$y_t = \Theta * s_t + \sum_{i=1}^{\infty} \psi_i \varepsilon_{t-i} \tag{14}$$

$$m_t = p_{t-1} + {}_{t-1}\bar{\pi}_t - \frac{\psi_2}{\psi_1}(\pi_{t-1} - {}_{t-2}\bar{\pi}^*_{t-1}) + \Theta \sum_{i=0}^{2} \psi_i s_{t-i} + \sum_{i=3}^{\infty} \psi_i \varepsilon_{t-i} \tag{15}$$

To gain intuition about equations (13)–(15), it is helpful to first recognize that the term $\Theta \cdot s_t$ is the central bank's best estimate of the current supply shock, ε_t. Since the central bank's objective is to accommodate real shocks while maintaining price stability (around target), it adjusts the money supply to reflect its best guess of the current supply shock. Since prices are fixed, an expansion of the money supply is first reflected in output, as desired, and not in prices. That is, the central bank uses the money supply to allow the real economy to react to its signal on the current supply shock, thus partially overcoming the nominal rigidities inherent in the economy.

In the following period, the private sector becomes informed about the realization of last period's supply shock and adjusts prices accordingly. Note that inflation deviates from the target inflation level only to the extent that the central bank's estimate of the real shock in the previous period was mistaken. In effect, by adjusting prices in response to the central bank's error, the private sector actually reveals to the central bank the extent of its past error. The reason that private agents react to past mistakes is that they foresee that the central bank will continue to accommodate the effects of a perceived shock until it becomes aware that it has made an error. Hence, the profit-maximizing, price-setting rule is to increase prices in response to past excessive expansion on the part of the central bank.

Correspondingly, once the central bank recognizes that it has made an error by observing a deviation of inflation from its target, it readjusts the money supply. This can be seen from equation (15), where the past deviation of inflation from the target level enters negatively in the money supply rule. Although monetary authorities never directly observe the supply shocks, within two periods they are able to perfectly infer their values from observing developments in the economy. This explains why the money supply rule can be written as a function of lagged values of the ε_t terms.[13]

We now turn our attention to the implications of the above model for the nature of the Phillips curve. For now, let us define the Phillips curve as a purely statistical object, as we did in section 1. In particular, let the slope of the Phillips curve be the slope of the relationship between the change in inflation and the deviation of output from trend. Since our model is in terms of deviations from trend, the theoretical analogue to this slope is the co-variance between the change in inflation and output, divided by the variance of output. The analytical expression for this slope is reported in equation (16) and is denoted by β.

---

13. The money supply rule (15) can be used to calculate expected and unexpected money and thus verify that equations (13) and (14) are consistent with private agents' optimal behaviour given by equations (9) and (10). It is only slightly more difficult to verify that the money supply rule given by equation (15) is optimal. Note that for both prices and output the deviation from target is simply the difference between the central bank's guess of ε_t, $\Theta \cdot s_t$, and its realization. Hence, since this difference is minimized by setting

$$\Theta = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\mu^2},$$

this confirms the optimality of the policy.

$$\beta = \frac{cov(\Delta\pi_{t+1}, y_t)}{var(y_t)} = \frac{(\psi_1^2\tau^2)}{\left(\sum_{i=1}^{\infty}\psi_i^2\right)(1+\tau^2)} > 0, \qquad \frac{\partial\beta}{\partial\tau^2} \geq 0 \quad (16)$$

Our model suggests that we focus on the relationship between the change in inflation and the lagged deviation of output from trend, since it is only after one period that prices in the model can react to demand disturbances. Recall from section 1 that in the data such a distinction (at the annual level[14]) does not make much difference. If we enriched the dynamics of the model to allow for an autoregressive component to the output gap, this distinction would not matter in the theory either.[15]

The first thing to note from equation (16) is that the model generates a statistical Phillips curve; that is, even though monetary policy is set optimally, the economy nevertheless exhibits a systematic, positive co-movement between inflation growth and output. Moreover, this co-movement actually represents causality running from money to output and then to inflation, as is usually thought to be the case in discussions of the Phillips curve.

The second aspect to note is that the slope of the Phillips curve is strictly increasing in τ^2 (the noise-to-signal ratio for s_t). In other words, equation (16) implies that when the central bank becomes more aware of real developments in the economy (perhaps by expending greater effort to gather information about these developments and thereby reducing τ^2), it will make fewer errors in conducting monetary policy, and this will lead to a flatter Phillips curve.

This is the first result we want to highlight from this model: a flat Phillips curve may be a reflection of a well-run monetary policy. In particular, if σ_μ^2 were to go to zero, monetary authorities would make no errors and the statistical Phillips curve would become perfectly horizontal. In such a case, monetary authorities would be able to stabilize prices, while allowing the economy to respond efficiently to real forces. In contrast, the

14. With quarterly data, we generally found the lagged output gap to be a better predictor of inflation than the contemporaneous output gap.

15. One of the limitations of the current model is that, because we have not included any state variables, there is no endogenous propagation mechanism. This explains why money expansions affect output only for one period. If we included adjustment costs, such as a convex cost of changing labour, monetary shocks would have persistent effects and, hence, the distinction between the co-variance of $\Delta\pi_t$ with either y_t or y_{t-1} would be, as in the data, rather minor. Given the small returns and added complexity associated with adding such elements, we do not pursue this generalization here.

Phillips curve would tend to be more steeply sloped in an environment with substantial variations in real shocks or a poorly informed central bank.

Before discussing the potential relevance of equation (16) for explaining the changing nature of the Phillips curve, it is interesting to note the difference between the statistical Phillips curve implied by this model and the short-run output-inflation trade-off faced by the central bank. Even in a situation where the slope of the statistical Phillips curve is almost zero, this model does not imply that the central bank should perceive the short-run trade-off between inflation and output to be close to zero. In effect, such a trade-off could still be quite large. To see this, we can use equations (13) and (14) to derive the short-run relationship between inflation, target inflation, output, and supply shocks. This relationship is given by equation (17).

$$\pi_t = {}_{t-1}\bar{\pi}_t^* + \psi_1 y_{t-1} - \sum_{i=1}^{\infty} \psi_i \varepsilon_{t-i} \tag{17}$$

The term $\psi_1 y_{t-1}$ in equation (17) represents the effect on inflation induced by the central bank stimulating (or contracting) output in a one-time deviation from the optimal monetary policy. This equation nicely captures the type of short-run output-inflation trade-off often used to discuss the short-run effect of monetary shocks.[16]

The distinction in this model between the statistical Phillips curve and the short-run output-inflation trade-off reflects the difference between the effect of a systematic policy rule and the effects of monetary shocks conditional on agents believing that the policy rule is being followed. In particular, the statistical Phillips curve tends to become horizontal precisely when monetary authorities do not try to exploit the short-run trade-off and instead try to allow output to adjust to the real shocks. This result is reminiscent of that derived in Lucas (1972, 1973), but there is an important difference. In the Lucas model, when the statistical Phillips curve is horizontal, the output-inflation trade-off is zero. Here, this does not arise since private agents are not confused between real and monetary shocks. If the central bank decides to arbitrarily stimulate (or contract) the economy, the agents recognize this and respond by adjusting prices. This property of the model is, we believe, quite interesting, since it can potentially explain why strong monetary contractions are often associated with faster declines in prices than would be predicted by the statistical Phillips curve.

16. The only major difference between equation (17) and the more standard structural Phillips curve is that the relevant term for expected inflation is the agents' expectation of the central bank's inflation target as opposed to agents' expectation of actual inflation.

Now that we have described the functioning of the model, let us return to our question: What insight does this model provide towards explaining the flattening of the statistical Phillips curve in Canada and the United States over the last 15 years? The answer suggested by the model is that the decline in slope may have arisen because monetary authorities have learned to better identify and properly respond to real developments in the economy, thereby allowing such real developments to take place without large price effects. In other words, the flattening may be a reflection of improvements in the manner in which monetary policy is executed. In the rest of the paper we present both empirical and anecdotal evidence to support this view and compare its merits with alternative explanations.

