Probing Potential Output: Monetary Policy, Credibility, and Optimal Learning under Uncertainty

by

James Yetman

Bank of Canada   **Banque du Canada**

Bank of Canada Working Paper 2000-10

May 2000

Probing Potential Output: Monetary Policy, Credibility, and Optimal Learning under Uncertainty

by

James Yetman

Research Department, Bank of Canada
Ottawa, Canada K1A 0G9
JYetman@bank-banque-canada.ca

Contents

iv

Acknowledgements

Abstract

The effective conduct of monetary policy is complicated by uncertainty about the level of potential output, and thus about the size of the monetary policy response that would be sufficient to achieve the targeted inflation rate. One possible response to such uncertainty is for the monetary authority to "probe," interpreted here as actively using its policy response to learn about the level of potential output.

Monetary authorities have put significant emphasis in recent years on attaining credibility for their policy objectives. These steps have anchored inflation expectations to the target of the monetary authority more firmly. I consider a simple calibrated model in the Canadian context and examine the relationship between credibility and optimal probing.

I find that, for plausible parameter values, the optimal amount of probing is small and varies little with credibility. Only for low levels of credibility or unrealistically large levels of uncertainty or volatility does the optimal policy with probing diverge significantly from a policy that ignores learning. Even then, the optimal amount of probing diminishes as credibility rises.

JEL classification: E52, E58
Bank classification: Credibility; Potential output; Uncertainty and monetary policy

Résumé

L'incertitude qui entoure le niveau de la production potentielle entrave l'efficacité de la politique monétaire car elle rend plus difficile la détermination du degré de resserrement nécessaire à la réalisation de la cible d'inflation visée. Face à l'incertitude, les autorités monétaires peuvent choisir de procéder par « tâtonnement », c'est-à-dire en tirant parti de leurs interventions pour se renseigner sur le niveau de la production potentielle.

Depuis quelques années, les autorités monétaires ont déployé beaucoup d'efforts en vue d'établir la crédibilité de leurs objectifs. Elles ont ainsi pu arrimer plus fermement à la cible visée les attentes en matière d'inflation. L'auteur fait appel à un modèle simple, étalonné en fonction des données canadiennes, pour examiner la relation entre la crédibilité et le degré optimal de tâtonnement.

Il constate que, pour des valeurs plausibles des paramètres, le degré optimal de tâtonnement est minime et varie peu selon la crédibilité. Ce n'est que lorsque cette dernière est faible ou que l'incertitude ou la volatilité atteint des niveaux tout à fait irréalistes que la politique optimale avec tâtonnement diffère sensiblement d'une politique qui ne repose pas sur

l'apprentissage. Même alors, le degré optimal de tâtonnement s'amenuise à mesure que la crédibilité augmente.

Classification JEL : E52, E58
Classification de la Banque : Crédibilité; Incertitude et politique monétaire; Production potentielle

A credible inflation target can help the Bank probe to find out what the limits of potential output really are.

Gordon G. Thiessen, Governor of the Bank of Canada, 1998. "The Canadian Experience with Targets for Inflation Control," *Canadian Public Policy* 24(4): 423.

Although no one knows exactly where the NAIRU is ... in testing the waters, we do not risk drowning. If need be, we can always reverse course. But by experimenting, and showing some hesitation about restraining the economy through higher interest rates or other methods as the NAIRU draws nigh, we might learn a little more about the depth of the waters and possibly become better swimmers in the process.

Joseph Stiglitz, 1997. "Reflections on the Natural Rate Hypothesis," *Journal of Economic Perspectives* 11(1): 10.

Although a closing of the output gap during the course of the next year cannot be excluded, given the uncertainties surrounding this measure monetary authorities should probe to see whether the economy can reach higher levels of output without inflation pressures.

"Developments in Individual OECD Countries: Canada," *OECD Economic Outlook* (June 1999): 65.

1.   Introduction

The effective conduct of monetary policy is complicated by uncertainty. There are many dimensions to this uncertainty: uncertainty about shocks, model parameters, data, and the "correct" model of the economy itself.[1] At a practical level, one of the key uncertainties facing policy-makers is the level of output that can be maintained without adding to inflation pressures (referred to as the level of potential output). While policy-makers can continue to refine and improve the measurement of potential output,[2] to a considerable degree uncertainty about potential output is fundamental. Thus, the challenge for policy-makers is how to deal with this uncertainty.

Three possible responses by the monetary authority to uncertainty about potential output that have been examined analytically are to (i) ignore the uncertainty and follow the "certainty equivalent" policy; (ii) act "conservatively," by which is meant moving interest rates by less than is implied by the certainty equivalent policy; or (iii) "probe" or experiment, which implies that the monetary authority actively uses its policy response to learn about the level of potential output.

To formalize probing within an economic model, one must understand what it means in terms of the behaviour of the monetary authority. However, there is no consensus on this. One interpretation of probing is that it entails optimal learning, that is, following a more aggressive policy to learn about the parameters of the economy. Probing of this type results in more precise estimates, and therefore smaller policy mistakes in future periods. Building on Wieland's (1998) analysis of this issue, I consider a simple calibrated model in the Canadian context and examine the relationship between this definition of probing and credibility.

Monetary authorities have put significant emphasis on attaining credibility for their policy objectives in recent years. Steps taken by the Bank of Canada have included announcing explicit inflation targets, publishing detailed accounts of inflation developments and the conduct of monetary policy, and issuing press releases explaining changes in the

---

1.   See Thiessen (1995) or Poole (1998) for discussions of the various dimensions of uncertainty facing monetary authorities.
2.   See Kuttner (1992), Laxton and Tetlow (1992), Butler (1996), St-Amant and van Norden (1997), or Dupasquier, Guay, and St-Amant (1999) for a discussion of the various ways potential output is measured.

Bank Rate.[3] These steps have increased the accountability of the monetary authority and that, together with the realized inflation record, has enhanced its credibility in the sense that expectations of inflation have become more firmly anchored to the inflation target.[4]

The question addressed here is whether an increase in credibility increases the desirability of probing. In other words, should a monetary authority that has increased its credibility follow a more aggressive policy in order to obtain more precise estimates of the parameters of the economy? I find that, for plausible parameter values, the optimal amount of probing is small and varies little with credibility. It is only for low levels of credibility or unrealistically large amounts of uncertainty or volatility that the optimal policy with probing diverges significantly from a policy that ignores learning. Even then, the optimal amount of probing diminishes as credibility rises.

At an intuitive level, the returns to probing decrease as credibility increases in the model I consider because credibility makes learning more difficult. As credibility increases, inflation becomes more firmly anchored to the inflation target; thus the out-turn for inflation is less informative about potential output. To illustrate this with an example, suppose that the monetary authority is underestimating potential output and, as a result, incorrectly believes that the economy is operating at potential. With low credibility, inflation will lie below the target, allowing the monetary authority to infer that its estimate of potential was incorrect. At higher levels of credibility, inflation is more firmly anchored to the target, so that inflation provides a weaker signal that potential output is higher than was previously believed.

The next section summarizes the literature supporting a conservative monetary policy in the face of uncertainty regarding the economy. Section 3 summarizes articles arguing for a more aggressive policy. An outline of the model is given in Section 4, followed by discussion of the parameter values in Section 5 and results in Section 6. Conclusions follow in Section 7.

2.   Uncertainty and conservatism

A number of authors, starting with Brainard (1967), argue that uncertainty is a motivator for a conservative monetary policy. Brainard considers a simple model given by

$$y = ap + u, \tag{1}$$

---

3.   See Amano, Coletti, and Macklem (1998) for more details.
4.   See Johnson (1997, 1998) or Perrier (1998) for evidence of this.

where the objective of the policy-maker is to choose the value of the policy variable p that minimizes the value of the policy-maker's loss function, $(y - y^*)^2$. Under certainty, the optimal policy takes the form of

$$p = (y^* - u)/a, \qquad (2)$$

and the policy-maker achieves the objective. Uncertainty can enter into this problem in two different ways: additive uncertainty, via the value of u; or multiplicative uncertainty, via the value of a.

In the presence of uncertainty, the policy-maker seeks to minimize the expected value of the loss function. Additive uncertainty has no effect on the optimal policy prescription, except that it is now a function of the expected, rather than the true, value of u:

$$p = (y^* - E(u))/a. \qquad (3)$$

This is referred to as the "certainty equivalent" policy, since the presence of uncertainty does not change the optimal policy response.

