What Happened to The Phillips Curve in the 1990s in Canada?[1]

Paul Beaudry [2] and Matt Doyle [3] [4]

June 2000

[1] We wish to thank Jeff Fuhrer, Steven James and Pierre Duguay for helpful comments. The responsibility for any remaining errors is our own.

[2] University of British Columbia, Canadian Institute for Advanced Research and N.B.E.R.

[3] University of British Columbia

[4] Please address comments to Beaudry & Doyle, Dept. of Economics, University of British Columbia, 1873 East Mall, Vancouver, B.C., V6T 1Z1, or paulbe@interchange.ubc.ca

Abstract

This paper begins by reviewing the empirical properties of the Phillips Curve in both Canada and the U.S over the last forty years. In particular, we document the extent to which the slope of the Phillips Curve has declined in both countries over the nineties. Then, building upon a commonly used macromodel, we attempt to explain this decline. The framework we develop focuses on the nature of the Phillips Curve when monetary authorities are imperfectly informed about real developments in the economy but nevertheless try to set monetary policy optimally. Our model explicitly recognizes two distinct activities performed by the central bank. On one hand, the central bank tries to provide sufficient liquidity to help private agents exploit gains from trade during periods in which prices are pre-set. On the other hand, the central bank also performs an information-gathering role as it continuously tries to infer the state of the economy. We show how this dual role gives rise to a Phillips Curve relationship which both exhibits causality running from output to prices and justifies a feedback from prices to the setting of monetary instruments. Based on this model, we argue that the observed flattening of the Phillips Curve may be the result of improvements in the manner in which central banks gather information regarding real forces affecting the economy, and that the flattening is not a reflection of a change in the output-inflation tradeoff faced by the central bank. Finally, we compare our proposed explanation of the flattening of the Phillips curve with leading alternative hypotheses.

Key Words: Optimal Monetary Policy, Phillips Curve, Output-Inflation tradeoff

1 Introduction

When examining observations on a country's output-inflation relationship (a country's Phillips Curve), there are at least two broad classes of interpretation. On one hand, there 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 the slope of the Phillips Curve 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, there is the view that the Phillips Curve is essentially a reduced form relationship, which mainly reflects rather than constrains the behavior of monetary authorities. In this paper we will argue that this second view helps explain recently observed changes in the Phillips Curve.

We begin the paper by reviewing the changing nature of the Phillips Curve relationship in both Canada and the U.S. from 1961-1999. We define the Phillips Curve 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 U.S the slope of the Phillips Curve has become much smaller over the last twenty years with a sharp reduction observed in the nineties. This observation raises two related issues: (1) what explains the decline in slope and (2) 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 seventies, central banks have continuously 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 sign of well executed, neutral stance, monetary policy.

Our explanation of the flattening of the Phillips Curve is presented in a simple model which 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 & Kiyotaki (1987) in a manner which 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 & Gordon (1983)), the information asymmetry we emphasize is such that the central bank is imperfectly informed regarding real developments in the economy and hence the central bank is constantly trying to infer the state of the economy while simultaneous 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.

Within 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 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 objective of monetary policy should be to simultaneously support the well functioning of the economy and maintain price stability. However, in the 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.

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

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

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 understand the observed flattening of the Phillips Curve over the last twenty years, as the central banks in both the U.S and Canada first became aware of the importance of real shocks in the seventies, and then learned to identify and react to them more appropriately throughout the eighties and nineties.

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

The remaining sections of the paper are structured as follows. In Section 2, we document the changing nature of the Phillips Curve for both the U.S and Canada over the period 1961-1999. In Section 3, 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 to one based on nominal wage rigidities. Finally, in Section 4 we offer concluding comments.

2 Overview of the Output-Inflation Relationship in Canada and the U.S

In this section we review the evidence related to the existence of a positively sloped Phillips Curve for both the U.S and Canada over the period 1961 to 1999. In particular, we present evidence to suggest that the Phillips Curve relationship is robust to various specifications and is roughly similar in Canada and the U.S.[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 20 years. We find that the reduction in slope, which has occurred in both the U.S and Canada, is quite substantive.

2.1 Basic Estimation and Results

In its simplest form, the Phillips Curve can be expressed as a relationship between inflation, lags of inflation, 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 which try 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 HP filters, structural VARs, structural macroeconomic models and simple time trends, to infer the trend level of output. We explored several alternatives and choose as our baseline measure the output gap series created by applying an HP 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]

---

[3]Fillion and Léonard(1997) present linear Phillips Curve estimates for Canada which resemble our own.

[4]We also used the CPI as an alternate measure of prices in order to check the robustness of our results.

[5]We have verified that the results are robust to various values of λ, the smoothing parameter of the HP 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 take our data for Canada from Cansim, and our U.S data from Basic Economics (formerly Citibase).

Table 1: Basic Phillips Curve Estimates for Canada: 1961-1999

	$\Delta\pi_t$	$\Delta\pi_t$	$\Delta\pi_t$
Constant	-0.0360	-0.0690	0.0021
(Std. Error)	(0.2259)	(0.2231)	(0.2219)
HP-GAP$_t$	0.2141	-	-
	(0.0828)		
HP-GAP$_{t-1}$	-	0.1890	-
		(0.0814)	
T-GAP$_{t-1}$	-	-	0.2599
			(0.0998)

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

$$\Delta\pi_t = \alpha + \beta \cdot GAP_t + \epsilon_t.$$

In Figures 1 and 2 we plot this relationship, along with regression line, for Canada and the U.S from 1961 to 1999. The slope of the estimated Phillips Curve for the U.S is 0.256, suggesting that a positive output gap of 1 percent is associated with an increase in inflation of around one quarter of one percent on average. The Canadian estimate of 0.214 is similar to that of the U.S. In both countries, we reject the hypothesis that the slope of the Phillips Curve is zero at conventional levels. To allow for the possibility that inflation responds to real developments with some delay, in column 2 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.

In order to illustrate the robustness of these results, we consider a variety of alternate 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 level of output into trend and gap components using an HP filter, we wish to repeat our analysis using alternate detrending

Table 2: Basic Phillips Curve Estimates for the U.S.: 1961-1999

	$\Delta\pi_t$	$\Delta\pi_t$	$\Delta\pi_t$
Constant	-0.0014	-0.0531	0.0133
(Std. Error)	(0.1567)	(0.1567)	(0.1576)
HP-GAP$_t$	0.2560	-	-
	(0.0650)		
HP-GAP$_{t-1}$	-	0.2404	-
		(0.0619)	
T-GAP$_{t-1}$	-	-	0.2719
			(0.0714)

methods. In column 3 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 :

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

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 column 1 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 U.S. 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 there is a positive and statistically significant co-movement between output and changes in inflation over the period 1961-

Table 3: Phillips Curve Estimates for Canada: 1961-1999

	$\Delta\pi_t$	$\Delta\pi_t$	π_t	π_t
Constant	-0.0787	-0.0184	1.2991	1.7040
	(0.2310)	(0.2319)	(0.4042)	(0.5207)
HP-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)

1999.

Respecifying the problem in terms of inflation rather than the change in inflation, as shown in columns 3 and 4 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 on 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 lags of inflation 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 supports the

[7]We define inflation in relative energy prices as the percentage change in the ratio of the CPI for energy to the all items CPI.

Table 4: Phillips Curve Estimates for the U.S.: 1961-1999

	$\Delta\pi_t$	$\Delta\pi_t$	π_t	π_t
Constant	-0.0559	-0.0299	0.4776	1.2633
	(0.1556)	(0.1474)	(0.3095)	(0.2969)
HP-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)

existence of a positively sloped Phillips Curve in both Canada and the U.S, and that this observation is robust to alternate specifications. Our estimates of the slope of the Phillips curve vary mostly between 0.2 and 0.3. In all cases, the estimated slope of the Phillips Curve is positive and significantly different from zero at conventional levels.

2.2 The Changing Slope of the Phillips Curve

Having reviewed the case for the existence of a positive co-movement between inflation and output in both Canada and the U.S from 1961 to 1999, we now turn our attention examining whether the Phillips Curve relationship may have changed over time. As we will show, the slope of the Phillips Curve in both Canada and the U.S has declined markedly from its peak in the late 1970's.

To examine the slope of the Phillips Curve relationship over time, we employ 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 years. For example, the estimates for 1983 are derived from observations over the period 1969-1983.

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 as it is easily tied in to the theoretical results presented in subsequent sections. Note that the estimated slope of the Phillips Curve peaks around 1982 in both countries. The slope then declines throughout the 1980's and 1990's. In the U.S 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 1990's. By the end of the 1990's, the slope of the Phillips Curve is not significantly different from zero in either Canada or the U.S.

As 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 1980's in both Canada and the U.S 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 also have been quite low in the mid to late 1970's. This implication, however, is not robust to the choice of estimation framework.

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

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

the chance that one or two observations might unduly influence the profile of our estimates. We found that this approaches yields similar results to that presented in Figures 3 to 6. That is, the slope of the Phillips Curve appears to decline substantially over the 1980's and 1990's.

The point estimates presented in the above figures 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 tradeoff: 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 alternate 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 1980's, 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-1980's, 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

[9] That is we regress the change in inflation on the lag of the output gap.

conventional levels. As can be seen from the figure, one outlier drives 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 U.S and Canada since 1985 to be merely 0.0212.

The evidence presented above leads us to believe that the Phillips Curve relationship has changed significantly in recent decades. Particularly, we find that the Phillips Curve has flattened substantially in both the U.S and Canada. Furthermore, at least in Canada, this flattening has occurred markedly in the 1990's.

3 Why is there is 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 upon a commonly used monopolistic competition macro model (see, for example, Blanchard & 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 preset 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 CRS production function given in (3.1).

$$Y_t = (\sum_{i=1}^{N} X_{it}^{\alpha})^{(\frac{1}{\alpha})} N^{(\frac{1-\alpha}{\alpha})} \tag{3.1}$$

Each intermediate goods firm has access to a production technology given by (3.2).

$$X_{it} = A_t^{(1-\gamma)} L_{it}^{\gamma} \tag{3.2}$$

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 (3.3).[10]

$$a_t = \sum_{j=0}^{\infty} \psi_j \epsilon_{t-j}, \qquad \psi_0 = 1, \qquad \sum_{i=1}^{\infty} \psi_i^2 < \infty \tag{3.3}$$

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.

In order 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 (3.4). 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 an exogenously fixed real wages.

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

$$U(C_t, \frac{M_t}{P_t}, L_t) = C_t^{\theta} \frac{M_t}{P_t}^{1-\theta} - \phi L_t \tag{3.4}$$

The household's budget constraint is given by (3.5), where P_t is the price of the final good, W_t is the nominal wage rate, $\bar{M}_t$ is money demanded and M_t is the money balances distributed by the central bank at the beginning of each period.

$$P_t C_t + \bar{M}_t = W_t L_t + M_t \tag{3.5}$$

In order to solve for the private sector's equilibrium behavior, we start by examining the household's decision problem. The representative household takes prices as given and chooses consumption, labour and money balances so as to maximize utility. The first order conditions associated with this maximization imply that money demanded satisfies equation (3.6), and that labour is supplied elastically at the real wage given by (3.7).

$$\bar{M}_t = P_t C_t \frac{(1-\theta)}{\theta} \tag{3.6}$$

$$\frac{W_t}{P_t} = (1-\theta)^{1-\theta} \theta^{\theta} \phi \tag{3.7}$$

Final good producers 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 (3.8), where P_{it} is the price of the ith intermediate good.

$$X_{it} = (\frac{P_{it}}{P_t})^{\frac{-1}{1-\alpha}} \frac{Y_t}{N} \tag{3.8}$$

The problem facing an intermediate good firm 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. An intermediate good producer's problem can therefore be expressed as follows:

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

$$s.t. \quad (3.2), (3.3), (3.7), (3.8)$$

Using the market clearing conditions for both the goods market and the money market, and imposing symmetry on the behavior of intermediate goods producers, one can easily derive Equations (3.9) and (3.10) which describe the behavior of the aggregate price level and aggregate output.[11] In these two equations, constant terms have been dropped.

$$p_t = E[m_t/\Omega_{t-1}] - \sum_{i=1}^{\infty} \psi_i \epsilon_{t-i} \tag{3.9}$$

$$y_t = m_t - p_t = (m_t - E[m_t/\Omega_{t-1}]) + \sum_{i=1}^{\infty} \psi_i \epsilon_{t-i} \tag{3.10}$$

Equations (3.9) and (3.10) represent the equilibrium behavior 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. It is important to note that the ϵ_t's in Equations (3.9) and (3.10) 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 preset prices macro model and generates

[11] In order to derive equation (3.9) from the intermediate good firm's problem, it is easiest to first use (3.2) and (3.7) to eliminate W_t and L_{it} from the firm's objective function. Then, using the market clearing conditions $C_t = Y_t$ and $M_t = \bar{M}_t$ in combination with (3.6) and (3.8), the demand facing the firm can be written as simply a function of P_{it}, P_t and M_t. Finally, by 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 (3.9).

