

Monetary Policy and Uncertainty about the Natural Unemployment Rate

Volker Wieland*

Board of Governors of the Federal Reserve System

April 1998

Abstract

Recent empirical research concerning the relationship between inflation and unemployment, a relationship that is central to the design of monetary policy, has been characterized by an active debate about the precision of relevant parameter estimates such as the estimated natural unemployment rate. This paper studies the optimal monetary policy in the presence of uncertainty about the natural rate and the short-run inflation-unemployment tradeoff in a simple macroeconomic model. Two conflicting motives drive the optimal policy. In the static version of the model, uncertainty provides a motive for the policymaker to move more cautiously than she would if she knew the true parameters. In the dynamic version, uncertainty also motivates an element of experimentation in policy. I find that the optimal policy that balances the cautionary and activist motives typically exhibits gradualism, i.e. it is less aggressive than a policy that disregards parameter uncertainty. Exceptions occur when uncertainty is very high and inflation close to target.

KEYWORDS: monetary policy, gradualism, parameter uncertainty, NAIRU estimation, experimentation, optimal control with unknown parameters.

JEL Classification System: E52, E24, D8, C61

*I appreciate the excellent research assistance of Kendrew Witt. Helpful comments by John Leahy, Andrew Levin, Athanasios Orphanides, Richard Porter, Peter von zur Muehlen, Carl Walsh, David Wilcox and by seminar participants at the Federal Reserve Board, the AEA Meeting in New Orleans, January 1997 and the NBER Summer Institute in Cambridge, July 1997 are gratefully acknowledged. The opinions expressed here are those of the authors and do not necessarily reflect views of the Board of Governors of the Federal Reserve System.

Correspondence: Volker Wieland, Division of Monetary Affairs, Board of Governors of the Federal Reserve System, Washington, D.C. 20551. tel: (202) 736-5620, email: vwieland@frb.gov.

1 Introduction

The short-run relationship between unemployment and inflation has long been central to the design of monetary policy. In implementing policy based on this relationship, policymakers have to rely on empirical estimates of the natural unemployment rate (or NAIRU)¹ and the slope of the short-run inflation-unemployment tradeoff. Estimates of these parameters have changed over time and their precision is the subject of a continuing active debate.² Indeed, Staiger, Stock and Watson (1997a, 1997b) investigate a variety of empirical specifications and find that a typical 95% confidence interval for the natural rate in 1990 was about 2.5 percentage points wide. The width of this confidence interval is closely related to the standard error of the slope of the short-run Phillips curve—most clearly in a linear framework, where estimates of the natural rate are obtained from the ratio of intercept and slope.

This paper investigates how a policymaker who believes in the existence of a short-run tradeoff between inflation and unemployment should adjust policy given her uncertainty about actual parameter values. Using a simple macroeconomic model I derive the optimal policy in the presence of uncertainty about the natural unemployment rate and the short-run inflation-unemployment tradeoff. Two conflicting motives drive the optimal policy. In the static version of the model, Phillips curve uncertainty provides a motive for the policymaker to move more cautiously than she would if she knew all the parameter values. In the dynamic version with learning on the part of the policymaker, uncertainty also motivates an element of experimentation in policy.

Analysis of the motive for cautionary policy due to multiplicative parameter uncertainty goes back to Brainard (1967) and has been used to justify a gradualist approach to monetary policy. For example, Alan Blinder (1995, p.13), when he was vice-chairman of the Board of Governors, argued that “*a little stodginess at the central bank is entirely appropriate*”,

¹An acronym for non-accelerating inflation rate of unemployment.

²See Gordon (1996), Staiger, Stock, and Watson (1997a) and (1997b), Blanchard and Katz (1997), Fuhrer (1995), Akerlof, Dickens and Perry (1996), Eisner (1996), Phelps and Zoega (1997), Debelle and Laxton (1996) and others.

and proposed in his Marshall lectures that “*central banks should calculate the change in policy required to get it right and then do less*”.³ However, there are a number of reasons to believe that such a Brainard-type analysis overstates the case for gradualism. For example, Caplin and Leahy (1996) show that in a game between a policymaker who attempts to stimulate the economy and potential investors, a cautious policy move may be ineffectual, because investors anticipate lower interest rates in the future. Another reason, investigated in this paper, is that a more aggressive policy move may generate more information, which would improve the precision of future estimates and thereby future policy performance. Policymakers have noted this link between policy and learning. For example, Stiglitz (1997), when Chairman of the Council of Economic Advisers, recognized that “*a fuller discussion (of NAIRU uncertainty) would take into account factors such as costs of adjustment and of variability in output and unemployment, and dynamic learning effects*” and then asked the question: “*are there policies that can affect the degree of uncertainty about the value of the NAIRU or of policy tradeoffs?*”

This paper makes the following three contributions. First, the static analysis of the cautionary motive extends Brainard’s analysis of optimal policy under multiplicative parameter uncertainty by including uncertainty about the natural rate. Specifically, I show that the covariance between estimates of the intercept and slope of the Phillips curve importantly affects the optimal policy and provide an economic interpretation of this effect. This effect also implies a form of gradualism, because it leads policy to gradually move out of disinflationary (or reflationary) episodes even at the expense of temporarily under-shooting (or over-shooting) the inflation target.

Second, the paper provides numerical results concerning the optimal policy in a dynamic model with learning. The reason why analytical results concerning optimal policies under parameter uncertainty are largely absent from the literature is the nonlinear nature of the dynamic learning problem. Furthermore, numerical analysis by dynamic programming methods is complicated, because modelling the policymaker’s beliefs increases the dimen-

³See also Blinder (1998) for a discussion of this strategy.

sion of the state space and the nonlinear updating equations induce multiple optima and discontinuities in the optimal policy.⁴ The analysis in this paper builds on earlier work on simple controlled regressions in Wieland (1996a,1996b) extending that framework to allow for a lagged dependent variable and a time-varying intercept. I find that even when parameter uncertainty is relatively low, the optimal policy incorporates a quantitatively significant degree of experimentation as indicated by a more aggressive policy response than under the cautionary Brainard-type policy. However, the optimal policy typically remains less aggressive than a certainty-equivalent policy which completely disregards parameter uncertainty. Thus, in most cases the recommendation for gradualist policymaking under parameter uncertainty still applies in the dynamic model with learning. Only, when uncertainty is very high and inflation close to target, does the optimal policy imply a more aggressive response than a policy that disregards parameter uncertainty.

Third, the policy functions derived in this paper can be related to the type of interest rate reaction functions, which have been the focus of the recent literature on monetary policy rules.⁵ So far, this literature has studied the performance of monetary policy rules in different macroeconomic models under the assumption that all structural parameters are known with certainty. Uncertainty about the parameters of the Phillips curve has important implications for the response coefficients in these rules. In the presence of parameter uncertainty optimal monetary policy would imply time-varying response coefficients on inflation and resource gaps that are typically *smaller* than for a policy that disregards parameter uncertainty. In most situations, the optimal interest rate rule would respond gradually to inflation and unemployment deviations with the degree of gradualism varying with the degree of parameter uncertainty.

The remainder of this paper is organized as follows. The next section first reviews

⁴The current state of analytical and numerical approaches to optimal learning in the literature is reviewed in the next section.

⁵See for example Taylor (1993), Henderson and McKibbin (1993), Bryant, Hooper and Mann (1993), Brayton, Levin, Tryon and Williams (1996), Rotemberg and Woodford (1997), McCallum (1997), Clarida, Gali and Gertler (1997,1998), Orphanides, Small, Wieland and Wilcox (1997), Williams (1997) and many others.