In the presence of multiplicative uncertainty, the optimal policy departs from the certainty equivalent policy, since the variance of a as well as the covariance of a and u now enter into the policy in the following way:

$$p = E(a)(y^* - E(u) - \sigma_{au})/(E(a)^2 + \sigma_a^2). \qquad (4)$$

In the special case that $E(u) = 0$ and $\sigma_{au} = 0$, the optimal policy rule reduces to

$$p = \frac{y^*}{E(a) + (\sigma_a^2/E(a))}. \qquad (5)$$

Since σ_a^2 is positive, the optimal policy response to shocks is smaller, or more conservative, than the certainty equivalent policy.

Other authors obtain similar results in a variety of frameworks. Aoki (1998) considers the effect of measurement errors on optimal monetary policy. He models the manner in which the central bank extracts information about economic shocks from noisy indicators using a dynamic sticky-price model. He shows that the central bank should respond to its forecasts of both the current output gap and current inflation, even if it is concerned only about inflation (as in the Taylor rule), although its response should be cautious due to the presence of measurement error.

Smets (1998) considers a simple model of the economy based on the Rudebusch and Svensson (1998) model in which the Taylor rule is non-optimal. He assumes that the output gap is measured with error, so that additive uncertainty is present in the model. As in the Brainard example, optimal central bank behaviour is not affected by this uncertainty. However, if the central bank were to restrict itself to using a Taylor rule to formulate policy, a conservative response to the estimated output gap would be desirable in the presence of output gap uncertainty.

Svensson (1997) finds that the optimal monetary policy under parameter uncertainty is more conservative than the certainty equivalent policy in a simple analytic model. Among other contributions, Srour (1999) extends his framework to an open economy context and obtains the same result, although the degree of conservatism is not great for plausible parameter values.

In some models, the NAIRU (Non Accelerating Inflation Rate of Unemployment) may serve the same role for monetary policy purposes as potential output. Estrella and Mishkin (1998) consider the impact of uncertainty in the NAIRU on optimal monetary policy in a simple linear model. They show that uncertainty of this type has no effect on the optimal policy, but uncertainty as to the trade-off between unemployment and inflation results in a more conservative optimal policy.

Bean (1999) studies the implications of a convex Phillips curve on the optimal policy under uncertainty. The optimal policy displays conservatism, and output is less than potential on average. In contrast to Brainard (1967), however, the presence of uncertainty here leads to a systematic bias in policy: policy should always be set tighter than it would be in the absence of uncertainty.[5]

Sack (1998) argues that the central bank is confident about the relationship between output and monetary policy if policy remains close to recent levels, but less confident as it moves away from levels implemented in the recent past. He assumed an I.S. curve given by

$$y_{t+1} = \alpha_{t+1} - \phi_{t+1} i_t, \tag{6}$$

---

5.   Alternatively, a systematically tight monetary policy may result from a linear Phillips curve if policy-makers think credibility (that is, the degree with which inflation expectations are anchored to the target) is difficult to attain but easy to lose. This has the effect of increasing the potential costs of expanding the economy too quickly relative to the costs of a recession, and so leads to a less expansionary policy than would be optimal without uncertainty. See Laxton, Ricketts, and Rose (1994) for an example of this.

where i_t is the policy instrument while α and ϕ, a measure of policy effectiveness, evolve through time. The variance of output is increasing in changes to the policy variable, so that the optimal policy entails gradual adjustment over time. These gradual changes provide informative observations about the effect of policy and the value of parameters in the economy and thereby reduce uncertainty about the impact of future policy.

In all of the above cases, uncertainty results in a more conservative optimal policy. The next section outlines frameworks in which uncertainty may lead to probing.

3.   Uncertainty and probing

A number of authors provide frameworks where the optimal policy of a central bank entails some probing or experimenting. For example, Caplin and Leahy (1996) suggest that policy-makers learn about the economy by observing the economy's response to policy shocks. When the economy is operating below potential, the aim of the central bank is to stimulate output via lowering interest rates to the point where some (but not all) planned investment projects will be undertaken. They argue that small decreases in the interest rate may result in little economic response, as agents will (correctly) infer that future reductions in interest rates are likely to follow. Profit-maximizing firms defer investment projects that are profitable at current interest rates until those rates fall further. As a result, both the length of recessions and the amount of policy adjustment required to attain potential output may be larger if the policy is changed gradually than if it is changed rapidly.

An alternative view of probing, and the one that is used here, assumes that policy-makers use the latest available data to estimate the parameters of the economy each period. These new estimates are then used in policy formulation. If policy-makers ignore the impact of their policy on this learning process, the policy-makers are said to be engaged in "passive learning." Alternatively, if the policy-maker explicitly takes account of the impact of their policy on the learning process, the policy-maker is engaged in "active learning" or "probing."

As a simple illustration, consider the example of Brainard given in (EQ 1) above. Suppose that the policy-maker regresses y on p each period and uses this regression to update the estimate of a. The optimal policy of the monetary authority will then take account of the amount of information generated by the policy. In general, the optimal policy that takes account of learning is more aggressive than the multiplicative uncertainty policy

(EQ 5), but less aggressive than the certainty equivalent policy (EQ 3). As will be outlined below, other authors obtain similar results with more general models.

Bertocchi and Spagat (1993) model the economy with the following equation:

$$y_t = \bar{y} + a_t + b_t M_t + \varepsilon_t, \tag{7}$$

where the policy-maker seeks to control y_t with M_t. The parameters a_t and b_t change every period and are randomly distributed with joint distribution F_{ab}. Policy-makers learn about this distribution by experimenting. The authors find that the optimal policy incorporates some experimentation.

Kendrick (1982) considers the potential for learning within a model that contained 10 unknown (constant) parameters. He finds that costly experimentation is desirable, and that increased model complexity increases the amount of costly experimentation that is optimal.

There have been examples in history where a major structural change in the economy has resulted in the central bank having little reliable data with which to inform their policy decisions. One such example was the German reunification in 1990. Wieland (1996) conducts dynamic simulations of monetary policy decisions in a model calibrated to the German economy at that time. Wieland shows that passive learning by the central bank could have resulted in persistent deviations from policy objectives since some policies yield little or no information about the state of the economy. In contrast, a policy that incorporates active learning eliminated persistent policy mistakes.

The basic premise behind these learning models is that the policy-maker lacks the data required to construct accurate estimates of the model parameters, despite the fact that the parameters remain constant over time. However, except when there are major structural changes to the economy (such as immediately following reunification in Germany), the main source of uncertainty facing the monetary authority is more likely to be related to the evolution of the economy than the lack of data, as Bean (1999, 15) notes:

> In practise the main source of uncertainty is ... not due to the imprecision with which parameters are estimated as a result of econometricians having limited sample information. Rather, a stochastic, or at best evolving, parameter model seems more appropriate in which learning about the value of today's parameters is of distinctly limited value for knowing their future value.

Further, these types of active learning models do not provide an explanation for the behaviour that may be observed in practice. For example, in the United States in the mid-1990s, monetary policy remained stimulative even after many economists believed that potential output had been attained. As a result, it was discovered that potential output was greater than had previously been believed. Probing in practice, if it takes place at all, appears to exist only at the point where the central bank perceives that it is getting close to full employment. Probing of the type discussed above should be just as valuable when the policy agency is far from full employment as when it is near.

Another class of learning models utilize an economy that evolves through time. For example, Balvers and Cosimano (1994) assume that the link between money growth and inflation is time-varying and uncertain. In particular, $\pi_t = \alpha_t + \beta_t m_t$, where both α_t and β_t follow an AR(1) process. Over time, the policy-maker learns about the parameters. The authors assume that anticipated inflation has little cost, while unanticipated inflation is costly. As a result, policy-makers seek to minimize the variability of inflation. High money growth leads to high inflation and also high inflation variability, since a large m_t implies a high multiplier on the unknown parameter β_t. The optimal policy is therefore one with zero money growth. Balvers and Cosimano use a dynamic programming framework to compute the optimal policy path. They assess the impact of taking into account learning with the "myopic" policy (when the benefits of learning about the parameters are ignored in the policy formulation process) and the "cold-turkey" policy (when money supply growth is immediately set to zero). They find that the optimal policy entails a significantly faster reduction in monetary growth than the myopic policy, but one that is slower than the cold-turkey policy.

