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

April 30, 2002

James Yetman*

Abstract

The effective conduct of monetary policy is complicated by uncertainty about the level of potential output. One possible response is for the central bank to "probe," interpreted here as actively using its policy to learn about the level of potential output. I consider a simple calibrated model in the Canadian context and examine the relationship between credibility and optimal probing. For plausible parameter values, the optimal amount of probing is small and diminishes slightly as credibility rises. Only for unrealistically low levels of credibility or unrealistically large levels of uncertainty or volatility does the optimal policy diverge significantly from a policy that ignores learning.

JEL codes: E52, E58

* School of Economics and Finance, University of Hong Kong, Pokfulam Road, Hong Kong; phone (852) 2857-8506; fax (852) 2548-1152; email jyetman@econ.hku.hk. The author thanks seminar participants at the Bank of Canada and the Society for Computational Economics meetings in Barcelona (2000), as well as an anonymous referee, for their comments. Frederic Beauregard-Tellier provided excellent research assistance. All remaining errors are the author's sole responsibility.

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 (see Thiessen (1995) and Poole (1998) for discussions). 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, (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), 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 central bank 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 central bank 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 behavior of the central bank. 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.

Many inflation targeting central banks have put significant emphasis on attaining credibility for their policy objectives in recent years (see Amano, Coletti, and Macklem (1999) for details of steps taken by the Bank of Canada, for example). These steps have increased the accountability of the central banks and that, together with the realized inflation record, has enhanced their credibility in the sense that expectations of inflation have become more firmly anchored to the inflation target (see Johnson (1997, 1998) or Perrier (1998) for evidence of this). The question addressed here is whether an increase in credibility increases the desirability of probing. In other words, should a central bank 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 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 central bank 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 central bank 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

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

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

$$p = (y^* - u)/a, \tag{2}$$

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:

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

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:

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

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

$$p = y^*/[E(a) + (\sigma_a^2/E(a))]. \tag{5}$$

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, 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 (1999) 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 behavior 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. Similarly, Svensson (1999) finds that the optimal monetary policy under parameter uncertainty is more conservative than the certainty equivalent policy in a simple analytic model. 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 (1999) 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.[1]

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 assumes an I.S. curve given by

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

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 policymakers 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 policymakers 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 (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 central bank will then take account of the amount of information generated by the policy. In general, the optimal policy that takes account of learning will be more aggressive than the multiplicative uncertainty policy (5), but less aggressive than the certainty equivalent policy (3), as other authors have argued.

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

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

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. Similarly, 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.

When an economy has undergone a major structural change, the central bank may have little reliable data with which to inform policy decisions. Wieland (2000) conducts dynamic simulations of monetary policy decisions in a model calibrated to the German economy at the time of reunification (1990) and 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. An alternative source of uncertainty is related to the evolution of the economy over time, as Bean (1999, 15) notes:

> In practice 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.

Balvers and Cosimano (1994) consider such a model where $\pi_t = \alpha_t + \beta_t m_t$, and both α_t and β_t follow an AR(1) process. Over time, the policy-maker learns about the parameters. If unanticipated inflation is costly, policy-makers will seek to minimize inflation variability. Because high money growth leads to high inflation variability (a large m_t implies a high multiplier on the unknown parameter β_t), the optimal policy entails 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 that will be examined below. He finds that in a static framework, a conservative policy is optimal. However, in a dynamic framework where the central bank 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.

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 central bank to better identify the (unknown, time-varying) NAIRU. They incorporate endogenous credibility, so that probing may result in long-term costs for the central bank, 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 central bank 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 central bank 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, 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,

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

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: $y_t^* = y_{t-1}^* + \eta_t$.[2] 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 central bank.

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

$$\pi_t - \pi_t^e = -\alpha_t + \beta y_t + \epsilon_t,$$

$$\alpha_t \equiv \beta y_t^* = \alpha_{t-1} + \nu_t. \tag{9}$$

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

$$\hat{y}_t^* = \hat{\alpha}_t / \hat{\beta}, \tag{10}$$

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

$$y_t = y_{t-1} - \gamma(r_t - r_{t-1}). \tag{11}$$

For simplicity, there is no uncertainty in this relationship: the central bank 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.[3]

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

$$\pi_t^e = \lambda \pi^* + (1 - \lambda)\pi_{t-1}, \tag{12}$$

where $\lambda \in [0, 1]$ is a measure of credibility. Inflation expectations are not rational in this model. Instead, this ad hoc specification captures the idea that as credibility increases, inflation expectations become more strongly anchored to the target of the central bank. Furthermore, agents do not try to learn about the central bank's preferences over time (see, for example, Cukierman and Meltzer (1986), Faust and Svensson (2001) and related papers). 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.[4,5] Given the observed persistence in inflation, realistic values of λ are likely to be significantly less than 1.