NAIRU estimation in the linear framework that is widely used in the empirical literature and then also discusses the current state of the literature on uncertainty and learning on the part of the policymaker. Section 3 introduces a simple macroeconomic model that forms the basis for the subsequent analysis of optimal policy under uncertainty about the parameters of the Phillips curve. In section 4 the cautionary policy, which is myopic because it disregards dynamic learning effects, is derived analytically. Section 5 extends the model to incorporate learning and a tradeoff between current stabilization and information. A quantitative comparison of the optimal, cautionary and certainty-equivalent policies is provided in section 6. In section 7, the asymptotic properties of beliefs and policies are discussed. Section 8 concludes.

2 Literature review

2.1 NAIRU estimation in the empirical literature

The standard framework for estimating the natural unemployment rate is the expectations-augmented Phillips curve. This relationship is often estimated in the following linear form:⁶

$$\pi_t = \pi_t^e + B(L)(u_t^* - u_t) + CX_t + \epsilon_t \quad (1)$$

where π_t^e is the public's expectation of inflation and may be approximated by survey expectations, or simply by lags of the inflation rate. This regression equation usually includes several lags of the unemployment rate and a vector X_t that contains proxy variables for supply shocks and dummies for the Nixon wage and price controls. Estimates of the natural unemployment rate can be computed as functions of the estimated parameters. For example, an estimate of a constant natural rate $u_t^* = u^* \forall t$, can be obtained from the ratio of the estimated regression constant and the sum of the coefficients on current and lagged

⁶In the past, the linear framework has been found to describe the historical process of U.S. inflation remarkably well (eg. Gordon (1977) and (1982)). Recently, Gordon (1996) has presented new estimates of a time-varying natural rate, which has declined somewhat over recent years, but is more precisely estimated than in Staiger et al. Eisner(1996), Debelle and Laxton (1996), and Akerlof et al. (1996) present evidence that various (perhaps conflicting) nonlinear specifications of the Phillips curve may fit the U.S. inflation process better than a linear specification.

unemployment rates $B(\cdot)$.⁷ Because the NAIRU estimate is computed as the ratio of the intercept and the slope estimates, its precision is importantly related to the precision of estimates of the slope of the Phillips curve.⁸ Thus any specification of the policymaker's decision problem, which addresses the issue of uncertainty about the natural rate must also consider uncertainty about the slope of the Phillips curve, i.e. about the (sum of) coefficient(s) on the unemployment rate(s).

Figure 1 illustrates the degree of uncertainty associated with the relationship between inflation and unemployment. It graphs year-to-year changes in inflation $\Delta\pi_t$ against the preceding year's unemployment rate u_{t-1} (annual observations (1955-1995)). The motivation for plotting the change of inflation is that last year's inflation rate is taken as a proxy for the public's expectation of inflation in the current year. All other dynamics and lags are neglected. For illustration, the chart contains a linear Phillips curve with a slope of -0.5 and a natural unemployment rate of 6 percent. These values are within the range of values discussed in the empirical literature and are quite close to the estimates one would obtain from an ordinary least squares regression based on this data. I will use these values later on in the numerical analysis of optimal and cautionary policies.

This chart suggests that unemployment variations only explain a fraction of the observed changes in inflation. As shown in the empirical literature, the fit of the regression improves when one includes more lags of inflation and unemployment, measures of inflation expectations and additional variables that account for wage and price controls, supply shocks and import prices. Furthermore, in a framework with a time-varying natural rate, some of the variation that is apparent in the cloud of observations scattered around the

⁷An approximate measure of the variance of the estimated NAIRU can be calculated by the delta method, which involves taking a first-order Taylor series approximation to the nonlinear function and computing the variance of this approximation.

⁸Note that while the estimates of the intercept and slope have a multivariate normal distribution under the assumption of Gaussian errors, the ratio of the intercept and the sum of slope coefficients has a doubly non-central Cauchy distribution with dependent numerator and denominator for which means and variances do not exist. Such a distribution may be skewed and heavy-tailed. Staiger et al. (1997b) point out that when the slope is estimated imprecisely, normality as implied by the delta method can provide a poor approximation to the distribution of this ratio. They provide an alternative method to calculate confidence intervals which are exact under the assumption of exogenous regressor and normal errors.

linear relationship in **figure 1** would be accounted for by variations in the natural rate. Nevertheless, the debate about the value of the natural rate is far from settled. The main thrust of the recent empirical literature has been to consider alternative explanations for the observed variation in inflation, such as structural shifts in the natural rate (eg. Staiger et al. and Gordon) and nonlinearities in the short-run or long-run Phillips curve (eg. Eisner, Akerlof et al. and Debelle and Laxton).

2.2 The literature on uncertainty and learning by the policymaker

The tradeoff between current stabilization and exploration for the sake of better control in the future has been the focus of a theoretical literature on optimal learning as well as an engineering-related literature on dual control. The theoretical optimal learning literature has primarily studied the asymptotic properties of beliefs and actions.⁹ The dual control literature instead has focussed on numerically computing payoffs and optimal policies based on linear approximation.¹⁰

In terms of monetary policy, recent studies of the implications of the tradeoff between current control and learning have been conducted by Bertocci and Spagat (1993), Balvers and Cosimano (1994) and Wieland (1996b).¹¹ Among these, Wieland (1996b) considers the most general learning framework, namely a simple linear regression equation with two unknown parameters, and numerically computes optimal policies. The asymptotic properties of beliefs and policies in this framework have been studied by Easley and Kiefer (1988) and Kiefer and Nyarko (1989), who have shown that incomplete learning may occur.¹² Wieland

⁹See for example Taylor (1974), Lai and Wei (1982), Easley and Kiefer (1988), Kiefer and Nyarko (1989), Nyarko (1991) and Aghion et al. (1993). The applicability of these results to the framework considered here is discussed in more detail in section 7.

¹⁰See for example Tse and Bar-Shalom (1973), Kendrick (1981), Kendrick (1982), Mizrach (1991), Amman and Kendrick (1995), Tucci (1996)).

¹¹Bertocci and Spagat have argued that a fixed money supply rule is not optimal in the context of uncertainty about the parameters of a Lucas-type supply curve, because learning leads to adjustments in monetary policy. However, they impose stringent simplifying restrictions on the type of uncertainty in the model and do not derive the optimal learning policy. Balvers and Cosimano consider a more general model with two unknown time-varying parameters and show that even when learning is taken into account the optimal transition from high to low inflation is still a gradual one. However, they impose simplifying restrictions on the type of uncertainty and the form of the objective function, which imply that the policymaker only cares about the variance but not the expectation of inflation.

¹²Using this incomplete learning result Kasa (1996) shows that a central bank that starts out thinking

(1996b) has evaluated the speed of learning under alternative policies as well as the frequency with which a persistent bias in money growth and inflation may arise due to such self-reinforcing incorrect beliefs subsequent a structural change such as German unification.

The reason, why analytical results concerning optimal policies under parameter uncertainty are largely absent from the literature and numerical results are rare, is the nonlinear nature of the dynamic learning problem. Compared to the simple regression framework considered in previous work, the problem studied in this paper is further complicated by the existence of lagged dependent variables. The associated nonlinear stochastic dynamic optimization problem is attacked using an extended version of the numerical algorithm developed in Wieland (1995).¹³ It is described in more detail in the appendix.