Our first argument in favour of this view is entirely anecdotal, since it reflects the change in macroeconomic thinking throughout the last 25 years, and in relation to the conduct of monetary policy. Prior to the 1970s, the importance of real shocks on the macroeconomy was perceived to be rather minimal. The substantial fluctuations in oil prices changed this view and led central banks to rethink the way they conducted monetary policy. The focus of macroeconomic research also changed over this period. In particular, the arrival of real business cycle theory showed that a well-functioning economy might optimally fluctuate around its steady-state growth path, and the rational expectations literature questioned the potential for monetary policy to have systematically large effects on the real economy. Correspondingly, it appears reasonable to think that central banks (at least in Canada and the United States) responded to these changes by focusing more on identifying the underlying real forces in the economy and on learning how to respond to them. In the context of the model, we believe that such a process would correspond to a reduction in τ^2, since $\frac{1}{\tau^2}$ captures the degree to which central banks are informed about changes in the fundamentals of the economy. As central banks focused more attention on understanding economic fundamentals throughout the 1980s and 1990s, and came to believe that market forces were appropriate for the short-run determination of economic activity, the quality of their economic indicators (captured by s_t) likely improved, and the degree to which central banks acted on these signals (captured by Θ in the money rule (15)) likely increased. These developments are of exactly the type that our model suggests would lead to a flattening of the Phillips curve.[17]

To examine the plausibility of the idea that improvements in the manner in which monetary policy is conducted could be the cause behind

17. It is interesting to note from Figures 3 and 4 that the period in which the statistical Phillips curve appears steepest is in the late 1970s and early 1980s, which is generally considered a period of high variation in real shocks and of substantial confusion.

Table 5
Rolling sample variances for Canada and the United States: 1983–99

	Canada		United States	
Year	**Var** $(\Delta\pi_t)$	**Var** gap_{t-1}	**Var** $\Delta\pi_t$	**Var** gap_{t-1}
1983	3.61	5.15	2.82	7.34
1984	4.04	7.40	2.79	8.07
1985	4.04	8.01	2.79	7.02
1986	3.65	7.62	2.82	6.71
1987	3.65	7.56	2.86	7.08
1988	3.24	7.56	2.72	6.86
1989	2.04	7.62	2.04	5.90
1990	2.04	7.62	1.99	6.45
1991	1.59	7.62	1.35	6.00
1992	2.19	8.41	1.37	6.15
1993	2.28	9.92	1.37	6.05
1994	2.28	10.76	1.37	6.35
1995	2.07	9.55	1.23	6.05
1996	1.88	9.06	1.12	5.20
1997	1.88	8.64	0.52	5.24
1998	1.34	8.06	0.31	3.57
1999	1.04	7.24	0.32	2.46

the observed flattening of the Phillips curve, it is useful to consider alternative explanations. One such potential explanation, often seen in the press, is that, over the 1980s and 1990s, monetary authorities began to disregard their role in controlling output fluctuations and conducted monetary policy with the sole aim of stabilizing prices. According to this view, greater price stability is achieved only at the cost of greater output instability. To help evaluate this view, Table 5 reports a series for the variance for the output gap and for the change in inflation for the United States and Canada since the early 1980s (the period over which we observe the decline in the slope of the Phillips curve). The variance reported for each year is calculated using the observations on the previous 15 years. As can be seen from the table, in both the United States and Canada, the variance of inflation growth has decreased quite substantially over the last 15 years. In contrast, the variance of output for Canada has remained about the same, while that for the United States appears to have declined. In particular, the variance of output in Canada was about 7.5 throughout much of the 1980s and was approximately at the same level by the end of the 1990s.

The main inference we draw from Table 5 is that the variance of output does not appear to have increased during the period in which the Phillips curve flattened. While such an observation is not inconsistent with our proposed explanation, it is somewhat at odds with the view that greater price stability was achieved at the cost of greater output variability.

2.1 The flattening Phillips curve: Evidence of optimal policy or downward nominal rigidities

One possible explanation for the observed flattening of the Phillips curve, as suggested by Akerlof, Dickens, and Perry (1996) and Fortin (1997), among others, is downward nominal-wage rigidity (DNWR). The reasoning is as follows: when inflation is very low, the unwillingness of workers to accept nominal-wage reductions prevents real wages from adjusting in response to excess supply in the labour market. In this case, if prices are a fixed markup on wages, then prices will not fall in response to a negative output gap. This causes the Phillips curve to be flatter in periods of lower inflation. The proponents of this explanation claim that this is the relevant difference between the experience of the 1990s relative to the early 1980s.

In this section, we attempt to differentiate between the theory of DNWR and our proposed explanation of the flattening of the Phillips curve, which relies on improved monetary policy. One implication of downward nominal rigidity, not shared by our explanation, is the prediction that the flattening of the Phillips curve should be associated with an increase in its degree of non-linearity. In particular, the hypothesis of downward nominal rigidity suggests that as inflation decreases, it is mainly the segment of the Phillips curve that relates to negative values of the output gap, which should flatten (because DNWRs are not relevant when the labour market is tight).

To explore this hypothesis empirically, we estimated several variants of the type of non-linear Phillips curve given by equation (18) and examined how the coefficients changed over time.[18]

$$\pi_{t+1} = \beta_0 + \pi_t + \beta_1(GAP_t) + \beta_2 PosGap_t + \varepsilon_{t+1} \quad (18)$$

In equation (18), the variable $PosGap_t$ takes the value of zero if the output gap is negative and is equal to the value of the output gap if the latter is positive.[19] Figures 9 and 10 report respectively values for β_1 and β_2 associated with successively estimating equation (18) based on pooled U.S. and Canadian data over 15-year periods. We present the results for the pooled estimates, since they are the most precise. However, we also estimated this equation for each country and for several different specifications, and obtained results similar to those represented in Figures 9 and 10.

18. Dupasquier and Ricketts (1998) provide a good entrance point to the literature on estimating non-linear Phillips curves for Canada.

19. Our approach is to estimate a Phillips curve that has a kink at a zero output gap. We adopt this simple approach to evaluate the presence of a non-linearity, although the nominal-wage-rigidity hypothesis does not precisely predict a kink at zero output gap.

Figure 9
Output gap coefficient, pooled sample
Change in inflation on lag of gap and lag of positive gap

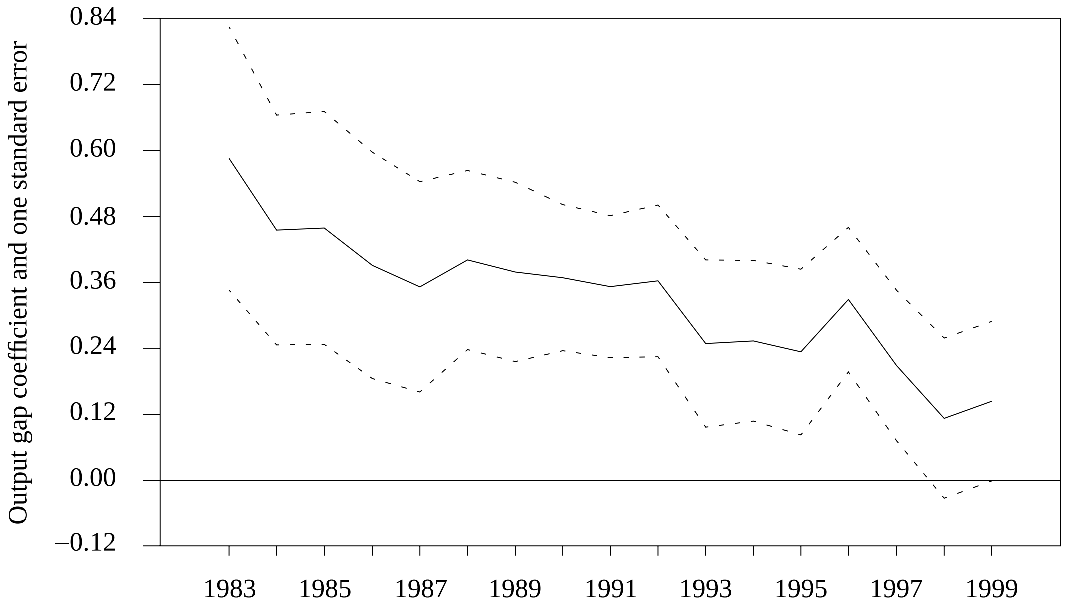

Figure 10
Positive gap coefficient, pooled sample
Change in inflation on lag of gap and lag of positive gap