The central bank seeks to minimize its loss given by

$$\underset{r_t}{Min} \sum_t \rho^t E_{t-1}[(\pi_t - \pi^*)^2 + \omega(y_t - y_t^*)^2], \tag{13}$$

where ρ is the discount rate. $\omega = 0$ represents a central bank that cares only about inflation deviations from target, while for $\omega \to \infty$, the central bank 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

$$r_t = r_{t-1} + \frac{1}{\gamma}(y_{t-1} - \hat{y}_{t-1}^*) + \frac{1}{\gamma}\Big(\frac{\hat{\beta}(1-\lambda)}{\hat{\beta}^2 + \omega}\Big)(\pi_{t-1} - \pi^*), \tag{14}$$

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

$$r_t \geq -\pi_t^e. \tag{15}$$

This is analogous to (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 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

$$r_t = r_{t-1} + \frac{1}{\gamma}(y_{t-1} - \hat{y}^*_{t-1}) + \frac{1}{\gamma}\Big(\frac{\hat{\beta}(1-\lambda)}{\hat{\beta}^2 + V(\hat{\beta}) + \omega}\Big)(\pi_{t-1} - \pi^*) - \frac{1}{\gamma}\Big(\frac{C(\hat{\alpha}_{t-1}, \hat{\beta}) - \frac{\hat{\alpha}_{t-1}}{\hat{\beta}}V(\hat{\beta})}{\hat{\beta}^2 + V(\hat{\beta}) + \omega}\Big), \tag{16}$$

again subject to the restriction that nominal interest rates cannot be negative (15). This is analogous to (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.[6] In general, (16) implies a more conservative policy response to shocks than (14).

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, $\omega \in \{0.0, \infty\}$), these policies are also optimal in a multi-period world where there is no learning. When the central bank 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 (14) for the certainty equivalent interest rate in the first period of a world that lasts for two periods and in which the central bank targets both inflation and output is given by

$$r_t = r_{t-1} + \frac{1}{\gamma}(y_{t-1} - \hat{y}^*_{t-1}) + \frac{1}{\gamma}\Big(\frac{[\hat{\beta}^2 + \omega + \rho\omega(1-\lambda)^2](1-\lambda)\hat{\beta}}{[\hat{\beta}^2 + \omega + \rho\omega(1-\lambda)^2]\hat{\beta}^2 + \omega[\hat{\beta}^2 + \omega]}\Big)(\pi_{t-1} - \pi^*). \quad (17)$$

This policy rule will be used later to see if varying ω impacts the optimal amount of probing that the central bank should undertake.

Consider a multi-period world in which the central bank 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:

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

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

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

$$F = \begin{pmatrix} 1 & -y_t \end{pmatrix} \Sigma_{t|t-1} \begin{pmatrix} 1 \\ -y_t \end{pmatrix} + \sigma^2_{\epsilon}. \quad (18)$$

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.[7] Here, the economy is simulated under varying degrees of policy credibility. This has the advantage of being less computationally intensive, allowing for a broader range of parameter values to be considered than has been elsewhere, at the cost of precision.

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 central bank 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 central bank 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 central bank follows a conservative policy and all learning is passive.

By allowing learning only in the first period and limiting the economy to 10 periods, there are two offsetting effects on the degree of experimentation in the first period. The benefits of learning last only 9 periods, which will reduce the amount of active learning. But the simulations have the effect of forcing central banks to endogenise the restriction on learning in later periods, so central banks will try to learn more in the first period than if they were solving the optimal dynamic policy problem. It is not clear which one of these will dominate. Nevertheless, examining both a 2-period and a 10-period economy offers insights on 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 central bank is characterized by the following parameters: an inflation target of 2 percent ($\pi^* = 0.02$); a rate of time preference of $\rho = 0.95$; and pure inflation targeting: $\omega = 0.0$. In reality, a central bank is likely to care about both the output gap and the inflation gap, an issue that will be addressed later.