As mentioned above, numerical approaches to solving this class of optimal learning problems have been studied extensively and for some time in engineering. Termed “dual control”, these methods have been further developed and applied to economic problems by Kendrick (1981) and (1982), Norman (1976), Mizrach (1991), Amman and Kendrick (1995) and others. There are several important differences between dual control and the dynamic programming algorithm used here: (i) while the dual control algorithm typically involves a first- or second-order linear approximation, the DP algorithm used here directly takes into account the nonlinearity of the updating equations which are at the center of the learning problem; (ii) while the dual control algorithm approximates ex-post payoffs for a given initial belief about the unknown parameters for alternative sequences of shocks as in Monte Carlo simulations, the DP algorithm used here approximates ex-ante payoffs and policies for a range of initial beliefs and any possible sequence of shocks.

Finally, it is important to note that the framework in this paper abstracts from the interaction between the public’s and the policymaker’s learning processes. By allowing the wage setters’ and the policymaker’s beliefs to depend on different information sets, it would

that there is a long-run tradeoff between inflation and unemployment, may be led to implement a policy that reinforces this belief even though the true Phillips curve is vertical.

¹³In this earlier work I have generalized the numerical DP approach used by Prescott (1972) for a one unknown parameter - one state variable problem to a problem with two unknown parameters and five state variables.

be possible to jointly address questions of credibility, time consistency and learning by the policymaker. While these questions have not yet been studied in a macroeconomic model with parameter uncertainty and rational learning such as the one considered in this paper, they have partly been addressed in two alternative frameworks. Caplin and Leahy (1996) study a game between a monetary policymaker who attempts to get the economy out of a recession and potential investors that are waiting for lower interest rates.¹⁴ Sargent (1998a) considers a framework with a policymaker and a public that are learning about the nature of the Phillips curve but are both 'boundedly rational'.¹⁵ Rather than learning optimally, both, the policymaker and the public, learn adaptively and use forecasting rules as in the least squares learning literature. There exist self-confirming equilibria that depend on the actual and perceived relationships between unemployment and inflation (not without some similarity to the self-reinforcing beliefs in the optimal learning literature).

3 A simple macroeconomic model

The starting point of the analysis is a standard linear Phillips curve:

$$\pi_t = \pi_t^e + \beta(u_t^* - u_t) + \epsilon_t \quad (2)$$

where π_t is the inflation rate, π_t^e is the inflation rate expected by the public (which may be different from the policymaker's expectation of inflation) and $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ are random shocks. The slope parameter β characterizes the short-run inflation-unemployment tradeoff.

The natural unemployment rate u_t^* is assumed to evolve according to:

$$u_t^* = u_{t-1}^* + \eta_t \quad (3)$$

where $\eta_t \sim N(0, \sigma_\eta^2)$. For $\sigma_\eta^2 > 0$, this specification implies that the natural rate follows a random walk. For $\sigma_\eta^2 = 0$, it implies a constant natural rate $u_t^* = u_0^* \forall t$. For simplicity, I

¹⁴The policymaker is uncertain about the investors' response to a specific value of the interest rate. She has to resolve the dilemma that an overly aggressive move may lead to overheating (too much investment) while an overly cautious move may be ineffectual, because investors anticipate lower interest rates in the future and delay investment.

¹⁵A related framework was studied by Sims (1988) and Chung (1990).

assume that the public forms expectations about inflation adaptively.¹⁶

$$\pi_t^e = \pi_{t-1} \quad (4)$$

The policymaker is endowed with a standard loss function which is quadratic in inflation and unemployment:¹⁷

$$L(\pi_t, u_t) = E_{t-1}[(\pi_t - \pi^*)^2 + \omega(u_t^* - u_t)^2] \quad (5)$$

where π^* is the target inflation rate.

Two additional equations are needed to close the model. First, according to “Okun’s Law” changes in the unemployment rate are related to deviations of real GDP growth Δy_t from its trend growth rate Δy_t^* .¹⁸

$$\Delta u_t = -\phi(\Delta y_t - \Delta y_t^*) \quad (6)$$

Second, monetary policy affects unemployment and inflation through the interest rate channel and aggregate demand. Deviations of the real interest rate r_t from its equilibrium value r^* are negatively related to the output gap:

$$y_t - y_t^* = \gamma(r_t - r^*) \quad (7)$$

where y_t and y_t^* denote the logarithms of actual and trend (potential) output. The policymaker is assumed to choose r_t based on her knowledge of the state of the economy and

¹⁶This assumption will simplify the subsequent numerical analysis of optimal policy under uncertainty. However, other than computational complexity there is no fundamental reason for excluding the possibility of rational expectations (or rational learning) on behalf of the price-setter. Alternatively, one could consider the “backward and forward-looking components” model of the Phillips curve (eg. Buiter and Miller (1985), Fuhrer and Moore (1995), and Clark, Goodhart and Huang (1996)):

$$\pi_t = \lambda\pi_{t-1} + (1 - \lambda)E_{t-1}\pi_t + \beta(u_t^* - u_t) + \epsilon_t$$

The backward-looking component reflects inertia in inflation due to overlapping wage contracts or adaptive expectations of a subset of price-setters and guarantees that monetary policy still has real effects. The public’s rational expectation of inflation $E_{t-1}[\pi_t]$ is a function of observable data as well as the policy rule and can be computed by the policymaker conditional upon her policy choice as long as they share the same information set.

¹⁷The quadratic specification is a natural starting point for analyzing optimal policy under uncertainty because of its prevalence in the policy literature. However, the numerical dynamic programming algorithm used here could be applied to non-quadratic specifications of the policymaker’s preferences. In this case higher-order moments may influence the optimal policy.

¹⁸For a textbook discussion of this empirical regularity see Dornbusch and Fischer (1990).

its structural parameters but before the shocks ϵ_t and η_t are realized.¹⁹ She can predict impending changes in inflation and unemployment based on the observations from the preceding period and respond appropriately.

As a first step, I consider a static version of this model, in which the policymaker minimizes the expected loss in the current period. If all parameters are known, including the natural rate u_{t-1}^* and the slope of the short-run Phillips curve β , the optimal policy is a linear feedback rule of the following form:

$$r_t = r_{t-1} + \frac{1}{\phi\gamma} \frac{\beta}{(\omega + \beta^2)} (\pi_{t-1} - \pi^*) + \frac{1}{\phi\gamma} (u_{t-1}^* - u_{t-1}) \quad (8)$$

This rule sets the current real interest rate based on past inflation and unemployment gaps. The optimal response depends on the structural parameters, the inflation target π^* and the weight ω which the policymaker assigns to unemployment deviations relative to inflation deviations. If inflation is above target, the policymaker raises the interest rate, if unemployment is above the NAIRU, she lowers the interest rate.

In solving her optimization problem, the policymaker balances the expected marginal loss from unemployment against the expected marginal loss from inflation. It is instructive to take a look at the associated first-order condition after substituting in the feedback rule (8):

$$E_{t-1}[u_t^* - u_t] = -\frac{\beta}{\omega + \beta^2} (\pi_{t-1} - \pi^*) = -\frac{\beta}{\omega} E_{t-1}[\pi_t - \pi^*] \quad (9)$$

Note that as long as the preceding period's inflation gap is zero, the optimal interest rate setting implies that both, the expected unemployment gap and the expected inflation gap for period t , are set to zero.

If the policymaker focuses exclusively on stabilizing inflation (i.e. $\omega = 0$), the optimal policy induces an expected inflation rate equal to the inflation target $E_{t-1}[\pi_t] = \pi^*$ and corresponds to a policy of “inflation forecast targeting” as defined by Svensson (1997a).