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 the constraints facing the central bank, and highlight the key elements of our model.

We assume that the objective of the central bank is to minimize deviations of output and prices from target levels y_t^* and p_t^*, as given by (3.11).

$$\sum_{i=0}^{\infty} \beta^i E[(y_t - y_t^*)^2 + \Phi \cdot (p_t - p_t^*)^2 | \bar{\Omega}_t] \tag{3.11}$$

In (3.11), Φ is the weight the central banker places on deviations of inflation from its target relative to output deviations.

With respect to the target for output, 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, $y_t^* = \sum_{i=0}^{\infty} \psi_i \epsilon_{t-i}$ (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 at hand. With respect to the target for the price level, we assume that it is driven by an exogenously given 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 as being either stochastic or deterministic; the key simplifying assumption being that it is exogenous. In order to allow for the possibility that the inflation target be 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 targets follow a known exogenous process, we are obviously sidestepping important issues related to the signaling 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

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

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 an informational disadvantage but a timing advantage relative to the private sector. The central bank's disadvantage is that it does not directly observe the ϵ_t's, and therefore must infer their values from past developments in the economy. Its advantage is that it has some information on the current state of the economy, which it can use during the period over which prices are preset.

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 (3.12), where μ_t is a normally distributed mean zero random variable with variance σ_μ^2. We denote by τ^2 the noise to signal ratio $\frac{\sigma_\mu^2}{\sigma_\epsilon^2}$.

$$s_t = \epsilon_t + \mu_t \tag{3.12}$$

The timing of moves is as follows. At the beginning of a period, intermediate good firms 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 different. Private agents know all past developments in the economy but do not know the realization of ϵ_t that is to arise during the period. In contrast, the central bank has past information only on output and prices (not the ϵ's), 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 $\Omega_{t-1} = \{\epsilon_{t-1}, \ldots, s_{t-1}, \ldots, p_{t-1} \ldots, y_{t-1} \ldots\}$.

Our justification for giving the central bank an informational advantage through s_t is meant to capture a timing advantage by which the central bank has more flexibility within a period to react to current shocks than the private sector (since the private sector has pre-set prices) but the central bank is nevertheless imperfectly informed regarding the right way to react.

The problem facing the central bank is to choose a monetary policy rule so as to minimize (3.11) subject to its informational restrictions and the optimizing behavior of the private economy, given by equations (3.9) and (3.10). 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 setup 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 which solves the central banks problem is given by equation (3.15), with the implied equilibrium solution for inflation (π_t) and output given by equations (3.13) and (3.14) receptively.

$$\pi_t =_{t-1} \bar{\pi}_t + \psi_1(\Theta s_{t-1} - \epsilon_{t-1}), \qquad \Theta = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_\mu^2} \tag{3.13}$$

$$y_t = \Theta \cdot s_t + \sum_{i=1}^{\infty} \psi_i \epsilon_{t-i} \tag{3.14}$$

$$m_t = p_{t-1} +_{t-1} \bar{\pi}_t - \frac{\psi_2}{\psi_1}(\pi_{t-1} -_{t-2} \bar{\pi}^*_{t-1}) + \Theta \sum_{i=0}^{2} \psi_i s_{t-i} + \sum_{i=3}^{\infty} \psi_i \epsilon_{t-i} \tag{3.15}$$

In order to gain intuition about equations (3.13)-(3.15), 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 so as to reflect its best guess of the current supply shock. Given that 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 only deviates from the target level of inflation 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 (3.15), where the past deviation of inflation from the target level enters negatively in the money supply rule. Although monetary authorities never directly observe the ϵ's, 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's. [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 we did in Section 2, as a purely statistical object. 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 covariance between the change in inflation and output, divided by the variance of output. The analytical expression for this slope is reported in equation (3.16) and is

[13] Using the money supply rule (3.15) to calculate expected and unexpected money, it is rather simple to verify that equations (3.13) and (3.14) are consistent with private agents optimal behavior given by Equations (3.9) and (3.10). It is only slightly more difficult to verify that the money supply rule given by (3.15) is in effect optimal. To see this, 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 it realization. Hence, since this difference is minimized by setting $\Theta = \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2+\sigma_\mu^2}$, this confirms the optimality of the policy.

denoted by β.

$$\beta = \frac{cov(\Delta\pi_{t+1}, y_t)}{var(y_t)} = \frac{\psi_1^2\tau^2}{(\sum_{i=1}^{\infty}\psi_i^2)(1+\tau^2)} > 0, \qquad \frac{\partial\beta}{\partial\tau^2} \geq 0 \tag{3.16}$$

Note that 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 2, 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 (3.16) is that the model generates a statistical Phillips Curve; that is, even though monetary policy is set optimally, the economy nevertheless exhibit 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 (3.16) 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 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

[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 only affect output 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.

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. The reason is that, 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 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 (3.16) 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 tradeoff faced by the central bank. In particular, 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 tradeoff could still be quite large. To see this, we can use Equations (3.13) and (3.14) to derive the short-run relationship between inflation, target inflation, output and supply shocks. This relationship is given by Equation (3.17).

$$\pi_t =_{t-1} \bar{\pi}_t^* + \psi_1 y_{t-1} - \sum_{i=1}^{\infty} \psi_i \epsilon_{t-i}, \tag{3.17}$$

In words, the term $\psi_1 y_{t-1}$ in Equation (3.17) represents the effect on inflation induced by the central bank stimulating (or contracting) output by deviating from its optimal monetary policy at one point in time. This equation nicely captures the type of short-run output-inflation tradeoff often used to discuss the short run effect of monetary shocks.[16]

The reason there is a distinction in this model between the statistical Phillips Curve and the short run output-inflation tradeoff 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

[16]The only major difference between Equation (3.17) 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.

particular, the statistical Phillips Curve tends to become horizontal precisely when monetary authorities do not try to exploit the short-run tradeoff and instead try to correlate output with 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 tradeoff 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 agent 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.

Now that we have described the functioning of the model, let us return to our question of interest: what insight does this model provide towards explaining the flattening of the statistical Phillips Curve in Canada and the U.S 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 so as to allow such real developments to take place without large price effects. In other words, the flattening of the Phillips Curve may be a reflection of improvements in the manner in which monetary policy is executed. In the rest of the paper we will present both empirical and anecdotal evidence in support of this view and compare its merits to alternative explanations.

Our first argument in favour of this view is entirely anecdotal as it reflects the change in macroeconomic thinking throughout the last 25 years and in relation to the conduct of monetary policy. Prior to the seventies, the importance of real shocks on the macroeconomy was perceived to be rather minimal. The substantial fluctuations in oil prices changed this view and lead 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 U.S) responded to these changes by focusing more effort on identifying the underlying real forces in the economy and 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 eighties and nineties, 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 upon these signals (captured by Θ in the money rule (3.15)) likely increased. These developments are exactly of the type that our model suggests would lead to a flattening of the Phillips Curve.[17]

In order to examine the plausibility of the idea that improvements in the manner in which monetary policy is conducted could be the cause behind the observed flattening of the Phillips Curve, it is useful to consider alternative explanations. One such potential explanation, often heard in the press, is that, over the eighties and nineties, 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 both the U.S and Canada since the early eighties (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 U.S and Canada, the variance of inflation growth has decreased quite

[17]In is interesting to note from Figures 3 and 4 that the period in which the statistical Phillips Curve appears steepest is in the late seventies and early eighties, which is generally thought as being a period of high variation in real shocks and substantial confusion.

Table 5: Rolling Sample Variances for Canada and the U.S: 1983-1999

	Canada		U.S	
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

substantially over the last fifteen years. In contrast, the variance of output for Canada has remained about the same while that for the U.S appears to have declined. In particular, the variance of output in Canada was about 7.5 throughout much of the eighties and is approximately at the same level by the end of the nineties.

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 slope of 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.

3.1 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 & Perry (1996) and Fortin(1997) among others, is downward nominal wage rigidity. 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 nineties relative to the early eighties.

In this section we attempt to differentiate between the downward nominal wage rigidity theory and our proposed explanation to the flattening of the Phillips Curve which relies on improved monetary policy. One implication of the downward nominal rigidity explanation, 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 downward nominal rigidity hypothesis suggests that as inflation decreases, it is mainly the segment of the Phillips Curve which relates to negative values of the output gap, which should flatten (because downward nominal wage rigidities are not relevant when the labour market is tight).

In order to explore this hypothesis empirically, we estimated several variants of the type of non-linear Phillips Curve given by Equation (3.18) and examined how the coefficients changed over time.[18]

$$\pi_{t+1} = \beta_0 + \pi_t + \beta_1(GAP_t) + \beta_2 PosGap_t + \epsilon_{t+1} \tag{3.18}$$

In Equation (3.18), the variable $PosGap_{t-1}$ takes the value of zero if the

[18] Dupasquier and Ricketts(1997) provide a good entrance point to the literature on estimating non-linear Phillips Curves for Canada.

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 (3.18) based on pooled U.S and Canadian data over periods of 15 years. We present the results for the pooled estimates since they are the most precise. However, it should be noted that we also estimated this equation for each country individually and for several different specifications, and obtained results similar to that represented in Figures 9 and 10.

As can be seen in Figure 9, the value of β_1 decreased substantially over the late eighties and nineties. Since this coefficient represents the slope of the Phillips Curve for negative values of the output gap, its decline is consistent with the hypothesis 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 which 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 downward nominal wage rigidity was the cause of the flattening of the Phillips Curve.

Our evidence against the downward nominal rigidity hypothesis can be inferred visually from the simple scatter plot presented in Figure 8. Since the mid 1980's, 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 out-lying 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 to the observed flattening of the Phillips Curve.

[19] Our approach is to estimate a Phillips Curve which has a kink at a zero output gap. We adopt this simple approach to evaluate the presence of a non-linearity even though the nominal wage rigidity hypothesis does not precisely predict a kink at zero output gap.

3.2 The Flattening of the Phillips Curve and the Ball, Mankiw & Romer hypothesis

A second potential explanation to 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 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) very often. 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 twenty years, the menu costs hypothesis predicts that the Phillips Curve should have become flatter over this period, which is exactly what we observe in the data.

The menu costs explanation and our model, however, have important differences regarding the effects of monetary surprises which arise due to their respective implications for the short-run output-inflation tradeoff 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 tradeoff, such a flattening implies that the output-inflation tradeoff has increased. In contrast, while our model predicts that the statistical Phillips Curve (whose slope is given by Equation (3.16)) becomes flatter when policy makers monitor the economy properly, this does not imply that the short-run output-inflation tradeoff changes. In effect, the relevant output-inflation tradeoff –that is, the tradeoff induced by a deviation from the perceived policy rule– is governed by the parameter ψ_1 in Equation (3.17) which is independent of the trend level of inflation.