The standard deviation of the error in the inflation process is $\sigma_\epsilon = 0.006$, or 0.6 percent on an annual basis, while the standard deviation of shocks to potential output is $\sigma_\eta = 0.004$ (Kichian 1999). Initial real interest rates, output, and inflation are given by $r_0 = 0.03$, $y_0 = 13.7$, and $\pi_0 = 0.015$ respectively.

It is assumed that the central bank believes the economy to be in excess supply at time 0,[8] 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 percent 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 percent of the upper tail being greater than Z. The 45 percent case may be thought of as high initial uncertainty as to the level of potential output, with the 1 percent case corresponding to low initial uncertainty. The initial level of potential output, y_0^*, satisfies $y_0^* = E_0(y_0^*) + \epsilon_{y^*}$ where $\epsilon_{y^*} \sim N(0, V_0(y_0^*))$ and $V_0(y_0^*) = (0.005)^2$. The impact of real interest rates on output is $\gamma = 1.0$ (Duguay 1994).

At time zero, the central bank 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 central bank has no credibility.[9] The value of α_0 is chosen to be consistent with this: $E_0(\alpha_0) = E_0(\beta)E_0(y_0^*)$. The true value of β satisfies $\beta = E_0(\beta) + \epsilon_\beta$, where $\epsilon_\beta \sim N(0, V_0(\beta))$, and $V_0(\beta) = (0.05)^2$. The central 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 $V_0(y_0^*)$:

$$
\begin{aligned}
V_0(\alpha_0) = (E_0(\beta))^2 V_0(y_0^*) + (y_0^*)^2 V_0(\beta) + V_0(y_0^*) V_0(\beta), \\
C_0(\alpha_0, \beta) = V_0(\beta) y_0^*.
\end{aligned}
\tag{19}
$$

The economy is simulated with varying degrees of central bank credibility ($0 \leq \lambda \leq 1$). In every period, the central bank 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 central bank'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.[10]

6. Results

The policy based on (14) is labeled the certainty equivalent policy, (16) 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 (2 and 10 periods) and Figure 1 (10 periods only). High initial uncertainty (simulation 1) refers to an economy where the initial point estimates of the central bank indicate a state of excess supply at time 0, but the initial variance estimates indicate a 45 percent 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 central bank has an extremely good initial information set and is almost certain that the economy is in excess supply. Under these circumstances, a central bank with little credibility should run a more aggressive monetary policy in order to learn optimally about the parameters of the economy.

Next, σ_ϵ, σ_η, $\sqrt{V_0(y_0^*)}$, and $\sqrt{V_0(\beta)}$ were increased by a factor of 5 in simulations 4 through 7 respectively to examine the robustness of this result to more extreme parameter variability (see Figure 2 for the 10 period results).[11]

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 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 central bank with low or moderate credibility if the economy is experiencing large inflation or potential output shocks or if the central bank 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.

Next, 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). The certainty equivalent policy for the first period is obtained from (17), and for the second period from (14). 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 (16).

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.

Finally, different levels of initial inflation are considered under low uncertainty (see Figure 4). In simulation 12, initial inflation is equal to the target ($\pi_0 = 0.02$), in simulation 13 $\pi_0 = 0.015$, while in simulation 14, initial inflation is far from the target ($\pi_0 = 0.00$). First note that the active learning policy is more aggressive, relative to the certainty equivalent policy, if initial inflation is closer to the inflation target. However, relative to the conservative policy, the degree of active learning is largely invariant to the initial level of inflation. Once again, optimal probing is small relative to the 25-basis-point policy increments observed in practice.

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 central bank 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 central bank's estimated equation (9) 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 central bank's estimates do not change very much over time whether the central bank 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 central bank'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 central bank.

There are several limitations to this analysis. First, credibility is assumed known by the central bank 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 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 central bank 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 central bank, 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.

Footnotes

1 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.

2 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.

3 It would be possible to include a demand shock term in (11), although in this model it is exactly equivalent to a shock to potential output. The central bank 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 (11) did not equal 1), then the effect of a demand shock would diverge from a potential shock.

4 λ is assumed known by the central bank. Srour (1999) shows that uncertainty about the propagation of inflation (in this framework, uncertainty about the level of credibility) leads to a more aggressive policy response being optimal.

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

6 See Wieland (1998, 15-18) for a more complete discussion.

7 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.