¹⁹Of course, in practice the central bank sets the nominal short-term interest rate rather than the real rate. However, the rigidity in inflation that is embodied in the short-run Phillips curve gives the policymaker short-run control over real interest rates.

The optimal setting of the nominal interest rate results in an unemployment gap that is expected to offset the inflation gap inherited from the preceding period.

4 Uncertainty as a motive for cautious policymaking

The preceding policy rule cannot be implemented directly if the natural unemployment rate u_t^* and the slope of the short-run inflation-unemployment tradeoff β are unknown. However the policymaker can obtain estimates from the following regression equation:

$$\pi_t - \pi_{t-1} = \alpha_t - \beta u_t + \epsilon_t \quad (10)$$

This is a simple version of the type of equation estimated in the empirical literature. The intercept is a function of both the unknown slope parameter and the unknown NAIRU, and may therefore vary over time.

$$\alpha_t = \beta u_t^* = \beta u_{t-1}^* + \beta \eta_t = \alpha_{t-1} + \nu_t \quad (11)$$

where $\nu_t \sim N(0, \sigma_\nu^2)$. I denote the means of the intercept and slope parameters based on information up to time t-1 as follows:

$$E_{t-1}\alpha_t = a_{t|t-1} = a_{t-1} \quad (12)$$

$$E_{t-1}\beta = b_{t-1}$$

Since $u_t^* = \frac{\alpha_t}{\beta}$, the empirical literature uses the ratio of the means as an estimator for the NAIRU.

$$\hat{u}_t^* = \frac{a_{t-1}}{b_{t-1}} \quad (13)$$

The degree of uncertainty about the unknown intercept and slope parameters based on information at t-1 is characterized by the following variance-covariance matrix:

$$\Sigma_{t|t-1} = \begin{pmatrix} v_{t-1}^a & v_{t-1}^{ab} \\ v_{t-1}^{ab} & v_{t-1}^b \end{pmatrix} = \begin{pmatrix} v_{t-1}^a + \sigma_\nu^2 & v_{t-1}^{ab} \\ v_{t-1}^{ab} & v_{t-1}^b \end{pmatrix} \quad (14)$$

The five variables $(a_{t-1}, b_{t-1}, \Sigma_{t|t-1})$ define a bivariate normal distribution which comprises all relevant information about the unknown parameters (α_t, β) at time t-1. This distri-

bution represents the policymaker's beliefs about the location and slope of the inflation-unemployment tradeoff.²⁰

Throughout the remainder of this paper I will focus exclusively on uncertainty about the relationship between inflation and output and maintain the assumption that all other parameters are known to the policymaker. The policy rule (8) can then be rendered implementable by replacing the actual values of u_{t-1}^* and β by the best available estimates $\frac{a_{t-1}}{b_{t-1}}$ and b_{t-1} .

$$r_t = r_{t-1} + \frac{1}{\phi\gamma} \frac{b_{t-1}}{(\omega + b_{t-1}^2)} (\pi_{t-1} - \pi^*) + \frac{1}{\phi\gamma} \left(\frac{a_{t-1}}{b_{t-1}} - u_{t-1} \right) \quad (15)$$

This ‘‘certainty-equivalent’’ policy rule will serve as a benchmark for comparison. An optimal rule in the presence of parameter uncertainty can be derived from the following one-period optimization problem:

$$\begin{aligned} \text{Min}_{r_t} \quad & E \left[(\pi_t - \pi^*)^2 + \omega (u_t^* - u_t)^2 \mid (\pi_{t-1}, u_{t-1}, a_{t-1}, b_{t-1}, \Sigma_{t-1}) \right] \\ \text{s.t.} \quad & \pi_t = \pi_{t-1} + \alpha_t - \beta u_t + \epsilon_t \\ & u_t = u_{t-1} + \phi\gamma (r_t - r_{t-1}) \end{aligned} \quad (16)$$

where Okun's law has already been substituted into the aggregate demand equation. The expected current period loss $L(\pi_t, u_t)$ is a function of the state variables at t-1, which comprise inflation, unemployment, the interest rate and the policymaker's beliefs about the unknown Phillips curve. It will be denoted by $L(\pi_{t-1}, u_{t-1}, r_{t-1}, a_{t-1}, b_{t-1}, \Sigma_{t-1})$ from here on.

The policy that minimizes expected one-period loss is best described as a cautionary policy, because it takes into account parameter uncertainty. Clearly, it is not necessarily optimal in a dynamic context. It is a ‘‘myopic’’ policy, because it disregards the effect

²⁰Note that for mathematical convenience, the variances of the normally distributed shocks $\sigma_\nu^2, \sigma_\epsilon^2$ are assumed to be known. This is a standard assumption in the optimal learning literature (see Easley and Kiefer (1988), Kiefer and Nyarko (1989)). It guarantees that given a normal prior, the posterior belief will also be a normal distribution.

of the current interest rate setting on future parameter estimates and policy performance. This cautionary policy takes the form of a feedback rule, which is linear in inflation and unemployment:

$$\begin{aligned}
r_t = & r_{t-1} + \frac{1}{\phi\gamma} \frac{b_{t-1}}{(b_{t-1}^2 + v_{t-1}^b + \omega)} (\pi_{t-1} - \pi^*) + \frac{1}{\phi\gamma} \left(\frac{a_{t-1}}{b_{t-1}} - u_{t-1} \right) \\
& + \frac{1}{\phi\gamma} \left(\frac{v_{t-1}^{ab} - v_{t-1}^b \frac{a_{t-1}}{b_{t-1}}}{b_{t-1}^2 + v_{t-1}^b + \omega} \right)
\end{aligned} \tag{17}$$

Neither the variance of the estimated intercept, v_{t-1}^a , nor the variance of the natural unemployment rate explicitly enter the feedback rule, although they affect the expected loss. This is not surprising. Since the model is linear and the loss function is quadratic, certainty-equivalence applies and additive uncertainty does not affect the optimal policy. However, the variance of the slope estimate v_{t-1}^b as well as the covariance v_{t-1}^{ab} of the intercept and slope enter the rule in two different places in a nonlinear way.

First, the variance v_{t-1}^b shows up in the denominator of the response coefficient on the inflation gap. Thus, as long as v^b is greater than zero, the cautionary policy rule recommends a more muted response to the preceding inflation gap than the certainty-equivalent rule (15). This muted policy response implies a gradualist approach to disinflation in the face of uncertainty about the slope of the Phillips curve. It mirrors the seminal result by Brainard (1967) concerning cautionary policy under multiplicative parameter uncertainty.²¹ As the slope estimate becomes more and more precise, policy will tend to respond to inflation more aggressively as prescribed by the certainty-equivalent rule.