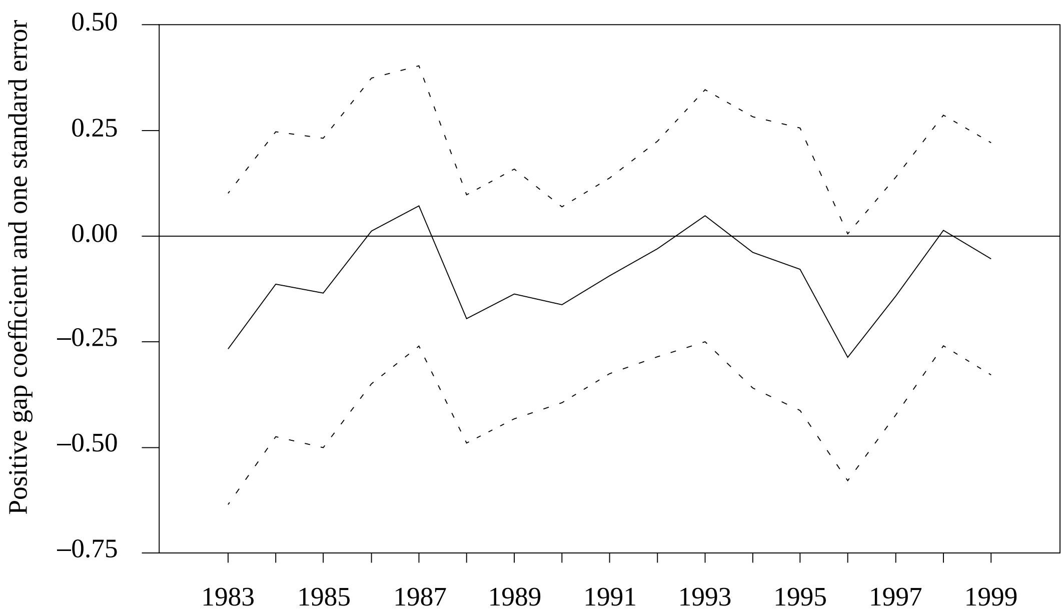

As can be seen in Figure 9, the value of β_1 decreased substantially throughout the 1980s and 1990s. Since this coefficient represents the slope of the Phillips curve for negative values of the output gap, its decline is consistent with the hypothesis that downward nominal rigidities may have caused the Phillips curve to flatten. However, our estimates of β_2, as shown in Figure 10, suggest that the degree of non-linearity of the Phillips curve has not increased over this period, an observation that is inconsistent with the nominal-wage-rigidity hypothesis. In effect, our estimates of β_2 suggest that the Phillips curve remained linear throughout the period, whereas an increase in β_2 would be expected if DNWR was the cause of the flattening.

Our evidence against the hypothesis of downward nominal rigidity can be inferred visually from the simple scatter plot presented in Figure 8. Since the mid 1980s, the Phillips curve has been very flat over the range of both positive and negative output gaps. In fact, the only evidence of non-linearity relates to the outlying observation of Canada in 1992. However, for this observation the output gap was negative and large, and inflation fell substantially. Hence, we take this evidence as contradicting the downward nominal rigidity hypothesis as an explanation of the observed flattening of the Phillips curve.

2.2 The flattening of the Phillips curve and the Ball, Mankiw, and Romer hypothesis

A second potential explanation for the flattening of the Phillips curve is the one proposed by Ball, Mankiw, and Romer (1988), based on menu costs. This theory suggests that in a period of low trend inflation, firms do not very often find themselves on the boundary of the set of acceptable prices (that is, the S,s boundary of acceptable prices defined by the size of the menu cost). Therefore, firms do not change their individual prices as frequently when trend inflation is low as when it is high. This greater sluggishness in individual prices increases the degree of overall nominal rigidity in the economy and therefore leads to a flatter Phillips curve. Since the trend level of inflation has fallen over the past 20 years, the menu-cost hypothesis predicts that the Phillips curve should have become flatter over this period, which is exactly what we observe in the data.

The menu-cost explanation and our model, however, have important differences regarding the effects of monetary surprises that arise because of their respective implications for the short-run output-inflation trade-off and the statistical Phillips curve. In the menu-cost explanation, when inflation is low, the Phillips curve is flat. Since there is no distinction in this story between the statistical Phillips curve and the short-run output-inflation trade-off, such a flattening implies that the output-inflation trade-off has increased. In contrast, while our model predicts that the statistical Phillips curve (whose slope is given by equation (16)) becomes flatter when policy-makers monitor the economy properly, this does not imply that the short-run output-inflation trade-off changes. In effect, the relevant output-inflation trade-off—that is, the trade-off induced by a deviation from the perceived policy rule—is governed by the parameter ψ_1 in equation (17), which is independent of the trend level of inflation.

In short, the difference between the two models is that the menu-cost story implies that the effect of a monetary shock varies inversely with the trend level of inflation, while our model predicts that the effect of a monetary shock on inflation is independent of the level of trend inflation. This difference indicates how the two models can be distinguished empirically. In effect, one can differentiate the two models by examining whether the co-movement of inflation and the output gap following a monetary shock differ in periods of high relative to low trend inflation.[20]

The major limitation of this strategy involves data. To compare these two competing theories, we need to observe monetary shocks in periods of both high and low trend inflation. Since monetary shocks are infrequent, we find ourselves confronted with the problem of having few observations. Nevertheless, using the Bank of Canada's Annual Reports as our source, we can identify two important disinflationary shocks in Canada between 1980 and 1999: the first occurred in 1982–83 and the second in 1991–92. In particular, the 1980 and 1981 Annual Reports suggest that the Bank was troubled by the high inflation of the late 1970s, but unable to act because of the need to respond to changes in U.S. interest rates and large capital flows out of the country. In 1982, then Governor Gerald Bouey wrote that "the Canadian economy has shown strong resistance to becoming less inflationary," (Bank of Canada 1982, 8) and noted that "inflation must sooner or later be fought." (Bank of Canada 1982, 7) A year later, he reflected on the "strong monetary medicine" that had been required to "beat the fever [of inflation]" (Bank of Canada 1983, 10) that gripped the Canadian economy in the late 1970s. We count the experience of 1982–83 as a disinflationary shock, since the Bank of Canada appeared focused on reducing what it regarded as an unacceptably high inflation rate, rather than responding to real developments in the economy.

The reports from 1984 through to 1990 portray a Bank of Canada on guard against a renewal of inflation, but not actively seeking to reduce the trend rate. In 1991, in response to the increased inflation of the late 1980s,

---

20. Like the menu-cost theory, the hypothesis of DNWR does not imply a distinction between the short-run output-inflation trade-off and the statistical Phillips curve; therefore, the evidence presented in this section also relates to that potential explanation.

the Bank of Canada and the Government of Canada jointly announced a set of inflation targets to take effect in 1992. The targets essentially mandated a reduction in inflation, which was then around 5 per cent, into a target band of 2 to 4 per cent. Governor Crow, in the 1991 Annual Report, wrote that the purpose of the inflation targets was "[t]o provide Canadians with a clear affirmation that price stability remains the goal of monetary policy." (Bank of Canada 1991, 8) Two years later, Governor Crow reflected that "a key purpose in establishing [the inflation targets] was to indicate as clearly as possible not a path for *sustaining* inflation, but a path for *reducing* inflation." (Bank of Canada 1993, 9) We regard the experience of 1991–92 as a second disinflationary shock.

After 1992, the Bank of Canada reduced the target band in 1994 and 1995, but since this had been announced in 1991, it is not clear that we would wish to count it as a monetary shock. After 1995, the Bank chose to maintain the target band at its 1995 level through 1998 and later, through 2001.

Therefore, we conclude that there have been two disinflationary shocks in Canada since 1980. The important difference between the two is that the 1982–83 shock occurred during a period of relatively high inflation, while the 1991–92 shock occurred while the trend rate of inflation was much lower. In principle, these two episodes provide an excellent opportunity to test the different theories. To make this comparison, Figure 11 plots the change in inflation against the output gap for Canada for the sample 1980 to 1999. The surprising and noticeable aspect is that the observations for 1983 and 1992 lie almost exactly on top of one another.[21] We see this as providing some, albeit limited, support for the view that the short-run output-inflation trade-off did not change as inflation decreased.