Wieland (1998) considers the impact on policy of uncertainty as to the natural unemployment rate, in a model very similar to the one we will examine below. The trade-off between inflation and unemployment follows a standard Phillips curve,

$$\pi_t = \pi_t^e + \beta(u_t^* - u_t) + \varepsilon_t, \tag{8}$$

where the natural unemployment rate u_t^* follows a random walk and $\pi_t^e = \pi_{t-1}$. The monetary authority faces the following minimization problem:

$$\underset{r_t}{Min} \quad L(\pi_t, u_t) = E_{t-1}[(\pi_t - \pi^*)^2 + \omega(u_t^* - u_t)^2]$$

$$s.t. \quad \Delta u_t = -\phi(\Delta y_t - \Delta y_t^*)$$
$$y_t - y_t^* = -\gamma(r_t - r^*). \tag{9}$$

The monetary authority does not know the values of u_t^* or β, but must estimate them with available data.

Wieland finds that in a static framework, a conservative policy is optimal. However, in a dynamic framework where the monetary authority takes explicit account of the impact of their policy on the amount of learning they can accomplish, the optimal policy lies between the static and the certainty equivalent policies. The only exception to this is when there is a very high degree of uncertainty, and inflation is close to the target. Then, the optimal policy with learning is more extreme than the static policy. This is consistent with recent experience in the United States, as outlined above.

Taking an entirely different approach, Isard and Laxton (1998) consider a model calibrated to the Australian economy in which experimentation only occurs when inflation is low in an attempt by the monetary authority to better identify the (unknown, time-varying) NAIRU. They incorporate endogenous credibility, so that probing may result in long-term costs for the monetary authority and a convex Phillips curve. While a probing policy may result in a slightly lower average rate of unemployment in their framework, this occurs at the expense of a rise in average inflation rates.

Finally, Stock (1999) argues that time-varying parameters make the use of robust control desirable. He considers a simple linear model of the United States where the parameters follow random walks, and the monetary authority chooses policy utilizing the minimax criterion. He finds that, for some types of uncertainty, policies should be more aggressive than point-estimates would suggest.

In general, the literature examined here suggests that the benefits to actively probing in a bid to determine the level of potential output are typically small. The only circumstance when the optimal "learning" policy is more aggressive than the certainty equivalent policy is when output is close to potential, and the monetary authority faces an extremely high amount of uncertainty (Wieland 1998).

In the remainder of this paper, the relationship between credibility and the benefits to probing is examined. In an economy in which there are explicit inflation targets, such as Canada's, the question addressed is whether probing is more desirable when people believe those targets will be attained than when they do not.

4. The model

The economy considered here is similar to that outlined in Wieland (1998), but with the Phillips curve defined in terms of output rather than unemployment,

$$\pi_t = \pi_t^e + \beta(y_t - y_t^*) + \varepsilon_t, \tag{10}$$

where ε_t is a price shock.

The central bank does not know the value of potential output, y_t^*, which follows a random walk: $y_t^* = y_{t-1}^* + \eta_t$.[6] They also do not know the slope of the Phillips curve, β (assumed constant), and so must learn about each of these over time. Clearly, there are also many other sources of uncertainty that enter into the problem of setting monetary policy that are ignored here; all other parameters are assumed known by the monetary authority.

Each period, the central bank uses all available information to estimate the following equations:

$$\pi_t - \pi_t^e = -\alpha_t + \beta y_t + \varepsilon_t,$$
$$\alpha_t \equiv \beta y_t^* = \alpha_{t-1} + v_t. \tag{11}$$

The estimates of α_t and β from this regression are then used to form an estimate of y_t^* given by

$$\hat{y}_t^* = \frac{\hat{\alpha}_t}{\hat{\beta}}, \tag{12}$$

which is used in the formulation of monetary policy in the following period. Monetary policy entails the setting of the real interest rate, which influences real output according to the relation

$$y_t = y_{t-1} - \gamma(r_t - r_{t-1}). \tag{13}$$

---

6. This provides the simplest possible case in which shocks to potential output are permanent. It would also be possible to consider alternative, more realistic characterizations of the evolution of potential output.

For simplicity, there is no uncertainty in this relationship: the monetary authority can always attain a desired level of output via an appropriate choice of r_t in this model, subject to the constraint that nominal interest rates cannot be negative.[7]

Inflation expectations are a weighted mean of the target and lagged inflation,

$$\pi_t^e = \lambda\pi^* + (1 - \lambda)\pi_{t-1}, \tag{14}$$

where $\lambda \in [0,1]$ is a measure of credibility. If $\lambda = 0$, then inflation expectations are equal to last period's inflation rate, while if $\lambda = 1$, inflation expectations are equal to the inflation target of the central bank.[8,9]

The monetary authority seeks to minimize its loss given by

$$\underset{r_t}{Min} \quad \sum_t \rho^t E_{t-1}[(\pi_t - \pi^*)^2 + \omega(y_t - y_t^*)^2], \tag{15}$$

where ρ is the discount rate. $\omega = 0$ represents a monetary authority that cares only about inflation deviations from target, while for $\omega \to \infty$, the monetary authority cares only about deviations of output from potential.

In a one-period world with certainty, the optimal real interest rate would be set according to the rule

$$r_t = r_{t-1} + \frac{1}{\gamma}(y_{t-1} - \hat{y}_{t-1}^*) + \frac{1}{\gamma}\left[\frac{\hat{\beta}(1-\lambda)}{\hat{\beta}^2 + \omega}\right](\pi_{t-1} - \pi^*), \tag{16}$$

subject to the restriction that the nominal interest rate cannot be negative

$$r_t \geq -\pi_t^e. \tag{17}$$

This is analogous to (EQ 3) in the Brainard case above, and will be referred to as the "certainty equivalent" policy for the remainder of the paper. An increase in central banker credibility (measured as an increase in λ) has the effect of reducing the optimal policy

---

7. It would be possible to include a demand shock term in (EQ 13), although in this model it is exactly equivalent to a shock to potential output. The monetary authority is concerned about the value of the output gap, and uncertainty as to either component of that gap is identical from their standpoint. If demand shocks were not permanent (that is, the coefficient on lagged output in (EQ 13) did not equal 1), then the effect of a demand shock would diverge from a potential shock.

8. λ is assumed known by the monetary authority. Srour (1999) showed that uncertainty about the propagation of inflation (in this framework, uncertainty about level of credibility) leads to a more aggressive policy response being appropriate.

9. With $\lambda = 0$, the model is equivalent to Wieland (1998).

response to a deviation of the inflation rate from target. Inflation expectations (and therefore future inflation rates) are less sensitive to current inflation at higher levels of credibility, so that the optimal policy is less aggressive in responding to current variation in inflation, all other things being equal.

If the central bank were to explicitly allow for the impact of uncertainty on the optimal policy in a static environment, that policy would be set according to the rule

$$
r_t = r_{t-1} + \frac{1}{\gamma}(y_{t-1} - \hat{y}_{t-1}^*) + \frac{1}{\gamma}\left[\frac{\hat{\beta}(1-\lambda)}{\hat{\beta}^2 + V(\hat{\beta}) + \omega}\right](\pi_{t-1} - \pi^*) - \frac{1}{\gamma}\left(\frac{C(\hat{\alpha}_{t-1}, \hat{\beta}) - \frac{\hat{\alpha}_{t-1}}{\hat{\beta}}V(\hat{\beta})}{\hat{\beta}^2 + V(\hat{\beta}) + \omega}\right)
$$

$$(18)$$

again subject to the restriction that nominal interest rates cannot be negative (EQ 17).

This is analogous to (EQ 4) in the Brainard example above and will be referred to as the "conservative" policy for the remainder of the paper. The additional term in the policy rule may be positive or negative and, for some economic shocks, may result in a more aggressive policy response than the certainty equivalent policy. Its presence is somewhat counterintuitive, as Wieland (1998, 15) explains:

> It implies that even in a situation where the observed inflation rate is on target and [output] equals [estimated potential output], the central bank would pursue a policy that drives [output] away from estimated [potential output] in expectation.

He goes on to explain that the final term is a function of estimates based on historical data and captures the idea that, with uncertainty, it is optimal for a central bank to lean towards the historical mean of output rather than seeking to end an inflationary or disinflationary period abruptly.[10] In general, (EQ 18) implies a more conservative policy response to shocks than (EQ 16).