Second, there is a new term in the feedback rule, which is a function of the variance v_t^b , the covariance v_{t-1}^{ab} and the point estimate of the natural rate $\frac{a_{t-1}}{b_{t-1}}$. This term may result in a negative or a positive adjustment to the interest rate. While the first result is intuitive, the second is somewhat puzzling. It implies that even in a situation where the observed inflation rate is on target and the unemployment rate equals the estimated

²¹Other papers that have looked at this effect recently are Svensson (1997b), Clarida, Gali and Gertler (1997) and Estrella and Mishkin (1998). Sack (1997) studies how parameter uncertainty can explain the high degree of serial correlation in interest rates.

natural rate, the central bank would pursue a policy that drives unemployment away from the estimated natural rate in expectation. One might have expected that the policymaker would abstain from any action, when unemployment is at the estimated natural rate and inflationary pressures are absent. In fact, this would still be the case if the covariance v_{t-1}^{ab} were to satisfy the following condition:

$$v_{t-1}^{ab} = v_{t-1}^b \frac{a_{t-1}}{b_{t-1}} = \hat{u}_t^* v_{t-1}^b \quad (18)$$

It may seem at first that nothing more can be said about this covariance effect. Fortunately however, one can make use of the fact that the covariance and variance of the intercept and slope estimates are functions of the historical data on inflation and unemployment. For example, if the natural rate is assumed to be constant, $u_t^* = u_0^* \forall t$, the estimates a_t and b_t of the parameters α and β can be obtained by ordinary least squares. The covariance of these least squares estimates is simply the product of the variance of the slope and the sample mean of the unemployment rate:

$$v_{t-1}^{ab} = \bar{u}_{t-1} v_{t-1}^b \quad (19)$$

where $\bar{u}_{t-1} = \sum_{i=0}^{t-1} u_i$ denotes the sample mean of the historical unemployment path.²² Whenever \bar{u}_{t-1} equals the estimated natural rate \hat{u}_{t-1}^* then (18) holds and the cautionary policy rule will keep unemployment as close as possible to the natural rate in the absence of inflationary pressures. The difference between \hat{u}_{t-1}^* and \bar{u}_{t-1} is a function of the sample mean of the change in inflation, which is equal to the trend in inflation over the data period:

$$\frac{a_{t-1}}{b_{t-1}} - \bar{u}_{t-1} = \frac{1}{b_{t-1}} \overline{\Delta \pi}_{t-1} = \frac{1}{b_{t-1}} \frac{\pi_{t-1} - \pi_0}{t-1} \quad (20)$$

This implies that the optimal approach to controlling inflation under uncertainty is to gradually ease out of inflationary or disinflationary episodes rather than attempt to end them abruptly. For example, having observed an upward trend in inflation, the policymaker

²²For a derivation see Greene (1993) pp 155-157. In the case of a time-varying natural rate, the covariance and variances of the parameter estimates would be related to an appropriately weighted average of the historical unemployment rate rather than the simple average. The Kalman filter generates the appropriate weighting.

concludes that in the past the unemployment rate has on average been below the natural rate. Rather than pushing unemployment aggressively above the natural rate to reverse the trend and bring inflation back to target as quickly as possible, policy will be leaning towards the historical mean of unemployment and respond more gradually.

To provide some economic intuition for this result, I turn to the case where the policymaker is exclusively concerned with inflation stabilization ($\omega = 0$). The expected loss can then be decomposed into two terms, the conditional expectation of inflation deviations from target squared and the conditional variance of inflation:

$$L(\pi_t, u_t) = (E_{t-1}[(\pi_{t-1} - \pi^*) + \beta(u_t^* - u_t)])^2 + Var_{t-1}[\pi_t - \pi^*] \quad (21)$$

If the preceding period's inflation gap is zero, the first term—the squared conditional expectation—will be minimized by a policy that drives the unemployment rate to the estimated natural rate $u_t = \frac{a_t}{b_t}$. However this is not true for the second term—the conditional variance of inflation:

$$Var_{t-1}[\pi_t - \pi^*] = v_{t-1}^a + v_{t-1}^b u_t^2 - 2v_{t-1}^{ab} u_t + \sigma_\epsilon^2 \quad (22)$$

The conditional variance is minimized for $u_t = \frac{v_{t-1}^{ab}}{v_{t-1}^b}$, which is equal to the sample mean and may well differ from the estimated natural rate. Thus, the covariance effect in the cautionary policy rule arises because the policymaker balances the loss from a deviation of the conditional expectation of inflation from target against an increase in the conditional variance of inflation. This tradeoff is illustrated graphically in **figure 2**. The solid line represents the estimated Phillips curve based on point estimates of the intercept and slope. It intersects the horizontal axis at the estimated natural unemployment rate \hat{u}^* and also goes through the point defined by the sample means $\overline{\Delta\pi}$ and \bar{u} when estimated by ordinary least squares. Since the parameter estimates are imprecise, the true Phillips curve may be steeper or flatter than the estimated one; and the inflation rate that results from a specific value of the unemployment rate is uncertain. The dotted lines in figure 2 are one standard deviation bands of inflation conditional on unemployment. As can be seen,

uncertainty about inflation is minimized when the unemployment rate equals the sample mean. Consequently, the unemployment rate which minimizes the loss function must lie between the sample mean and the natural rate whenever the preceding period's inflation gap is zero. Its value can easily be derived from the first-order condition:

$$E_{t-1}[u_t^* - u_t] = -\frac{b_{t-1}}{b_{t-1}^2 + v_{t-1}^b}(\pi_{t-1} - \pi^*) - \frac{(v_{t-1}^{ab} - v_{t-1}^b \frac{a_{t-1}}{b_{t-1}})}{(b_{t-1}^2 + v_{t-1}^b)} \quad (23)$$

If $(\pi_{t-1} - \pi^*)$ is zero, the unemployment rate that is consistent with the optimal policy in the static model is a weighted average of the sample mean and the estimated natural rate:

$$u_t(r_t) = \frac{b_{t-1}^2}{b_{t-1}^2 + v_{t-1}^b} \hat{u}_{t-1}^* + \frac{v_{t-1}^b}{b_{t-1}^2 + v_{t-1}^b} \bar{u}_{t-1} \quad (24)$$

Finally, it can be shown that the cautionary policy rule (17) formalizes the recipe for gradualist policymaking provided by Blinder (1995). For example, suppose that inflation exceeds the target rate by x percentage points, the unemployment rate equals the estimated natural rate and the sample mean of unemployment is z percentage points below the estimated natural rate (which would be consistent with a past upward trend in inflation). The change in policy required to “get it right” given the parameter estimates and the policymaker's objective function would be to raise the interest rate by $\frac{b_{t-1}}{(\phi\gamma)(\omega + b_{t-1}^2)}x$. However, under uncertainty about the parameters of the Phillips curve the policymaker will “do less”. How much less? As shown above, there are two adjustments, the first is a function of the inflation gap x and the second is a function of the difference between the sample mean of unemployment and the natural rate z . The first adjustment reduces the prescribed interest rate change by $x \frac{b_{t-1}v_{t-1}^b}{(\phi\gamma)(\omega + b_{t-1}^2)(\omega + b_{t-1}^2 + v_{t-1}^b)}$. The second adjustment further reduces the interest rate prescription by an amount equal to $z \frac{v_{t-1}^b}{(\phi\gamma)(\omega + b_{t-1}^2 + v_{t-1}^b)}$.

As a consequence of these adjustments, inflation will decline more gradually towards the target rate. This is best illustrated by computing the expected disinflation path when the inflation rate is above target. An example is shown in **figure 3**. The associated initial conditions for inflation and unemployment are $(x = 1, z = 0.4, \bar{u} = 5.6, \hat{u} = 6, u_{t-1} = \hat{u})$, the prior belief is $(a = 3, b = 0.5, v^a = 5.22, v^b = 0.16, v^{ab} = 0.896, \bar{u} = 5.6, \hat{u} = 6)$ and

the known parameters are set to $(\phi = 0.5, \gamma = 1, \omega = 0, \sigma_\nu^2 = 0, \sigma_\epsilon^2 = 1)$. Thus, initially inflation is one percentage point above target, unemployment equals the estimated natural rate, uncertainty about the Phillips curve is very high and both, the effect due to slope uncertainty and the covariance effect, come into play. The covariance effect is related to the difference between the sample mean of unemployment and the estimated natural rate, z , which equals 0.4 percentage points and is consistent with a five-year upward trend in inflation of one percentage point.