8 This assumption does not limit the applicability of the results, since monetary policy is symmetric in this model.

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

10 By definition, the optimal policy in a one-period world (when there are no benefits to probing) is given by (16). 10,000 runs were sufficient to ensure that the simulated optimal policy matches the theoretical optimal policy to six decimal places (1/100th of a basis point).

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

Notation.

π_t	inflation
π_t^e	inflation expectations
y_t	output
y_t^*	potential output
r_t	interest rate
ϵ_t	inflation shock
η_t	potential shock
α_t	Phillips Curve intercept
β	Phillips Curve slope
ν_t	Phillips Curve error
π^*	inflation target
ω	weight on output in loss function
ρ	discount rate
λ	credibility

Table 1. Real Interest Rate.

λ	Certainty Equivalent	Conservative	Active Learning 10 period	Active Learning 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

References

Amano, R., D. Coletti, and T. Macklem. "Monetary Rules When Economic Behaviour Changes." Bank of Canada Working Paper: 99-8. 1999.

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

Balvers, R. J. and T. F. Cosimano. "Inflation Variability and Gradualist Monetary Policy." Review of Economic Studies 61 (1994): 721-38.

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

Bertocchi, G. and M. Spagat. "Learning, Experimentation, and Monetary Policy." Journal of Monetary Economics 32 (1993): 169-83.

Brainard, W. "Uncertainty and the Effectiveness of Policy." American Economic Review 57 (1967): 411-25.

Butler, L. "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." Bank of Canada Technical Report: 77. 1996.

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

Cukierman, A. and A. H. Meltzer. "A Theory of Ambiguity, Credibility, and Inflation Under Discretion and Asymmetric Information." Econometrica 54 (1986): 1099-1128.

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

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

Dupasquier, C., A. Guay, and P. St-Amant. "A Survey of Alternative Methodologies for Estimating Potential Output and the Output Gap." Journal of Macroeconomics 21 (1999): 577-95.

Estrella, A. and F. Mishkin. "Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty." In Monetary Policy Rules: NBER Conference Report Series, J. B. Taylor, ed, 405-30. Chicago: University of Chicago Press. 1999.

Faust, J. and L. E. O. Svensson. "Transparency and Credibility: Monetary Policy with Unobservable Goals." International Economic Review 42 (2001): 369-97.

Fillion, J. F. and A. Leonard. "La courbe de Phillips au Canada: un examen de quelques hypotheses." Bank of Canada Working Paper: 97-3. 1997.

Isard, P. and D. Laxton. "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. 1998.

Johnson, D. R. "Expected Inflation in Canada 1988-1995: An Evaluation of Bank of Canada Credibility and the Effect of Inflation Targets." Canadian Public Policy 23 (1997): 233-58.

Johnson, D. R. "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. 1998.

Kendrick, D. "Caution and Probing in a Macroeconomic Model." Journal of Economic Dynamics and Control 4 (1982): 149-70.

Kichian, M. "Measuring Potential Output within a State-Space Framework." Bank of Canada Working Paper: 99-9. 1999.

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

Laxton, D., N. Ricketts, and D. Rose. "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. 1994.

Laxton, D. and R. Tetlow. "A Simple Multivariate Filter for the Measurement of Potential Output." Bank of Canada Technical Report: 59. 1992.

Perrier, P. "Un examen de la Credibilite de la Politique Monetaire au Canada." Bank of Canada Working Paper: 98-12. 1998.

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

Rudebusch, G.D. and L.E.O. Svensson. "Policy Rules for Inflation Targeting." In Monetary Policy Rules: NBER Conference Report Series, J. B. Taylor, ed, 203-46. Chicago: University of Chicago Press. 1999.

Sack, B. "Uncertainty, Learning, and Gradual Monetary Policy." Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, Discussion Paper: 98-34. 1998.

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

Srour, G. "Inflation Targeting under Uncertainty." Bank of Canada Technical Report: 85. 1999.

St-Amant, P. and S. van Norden. "Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada." Bank of Canada Technical Report: 79. 1997.

Stock, J. H. "Monetary Policy in a Changing Economy: Indicators, Rules, and the Shift Towards Intangible Output." Bank of Japan IMES Discussion Paper: 99-E-13. 1999.

Svensson, L. E. O. "Inflation Targeting: Some Extensions." Scandinavian Journal of Economics 101 (1999): 337-61.