In short, the difference between the two models is that the menu costs 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 differs in periods of high relative to low trend inflation.[20]

The major limitation of this strategy involves data. In order 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 of Canada was troubled by the high inflation of the late 1970's, 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", and noted that "inflation must sooner or later be fought". A year later, he reflected on the "strong monetary medicine" which had been required to "beat the fever of inflation" that gripped the Canadian economy in the late 1970's. We count the experience of 1982-83 as a disinflationary shock, since the Bank of Canada appeared to be focused on reducing what it regarded as an unacceptably high rate of inflation 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 1980's, the Bank of Canada jointly with the Government of Canada announced a set of inflation targets which were to take effect starting in 1992. The targets essentially mandated a reduction in inflation, which was then around 5 percent, into a target

[20]Like the menu cost theory, the hypothesis of downward nominal wage rigidity does not imply a distinction between the short run output-inflation tradeoff and the statistical Phillips Curve, therfore, the evidence presented in this section also relates to that potential explanation.

band of 2 to 4 percent. Governor Crow, in the 1991 annual report, wrote that the purpose of the inflation targets was "to provide Canadians with a clear affirmation that price stability remains the goal of monetary policy". Two years later, Governor Thiessen 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". We regard the experience of 1991-92 as a second disinflationary shock.

Subsequent to 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 being 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. In order 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 tradeoff did not change as inflation decreased.

We also find it informative to contrast the inferred size of the output-inflation tradeoff 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 tradeoff, Figure 11 implies that the cost of reducing inflation by one percent is a negative output gap of approximately 1.3 percent (the slope implied by the 1983 and 1992 observations). If, on the other hand, the

[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.

menu costs theory is correct then the slope of the Phillips Curve is the proper estimate of the output-inflation tradeoff. 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 tradeoff, the negative output gap induced by a one percent reduction in inflation would be on the order of ten percent. Clearly, the two interpretations differ by orders of magnitude and hence suggest the need to provide further evidence to differentiate the two views more convincingly.

4 Conclusion

Our answer to the title of the paper "What has happened to the Phillips Curve over the 1990 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 U.S. Since we are interested in interpreting these observations for policy discussion, we have used a prototypical macro model to attempt to understand why the slope of the Phillips curve may have changed over time and what implication this may have for the output-inflation tradeoff 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 likely incorporated the improved knowledge into their behavior.[22] The second insight drawn from the model is that a flatter Phillips curve does not necessarily imply a change in the output-inflation tradeoff faced by the cental

[22] One possibility we have not addressed in the paper is that monetary authorities may have learned how to properly unlink real and nominal developments in the economy over the last twenty 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 preset prices, it may be socially optimal for monetary authorities to favor 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.

bank. In effect, we showed why the Phillips curve can become flatter while the relevant output-inflation tradeoff 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 explained in the paper, 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 eighties and nineties. In other words, we believe that the evidence on inflation and output over the last twenty years supports the view that the costs associated with reducing inflation have likely neither increased nor decreased over the last twenty years even if the statistical Phillips Curve appears to have flattened.

References

1. Akerlof, George A., William T. Dickens, and George L. Perry, *The Macroeconomics of Low Inflation*, Brookings Papers on Economic Activity, 1996, 1, 1-76.

2. Ball, Laurence, Gregory N. Mankiw and David Romer, *The New Keynesian Economics and the Output Inflation Tradeoff*, Brookings Papers on Economic Activity, 1988, 1, 1-65.

3. Blanchard, O. and N. Kiyotaki, *Monopolistic Competition on the Effect of Aggregate Demand*, American Economic Review, 1987, 77(4), 647-666.

4. Caplin, Andrew, and John Leahy, *Monetary Policy as a Process of Search*, American Economic Review, 1996, 86(4), 689-702.

5. Dupasquier, Chantal, and Nicholas Ricketts, *Non-Linearities in the Output-Inflation Relationship*, In Price Stability, Inflation Targets and Monetary Policy. Proceedings of a conference held by the Bank of Canada, May 1997, Bank of Canada.

6. Fillion, Jean-Francois and André Léonard, *La Courbe de Phillips au Canada: Un Examen de Quelques Hypothèses*, Bank of Canada Working Paper 97-3, 1997.

7. Fortin, Pierre, *Canada's Job Growth Potential*, Paper presented at the June 1997 meetings of the Canadian Economics Association.

8. Goodfriend, Marvin and Robert G. King *The New Neoclassical Synthesis and the Role of Monetary Policy*, NBER Macroeconomics Annual 1997, 231-283.

9. Lucas, Robert E,. *Some International Evidence of Output-Inflation Trade-offs*, American Economic Review, 1973, 63(3), 326-334.