The three panels in the top row of **figure 3** show the expected paths of inflation, unemployment and interest rates under the certainty-equivalent and cautionary policy respectively. The certainty-equivalent policy raises real interest rates and unemployment by 4 and 2 percentage points respectively, so that inflation is expected to return to target within a year. The cautionary policy responds more slowly and initially raises interest rates and unemployment by less than the certainty-equivalent rule. In terms of the interest rate, the total downward adjustment in the first period is 1.9 percentage points with 1.6 percentage points due to the slope uncertainty effect and 30 basis points due to the covariance effect. As a consequence, inflation, unemployment and interest rates are expected to return more gradually towards their steady-state values over the course of the next four years. Note that under different initial conditions, specifically when the initial sample mean of unemployment is very low, unemployment and the interest rate may first increase gradually over a few periods and then again decrease gradually.

The dynamic simulation of the expected disinflation path in **figure 3** takes into account that the policymaker will be able to reestimate the Phillips curve every time new inflation and unemployment observations become available. This raises the question whether the policymaker should expect her beliefs about the unknown parameters to change based on the expected path of inflation and unemployment. It turns out that the point estimates of the intercept and slope are martingales; and the policymaker would not expect that they will change. This property of the point estimates, which has important implications for their asymptotic behavior, will be discussed in more detail in section 7. The variances and

covariances however, are deterministic functions of the sum of squared deviations of the unemployment rate from its sample mean, and the policymaker would expect that they change along the expected path for unemployment. This is reflected in the bottom row of panels in **figure 3**. All three elements of the covariance matrix are expected to decline monotonically during the disinflation.²³ The temporary increase in the unemployment rate generates information and is expected to raise the precision of the parameter estimates. The more aggressive the initial increase in interest rates and unemployment, the more rapid will be the reduction in uncertainty. While the certainty-equivalent policy implies a higher conditional variance of inflation than the cautionary policy in the current period, it may well lead to a lower variance of inflation in future periods due to the rapid reduction in parameter uncertainty. A forward-looking policymaker may want to take this dynamic learning effect into account and adjust policy in an optimal manner. It remains to be seen whether the recommendation for gradualist monetary policy survives in a dynamic setting with learning.

5 The optimal balance between caution and experimentation

As new observations on inflation and unemployment become available the policymaker may reestimate the regression and update her estimates of the unknown parameters. The relevant updating equations for the policymaker's beliefs $(a_{t-1}, b_{t-1}, \Sigma_{t|t-1})$ can be cast in form of the Kalman filter:

$$\begin{aligned} \begin{pmatrix} a_{t|t} \\ b_{t|t} \end{pmatrix} &= \begin{pmatrix} a_{t|t-1} \\ b_{t-1} \end{pmatrix} + \Sigma_{t|t-1} \begin{pmatrix} 1 \\ u_t \end{pmatrix} F^{-1} (\pi_t - \pi_{t-1} - a_{t|t-1} + b_{t-1} u_t) \\ \Sigma_{t|t} &= \Sigma_{t|t-1} - \Sigma_{t|t-1} \begin{pmatrix} 1 \\ u_t \end{pmatrix} F^{-1} \begin{pmatrix} 1 & u_t \end{pmatrix} \Sigma_{t|t-1} \end{aligned} \quad (25)$$

where $F = \begin{pmatrix} 1 & u_t \end{pmatrix} \Sigma_{t|t-1} \begin{pmatrix} 1 \\ u_t \end{pmatrix} + \sigma_\epsilon^2$

Under the assumption that the error terms are normally distributed with known variances σ_ϵ^2 and σ_ν^2 , (25) is equivalent to Bayesian updating of the bivariate normal distribution that

²³If the NAIRU is assumed to be time-varying, then the variance of the intercept would be expected to be renewed every period by σ_η^2 .

represents the policymaker's belief about α_t and β . Under the assumption of a constant natural unemployment rate, these updating equations are also equivalent to recursive least squares.²⁴ The current choice of the interest rate will affect the precision of the point estimates as well as the point estimates themselves through its impact on unemployment and inflation in (25). By choosing the interest rate appropriately, the policymaker can raise the precision of parameter estimates and improve future performance, albeit at the expense of higher current variability of inflation and unemployment. The optimal policy in this model solves the following dynamic optimization problem:

$$\begin{aligned} \text{Min}_{\{r_t\}_{t=0}^{\infty}} \quad & E \left[\sum_{t=0}^{\infty} \delta^t \left((\pi_t - \pi^*)^2 + \omega(u_t^* - u_t)^2 \right) \mid (\pi_0, u_0, a_0, b_0, \Sigma_0) \right] \\ \text{s.t.} \quad & \pi_t = \pi_{t-1} + \alpha_t - \beta u_t + \epsilon_t \\ & u_t = u_{t-1} + \phi \gamma (r_t - r_{t-1}) \\ \text{and s.t.} \quad & (25) \end{aligned} \tag{26}$$

This is a dynamic discrete-time stochastic control problem, which can be rewritten as a dynamic program. A nonstandard feature of this dynamic problem is that decisions affect the expectations operator itself. However, one can still use a standard contraction mapping argument as in Kiefer and Nyarko (1989) to show that a unique value function exists, which solves the dynamic program and corresponds to the infimum of the sum of expected current and discounted future losses in (26). To simplify notation, I will use the vector θ_{t-1} to denote the policymaker's beliefs about the unknown parameters at time t-1, $\theta_{t-1} = (a_{t-1}, b_{t-1}, \Sigma_{t-1})$. The value function is then denoted as $V(\pi, u, \theta)$ and the associated Bellman equation is:

²⁴This dynamic learning model extends earlier analysis of optimal learning by Wieland (1996a) and (1996b) in a simple regression framework in two ways. It allows for a time-varying intercept and, more importantly, it includes a lagged dependent variable, which raises the number of state variables and increases computational complexity.

$$\begin{aligned}
V(\pi_{t-1}, u_{t-1}, r_{t-1}, \theta_{t-1}) &= \underset{r_t}{\text{Min}} && L(\pi_{t-1}, u_{t-1}, r_{t-1}, \theta_{t-1}) \\
&&& + \delta \int V(\pi_t, u_t, r_t, \theta_t) f(\pi_t | \pi_{t-1}, u_{t-1}, \theta_{t-1}) d\pi \\
&&& (27) \\
&= \underset{r_t}{\text{Min}} && L(\pi_{t-1}, u_{t-1}, r_{t-1}, \theta_{t-1}) \\
&&& + \delta \int V(\alpha_t, \beta, \epsilon_t, r_t, \pi_{t-1}, u_{t-1}, r_{t-1}, \theta_{t-1}) \\
&&& p(\alpha_t, \beta | \pi_{t-1}, \theta_{t-1}) q(\epsilon) d\alpha d\beta d\epsilon
\end{aligned}$$

First focus on the upper equation in (27) and note that (π, u, r, θ) are state variables. The two terms on the right-hand side of this equation characterize the tradeoff between current control and estimation. $L(\cdot)$ is the expected current loss, while the second term denotes next period's expected loss and is multiplied by the discount factor δ . This term incorporates the value of information. Note that θ_t , the vector of beliefs at time t , is stochastic and can only be calculated once time t unemployment and inflation observations become available. $f(\pi_t | \cdot)$ is the corresponding predictive distribution of inflation.