We also find it worthwhile to contrast the inferred size of the output-inflation trade-off under the two views. Under the assumptions that our model is correct and that the 1983 and 1992 points are representative of the short-run output-inflation trade-off, Figure 11 implies that the cost of reducing inflation by 1 per cent is a negative output gap of approximately 1.3 per cent (the slope implied by the 1983 and 1992 observations). If, on the other hand, the menu-cost theory is correct, then the slope of the Phillips curve is the proper estimate of the output-inflation trade-off. In this case, using the final estimate of the slope of the Phillips curve from the rolling regressions for Canada (which is around 0.1) as the measure of the trade-off, the negative output gap induced by a 1 per cent reduction in inflation would be on the order of 10 per cent. Clearly, the two interpretations differ by

21. While we do not include 1994 as a shock, it is interesting to note that it does lie along the same line as the two stronger shocks.

Figure 11
Phillips curve, Canada 1980–99
Regression includes dummy variables for 1983 and 1992

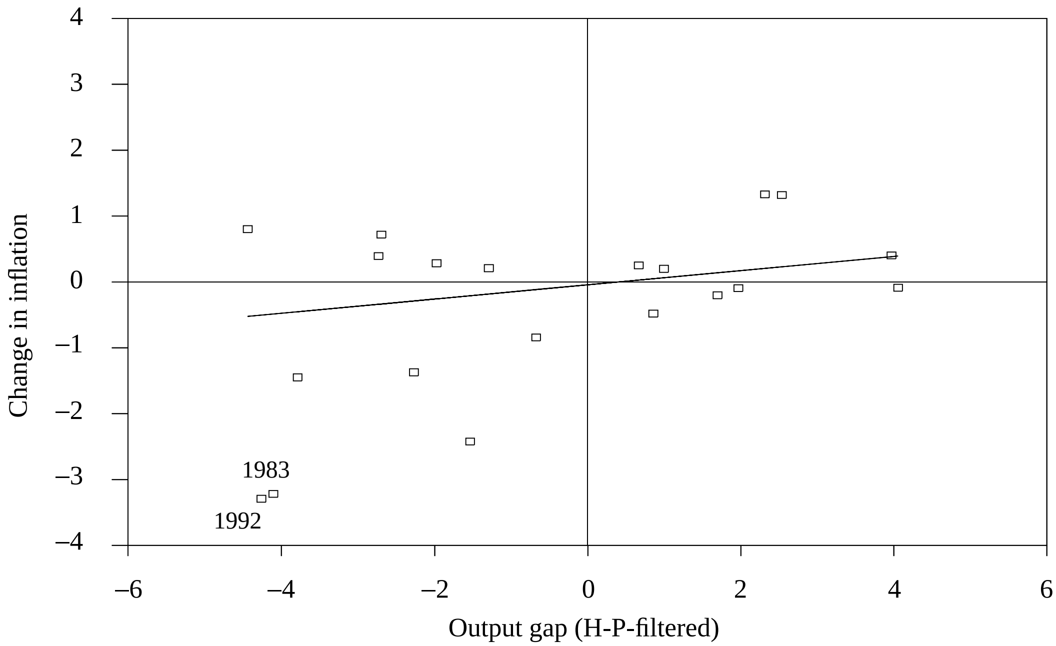

orders of magnitude and hence suggest the need to provide further evidence to differentiate the two views more convincingly.

Conclusion

Our answer to the title of the paper, "What Happened to the Phillips Curve in the 1990s in Canada?," is both empirical and theoretical. From a statistical point of view, we have shown that the slope of the Phillips curve in Canada has decreased substantially over the period. We also document that the same phenomenon is observed in the United States. Since we are interested in interpreting these observations for policy discussion, we have used a prototypical macromodel to attempt to understand why the slope may have changed over time and what implication this may have for the output-inflation trade-off faced by the central bank. In particular, we have shown why a change in monetary policy, which incorporates a better understanding of the real side of the economy, will lead to a flatter Phillips curve. The reason we believe that the conduct of monetary policy may have changed in this direction is that, after the oil shocks of the seventies, central banks appear to have devoted more effort towards tracking the real forces affecting aggregate output, and have probably incorporated the improved knowledge into their behaviour.[22] The second insight drawn from the model is that a flatter Phillips curve does not necessarily imply a change in the output-inflation trade-off faced by the central bank. In effect, we show why the Phillips curve can become flatter while the relevant output-inflation trade-off remains constant.

Based on several pieces of evidence, we have argued that our model provides a reasonable framework for interpreting recent observations on the Phillips curve. As we have explained, the main implication of this view for policy is that the best guess of the potential costs associated with a disinflation undertaken today is that inferred from the disinflationary episodes of the early 1980s and 1990s. In other words, we believe that the evidence on inflation and output over the last 20 years supports the view that the costs associated with reducing inflation have likely neither increased nor decreased over that period, even if the statistical Phillips curve appears to have flattened.

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.

Ball, L., N.G. Mankiw, and D. Romer. 1988. "The New Keynesian Economics and the Output-Inflation Trade-off." *Brookings Papers on Economic Activity* 1: 1–65.

Bank of Canada. 1982. *Annual Report*. Ottawa: Bank of Canada.

———. 1983. *Annual Report*. Ottawa: Bank of Canada.

———. 1991. *Annual Report*. Ottawa: Bank of Canada.

———. 1993. *Annual Report*. Ottawa: Bank of Canada.

Barro, R.J. and D.B. Gordon. 1983. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy* 91 (4): 589–610.

Blanchard, O.J. and N. Kiyotaki. 1987. "Monopolistic Competition on the Effects of Aggregate Demand." *American Economic Review* 77 (4): 647–66.

Caplin, A. and J. Leahy. 1996. "Monetary Policy as a Process of Search." *American Economic Review* 86 (4): 689–702.

Dupasquier, C. and N. Ricketts. 1998. "Non-Linearities in the Output-Inflation Relationship." In *Price Stability, Inflation Targets, and Monetary Policy*, 131–73. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Fillion, J.-F. and A. Léonard. 1997. "La courbe de Phillips au Canada : un examen de quelques hypothèses." Bank of Canada Working Paper No. 97–3.

---

22. One possibility we have not addressed is that monetary authorities may have learned how to properly unlink real and nominal developments in the economy over the last 20 years, which would explain the flattening of the Phillips curve, but that this simple unlinking is not necessarily the best policy to follow. In effect, if there are other imperfections in the market besides those associated with pre-set prices, it may be socially optimal for monetary authorities to favour lower output variability than that which would arise by letting market forces work freely. We believe that this possibility is very relevant and should be examined in future research.

Fortin, P. 1997. “Canada’s Job Growth Potential.” Paper presented at the June 1997 Meetings of the Canadian Economics Association, Memorial University, St. John’s, Newfoundland.

Goodfriend, M. and R.G. King. 1997. “The New Neoclassical Synthesis and the Role of Monetary Policy.” *1997 NBER Macroeconomics Annual*, 231–83. Cambridge and London: MIT Press.

Lucas, R.E., Jr. 1972. “Expectations and the Neutrality of Money.” *Journal of Economic Theory* 4 (2): 103–24.

———. 1973. “Some International Evidence on Output-Inflation Tradeoffs.” *American Economic Review* 63 (3): 326–34.

Townsend, R.M. 1983. “Forecasting the Forecasts of Others.” *Journal of Political Economy* 91 (4): 546–88.

Discussion

Steven James

In their paper, Beaudry and Doyle argue that the short-run Canadian Phillips curve has flattened over the last 20 years. This conclusion is based on 15-year rolling regressions of the change in inflation on an H-P-filter-derived output gap and on the growth of relative energy prices. They find similar results employing an output gap constructed using a cubic time trend. Based on results from a theoretical model that they construct, the authors argue that this flattening stems from an improved understanding by the monetary authorities of the role played by supply shocks. In their theoretical model, such an improved understanding would lead to a flattening of the statistical Phillips curve but would not affect the exploitable short-run trade-off between inflation and unemployment. The sacrifice ratio would thus be unaffected. They note that a competing explanation for such a flattening is the view that DNWR has tended to become more binding at low rates of inflation. The authors assert that this would imply a kinked non-linearity at potential, that is, the flattening would occur only in the excess supply range of the short-run Phillips curve. They cannot reject the null hypothesis that the Phillips curve is linear, leading them to prefer their own model-based explanation.