10. Lucas, Robert E., *Expectations and the Neutrality of Money*, Journal of Economic Theory, 1972, 4(2), 103-124.

11. Townsend, Robert M., *Forecasting the Forecasting of Others*, Journal of Political Economy, 1983, 91(4), 546-588.

Fig. 1: Phillips Curve, Canada 1961-1999

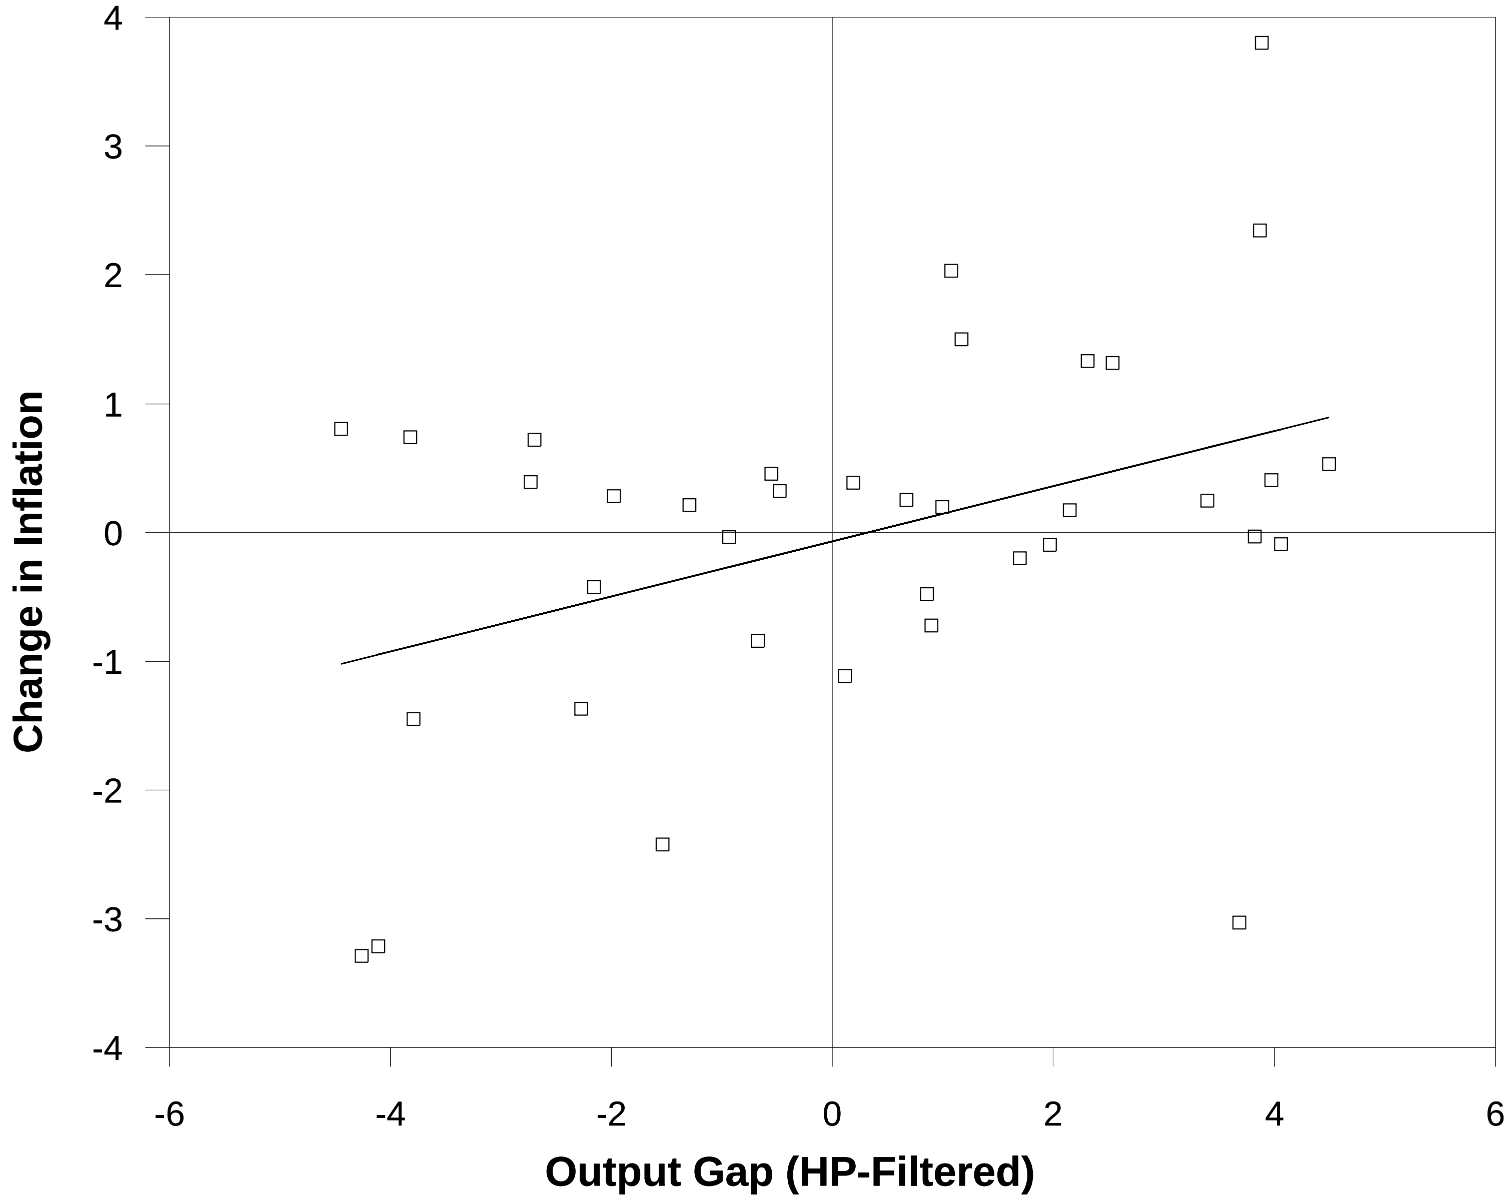

Fig. 2: Phillips Curve, U.S 1961-1999