Note that the difference between these two policies diminishes as the weight on output increases in the central bank's loss function and in the limit, as $\omega \to \infty$, the policies converge and do not vary with credibility. For extreme values of ω (that is, $\omega \in \{0.0, \infty\}$),

---

10. See Wieland (1998, 15–18) for a more complete discussion.

these policies are also optimal in a multi-period world where there is no learning. When the monetary authority targets only inflation or output, there is no trade-off between meeting the target this period and next. For $0 < \omega < \infty$, the extent to which meeting the inflation and output targets this period precludes meeting the inflation target next period varies with credibility. As a result, the optimal dynamic policy without learning diverges from the optimal static policy. As an example of the impact of this, the analogue of (EQ 16) for the certainty equivalent interest rate in the first period of a world that lasts for two periods and in which the monetary authority targets both inflation and output is given by

$$r_t = r_{t-1} + \frac{1}{\gamma}(y_{t-1} - \hat{y}_{t-1}*) + \frac{1}{\gamma}\left[\frac{[\hat{\beta}^2 + \omega + \rho\omega(1-\lambda)^2](1-\lambda)\hat{\beta}}{[\hat{\beta}^2 + \omega + \rho\omega(1-\lambda)^2]\hat{\beta}^2 + \omega[\hat{\beta}^2 + \omega]}\right](\pi_{t-1} - \pi*).$$

(EQ 19)

We will use this policy rule later to see if varying ω impacts the optimal amount of probing that the monetary authority should undertake.

We now consider a multi-period world in which the monetary authority learns over time. Each period, their estimates of α_t and β are updated optimally using the new data obtained. In a world with constant parameters, this would involve Bayesian updating. Because α_t is time-varying here, the appropriate analogue to Bayesian updating that results in efficient, unbiased estimates may be cast in the form of the Kalman filter:

$$\Sigma_{t|t-1} = \begin{bmatrix} v^{\alpha}_{t|t-1} & v^{\alpha\beta}_{t-1} \\ v^{\alpha\beta}_{t-1} & v^{\beta}_{t-1} \end{bmatrix} = \begin{bmatrix} (v^{\alpha}_{t-1|t-1} + \beta^2\sigma^2_{\eta}) & v^{\alpha\beta}_{t-1} \\ v^{\alpha\beta}_{t-1} & v^{\beta}_{t-1} \end{bmatrix},$$

$$\begin{bmatrix} \alpha_{t|t} \\ \beta_{t|t} \end{bmatrix} = \begin{bmatrix} \alpha_{t|t-1} \\ \beta_{t|t-1} \end{bmatrix} + \Sigma_{t|t-1}\begin{bmatrix} -1 \\ y_t \end{bmatrix}F^{-1}(\pi_t - \pi^e_t + \alpha_{t|t-1} - \beta_{t|t-1}y_t),$$

$$\Sigma_{t|t} = \Sigma_{t|t-1} - \Sigma_{t|t-1}\begin{bmatrix} 1 \\ -y_t \end{bmatrix}F^{-1}\begin{bmatrix} 1 & -y_t \end{bmatrix}\Sigma_{t|t-1},$$

$$F = \begin{bmatrix} 1 & -y_t \end{bmatrix}\Sigma_{t|t-1}\begin{bmatrix} 1 \\ -y_t \end{bmatrix} + \sigma^2_{\varepsilon}.$$

(EQ 20)

The optimal policy that takes account of the learning process and optimizes the amount of learning is now the solution to a highly non-linear problem that cannot be solved analytically. Other authors resort to computationally intensive techniques in order to

approximate the optimal policy.[11] Here, the economy is simulated under varying degrees of policy credibility to determine the degree of aggressiveness that is warranted in the face of a desire by the central bank to minimise loss when learning about the economy is endogenous to the model.

For tractability reasons, the economy is assumed to have a finite life. In an economy with only one period, the optimal policy with active learning coincides with the conservative policy, since there is no time for the monetary authority to benefit from information obtained in the first period. With two periods, these policies differ only in the first period. In reality, the benefits from learning accrue in all future periods, and not just the period immediately following. Therefore, an economy with a life of two periods provides a lower bound on the benefits of active policy.

To further examine the benefits of active learning, an economy with a life of 10 periods is also considered. An optimizing monetary authority may be expected to undertake active learning in every period except the final one. However, because a grid search is used to determine this policy, it would be computationally demanding to allow for active learning in more than one period. Active learning is therefore restricted to the first period only; thereafter, the monetary authority follows a conservative policy and all learning is passive. Examining both a 2-period and a 10-period economy allows us to assess the sensitivity of the results to the length chosen.

5.   Parameter values

Clearly the results obtained from this exercise are somewhat dependant upon the choice of parameter values. Here the values chosen are outlined, as well as the reasons for choosing them. In general, parameter values are consistent with recent studies using Canadian data, interpreting the model at an annual frequency. Further, it is assumed that the central bank knows how much it does not know. That is, if the bank does not know a parameter value, then it knows the distribution from which that parameter is drawn.

The loss function of the monetary authority is characterized by the following parameters: an inflation target of 2 per cent (that is, $\pi^* = 0.02$); a rate of time preference of 0.95; and pure inflation targeting: $\omega = 0.0$. The optimal policy with pure output targeting

---

11. For example, Wieland (1998) uses a dynamic programming algorithm that provides numerical approximations to the solution for the special case when potential output is constant.

($\omega = \infty$) was also considered, but this was identical to the policy that ignores learning. In reality, a central bank is likely to care about both the output gap and the inflation gap, an issue that will be addressed a little later.

The standard deviation of the error in the inflation process is $\sigma_\varepsilon = 0.006$, or 0.6 per cent on an annual basis. This is consistent with total variability of inflation over the past 10 years. The standard deviation of shocks to potential output is taken from Kichian (1999), who measures potential output in a state-space framework: $\sigma_\eta = 0.004$, or 0.4 per cent of the level of potential output.

Real interest rates at time zero are taken to be consistent with a nominal interest rate of 5 per cent and inflation expectations of 2 per cent: $r_0 = 0.03$; initial real output is the log of output in millions of dollars: $y_0 = 13.7$; and inflation at time zero is $\pi_0 = 0.015$.

It is assumed that the central bank believes the economy to be in excess supply at time 0,[12] with $E_0(y_0^*)$ chosen consistent with the belief of the central bank being incorrect (and the economy actually being in excess demand) 45, 15, and 1 per cent of the time respectively. That is, $E_0(y_0^*) = y_0 + Z(\sqrt{V_0(y_0^*)})$ where (for the former case) Z corresponds to the score in the standard normal distribution associated with 45 per cent of the upper tail being greater than Z. The 45 per cent case may be thought of as high initial uncertainty as to the level of potential output, with the 1 per cent case corresponding to low initial uncertainty.

The initial level of potential output, y_0^*, is chosen at random from the normal distribution that is centred on the monetary authority's expectation of potential output, and has a variance consistent with the belief of the central bank. That is, $y_0^* = E_0(y_0^*) + \varepsilon_{y^*}$ where $\varepsilon_{y^*} \sim N(0, V_0(y_0^*))$. The variance is chosen consistent with the variance found in estimates of the level of potential output in recent years: $V_0(y_0^*) = (0.005)^2$. The impact of real interest rates on output is consistent with estimates obtained by Duguay (1994): $\gamma = 1.0$.

At time zero, the monetary authority believes that the slope of the Phillips curve is $E_0(\beta) = 0.5$, which is consistent with a sacrifice ratio of 2 when the monetary authority has no credibility.[13] The value of α_0 is chosen to be consistent with this:

---

12. This assumption does not limit the applicability of the results. The optimal policy when the monetary authority believes it is facing an excess demand will be the mirror image of that obtained here.

13. Recent estimates of the sacrifice ratio for Canada include 1.5 (Dupasquier and Girouard 1992), 1.7 (Duguay 1994), and 2.2 (Fillion and Léonard 1997).