In the lower equation in (27), time t values of inflation, unemployment and beliefs have been substituted out using the respective transition equations, including (25). They are functions of the previous period's inflation rate, unemployment and beliefs, and also of the unknown parameters and random shock ϵ_t . Expectations are taken with respect to the unknown parameters and the random shock. $p(\alpha_t, \beta | \cdot)$ is the bivariate normal distribution that describes the policymaker's beliefs about α_t and β . $q(\epsilon)$ refers to the normal density function of the price shocks in the Phillips curve.

Associated with this functional equation is a stationary optimal policy. This policy function maps the state variables into a value for the real interest rate:

$$r_t = r_{t-1} - \frac{1}{\phi\gamma} u_{t-1} + \frac{1}{\phi\gamma} H(\pi_{t-1} - \pi^*, \theta_{t-1}) \quad (28)$$

It is the dynamically optimal counterpart of the certainty-equivalent and cautionary feedback rules (15), (17) that were derived analytically in the preceding section. As long as

the aggregate demand parameters are considered known to the decision maker, the current unemployment rate is a deterministic function of the lag of unemployment as well as the lag and the current value of the interest rate. Thus, determining the optimal response of the current interest rate to the lagged unemployment rate is straightforward. The difficulty lies in determining the optimal response of the interest rate to lagged inflation and the policymaker’s beliefs as denoted by the nonlinear function $H(\cdot)$. Unfortunately analytical solutions for $H(\cdot)$ are not available due to the nonlinear nature of the dynamic decision problem. However, one can use numerical dynamic programming methods to approximate the value function and the optimal policy.

The functional equation (27) defines a contraction mapping with a unique fixed point, which is the value function. Starting from an initial guess of the value function, one can obtain successively better approximations by repeatedly solving the optimization problem on the right-hand side of (27). As is well known, this iterative method can be implemented numerically. Its application is hampered by the “curse of dimensionality” which implies that the number of necessary computations increases geometrically with the number of state variables. The numerical algorithm used here combines such value function iterations with policy iterations to speed up convergence. Nevertheless, this optimal learning problem, which has at least six state variables,²⁵ is at the borderline of what is computationally feasible even with very powerful workstations.

6 Numerical Results

In the following I investigate the properties of the optimal policy rule. Two questions arise naturally:

²⁵Note that the lag of the unemployment rate, and the lag of the interest rate are not crucial for approximating the optimal decision rule. Since the coefficients of the aggregate demand equation are treated as known, the policymaker can always adjust the interest rate setting by the appropriate amount for any observed value of the lagged unemployment rate and lagged interest rate. Effectively, one can treat the current unemployment rate as the control variable in this problem and does not need to consider the lagged unemployment and interest rates in computing the optimal policy.

1. To what extent does the optimal policy incorporate experimentation, or in other words, to what extent does it differ from the cautionary policy rule?
2. Does the optimal policy exhibit gradualism, or in other words, does it imply a less activist policy response than the certainty-equivalent rule?

Answers to these questions can be provided based on numerical approximations of the $H(\cdot)$ mapping in (28). I have obtained such approximations under the simplifying assumptions that the natural unemployment rate is constant, $\sigma_{\eta}^2 = 0$, and the policymaker focuses exclusively on stabilizing inflation, $\omega = 0$. I consider two different values for the discount factor, $\delta = 0.75$ and $\delta = 0.95$, and set the variance of price shocks to $\sigma_{\epsilon}^2 = 1.0$.

The dynamic programming algorithm provides numerical approximations for the value and policy functions over a wide range of the state space. However, to keep the number of charts in this section manageable I restrict attention to a comparison of different policies at a few points in the state space. As initial estimates of the intercept and slope of the Phillips curve I consider $a_{t-1} = 3.0$ and $b_{t-1} = 0.5$. The implied estimate of the natural unemployment rate is 6%. These values correspond to the line drawn in **figure 1** and are reasonable values given the different estimates in the recent empirical literature.²⁶

As to the degree of uncertainty, I will first consider four scenarios with very low, low, moderate and high uncertainty. The four sets of variances would imply t-statistics in the Phillips curve regressions of about 2.5 (very low uncertainty), 2 (low uncertainty), 1.875 (moderate uncertainty) and 1.5 (high uncertainty). In all four scenarios, the respective covariances are chosen to be equal to the product of the variance of the slope and the estimated natural rate as in (18). As established earlier, under these conditions of uncertainty, the covariance effect discussed in section 4 does not come into play. Also, the comparison of the response coefficients in (15) and (17) shows that under such conditions of uncertainty the cautionary policy rule implies a more muted response to the observed inflation gap than the certainty-equivalent rule.

²⁶Results for alternative values of the slope and natural rate can easily be provided upon request. While the qualitative nature of these results is the same, there are of course quantitative differences.

Figure 4 compares the policy response to the inflation gap under the optimal policy with the cautionary and certainty-equivalent policies. The four panels correspond to the four different degrees of uncertainty. The horizontal axis in each panel measures the size of the observed inflation gap. The vertical axis measures the unemployment rate which occurs in the current period as a consequence of the interest rate set by the policymaker in response to the observed inflation gap. For illustrative purposes, these panels focus on the policy response when inflation is above target and show unemployment rates equal or above the estimated natural rate of 6%. The relationship between unemployment and inflation when inflation is below target is symmetric. In each panel the solid line is associated with the certainty-equivalent policy, the dashed line with the cautionary policy and the dashed-dotted line with the optimal policy. In case of the certainty-equivalent and cautionary policies the relationship between unemployment and inflation shown in the panels corresponds to the first-order conditions (9) and (23). The slope of the lines is given by the respective coefficient on the inflation gap in (9) and (23), which is a function of the reaction coefficients in the certainty-equivalent and cautionary policy rules. The slope under the certainty-equivalent policy is 2 in all four panels, since the policy reaction coefficient is not affected by uncertainty. The relation between unemployment and inflation under the cautionary policy is not as steep because the reaction coefficient for the cautionary rule is smaller. The four panels illustrate that the cautionary policy reacts less and less aggressively as uncertainty increases, opening up an increasingly wider wedge between the certainty-equivalent and the cautionary policies. The cautionary policy at first only responds partially to an increase in inflation and results in gradualist approach to disinflation.

Figure 4 provides several new results concerning the optimal policy response (dashed-dotted line). First, as one would expect, the optimal policy turns out to be more aggressive than the cautionary policy due to the incentive for experimentation. In all four panels, the unemployment rate that results from the optimal policy response to a rise in inflation lies above the unemployment rate resulting from the cautionary policy. The most interesting point to notice in **figure 4** is that the optimal policy typically lies inside the wedge made up

of the cautionary and certainty-equivalent policies. To the extent that the optimal policy implies active exploration for the purpose of gathering information, it is more aggressive than the cautionary policy but typically remains less aggressive than the certainty-equivalent policy. In other words, the policy recommendation of gradualism under parameter uncertainty is still appropriate.

However, there are exceptions as can be seen in the lower right panel. When uncertainty is high and inflation relatively close to target, the gains from experimentation are large enough to justify a more aggressive increase in unemployment than would be required to keep the expected inflation rate on target. Even though current inflation is close to target, the policymaker realizes that shocks are likely to drive the inflation rate away from the target in the future and future stabilization policy will be more effective under reduced uncertainty. Note that this relationship is symmetric and that when inflation is below target the optimal policy will lower the unemployment rate more aggressively than the cautionary policy. Finally, in all four panels the relative importance of experimentation declines as the observed inflation gap increases. If inflation is substantially above target then even the cautionary policy will result in a substantial increase in unemployment. Such an increase is expected to generate quite a bit of information about the inflation-unemployment tradeoff and the location of the natural rate. Consequently, in relative terms there is less incentive to perturb unemployment further for the purpose of improving the quality of the parameter estimates.