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

Wieland, V. "Monetary Policy, Parameter Uncertainty and Optimal Learning." Journal of Monetary Economics 46 (2000): 199-228.

Wieland, V. "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. 1998.

Figure 1

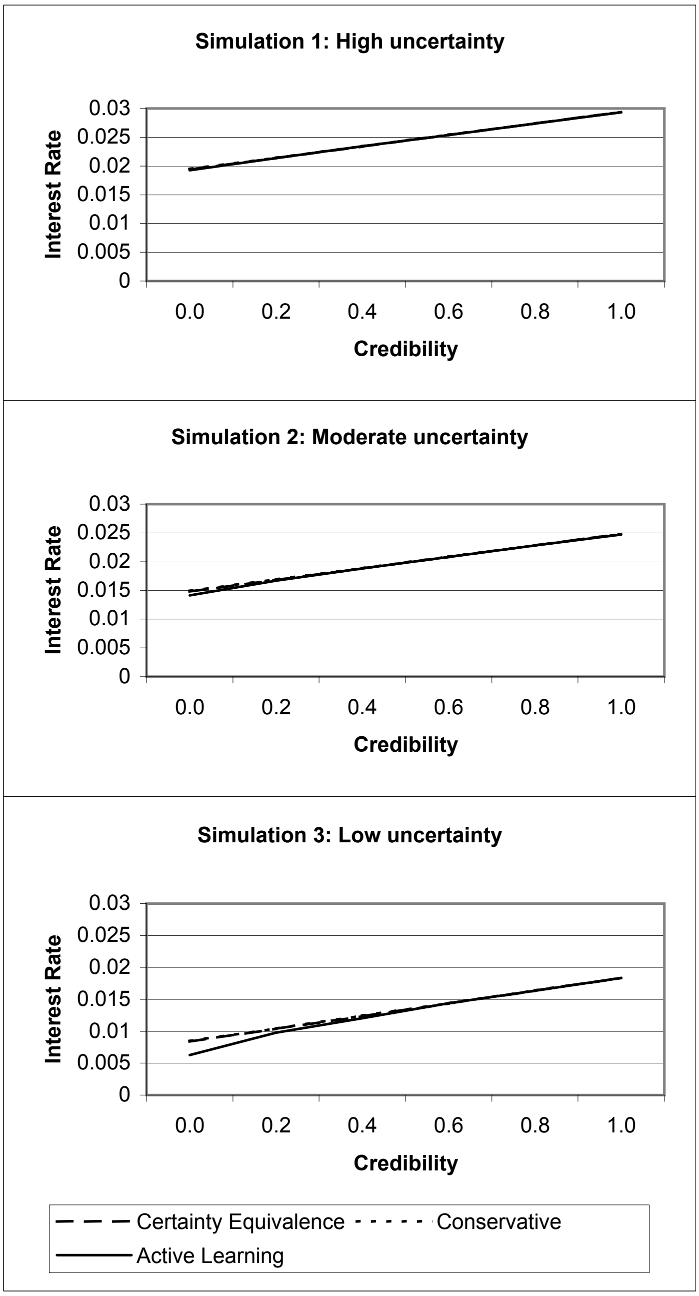

Simulation 1: High uncertainty
Interest Rate
0.03
0.025
0.02
0.015
0.01
0.005
0
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Simulation 2: Moderate uncertainty
Interest Rate
Credibility
Simulation 3: Low uncertainty
Interest Rate
Credibility
Certainty Equivalence
Conservative
Active Learning