Beaudry and Doyle's theoretical model provides a useful reminder that the statistical slope of the Phillips curve need not map directly to the short-run output-inflation trade-off. If the central bank credibly follows an inflation-control policy rule and ensures that supply shocks affect real output rather than prices, then the statistical Phillips curve will be flatter than if the bank does not discern supply shocks or does not seek to neutralize their effect on prices. In their model, the short-run trade-off reflects the cost of a deliberate disinflation, which may be positive even if the statistical Phillips curve is flat. The distinction between the statistical slope and the trade-off stems from their model assumption that all fluctuations are driven by supply and monetary policy shocks. The assumption of a greater relative role for non-monetary demand shocks would reduce the distinction between the statistical slope and the short-run trade-off and would weaken the link between flattening of the slope and the ability of the central bank to distinguish supply shocks.

A similar Phillips-curve methodology was used by Cozier and Wilkinson (1991) to estimate the Canadian sacrifice ratio. In this approach, the sacrifice ratio is a function of the coefficient on the output gap and the adjustment speed of inflation expectations to inflation shocks. Lipsett and James (1995) extended this approach by adding a role for the change in the gap and by conducting 13-year rolling regressions to estimate and compare the evolution of international sacrifice ratios over time. They found that the Canadian sacrifice ratio increased from 0.6 to 2.3 as the fixed-sample size end point varied between 1981 and 1993. The increase stemmed particularly from the decline in the coefficient on the output gap rather than from changes in the adjustment speed of inflation expectations. This decline was particularly marked as information from the 1981–82 recession was added, a result that is not consistent with the Beaudry-Doyle model feature that arbitrary contractions lead to quicker price declines than implied by the statistical Phillips curve. Lipsett and James found unstable and frequently statistically insignificant sacrifice ratios for seven other industrialized countries, including the United States. Output-gap coefficients tended to be considerably more stable when persistence was captured through the inclusion of a change in the gap variable. It would be interesting to see how the same extension would affect the Beaudry-Doyle results.

There are a number of shortcomings with these types of Phillips-curve-based estimates. The first stems from the inherent arbitrariness of the H-P-filter and time-trend-based output gaps. Actual gaps will be under-estimated in the presence of hysteresis or significant persistence of demand shocks. Persistent gaps will be defined away, since such shocks will be treated as changes in potential. One alternative is to regress on a production-function-based gap using a structurally determined natural rate or on the unemployment rate and structural variables such as a measure of EI generosity. An SVAR approach can also allow for more structurally based estimates. Cecchetti and Rich (1999) argue that econometric studies of the sacrifice ratio should try to distinguish movements in a policy variable associated with a shift in policy (e.g., a decision to disinflate) from those reflecting a systematic response to the state of the economy (e.g., an inflation-control rule). The SVAR approach that they adopt allows monetary policy to be decomposed into systematic (policy-reaction) and random (policy-shift) components. Such an approach is closer in spirit to the Beaudry-Doyle theoretical model than is the Phillips-curve approach the authors use. Unfortunately, the Cecchetti and Rich methodology does not generate robust estimates. In particular, results tend to be very sensitive to the choice of identifying restrictions.

Second, rolling regressions, while useful, must be interpreted carefully. It is tempting to view the sudden flattening of a Phillips curve in a particular year as the consequence of a particular event or policy-regime change. In reality, the rolling regressions do not provide end point estimates. The removal of information 15 years earlier is as important as new information being added.

Beaudry and Doyle note that the existence of DNWR could also lead to a flattening of the Phillips curve. This is correct; however, other factors could also affect the slope, including a reduction in average inflation and its associated variability, changes in labour market rigidities, changes in the credibility of discretionary bank actions, and the menu-cost explanation of Ball, Mankiw, and Romer (1988), noted by the authors. It is not possible to distinguish among these factors in the empirical framework used.

The authors assert that "... the hypothesis of downward nominal rigidity suggests that as inflation decreases, it is mainly the segment of the Phillips curve that relates to negative values of the output gap, which should flatten (because DNWRs are not relevant when the labour market is tight)" (page 75). This is not correct. The Phillips curve in Akerlof, Dickens, and Perry (1996) has the form:

$$\pi = \pi^{\varepsilon} + c - a \cdot u + S \cdot M,$$

where π is inflation, c is a constant, u is the unemployment rate, S is the shift in unit labour costs resulting from DNWR, and M is a markup.

The S variable may be expressed as:

$$S_t = \sigma_0 \cdot \phi(v_t / \sigma_0) + \Phi(v_t / \sigma_0) \cdot v_t,$$

where ϕ and Φ are the probability density function and cumulative distribution function, respectively, of the standard normal distribution. The parameter σ_0 is the standard deviation of the gap between notional wages and lagged actual wages and v is the expected value of this gap, deflated by expected nominal productivity. The variable v may be expressed as:

$$v_t = \frac{S_{t-1} - \left[\frac{\beta - 1}{\beta}\right] \cdot [\pi_t^{ee} + g_t - a(u_t - u_{t-1})]}{1 + \pi_t^{ee} + g_t} + d(r_t - r_{t-1}),$$

where *g* is labour productivity growth. Lower combined inflation and productivity growth and higher unemployment raise *v* and hence *S*. The presence of the change in unemployment term causes the short-run Phillips curve to flatten as DNWR binds. It also introduces a non-linearity; however, there is no kink at a zero output gap or at any other point. The degree of non-linearity depends on the size of the parameter *a* and need not, in principle, be large. The recursive nature of *S* also means that there is no direct mapping between the degree of flattening/non-linearity of the short-run trade-off and the rate of inflation. Sustained periods of low nominal-labour productivity growth (not low inflation per se) are the key to driving up *S* in the Akerlof, Dickens, and Perry Phillips curve.

Figure 1 shows a stylized Akerlof, Dickens, and Perry Phillips curve. The long-run Phillips curve is labelled LR. It is vertical if nominal-labour productivity growth is greater than b. Below this point there is a trade-off between the long-run NAIRU and nominal-labour productivity growth.

The short-run Phillips curve is flatter in the binding DNWR zone (SR2) than in the non-binding zone (SR1) and is somewhat non-linear, but not kinked at the zero output gap, point C. Flattening occurs at point A even though the labour market is in excess demand. Testing for a kink at point C or for a flattening that is limited to the zone of excess supply does not test whether DNWR is binding. While the authors cannot reject linearity, we have no way of knowing whether they can reject the degree of non-linearity implied by binding DNWR or how pronounced this non-linearity should be.

Figure 1
Akerlof, Dickens, and Perry Phillips curve

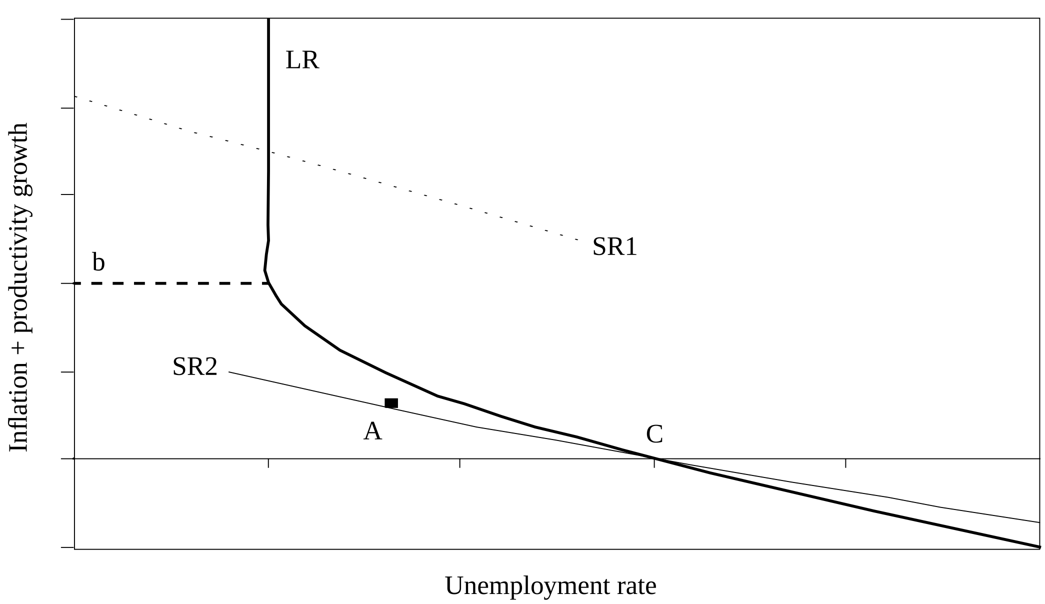

Pierre Fortin, in his 1991 Innis lecture on the Phillips curve, stated that, "detecting parameter instability and explaining its sources remains a major task of macroeconomic research." This is doubtless still true today, and the Beaudry-Doyle paper provides a useful contribution to advancing this research, particularly in providing an interesting theoretical underpinning for distinguishing between statistical and exploitable trade-offs between inflation and output. It is doubtful, however, that their empirical work can allow us to distinguish between a variety of competing explanations for this instability. The very existence of this instability has led authors such as Cecchetti and Rich (1999) to conclude that "while a better understanding of the true costs of disinflation would be of particular interest and importance to policymakers, we are skeptical that current data and econometric techniques can provide a meaningful set of estimates." This pessimism will doubtless not deter other researchers in the future.