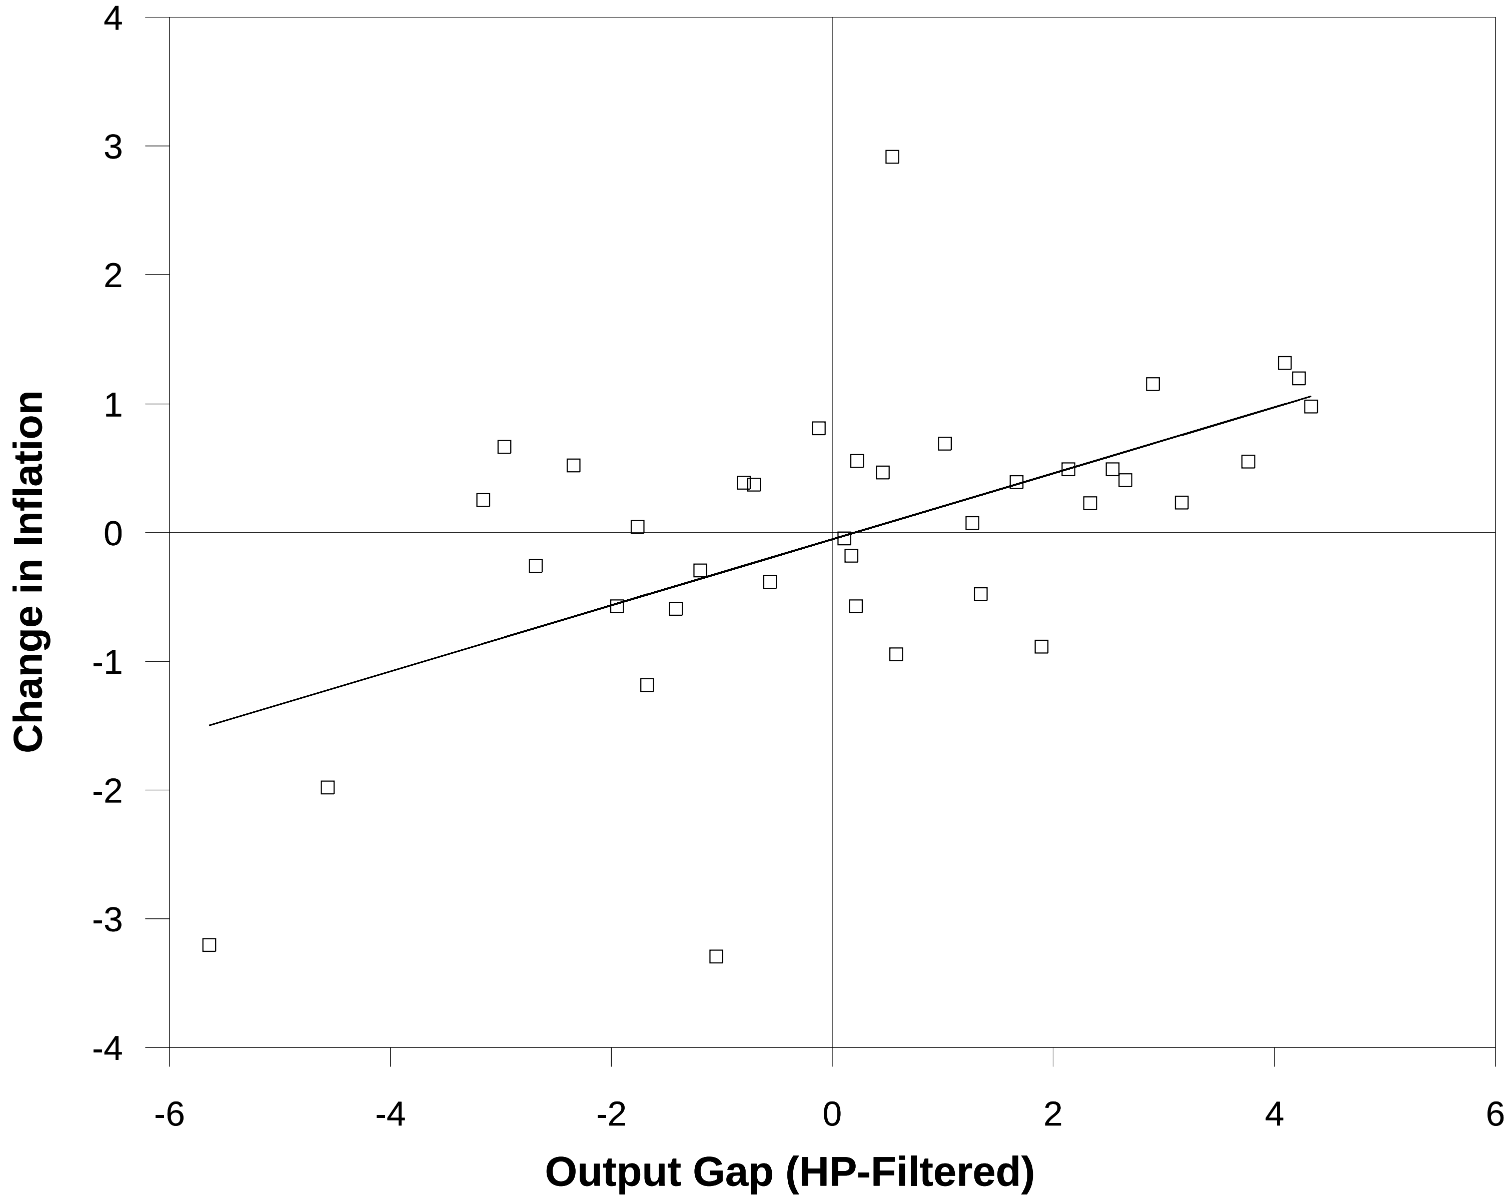

Fig. 3: Slope of Canadian Phillips Curve Over Time

15 Year Rolling Regression of Change in Inflation on Lag of Output Gap

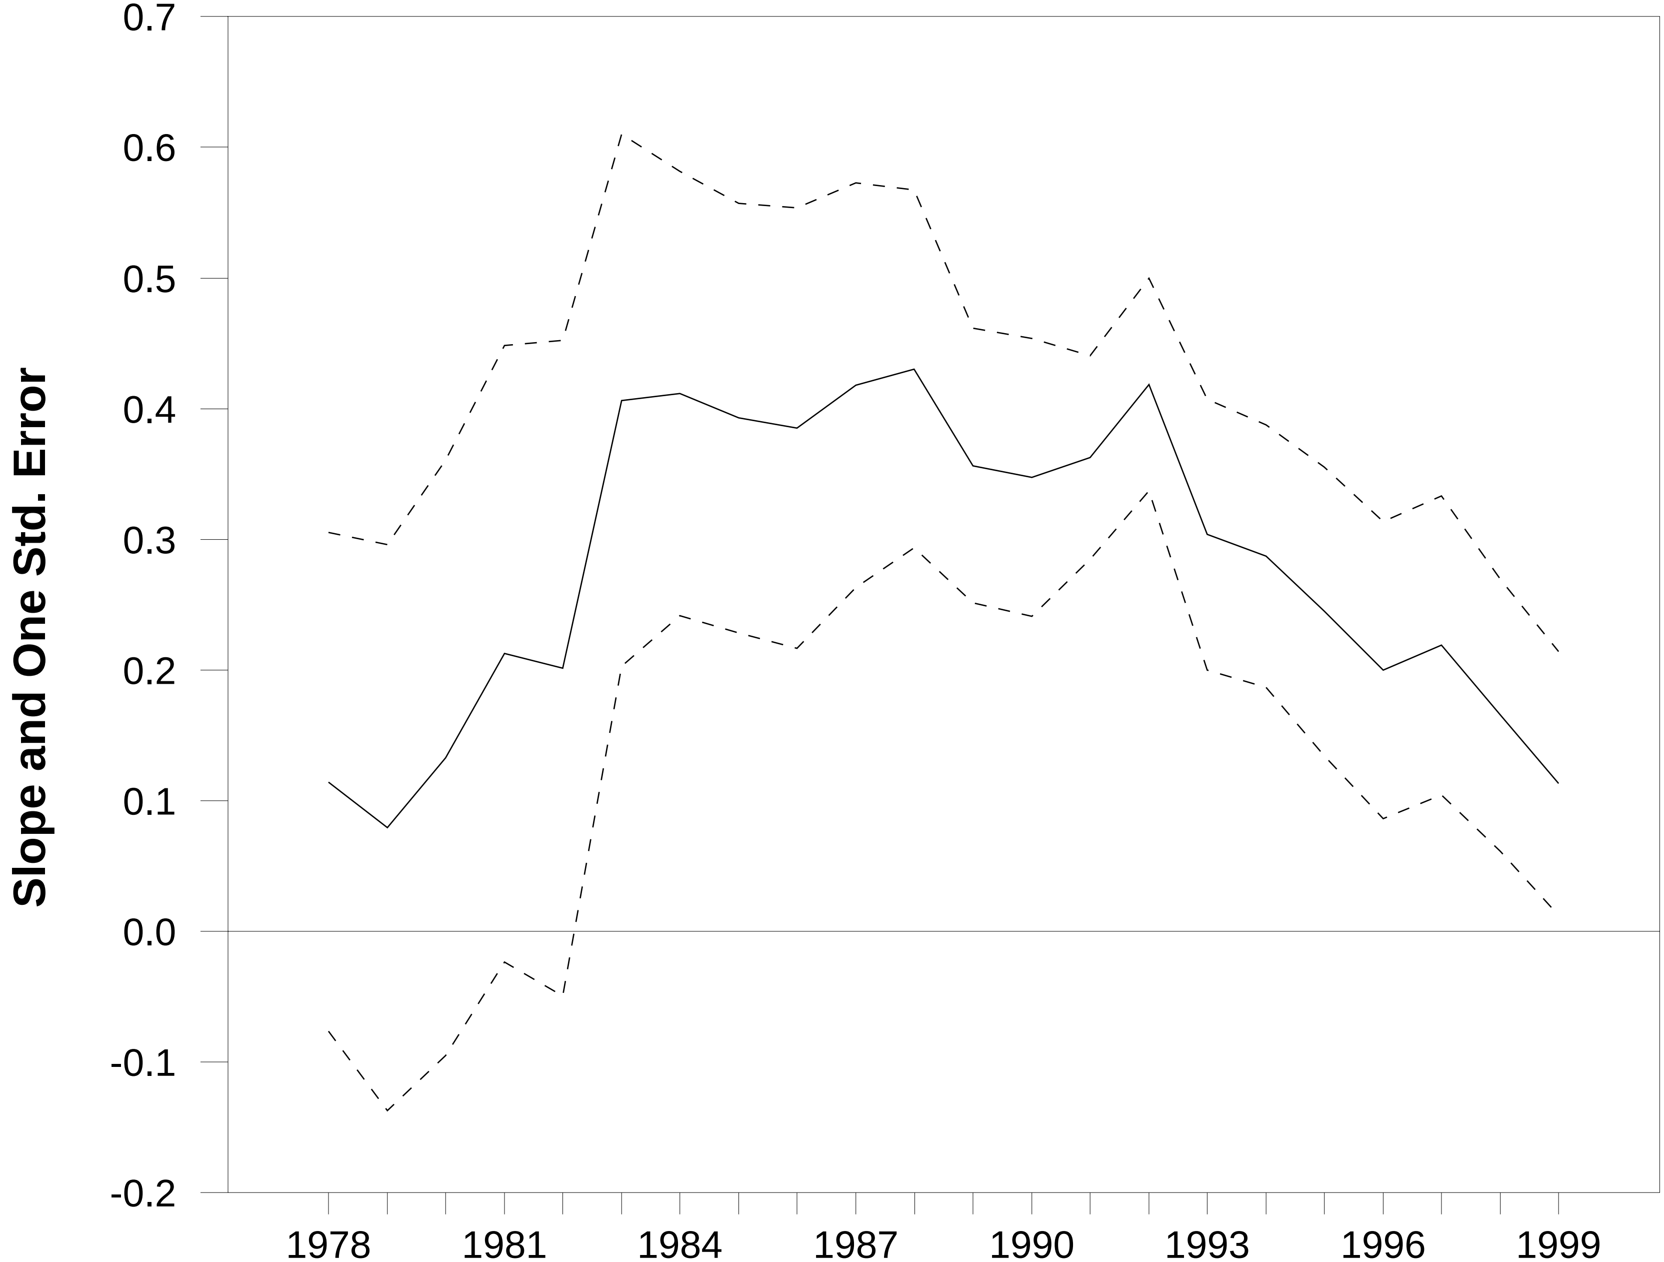

Fig. 4: Slope of U.S Phillips Curve Over Time

15 Year Rolling Regression of Change in Inflation on Lag of Output Gap