$E_0(\alpha_0) = E_0(\beta)E_0(y_0^*)$. The true value of β is drawn from a distribution that is centred on the monetary authority's expectations, so that $\beta = E_0(\beta) + \varepsilon_\beta$, where $\varepsilon_\beta \sim N(0, V_0(\beta))$, and $V_0(\beta) = (0.05)^2$. The Bank's initial estimates of $V_0(\alpha_0)$ and $C_0(\alpha_0, \beta)$ are chosen to be consistent with $V_0(\beta)$ and $V_0(y_0^*)$:

$$V_0(\alpha_0) = (E_0(\beta))^2 V_0(y_0^*) + (y_0^*)^2 V_0(\beta) + V_0(y_0^*)V_0(\beta),$$
$$C_0(\alpha_0, \beta) = V_0(\beta)y_0^*. \tag{21}$$

The economy is simulated with varying degrees of central bank credibility ($0 \le \lambda \le 1$). In every period, the monetary authority updates their estimates of α_t, β, their variances and covariance, and uses these new estimates in the selection of policy. Certainty equivalent and conservative "passive learning" policies are constructed for all periods. A grid search is then used to find the first period interest rate that minimizes the expected value of the monetary authority's losses over 10,000 artificial runs of the future, assuming a conservative policy for all periods following the first period. This will converge to the optimal active learning policy as the sample size increases.[14]

6.   Results

In analyzing the results, there are several important matters to bear in mind. First, the initial real interest rate is 3 per cent, so all interest rates should be compared with this. Second, the central bank believes that the economy is initially in excess supply and knows the probability with which that belief is correct. And third, only two sources of parameter uncertainty have been incorporated in the model: uncertainty as to the level of potential output, and uncertainty as to the slope of the Phillips curve. There are many other sources of uncertainty, which would lead to a greater difference between the certainty equivalent and conservative policies and would also likely increase the potential gains to probing.[15]

---

14. By definition, the optimal policy in a one-period world (when there are no benefits to probing) is given by (EQ 17). Experimentation revealed that 10,000 runs were sufficient to ensure that the simulated optimal policy equals the theoretical optimal policy to six decimal places for the formulations of the model considered here.

15. For example, the monetary authority may face uncertainty as to whether the Phillips curve is linear or not, or be unsure of the value of other parameters in the economy, which may also be evolving over time.

The policy based on (EQ16) is labelled the certainty equivalent policy, (EQ18) the conservative policy, and the simulated policy that incorporates an optimal amount of learning the active learning policy.

First consider simulations 1 to 3, the results of which are given in Table 1 and Figure 1. High initial uncertainty (simulation 1) refers to an economy where the initial point estimates of the monetary authority indicate a state of excess supply at time 0, but the initial variance estimates indicate a 45 per cent probability of being wrong. For moderate and low uncertainty (simulations 2 and 3), these percentages are 15 and 1 respectively.

The certainty equivalent policy is slightly more aggressive than the conservative policy, although as credibility increases, the extent to which they differ diminishes. The optimal policy with active learning is more aggressive than the alternative policies at low levels of credibility, but becomes less aggressive as credibility rises. In the recent past, monetary policy in Canada has generally adjusted in 25-basis-point increments. Except at low levels of credibility, the effect of active learning on policy is always much less than one increment. Also, the impact of uncertainty (that is, the difference between the certainty equivalent policy and the conservative policy) is not large. Even the difference between an economy with a life of two periods and one with 10 periods is negligible, except at very low levels of credibility.

The results of simulation 3 yield the greatest difference in policy caused by active learning. This is the situation where the monetary authority has an extremely good initial information set and is almost certain that the economy is in excess supply. Under these circumstances, a monetary authority with little credibility should run a more aggressive monetary policy in order to learn optimally about the parameters of the economy.

Next, more extreme parameter variability was considered under moderate uncertainty, to examine the robustness of this result (see Figure 2). In simulation 4, the standard deviation of innovations in the Phillips curve (σ_ε) was increased by a factor of 5, while in simulations 5 through 7, σ_η (the standard deviation of innovations to potential output), $\sqrt{V_0(y_0{}^*)}$ (the initial standard deviation of potential output), and $\sqrt{V_0(\beta)}$ (the initial standard deviation of the slope of the Phillips curve) were increased by factors of 5 respectively.[16]

---

16. To avoid the nominal interest bound ($r_t \geq -\pi_t^e$), simulation 6 was conducted assuming an initial real interest rate of 10 per cent.

Increased inflation shock volatility (simulation 4) results in an optimal learning policy that is substantially more aggressive than alternative policies, although the extent of this declines as credibility rises. At low levels of credibility, the optimal nominal interest rate is equal to close to its lower bound, and even with moderate levels of credibility, the differences are still significantly greater than 25 basis points. Very similar results are also obtained for increased potential output shock volatility (simulation 5). Increased uncertainty about the initial level of potential output (simulation 6) produces qualitatively similar results, although the magnitude of the difference in policies is smaller. In contrast, increased uncertainty as to the value of β (simulation 7) drives a wedge between the certainty equivalent and conservative policies, with the optimal learning policy lying between these.

These results indicate that probing may be beneficial for a monetary authority with low or moderate credibility if the economy is experiencing large inflation or potential output shocks or if the monetary authority has very poor information about the level of potential output. However, increased uncertainty about the slope of the Phillips curve does not warrant much change in the optimal policy in order to learn.

Finally, a two-period world in which both output and inflation are targeted was considered, with parameter values set equal to those considered in simulation 1 (see Figure 3). As noted in Section 4, this adds significant complexity to the problem. The certainty equivalent policy for the first period is obtained from (EQ 19), and for the second period from (EQ 16). The conservative and active learning policies for the first period are both obtained using simulation methods. For the conservative policy, the policy that minimizes first period loss is appropriate, while for the active learning policy, the policy that minimizes combined first and second period losses is appropriate, where the interest rate in the second period is set according to (EQ 18).

With a small weight on deviations of output from potential, there is little change from the results already described. However, if deviations of output from potential are weighted equally in percentage terms with deviations of inflation from target ($\omega = 1$; simulation 9), there is a divergence between the conservative policy and the certainty equivalent policy, with the optimal learning policy almost indistinguishable from the latter. Once again, as credibility rises, this divergence diminishes. Finally, if the monetary authority places very little weight on inflation deviations from the target (simulation 11), credibility has little bearing on the optimal policy and all three policies are quantitatively very similar.

7. Conclusions

Simulations have been conducted here on an artificial economy calibrated to reflect a simple model of the Canadian economy, where probing is interpreted as following a more aggressive policy in order to learn about the parameters of the economy. The optimal amount of probing for a monetary authority that seeks to target inflation has been shown to be generally small, and to vary little with credibility. Only with low levels of credibility or unrealistically large levels of uncertainty or volatility does the optimal policy with probing diverge by more than one policy increment (25 basis points) from a policy that ignores learning. Even then, for most forms of uncertainty, the optimal amount of probing diminishes as credibility rises.

The results also suggest that the optimal amount of probing decreases with credibility because of the positive impact increased credibility plays in reducing output and inflation volatility in the economy. The monetary authority's estimated equation (EQ 11) effectively equates the inflation gap (inflation less expectations) with the output gap. At higher levels of credibility, these gaps are small on average and increasingly indistinguishable from the shock terms. The information contained in a new observation is small under such circumstances, and the monetary authority's estimates do not change very much over time whether the monetary authority chooses to probe or not. In contrast, at lower levels of credibility, there will generally be significant inflation and output gaps, with larger improvements in the precision of the monetary authority's estimates from one period to the next. The informational benefits from probing are therefore greatest at low levels of credibility, resulting in a negative relationship between the optimal amount of probing and the level of credibility enjoyed by the monetary authority.

There are several limitations to this analysis. First, credibility is assumed to be known by the monetary authority and is independent of policy. In reality, a monetary authority cannot be sure of the amount of credibility it enjoys, and the act of probing may result in reduced credibility.

Second, the scope of uncertainty is this model is very limited. The monetary authority is uncertain only about the level of potential output and the slope of the Phillips curve. The reality facing policy-makers is that uncertainty is considerably more pervasive. In this setting, the benefits to probing may be larger, although this remains to be established.

Third, the initial estimation errors the monetary authority makes in estimating β and y_0^* are independent of each other. In general, these estimates will be negatively correlated

so that an overestimate of the slope of the Phillips curve will imply an underestimate of the output gap.

Fourth, inflation has no impact on output in this model, and so aside from entering the loss function of the monetary authority, has no cost to the economy. Therefore, the level of credibility does not influence the policy that is required to attain an output target, although it has a substantial effect on the policy that is required to attain an inflation target.

Finally, these results may be sensitive to the interpretation of probing considered. Isard and Laxton (1998) develop an alternative view of probing in which probing only occurs when inflation is low, and credibility is endogenous.