References

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

Ball, L., N.G. Mankiw, and D. Romer. 1988. "The New Keynesian Economics and the Output-Inflation Trade-off." *Brookings Papers on Economic Activity* 1: 1–65.

Cecchetti, S. and R. Rich. 1999. "Structural Estimates of the U.S. Sacrifice Ratio." Federal Reserve Bank of New York, March.

Cozier, B. and G. Wilkinson. 1991. *Some Evidence on Hysteresis and the Costs of Disinflation in Canada*. Technical Report No. 55. Ottawa: Bank of Canada.

Fortin, P. 1991. "Innis Lecture: The Phillips Curve, Macroeconomic Policy, and the Welfare of Canadians." *Canadian Journal of Economics* November XXIV (4): 774–803.

Lipsett, B. and S. James. 1995. "Interpreting Sacrifice Ratios Across Countries and Over Time." Department of Finance Working Paper No. 95–06.

Discussion

Jeffrey Fuhrer

Overview

This paper tackles two perplexing questions, the first of fairly recent origin, the second of long standing. The first question is largely empirical: What has happened to the Phillips curve in the 1990s? The second is theoretical: Where does the statistical regularity that we call the Phillips curve come from? And can an understanding of its underpinnings help us explain its recent behaviour?

The paper first develops empirical evidence documenting a decline in the slope of the Phillips curve in the 1990s. It then explores a model in which the decline in the slope of the Phillips curve arises because the central bank's noise-to-signal ratio in identifying productivity shocks has declined. As a result, the Fed is better able to accommodate such real shocks, and their impact on inflation is reduced, attenuating the statistical link between output and the change in inflation.

I will take some issue with both the empirical and theoretical points made in the paper.

Empirical Issues

Beaudry and Doyle present a number of rolling-regression estimates of Phillips-curve slopes with standard-error bands. They examine Canadian and U.S. data separately, as well as a data set that pools both countries' inflation and output data. In most cases, the estimated slope coefficient declines in the 1990s.

But nowhere does the paper conduct formal tests of the hypothesis of stability of the Phillips-curve coefficients. In addition, while the focus of this paper is on the slope of the Phillips curve, a good deal of recent literature suggests that the *intercept* of the Phillips curve may also have shifted. Recent examples include the time-varying NAIRU papers of Staiger, Stock, and Watson (1997), Gordon (1997), and Brainard and Perry (1999).

I examine U.S. data on inflation, output, and relative oil prices to see if formal tests reject stability for the Phillips curve. I look at both simple, known breakpoint tests that presume a breakpoint starting in 1990, and more rigorous, unknown breakpoint tests that test for multiple breakpoints at all possible locations in the data set. These tests are used in Estrella and Fuhrer (2000), and draw on work in Bai (1999).

Table 1 presents results for U.S. quarterly and annual data, for two measures of prices and output. I compute the tests for both the GDP deflator and the core CPI, and for an H-P-filtered output gap ($\lambda = 1600$) and the unemployment rate. In the known breakpoint tests, I test separately for a break in the slope or the intercept. In the unknown breakpoint tests, I test jointly for a break in the slope and intercept. I consider both the "simple" Phillips correlation model, as well as the "dynamic" model that allows for more lags of inflation and output, and accounts for the influence of relative oil prices.

Interestingly, the unknown breakpoint test for Canadian data (the bottom panel) finds *no* evidence of a breakpoint at any date.

The results suggest the following:

- The known breakpoint tests develop no evidence of behaviour that differs in the 1990s versus the pre-1990 period, for either slope or intercept.
- The unknown breakpoint tests find evidence of breaks in both the simple and dynamic Phillips curves, but none later than the early 1980s.
- There is little support in U.S. data for a significant shift downwards in either the slope or the intercept of the Phillips curve.

These tests do not rule out the possibility of a shift in Phillips parameters in the last few years of the 1990s, although developing statistical significance for such a recent shift would likely be difficult.

The Unit Sum Constraint

In their more sophisticated Phillips regressions, Beaudry and Doyle do not impose the constraint that the sum of the lagged inflation coefficients be unity. I find that rolling-regression results behave in a qualitatively different manner when the unit sum constraint is imposed. Figure 1 shows my rolling-regression results for U.S. data (the GDP deflator and the H-P-filtered output

Table 1
Breakpoint test results

Known breakpoint tests (U.S. data)		
	***t*-statistic for break in 1990s**	
Variable	**Slope**	**Intercept**
Annual data		
Simple model: $\Delta\pi_t = \beta y_t$		
GDP deflator, H-P-filtered output	–0.2	–0.5
Core CPI, unemployment	0.8	–0.9
Dynamic model: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma X_t$		
GDP deflator, H-P-filtered output	–0.9	–1.6
Core CPI, unemployment	0.8	–1.0
Quarterly data		
Simple model: $\Delta\pi_t = \beta y_t$		
GDP deflator, H-P-filtered output	0.3	0.0
Core CPI, unemployment	0.5	–0.6
Dynamic model: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma X_t$		
GDP deflator, H-P-filtered output	0.4	–1.1
Core CPI, unemployment	0.3	–0.5
Unit sum imposed: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma \mathrm{X}_t$		
GDP deflator, H-P-filtered output	–0.3	–0.2
Core CPI, unemployment	0.2	–0.4

Unknown breakpoint tests (U.S. data)		
Variable	**Number of breakpoints**	**Dates**
Annual data		
Simple model: $\Delta\pi_t = \beta y_t$		
GDP deflator, H-P-filtered output	0	—
Core CPI, unemployment	0	—
Dynamic model: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma X_t$		
GDP deflator, H-P-filtered output	1	78
Core CPI, unemployment	2	74, 82
Quarterly data		
Simple model: $\Delta\pi_t = \beta y_t$		
GDP deflator, H-P-filtered output	0	—
Core CPI, unemployment	0	—
Dynamic model: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma X_t$		
GDP deflator, H-P-filtered output	4	67:2, 72:4, 77:1, 81:1
Core CPI, unemployment	4	70:1, 74:1, 79:4, 83:4
Unit sum imposed: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma \mathrm{X}_t, \Sigma_i \alpha_i = 1$		
GDP deflator, H-P-filtered output	2	75:1, 79:2
Core CPI, unemployment	3	73:3, 80:2, 84:2
Quarterly data: Canada		
Dynamic model: $\pi_t = \Sigma_i \alpha_i \pi_{t-i} + \Sigma_k \beta_k y_{t-k} + \gamma X_t$		
GDP deflator, H-P-filtered output	0	—