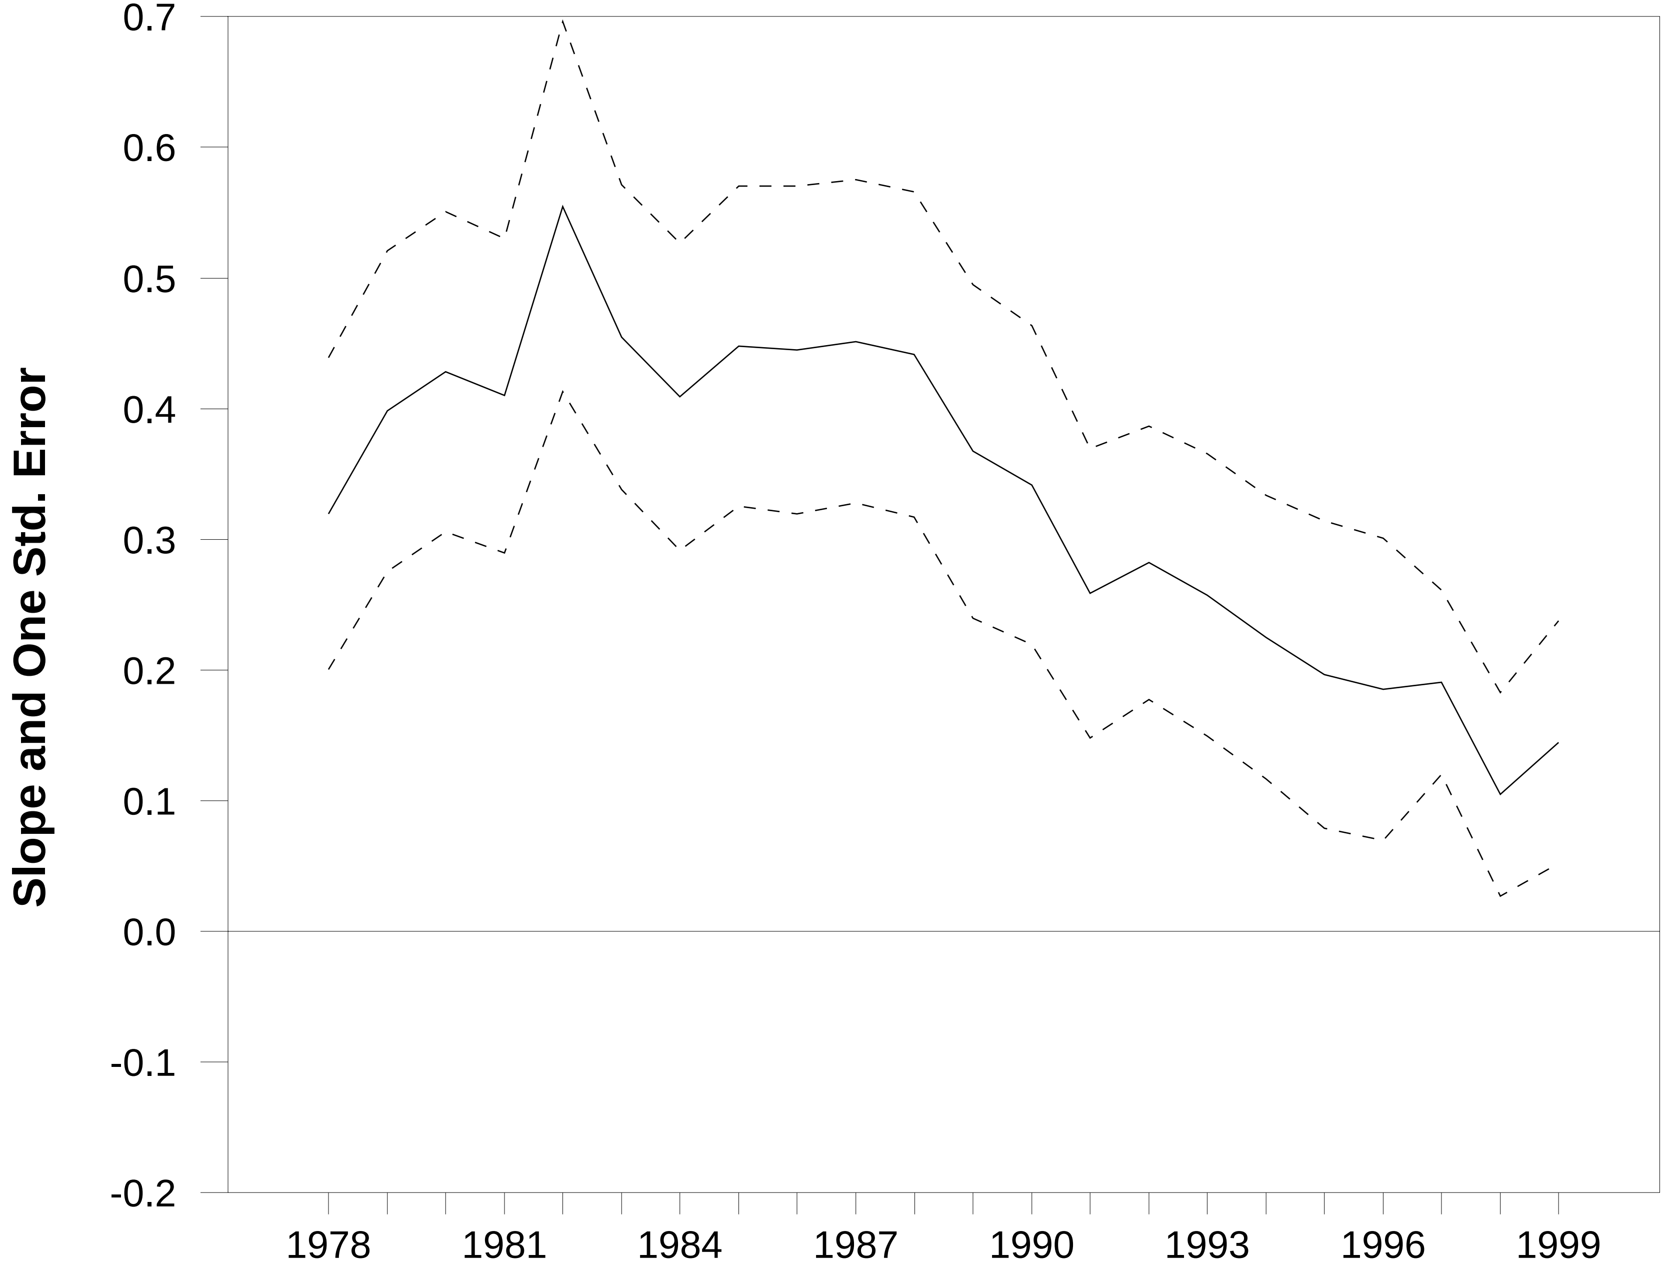

Fig. 5: Slope of Canadian Phillips Curve Over Time

Change in Inflation on Lag of Output Gap and Inflation in Relative Energy Prices

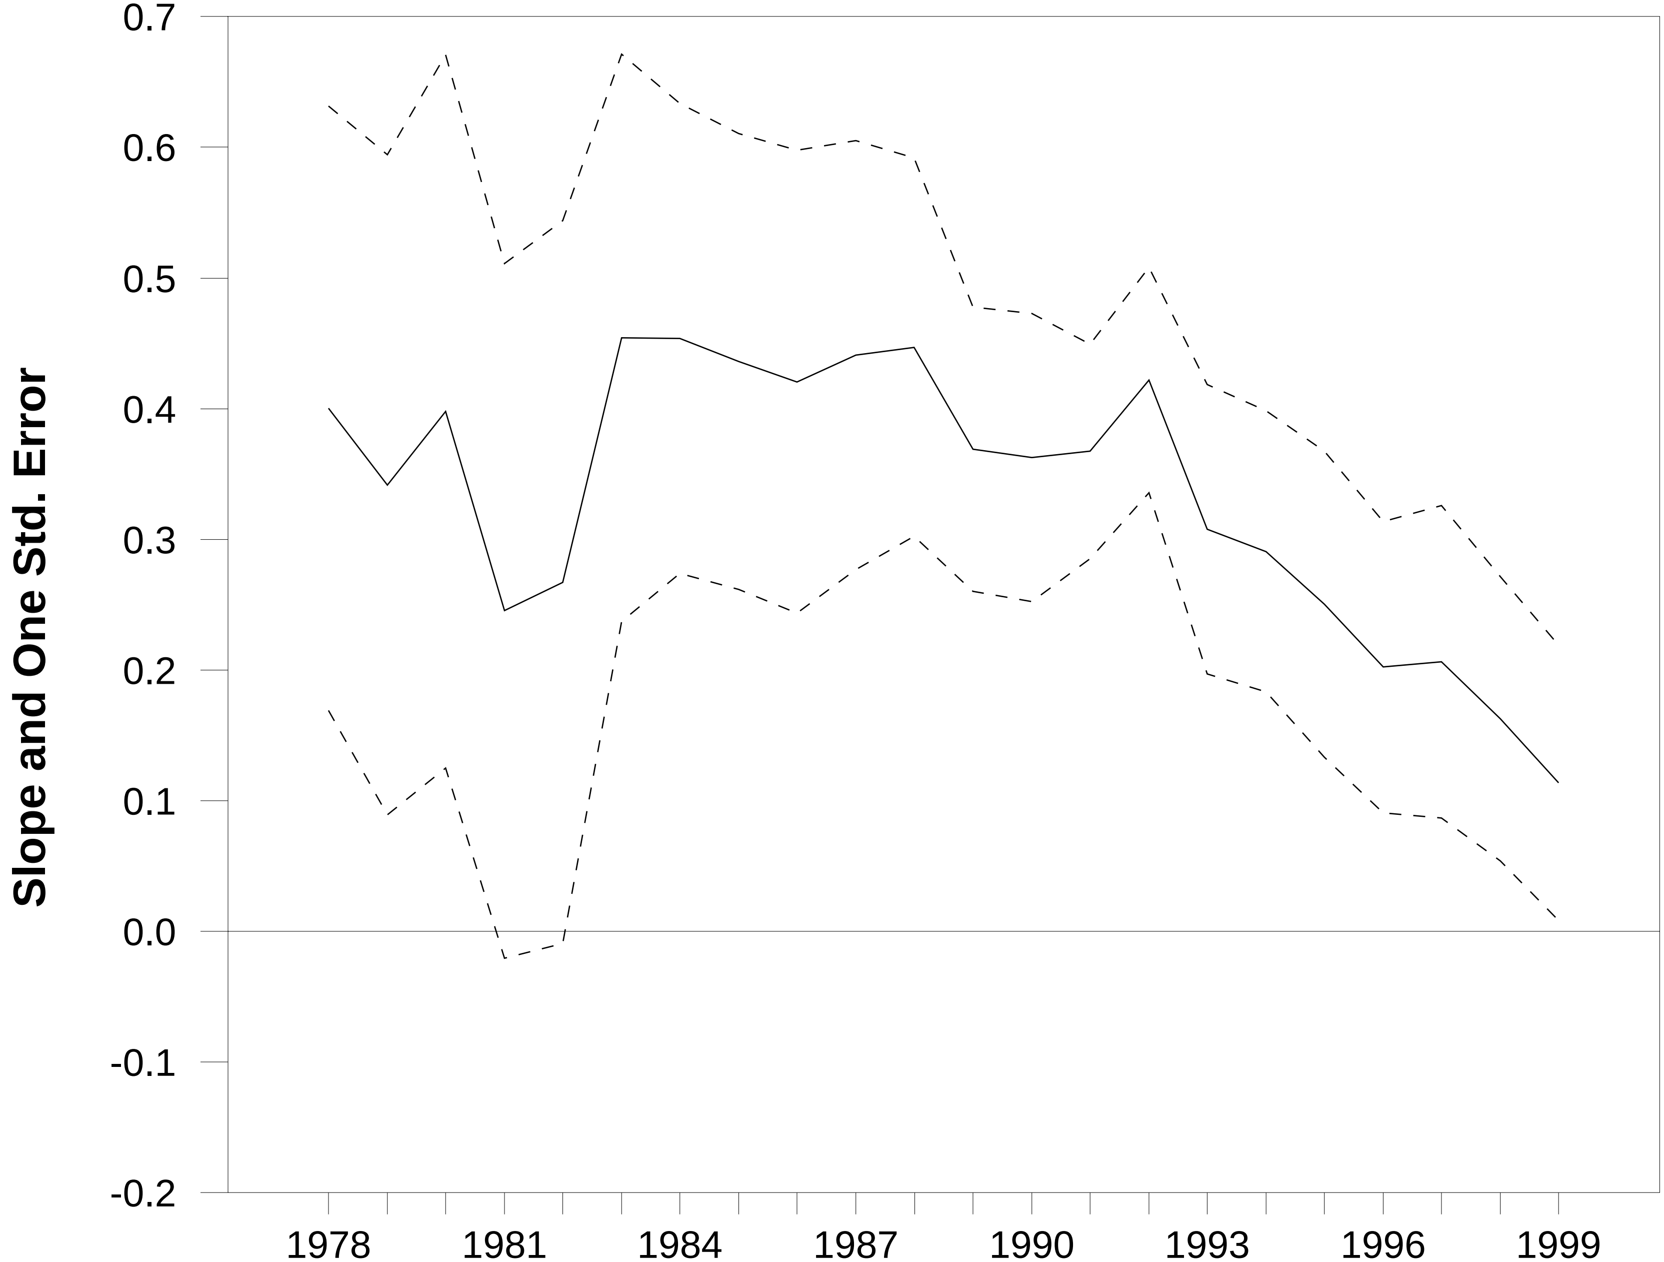

Fig. 6: Slope of U.S Phillips Curve Over Time

Change in Inflation on Lag of Output Gap and Inflation in Relative Energy Prices