Figure 2

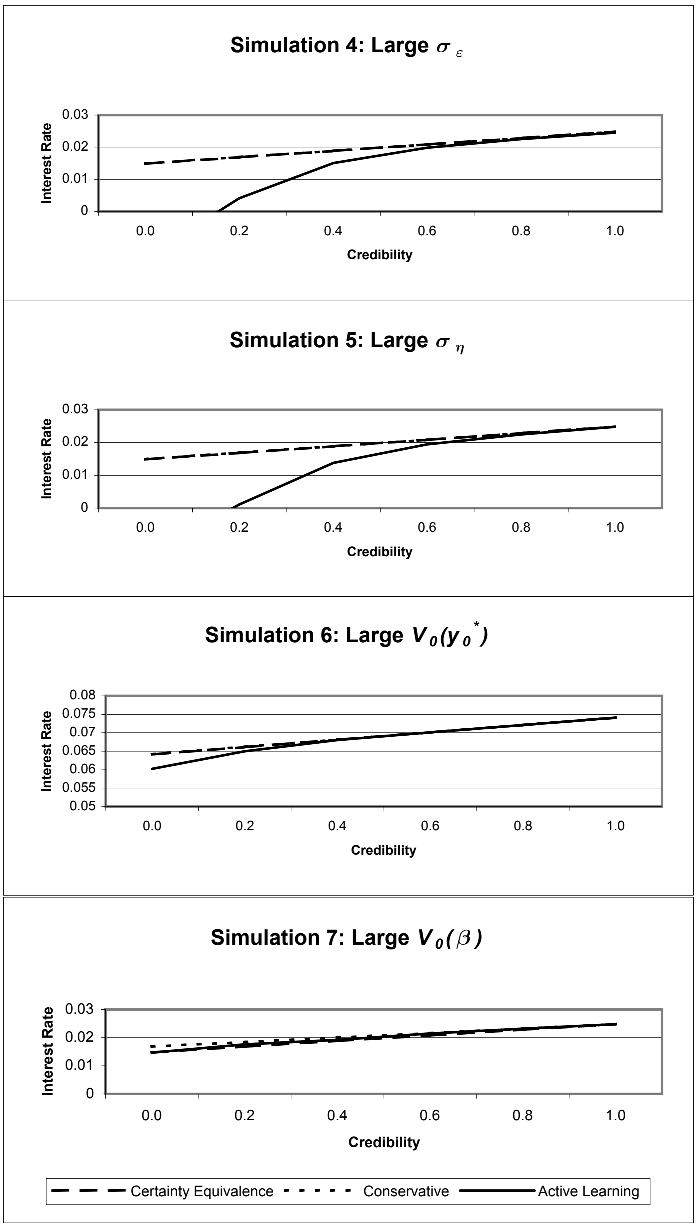
Simulation 4: Large σ_{ε}
Interest Rate
0.03
0.02
0.01
0
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Simulation 5: Large σ_{η}
Interest Rate
0.03
0.02
0.01
0
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Simulation 6: Large $V_0(y_0^*)$
Interest Rate
0.08
0.075
0.07
0.065
0.06
0.055
0.05
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Simulation 7: Large $V_0(\beta)$
Interest Rate
0.03
0.02
0.01
0
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Certainty Equivalence
Conservative
Active Learning