Figure 1
Rolling regressions, U.S. Phillips curves

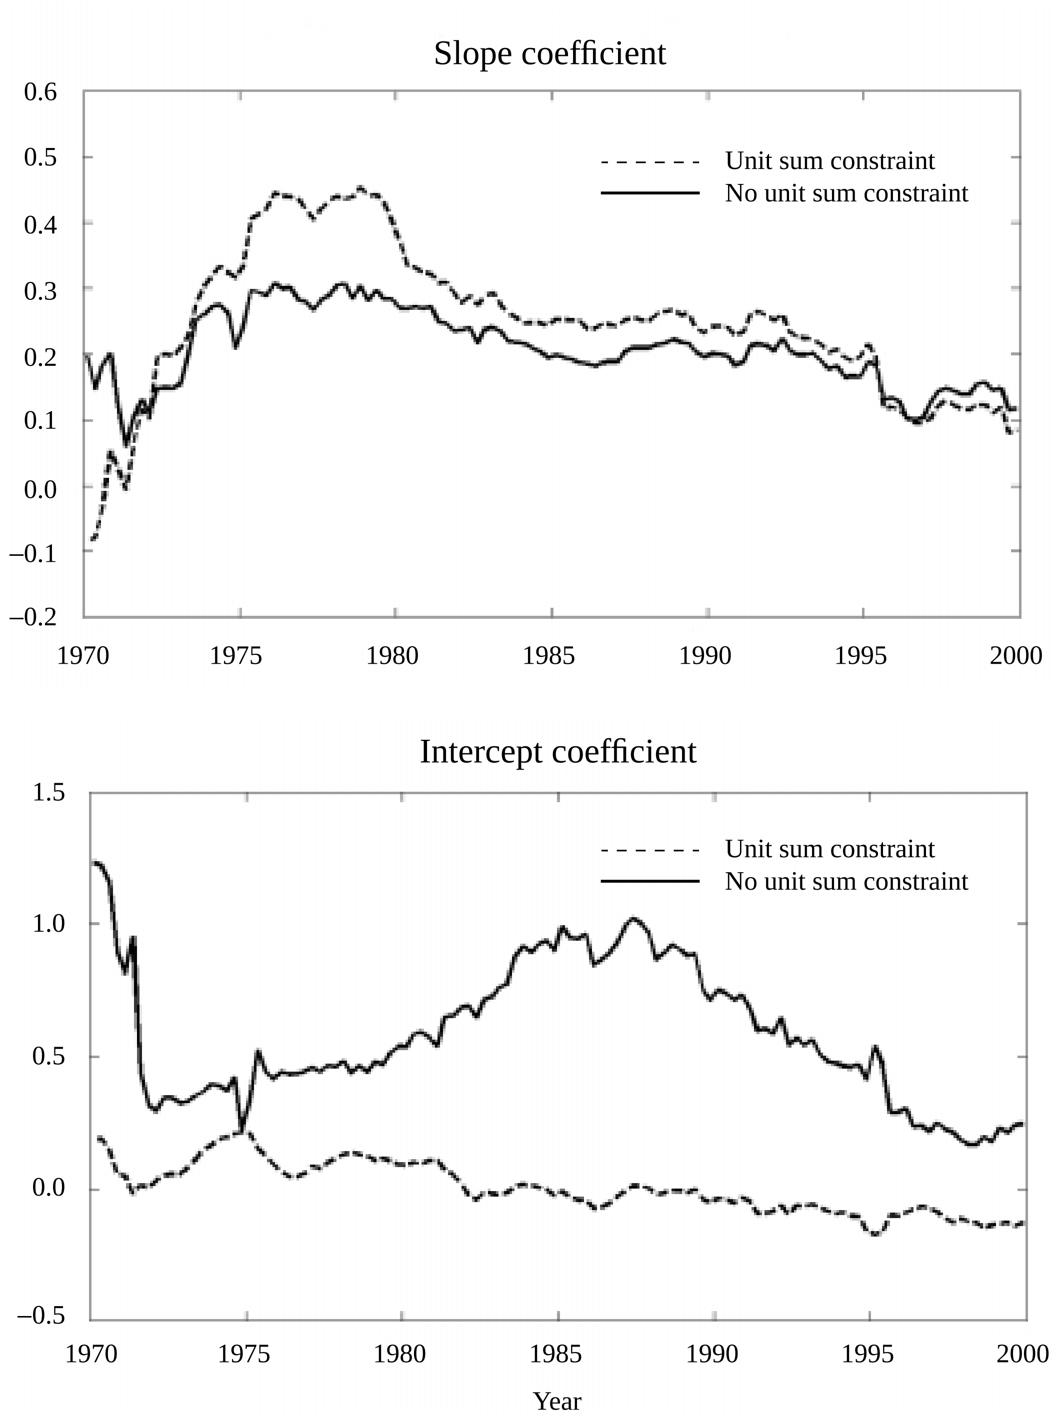

gap), with and without the unit-sum constraint (the broken and solid lines, respectively). The top panel shows the slope estimates, and the bottom panel shows the intercept.

Without the constraint, as in Beaudry and Doyle, most of the variation is in the *intercept*, not the slope. The reverse is true for the constrained model. This points to my first concern about the empirical results, which is that they may not be robust to the imposition of this relatively standard Phillips-curve constraint.

Second, I am concerned about the interpretation of unconstrained lags of inflation in the Phillips curve. These lags may simply reflect the pure reduced-form time-series properties of inflation. But the sum of the unconstrained lags can have little to do with whether the Phillips curve is vertical or not in the long run. Instead, they may simply reflect the reduced-form behaviour of inflation that arises from the joint behaviour of a long-run vertical Phillips curve and monetary policy. For example, in a model in which the Phillips curve imposes the unit constraint, if the central bank targets the *change* in inflation, then the unconstrained lags will sum to close to one. But if instead the central bank aggressively targets inflation, then the sum of the unconstrained lags on inflation can be well below one. In neither case does the sum of the unconstrained lags tell you whether the unit-sum constraint should be imposed.

Theoretical Issues

How Does the Model Work? The Simple Version

A simple way to understand the Beaudry and Doyle model is to start with the quantity equation: $M\,V = P\,Y$. For simplicity, assume a constant V (there is nothing in the model to change it). The only shocks in the model are transitory (around long-run trend) but persistent productivity shocks.

In a perfect world, a positive productivity shock would expand real supply, and the central bank would completely accommodate it, raising M one-for-one with Y (this is the definition of an “elastic currency”). In the Beaudry-Doyle world, the central bank can’t see the productivity shock perfectly. So it filters the noisy measure of the productivity shock in the standard way and interprets the observed shock as part productivity, part noise.

As a result, it moves M up by less than one-for-one with the productivity shock. The quantity equation implies that prices must *fall*. In fact, this scenario sounds much like the behaviour of the economy following the large productivity shocks of the nineteenth century, which were imperfectly accommodated. The result was massive deflation.

So how does this model generate a positive Phillips-curve correlation between the change in inflation and output? In the model, as the price level drops, inflation is negative for one period. As soon as the expectation error is revealed (one period later), inflation rises to its long-run target (see equation (13)). So while output is rising, the *change* in inflation is positive (for some of the time).

In Figure 2, a simulation of the model (Beaudry-Doyle equations (13–15)) confirms this simple intuition. With a positive productivity shock, output (the dotted line) rises with some persistence (given the MA process that drives output, see equation (14)). Here I have chosen an MA(4) with stationary coefficients, and set θ to 0.5. The inflation target is deterministic, and is arbitrarily set to zero. The price level (the solid line) drops after one period and stays down (the size of the drop depends on θ). Inflation (the dashed line), which is always at its long-run target but for a one-period expectation error, drops for one period, and then rises back to its target. The change in inflation, shown in the dash-dot line, turns from negative to positive. Because the Phillips-curve correlation is defined as the correlation between $\Delta\pi_{t+1}$ and y_t, the model predicts a mildly positive correlation between output and the change in inflation. In this simulation, the correlation is about 0.15.

While this mechanism for producing a Phillips-curve correlation "works," it seems both counterintuitive and counterfactual. While the Phillips-curve correlation remains positive through most of the post-war period, we have never seen a sustained episode during which (i) productivity surged, boosting output, but (ii) the monetary authority imperfectly accommodated, so (iii) the price level fell, but (iv) at a declining rate, so that (v) the second difference in inflation was positive while output rose. These statements are true whether they apply to raw or detrended data. While the model gets the $\Delta\pi$–y correlation right, it does so by implying an unusual set of price, inflation, and output dynamics that we have yet to observe over the past 60 years.

Finally, I am surprised that the authors suggest that a *lower* noise-to-signal ratio is the explanation for the 1990s decline in the slope of the Phillips curve. This is a direct implication of their model, because the less noise present in real shocks, the more perfectly the central bank can accommodate real shocks, and the less they will show up in price movements, all of which implies a flatter Phillips curve.