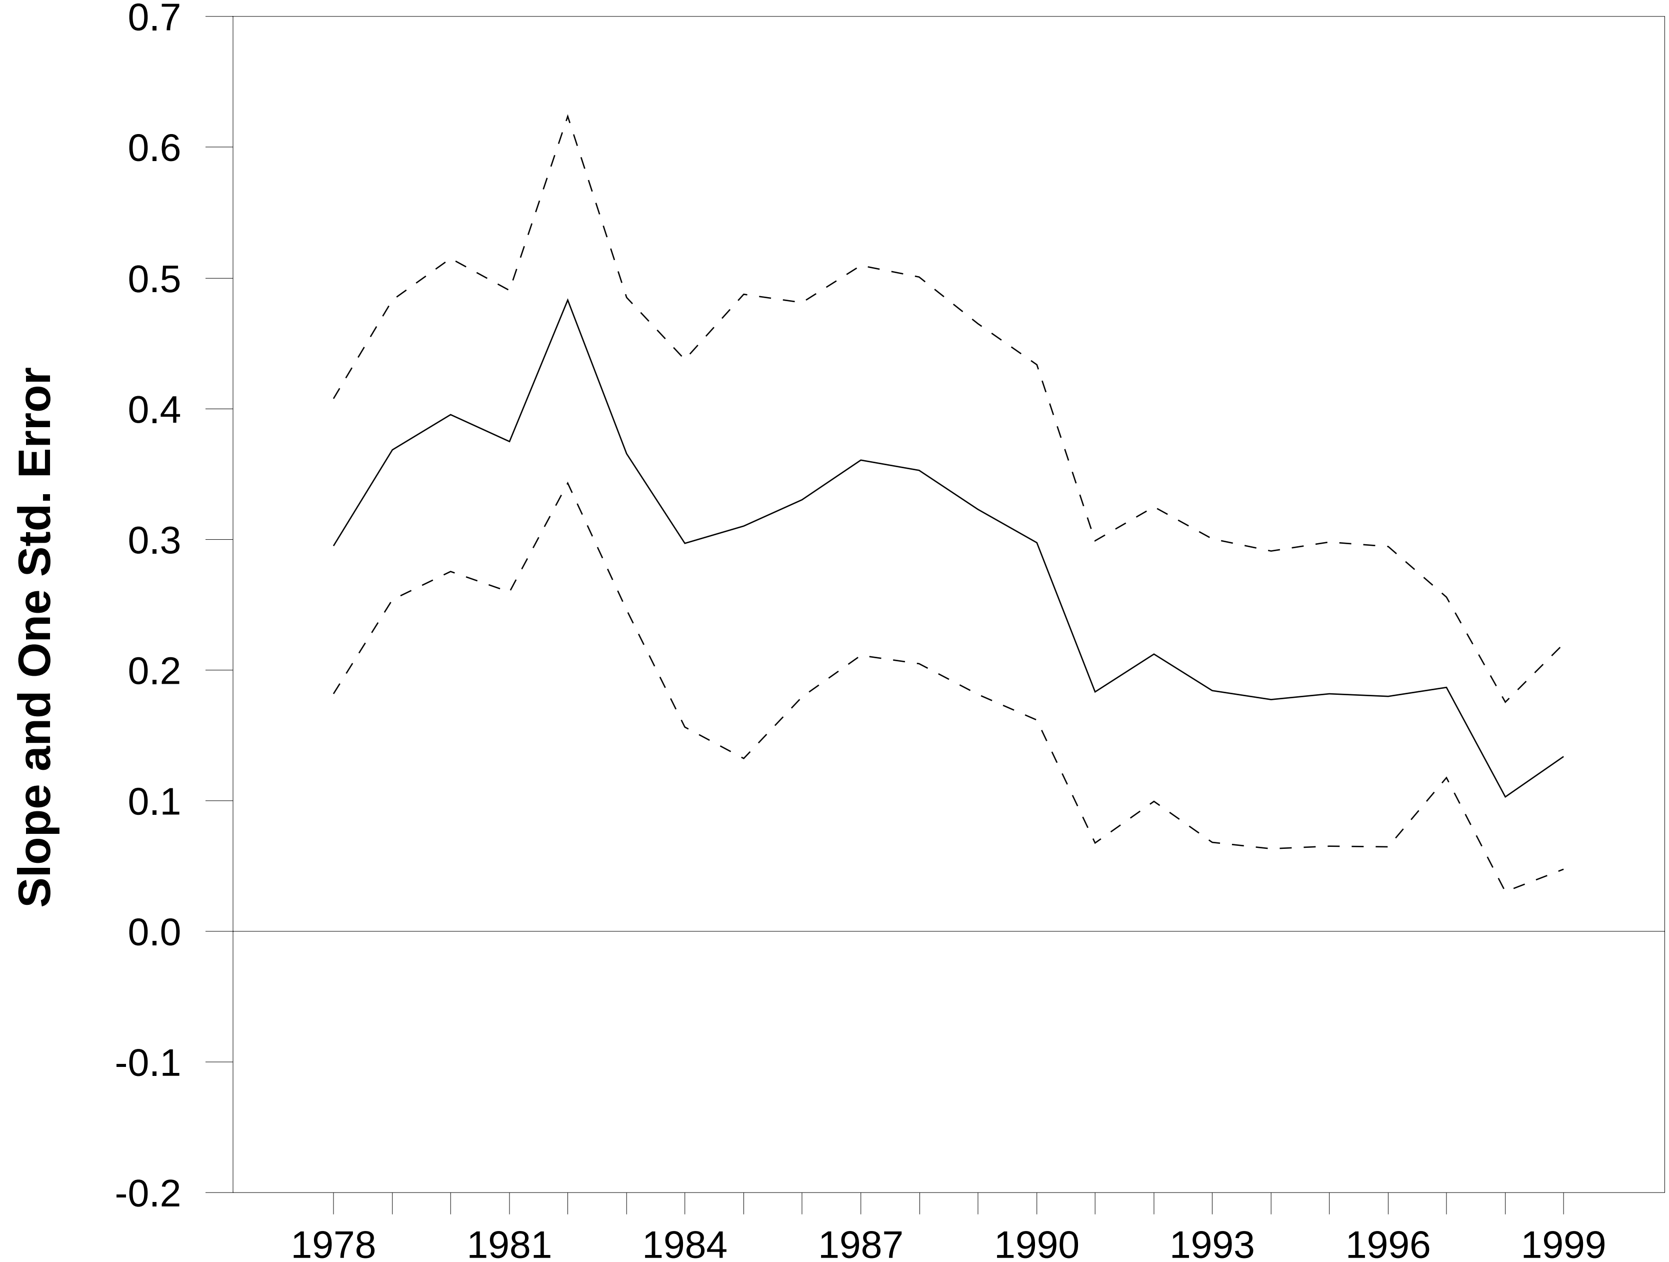

Fig. 7: Slope of Phillips Curve Over Time, Pooled Sample

Rolling Regression of Change in Inflation on Lag of Output Gap

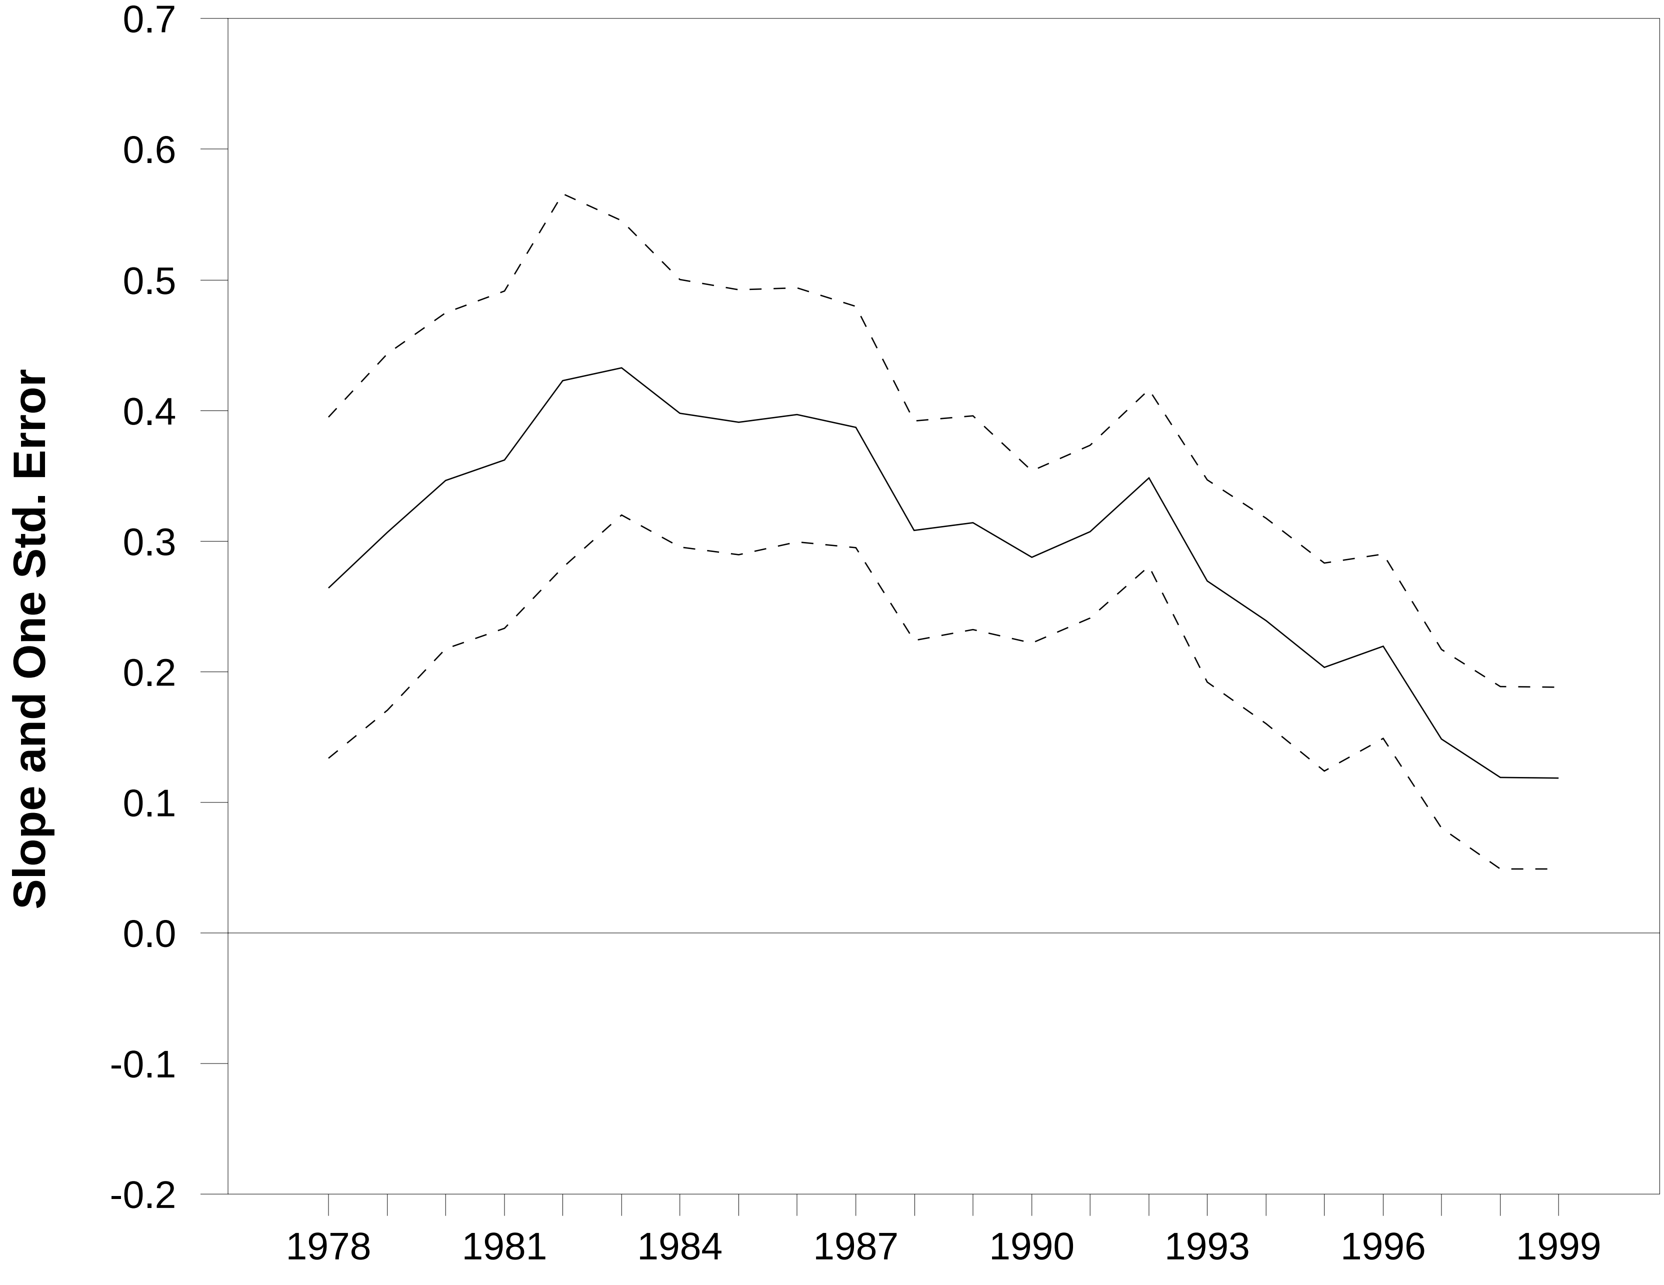

Fig. 8: Change in Inflation vs. Gap 1985-1999, Pooled Sample

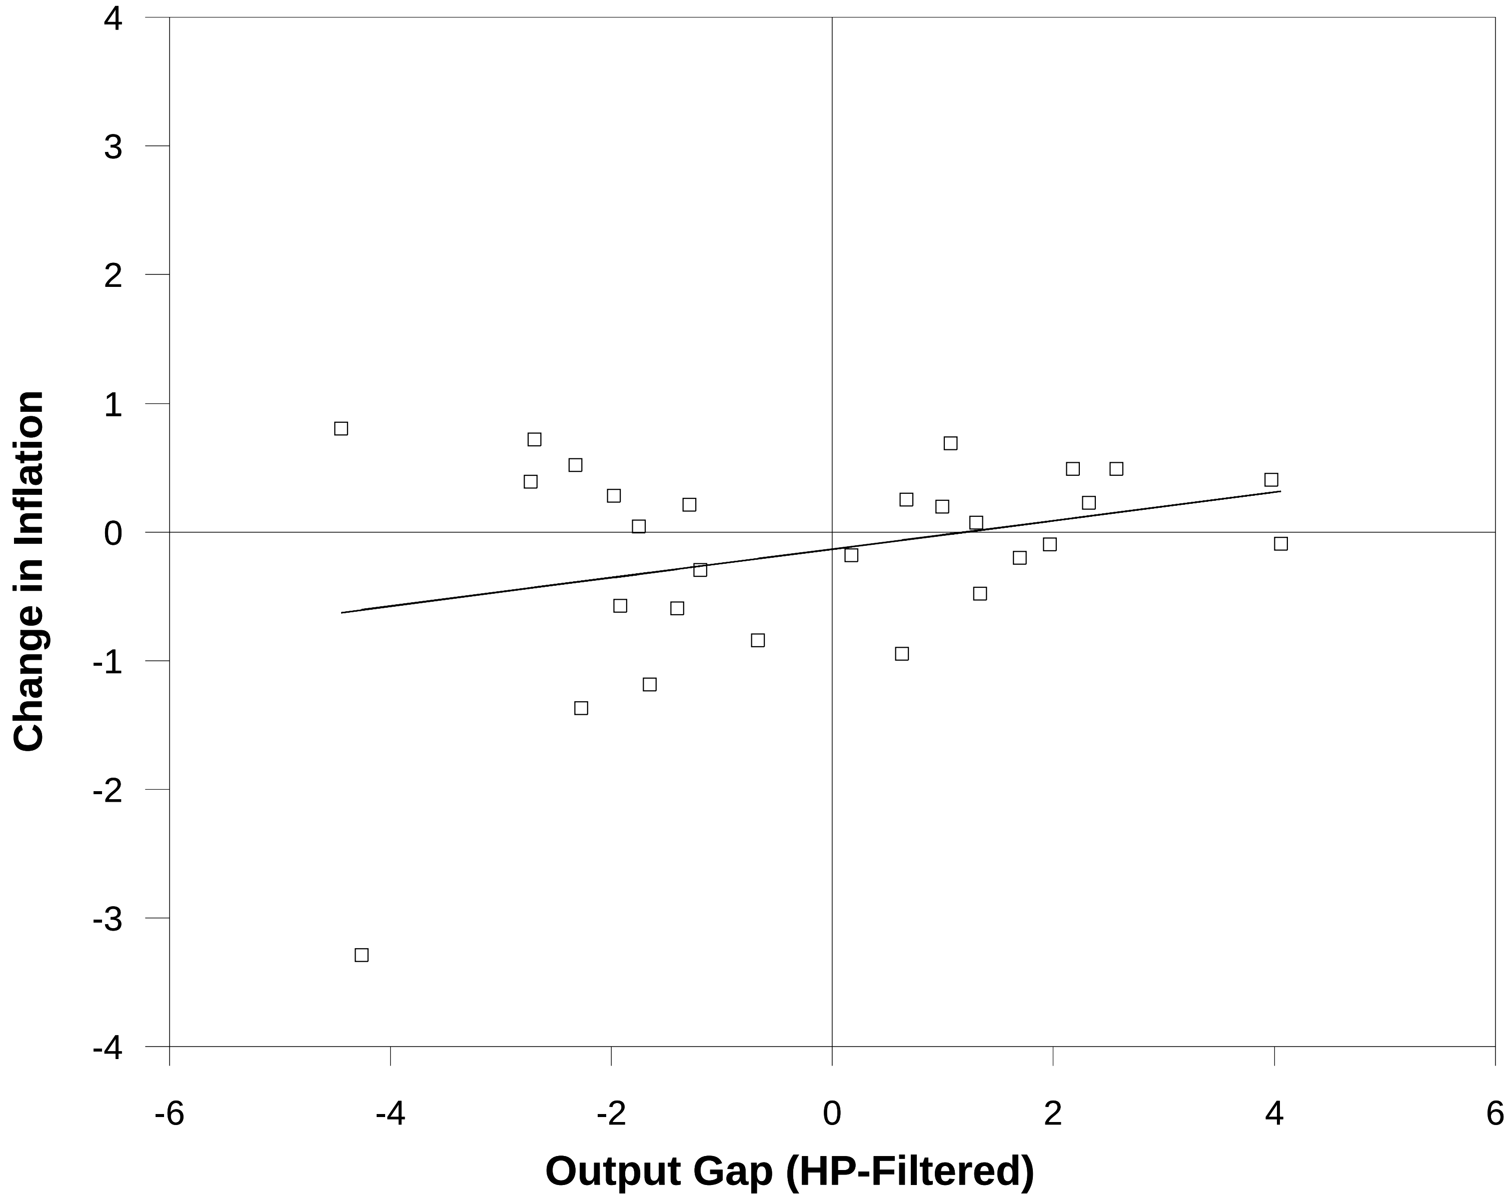

Fig. 9: Output Gap Coefficient, Pooled Sample