Table 1: Real interest rate (%)

λ	Certainty Equivalent	Conserva- tive	Active Learning	
			10 Period	2 Period
Simulation 1 High Initial Uncertainty About Potential				
0.0	0.01937	0.01947	0.01926	0.01911
0.2	0.02138	0.02145	0.02139	0.02131
0.4	0.02338	0.02343	0.02342	0.02336
0.6	0.02538	0.02541	0.02539	0.02534
0.8	0.02738	0.02739	0.02737	0.02731
1.0	0.02938	0.02938	0.02935	0.02929
Simulation 2 Moderate Initial Uncertainty About Potential				
0.0	0.01482	0.01492	0.01416	0.01442
0.2	0.01682	0.01690	0.01670	0.01665
0.4	0.01882	0.01880	0.01891	0.01879
0.6	0.02082	0.02086	0.02083	0.02078
0.8	0.02282	0.02284	0.02281	0.02276
1.0	0.02482	0.02482	0.02472	0.02474
Simulation 3 Low Initial Uncertainty About Potential				
0.0	0.00837	0.00846	0.00629	0.00740
0.2	0.01037	0.01044	0.00981	0.00995
0.4	0.01237	0.01242	0.01205	0.01231
0.6	0.01437	0.01440	0.01438	0.01432
0.8	0.01637	0.01638	0.01636	0.01630
1.0	0.01837	0.01837	0.01834	0.01828