Figure 3

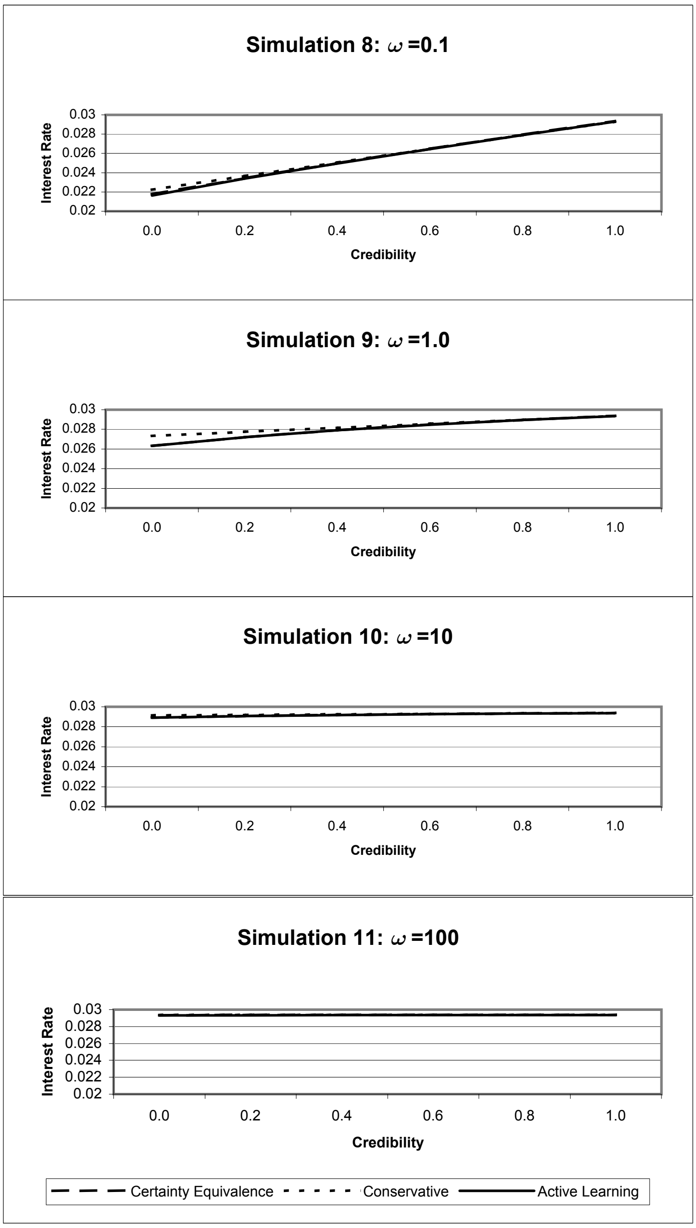

Simulation 8: ω =0.1
Simulation 9: ω =1.0
Simulation 10: ω =10
Simulation 11: ω =100
Interest Rate
0.03
0.028
0.026
0.024
0.022
0.02
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Certainty Equivalence
Conservative
Active Learning

Figure 4

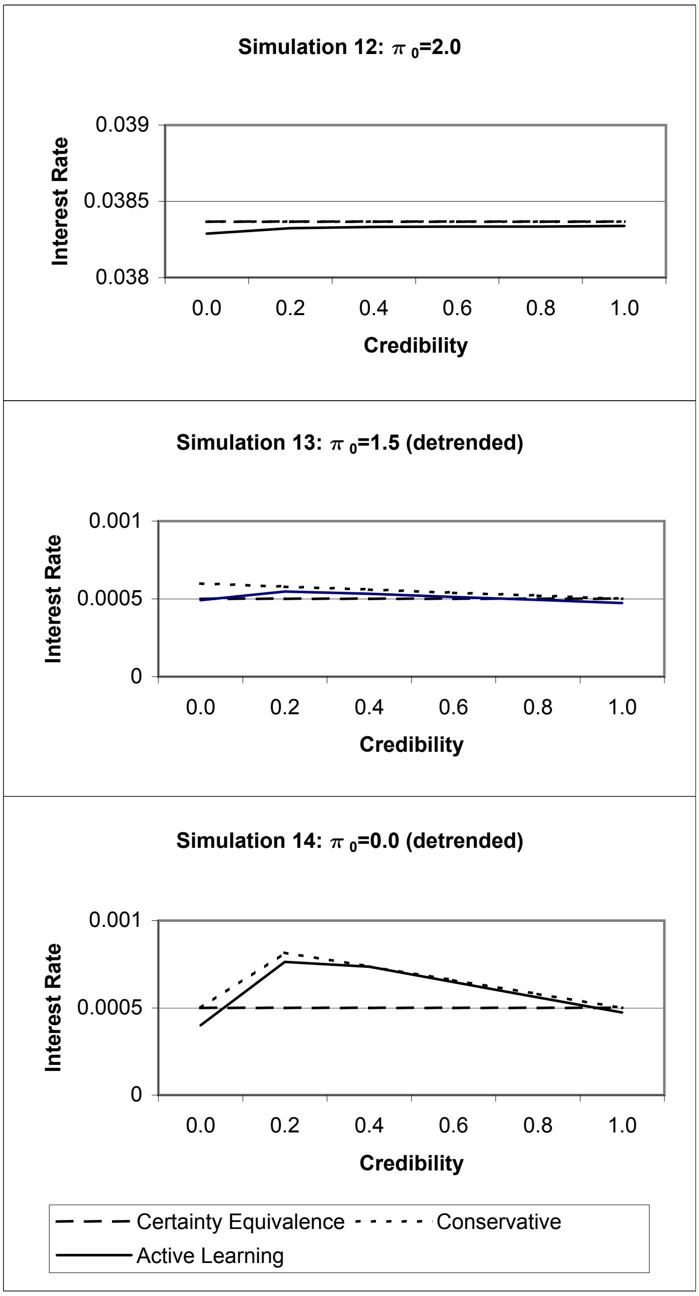
Simulation 12: π 0=2.0
Interest Rate
0.039
0.0385
0.038
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Simulation 13: π 0=1.5 (detrended)
Interest Rate
0.001
0.0005
0
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Simulation 14: π 0=0.0 (detrended)
Interest Rate
0.001
0.0005
0
0.0
0.2
0.4
0.6
0.8
1.0
Credibility
Certainty Equivalence
Conservative
Active Learning