Change in Inflation on Lag of Gap and Lag of Positive Gap

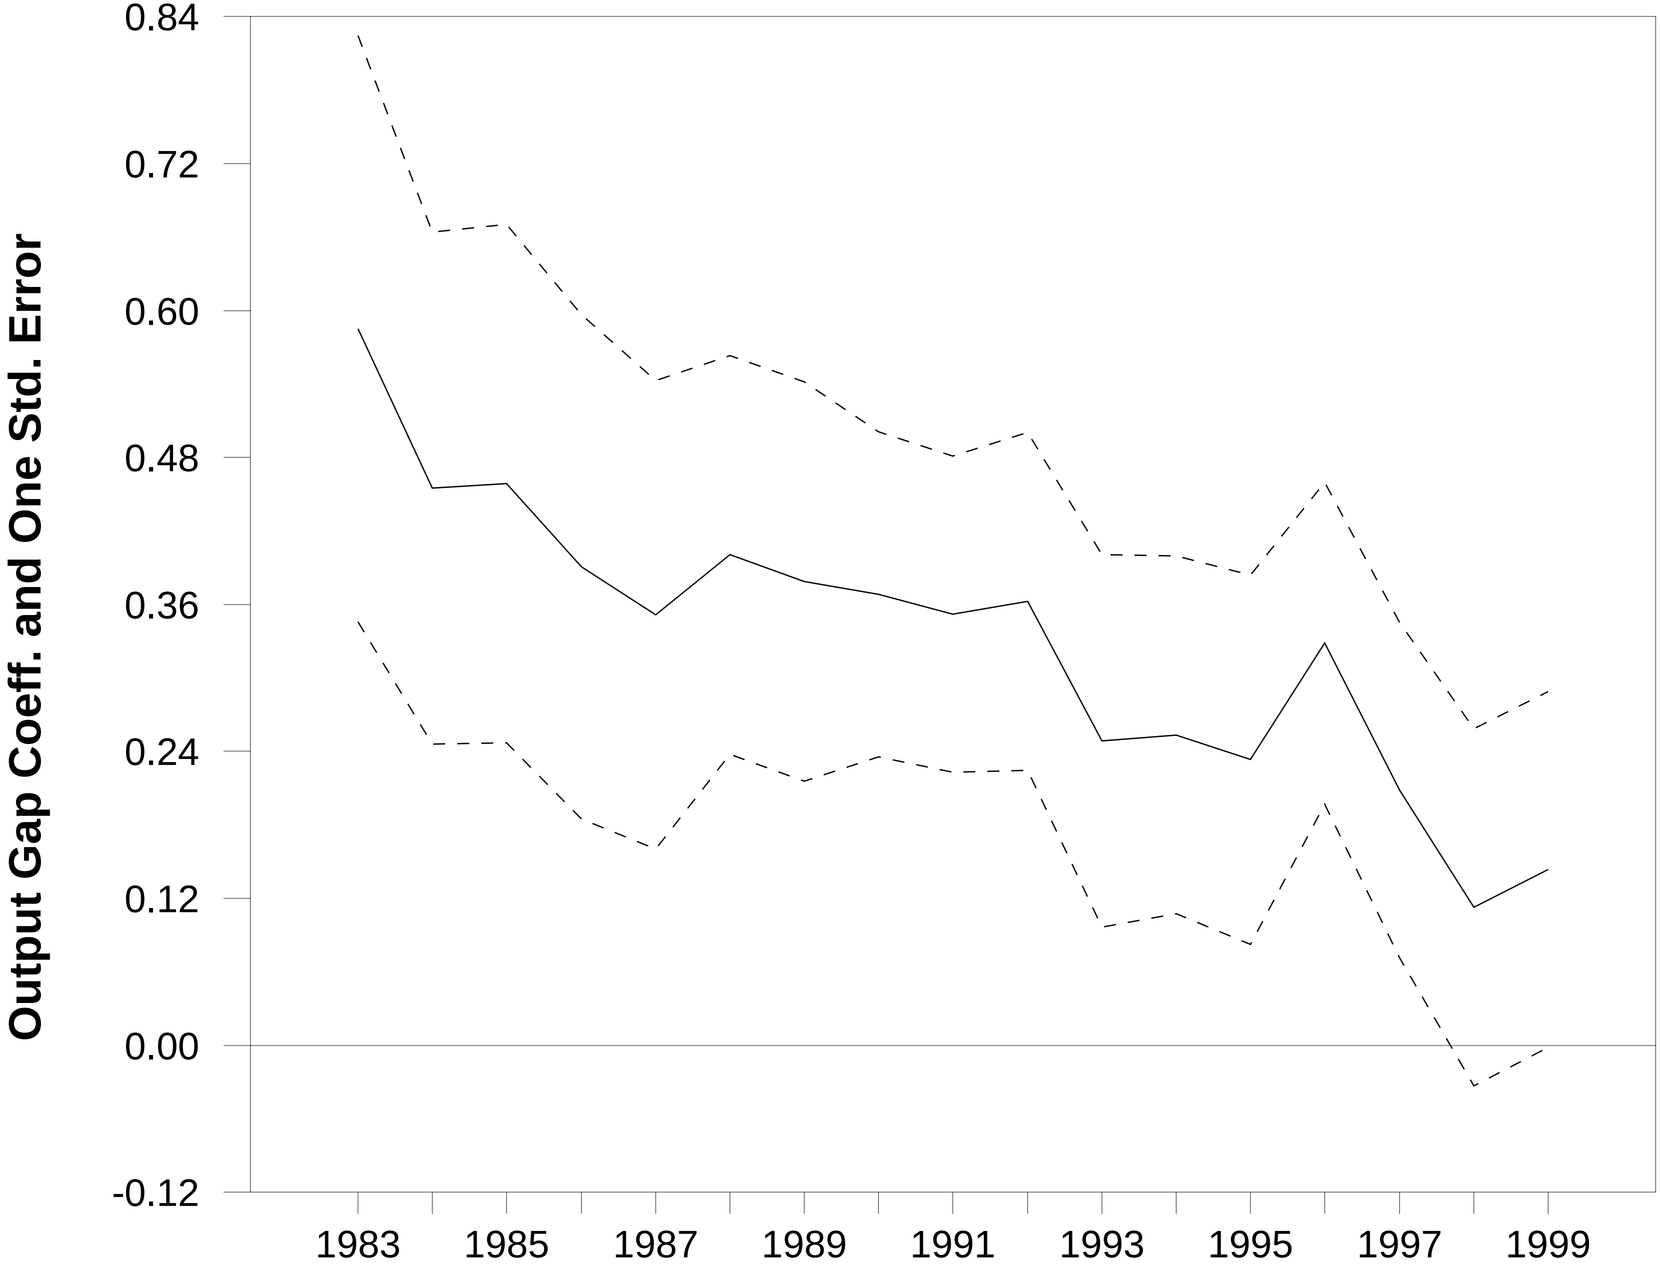

Fig. 10: Positive Gap Coefficient, Pooled Sample

Change in Inflation on Lag of Gap and Lag of Positive Gap

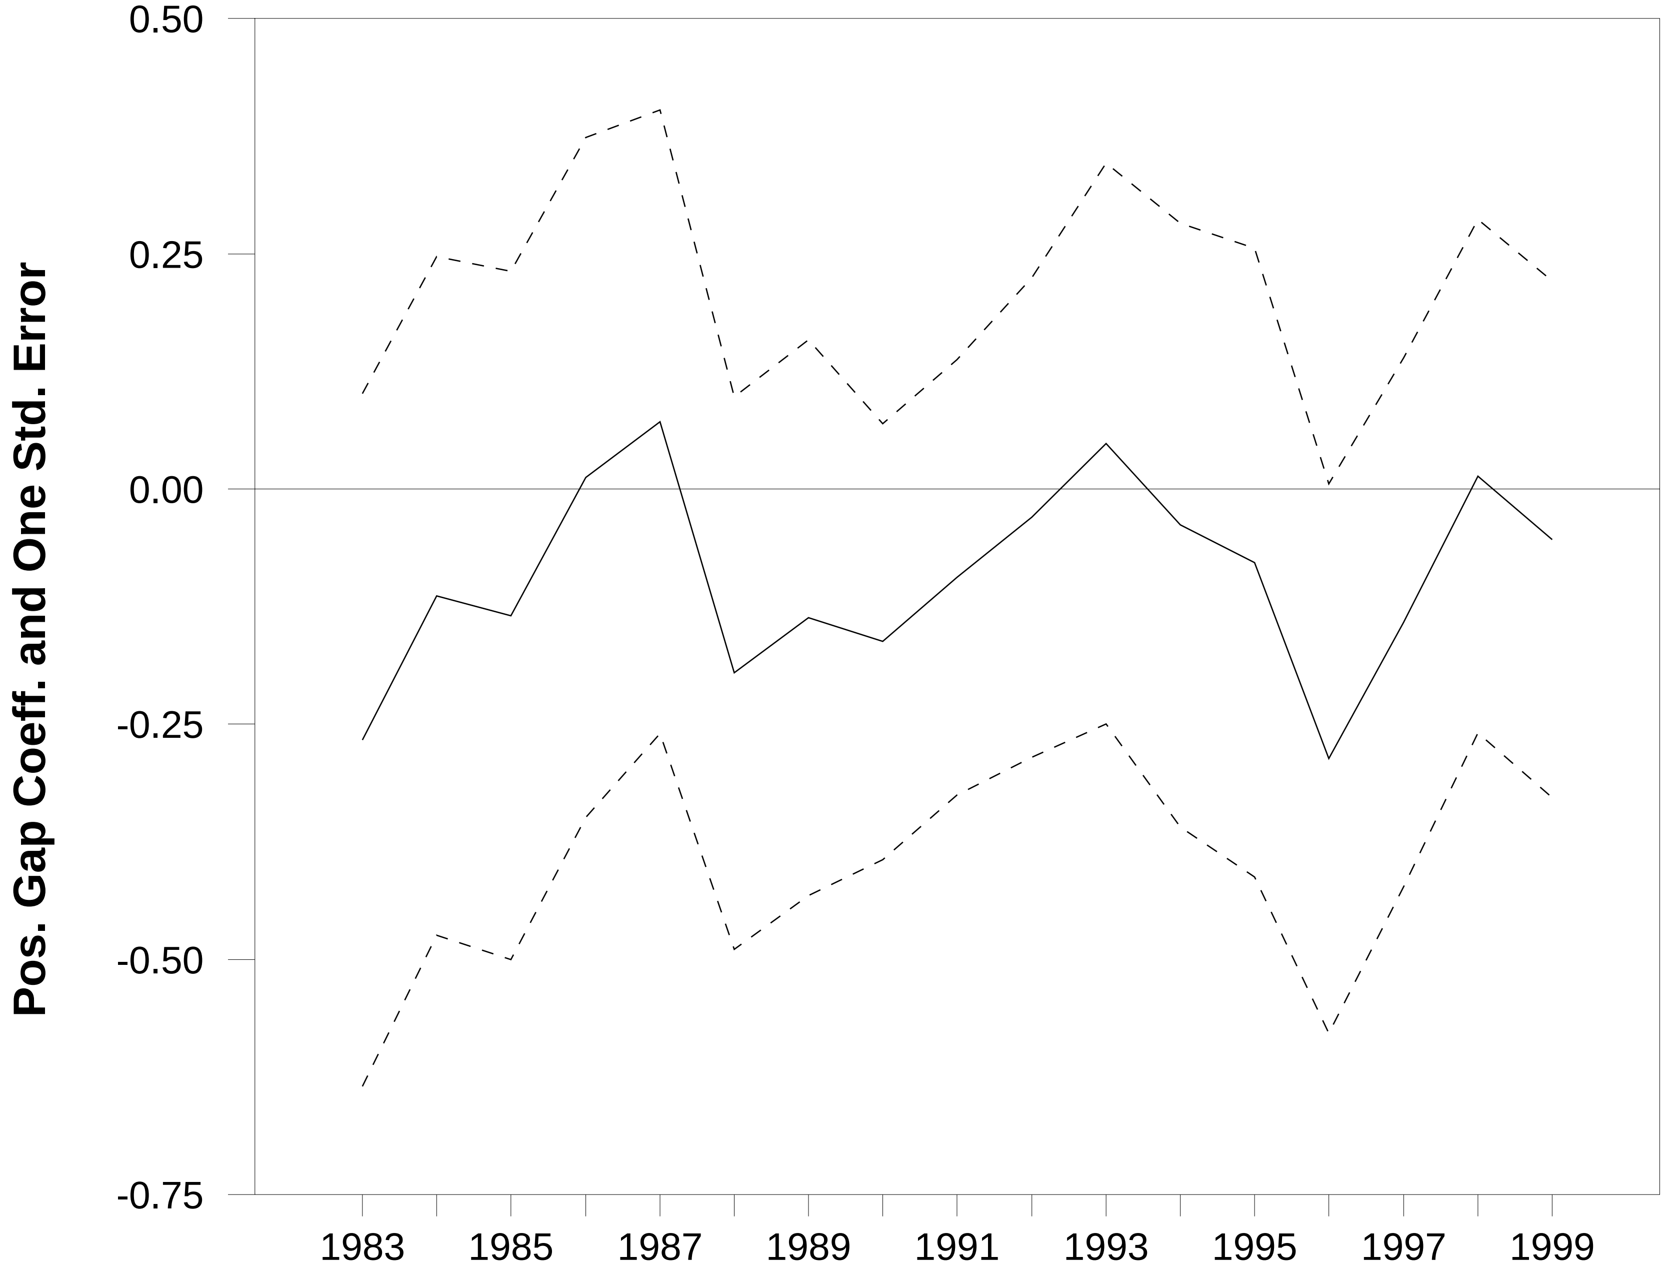

Fig. 11: Phillips Curve, Canada 1980-1999

Regression includes dummy variables for 1983 and 1992