Figure 1

Simulation 1
10 period

Simulation 1
2 period

Simulation 2
10 period

Simulation 2
2 period

Simulation 3
10 period

Simulation 3
2 period

Figure 2

Figure 3

Simulation 8

2 period,omega=0.1

Simulation 9

2 period,omega=1

Simulation 10

2 period,omega=10

Simulation 11

2 period,omega=100
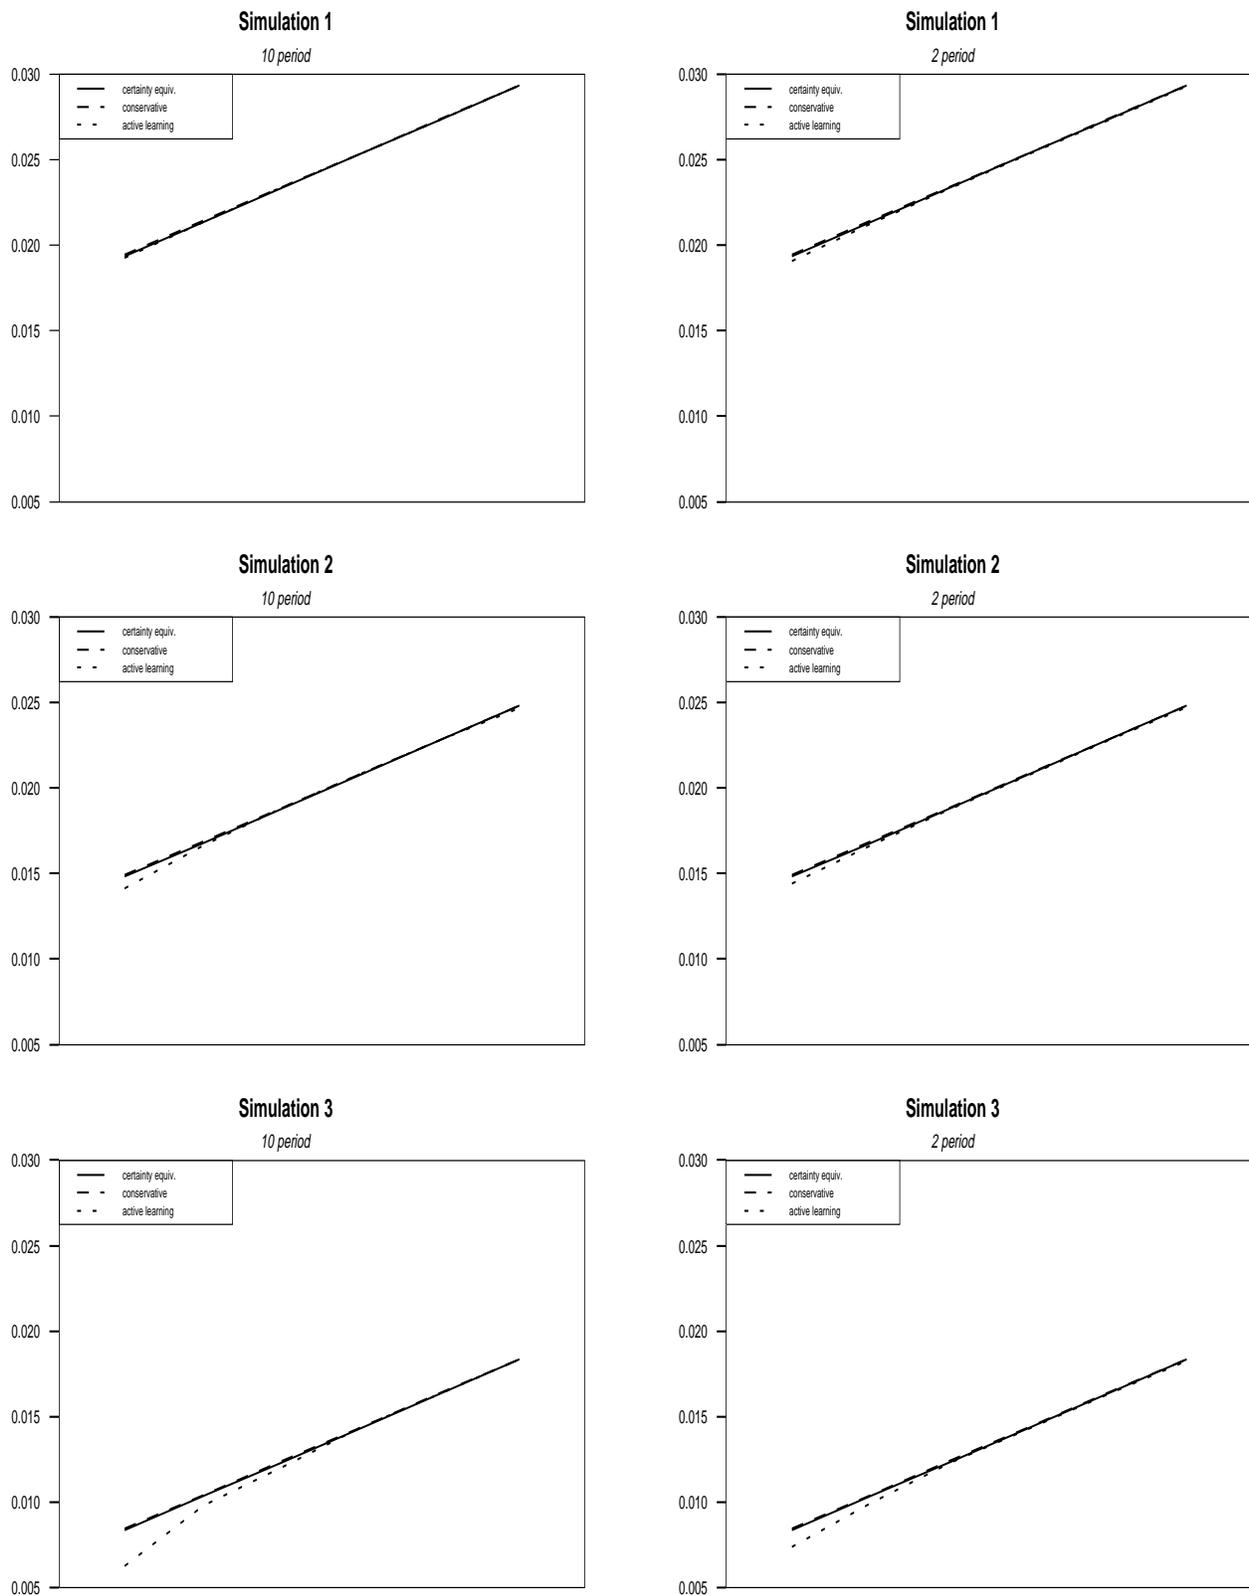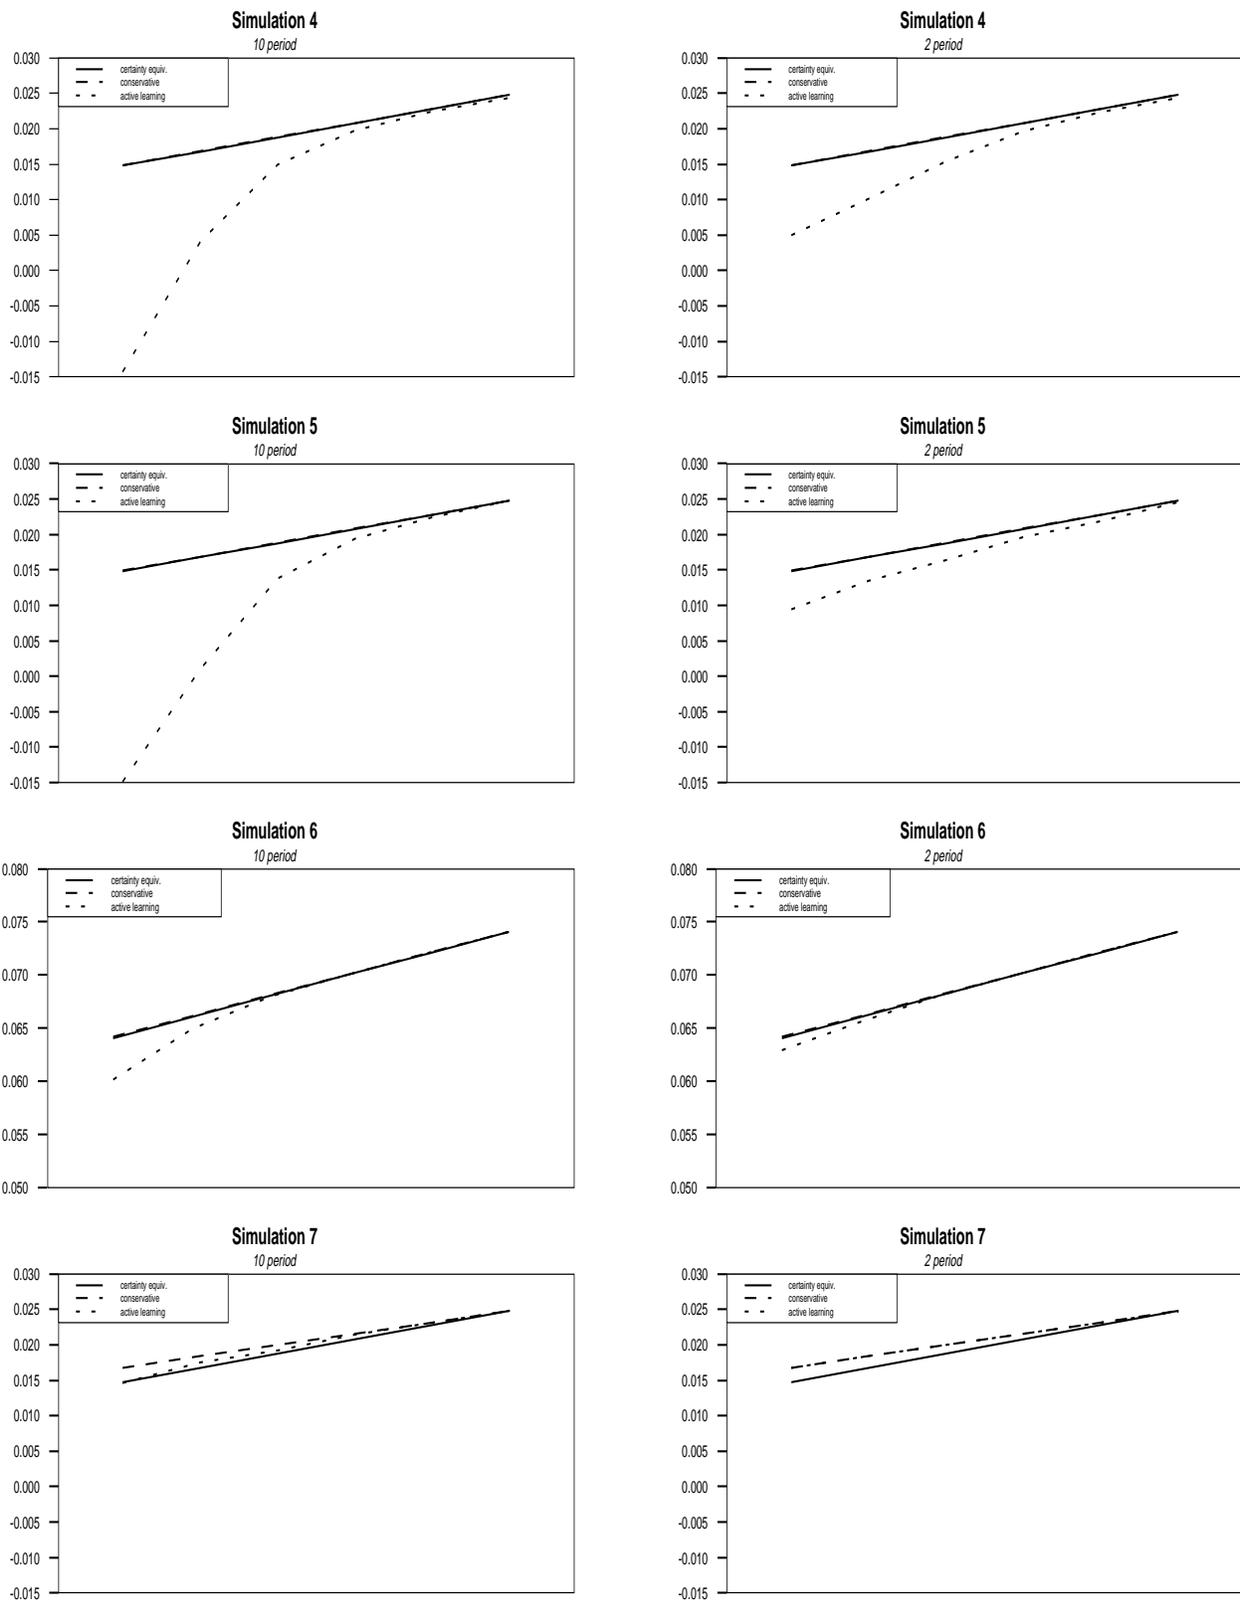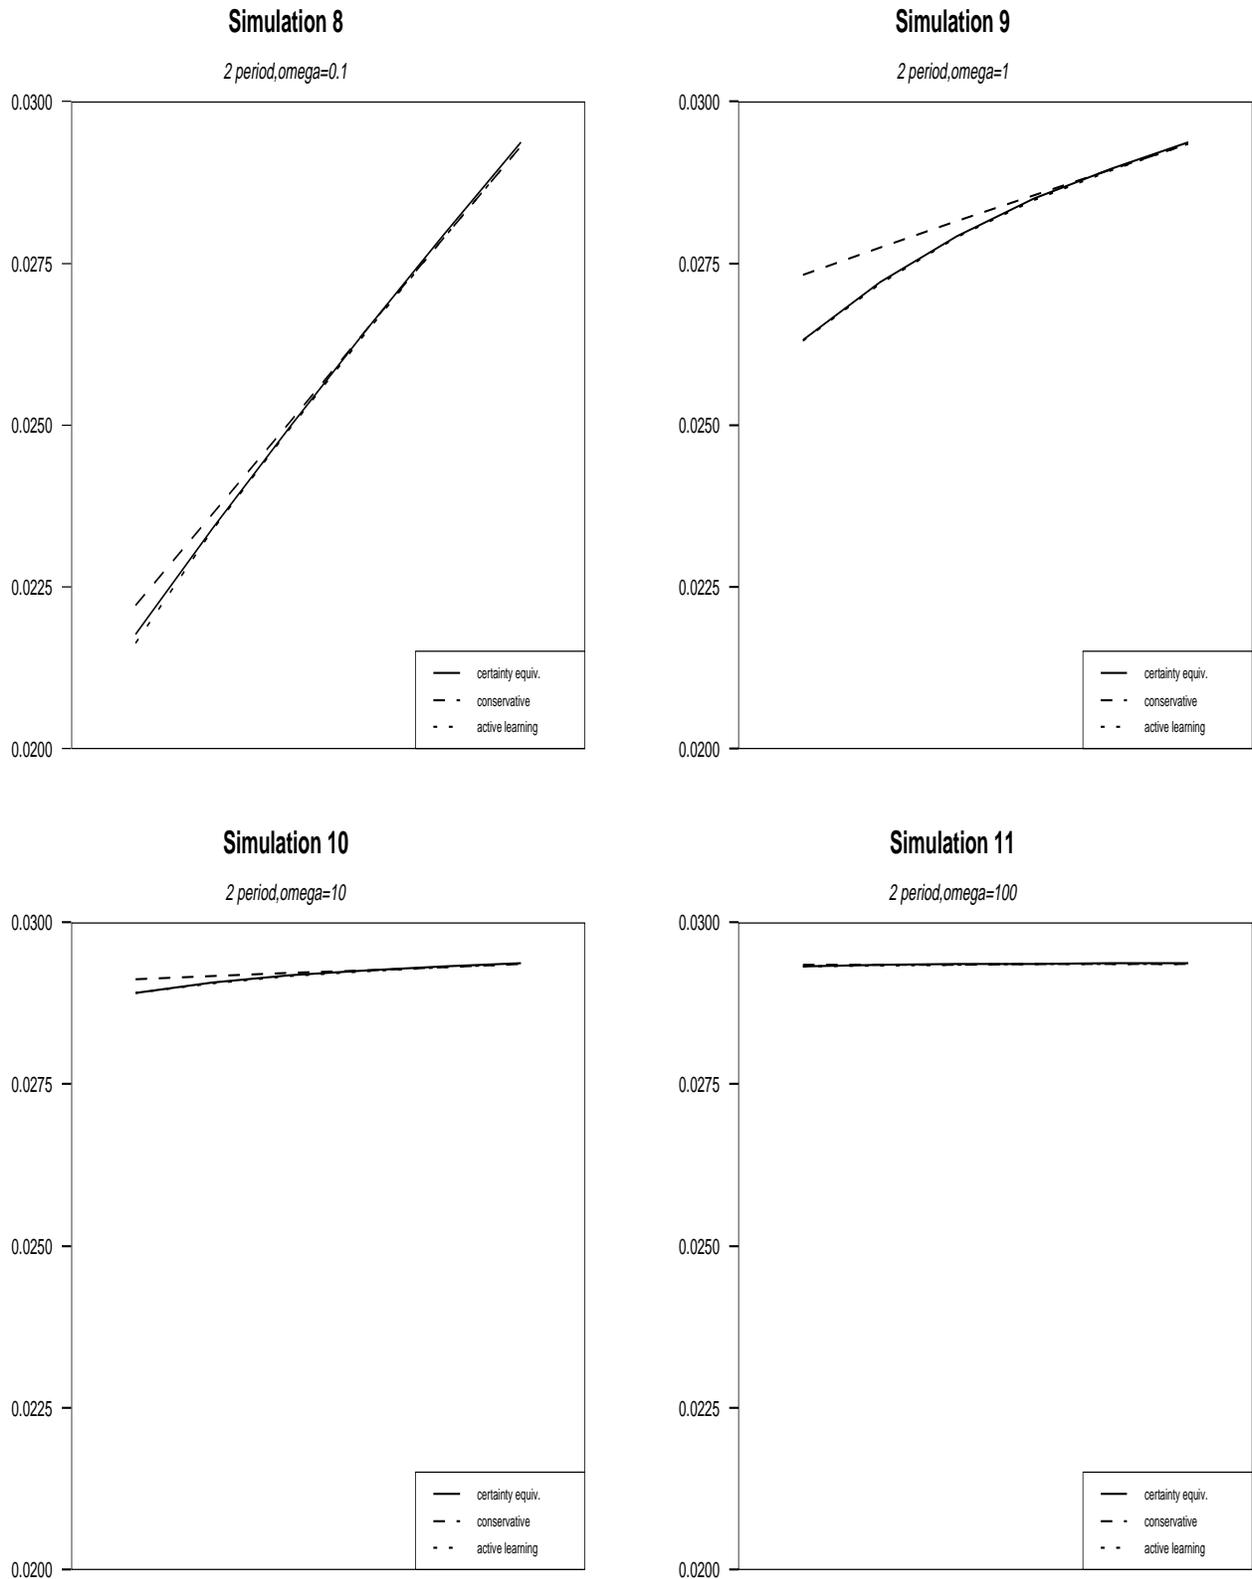
References