My impression is that, if anything, the noise-to-signal ratio, particularly for productivity shocks, is *higher* now. At best, it is no lower than it has been over the past 20 years. Certainly a good deal of (public) Fed discussion has centred on the sustainability and sources of productivity growth. But this, of course, goes in the opposite direction from the Beaudry and Doyle

Figure 2
Simulation of Beaudry-Doyle model, shock to ε
θ = 0.5, productivity shocks ~ MA(4), $\bar{\pi}$ = 0

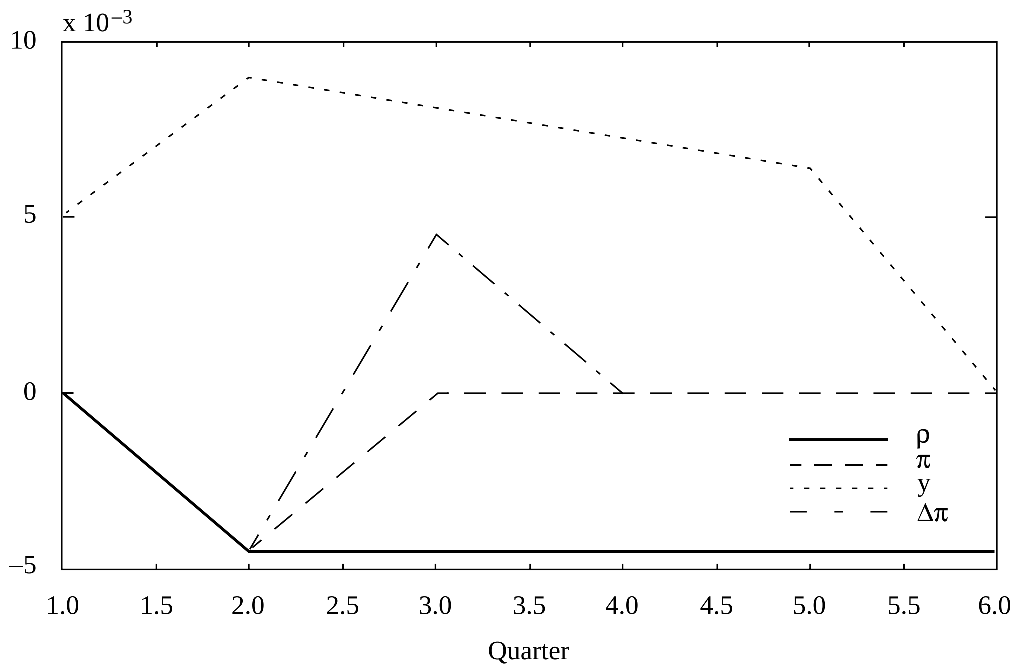

Simulation of equations 13–15 with parameters as indicated.

conclusion. Putting aside concerns over the model, a higher or unchanged noise-to-signal ratio for productivity shocks would argue that we should see a *steeper*, not flatter, Phillips curve.

A Few Minor Points

The authors discuss and dismiss alternative explanations for the recent benign behaviour of inflation. One possibility is heightened central bank attention to price stability. But an increased responsiveness to inflation should, in conventional models, produce more output variability. As Beaudry and Doyle point out, we don't see any such increase.

But this point holds only if the variance and composition of shocks remain constant across time. An alternative explanation for the pattern of variances that we have observed recently is that the variance of "price" shocks has declined, leading to low variance in both prices and output. Central banks were fortunate in the 1990s in facing relatively few and small price shocks, resulting in low variances for both output and inflation.

Beaudry and Doyle attribute central banks' improved knowledge about productivity shocks to the contributions of the real business cycle literature, and to the rational expectations literature that "questioned the potential for monetary policy to have systematically large effects on the real economy" (p. 73). I will not comment on the first premise. But as for the second, rational expectations per se do not rule out large and persistent effects of monetary policy on the real economy. That result depends critically on the model structure within which rational expectations are embedded.

Summary

This paper raises some serious issues. I am quite sympathetic to the possibility that the Phillips-curve slope or intercept may have shifted of late. But proof of that point will be found in more serious empirical testing of shifts in Phillips relationships. My first pass at such tests suggests that it will be hard to find statistically significant shifts in the 1990s. If such a shift occurred more recently, we may need more time and data before we can detect a statistically significant shift.

Similarly, I have no doubt that the Fed—like everyone else—finds it difficult to identify persistent productivity shocks. This may well be an important aspect of the historical monetary policy response, as Orphanides (1999) has suggested. While this model can produce a positive Phillips correlation from misperceptions about productivity shocks, it does so through a set of price dynamics that are difficult to reconcile with the historical behaviour of prices, productivity, and output.

Finally, the model can produce the required shift in the Phillips slope only if the noise-to-signal ratio has declined. Given the recent data on productivity in the United States and the ensuing public policy discussions, a decline in the noise-to-signal ratio seems counterfactual.

References

Bai, J. 1999. "Likelihood Ratio Tests for Multiple Structural Changes." *Journal of Econometrics* 91 (2): 299–323.

Brainard, W. and G.L. Perry. 1999. "Making Policy in a Changing World." In *Economic Events, Ideas, and Policies: The 1960s and After,* edited by G.L. Perry and J. Tobin, Chapter 2. Brookings Press.

Estrella, A. and J. Fuhrer. 2000. "Are 'Deep' Parameters Stable? The Lucas Critique as an Empirical Hypothesis." Federal Reserve Bank of Boston Working Paper No. 99–4, May 1999, revised September 2000.

Gordon, R.J. 1997. "The Time-Varying NAIRU and Its Implications for Economic Policy." *Journal of Economic Perspectives* 11 (1): 11–32.

Orphanides, A. 1999. "The Quest for Prosperity Without Inflation." Working Paper No. 93, Board of Governors of the Federal Reserve System.

Staiger, D., J.H. Stock, and M.W. Watson. 1997. "The NAIRU, Unemployment and Monetary Policy." *Journal of Economic Perspectives* 11 (1): 33–49.

General Discussion*

Paul Beaudry thanked the discussants for their comments and responded first by indicating that even though he and Doyle did not estimate the decline in the slope of the Phillips curve with precision, this decline is, nevertheless, apparent in public debate and in the Bank of Canada's own work, making it a reasonable null hypothesis. In addition, the nature of this break, whether gradual or sharp, was not obvious to him. Second, he argued that, for the recent past, a declining signal-to-noise ratio is more plausible than an increasing one, since the distance between central bank and private sector perceptions is smaller. Third, he mentioned that the authors investigated for the presence of a kink in the Phillips curve (the Akerlof, Dickens, and Perry hypothesis) as one possible nonlinearity and found no support for it, and that it is difficult to find other types of nonlinearities in this curve for the period covering the last 15 years. Fourth, he wondered whether using other shocks in the model would have been constructive, since, for example, the response of the Bank to the negative shocks of the 1970s by policy tightening was not what they observed. Finally, Beaudry agreed that the sacrifice ratio was indeed relevant, that it was strongly related to the other inflation equation in their model, and that it should be examined carefully.

Charles Freedman pointed out that inflation expectations are partially forward-looking instead of entirely backward-looking, as Beaudry and Doyle assume. In this case, the change in the Phillips curve might have arisen from changing expectations; that is, as credibility increased, expectations became more anchored. He then raised a theoretical point: in the model the Bank considers only supply shocks, yet many other shocks (fiscal,

* Prepared by Maral Kichian.

restructuring of the private sector, Asian crisis) have also been important over the period examined. He added that while productivity has been a significant factor in the United States in the last three years, there is little evidence of increased productivity in Canada, and that demand shocks have played an important role for the Bank.

Peter Howitt suggested an alternative interpretation in the observed trade-off between inflation and output, saying that it is not that the Bank interprets shocks better, but rather that agents now know what works and what does not. It is the gradual evolution of how central banks are picking up points on the variance of the Phillips curve, which is not related to the noise-to-signal ratio.

Nicholas Rowe indicated that if the Phillips curve is described as a supply curve and the Bank's reaction function as a demand curve, the change in the Bank's preferences (i.e., the θ parameter) is what leads to a flatter estimated Phillips curve.

Beaudry responded to Freedman's comments saying that, by looking at the statistical Phillips-curve slope, defined with respect to output only, the authors were trying to be consistent with their theoretical model, adding that they dealt with the expectations in a separate equation in their model. As for the comment about demand shocks, he agreed that other shocks have occurred and that the model shock could be interpreted as a composite shock term that could be made more explicit.

Responding to Howitt's comment, Beaudry said that good monetary policy leads to a better economy, and while this was one possible explanation for a flatter Phillips curve, it is not too far from other interpretations.