Amano, R., D. Coletti, and T. Macklem. 1998. "Monetary Rules When Economic Behaviour Changes." Bank of Canada Working Paper No. 99-8.

Aoki, K. 1998. "On the Optimal Monetary Response to Noisy Indicators." Unpublished.

Balvers, R.J. and T.F. Cosimano. 1994. "Inflation Variability and Gradualist Monetary Policy. *Review of Economic Studies* 61: 721–738.

Bean, C. 1999. "The Convex Phillips Curve and Macroeconomic Policy Making Under Uncertainty." Unpublished.

Bertocchi, G. and M. Spagat. 1993. "Learning, Experimentation, and Monetary Policy." *Journal of Monetary Economics* 32: 169–183.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57: 411–425.

Butler, L. 1996. *The Bank of Canada's New Quarterly Projection Model, Part 4. A Semi-Structural Method to Estimate Potential Output: Combining Economic Theory with a Time-Series Filter.* Technical Report No. 77. Ottawa: Bank of Canada.

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

Duguay, P. 1994. "Empirical Evidence on the Strength of the Monetary Transmission Mechanism in Canada." *Journal of Monetary Economics* 33: 39–61.

Dupasquier, C. and N. Girouard. 1992. "Une nouvelle courbe de Phillips au Canada." Bank of Canada. Unpublished.

Dupasquier, C., A. Guay, and P. St-Amant. 1999. "A Survey of Alternative Methodologies for Estimating Potential Output and the Output Gap." *Journal of Macroeconomics* 21: 577–595.

Dupasquier, C. and N. Ricketts. 1998. "Non-Linearities in the Output-Inflation Relationship: Some Empirical Results for Canada." Bank of Canada Working Paper No. 98-14.

Estrella, A. and F. Mishkin. 1998. "Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty." NBER Working Paper No. 6518.

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.

Isard, P. and D. Laxton. 1998. "Monetary Policy with NAIRU Uncertainty and Endogenous Credibility: Perspectives on Policy Rules and the Gains from Experimentation and Transparency." Reserve Bank of New Zealand Conference on Monetary Policy Under Uncertainty.

Johnson, D.R. 1997. "Expected Inflation in Canada 1988-1995: An Evaluation of Bank of Canada Credibility and the Effect of Inflation Targets." *Canadian Public Policy* 23: 233–258.

———. 1998. "The Credibility of Monetary Policy: International Evidence Based on Surveys of Expected Inflation." In *Price Stability, Inflation Targets, and Monetary Policy,* 361–395. Proceedings of a conference held at the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Kendrick, D. 1982. "Caution and Probing in a Macroeconomic Model." *Journal of Economic Dynamics and Control* 4:149–170.

Kichian, M. 1999. "Measuring Potential Output within a State-Space Framework." Bank of Canada Working Paper No. 99-9.

Kuttner, K. 1992. "Monetary Policy with Uncertain Estimates of Potential Output." *Federal Reserve Bank of Chicago Economic Perspectives* 16: 2–15.

Laxton, D., N. Ricketts, and D. Rose. 1994. "Uncertainty, Learning and Policy Credibility." In *Economic Behaviour and Policy Choice Under Price Stability,* 173–226. Proceedings of a conference held at the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Laxton, D. and R. Tetlow. 1992. *A Simple Multivariate Filter for the Measurement of Potential Output.* Technical Report No. 59. Ottawa: Bank of Canada.

Perrier, P. 1998. "Un examen de la Crédibilité de la Politique Monétaire au Canada." Bank of Canada Working Paper No. 98-12.

Poole, W. 1998. "A Policymaker Confronts Uncertainty." *Federal Reserve Bank of St. Louis Review* (September/October): 3–8.

Rudebusch, G.D. and L.E.O. Svensson. 1998. "Policy Rules for Inflation Targeting." NBER Working Paper No. 6512.

Sack, B. 1998. "Uncertainty, Learning, and Gradual Monetary Policy." Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, Discussion Paper No. 98-34.

Smets, F. 1998. "Output Gap Uncertainty: Does it Matter for the Taylor Rule?" Bank of International Settlements Working Paper No. 60.

Srour, G. 1999. *Inflation Targeting under Uncertainty*. Technical Report No. 85. Ottawa: Bank of Canada.

St-Amant, P. and S. van Norden. 1997. *Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada.* Technical Report No. 79. Ottawa: Bank of Canada.

Stock, J.H. 1999. "Monetary Policy in a Changing Economy: Indicators, Rules, and the Shift Towards Intangible Output." Bank of Japan IMES Discussion Paper No. 99-E-13.

Svensson, L.E.O. 1997. "Inflation Targeting: Some Extensions." NBER Working Paper No. 5962.

Thiessen, G. 1995. "Uncertainty and the Transmission of Monetary Policy in Canada." *Bank of Canada Review* (Summer): 41–58.

van Norden, S. 1995. "Why is it So Hard to Measure the Current Output Gap?" Bank of Canada. Unpublished.

Wieland, V. 1996. "Monetary Policy, Parameter Uncertainty and Optimal Learning." Board of Governors of the Federal Reserve System. Unpublished.

———. 1998. "Monetary Policy and Uncertainty about the Natural Unemployment Rate." Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, Discussion Paper No. 98-22.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. *Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.*

2000

1999