Clearly, the extent of optimal experimentation depends on how forward-looking the policymaker is. The results shown in **figure 4** are for a discount factor $\delta = 0.95$, which implies a reasonable discount rate for an annual model. **Figure 5** provides a comparison with a more short-term oriented policymaker with a discount factor of 0.75. The two top panels apply to the low uncertainty scenario and the two bottom panels to the high uncertainty scenario. In both cases the optimal policy incorporates a lesser degree of experimentation with a smaller discount factor. The optimal policy remains well inside the wedge created by the cautionary and certainty-equivalent policies.

As noted above, the recommendation of gradualism fails to apply when uncertainty is high and inflation is close to target. Under extreme uncertainty the optimal policy may even exhibit a discontinuity at an inflation gap of zero as shown in **figure 6**. In this case, policy allows unemployment to fall below (or rise above) the estimated natural rate, even though current inflation is close to target, solely for the purpose of reducing uncertainty and improving policy performance in the future. The policymaker knows that in a stochastic economy there will be a need for stabilization policy in the future and desires more precise knowledge about the nature of the inflation-unemployment tradeoff and the location of the natural rate.

So far I have focussed on comparing the different policies in states where the covariance satisfies (18) and does not directly affect the optimal policy. As shown earlier, in a world with a constant natural rate, this state is equivalent to a situation in which the estimated natural rate is equal to the sample mean of the unemployment rate. Of course, this does not need to be the case in practice. In fact an observed up-drift (decline) in inflation may well have been the result of a period during which unemployment remained persistently below (above) the natural rate. This means the sample mean of unemployment would have been below (above) the estimated natural rate. The cautionary policy derived in the static analysis in section 3 responds to this situation by keeping unemployment on average somewhat lower (higher) than it otherwise would. This is apparent in **figure 7**, which shows two scenarios with low and very low uncertainty under alternative values of the covariance. The lower covariances in the two right-hand side panels imply that the sample mean of unemployment is 1.5 percentage points below the estimated natural rate of 6%. The covariance effect is clearly apparent in the cautionary policy (dashed line) which has shifted downwards compared to the left-hand side panels. The optimal policy again lies well in between the cautionary and certainty-equivalent policies. When inflation is close to target the covariance effect almost completely disappears under the optimal policy. However, when inflation is further away from target, the optimal policy is again closer to the cautionary policy.

7 Convergence of Beliefs and Policies

A remaining question is whether the policymaker will eventually learn the true parameter values as more and more data becomes available. This question has been the focus of a theoretical literature on optimal learning in a controlled regression framework (eg. Easley and Kiefer (1988) and Kiefer and Nyarko (1989)). The learning problem considered here differs from the regression framework studied in that literature in two important ways: (i) the intercept may be time-varying, and (ii) the regression includes a lagged dependent variable.

In a framework where the natural unemployment rate is time-varying, the need for learning and adjusting policy in response to changes in parameter estimates persists through time. A policymaker who considers that the NAIRU may change, will always attach a positive variance to her beliefs about the unknown intercept and adjust policy accordingly. Uncertainty about the intercept is renewed in every period and the policymaker will never learn the true natural rate because it will keep changing.

In the case of a constant natural rate, one can bring some of the convergence results obtained by Kiefer and Nyarko (1989) (KN) to bear on this problem. In this case, the parameters of the accelerationist Phillips curve can be estimated by means of a simple regression

$$\Delta\pi_t = \alpha - \beta u_t + \epsilon_t \tag{29}$$

with the change rather than the level of inflation as dependent variable. The parameter estimates and covariance are updated according to (25), which corresponds to Bayesian updating of bivariate normal beliefs. KN provide a general convergence result that applies to this class of regression equations. They show that under general assumptions concerning the form of beliefs and the shock process, the process of posterior beliefs always converges with probability 1 (Theorem 4.1., p. 577). However, the limiting belief may or may not be centered on the true values. The proof of this theorem relies on an application of the martingale convergence theorem. It is straightforward to confirm that the point estimates

a_t and b_t in (25) follow a martingale relative to the decisionmaker's information. Since $E_{t-1}[\Delta\pi_t - a_{t-1} + b_{t-1}u_t] = 0$,²⁷ it follows that $E_{t-1}[a_t] = a_{t-1}$ and $E_{t-1}[b_t] = b_{t-1}$. This property of the point estimates was previously alluded to in the discussion of the expected disinflation shown in figure 3. The policymaker does not expect her estimate of the natural unemployment rate or the slope of the Phillips curve to change during the disinflation.

Whether the process of posterior beliefs converges to the truth or not, depends on the behavior of the series of unemployment rates u_t . KN provide two results that hold for simple regressions. First, if u_t does not converge, then the process of posterior beliefs converges to the point mass on the true parameter values (Theorem 4.2., p. 577). Second, if u_t does converge to a limit value, then the posterior beliefs may converge to a limit belief that does not coincide with the true parameter values. This introduces the possibility of incomplete learning. At a minimum however the decision maker learns the mean of the dependent variable that corresponds to the limiting value of u_t (Theorem 4.3., p. 578). KN then characterize the set of possible (including incorrect) limit beliefs and policies. However, without solving for the optimal policy, KN cannot determine the frequency with which incomplete learning may occur. This question has been addressed in Wieland (1996a) and (1996b).²⁸

With a constant natural rate the model considered in this paper generates complete learning of the unknown parameters under all of the three policy feedback rules. Because the centerpiece of the model—the accelerationist Phillips curve—contains a unit root, any policy that attempts to permanently lower (raise) the unemployment rate below (above) the natural rate, would imply that the rate of inflation goes towards $+(-)$ infinity. Furthermore, a policy that stabilizes unemployment exactly at its natural rate, would render the inflation

²⁷This is true because u_t is a deterministic function of u_{t-1}, r_{t-1} and the policy instrument r_t and is effectively part of the policymaker's information set at $t-1$.

²⁸Wieland (1996a) using numerical methods has characterized the value function and optimal policy for controlling a simple regression with two unknown parameters as in KN. Optimal experimentation was found to be most pronounced in the neighborhood of potentially self-reinforcing incorrect beliefs. Wieland (1996b) has shown that a myopic, passive-learning policy in a model with unknown money demand may frequently be uninformative and induce a long-lasting bias in the setting of the policy instrument that would not emerge under the optimal policy.

process a random walk. Inflation only remains under control if the policymaker pursues an active stabilization policy that responds to past values of the inflation rate. Using this property of the model, one can appeal to theorems 4.1. and 4.2. in KN to prove that complete learning will occur. First, theorem 4.1. implies that the process of posterior beliefs (a_t, b_t, Σ_t) about the unknown parameters α and β in (29) converges with probability one to a limit belief $(a_\infty, b_\infty, \Sigma_\infty)$ as $t \rightarrow \infty$. For any given belief, the unemployment that obtains under the cautionary policy, is a function of the means, the variance of the slope, the covariance, and the preceding period's inflation rate. For example, in any time period t the unemployment rate associated with a given belief $(a_\infty, b_\infty, \Sigma_\infty)$ would be:

$$u_t = \frac{a_\infty}{b_\infty} + \frac{b_\infty}{b_\infty^2 + v_\infty^b} (\alpha + \beta u_{t-1} + \epsilon_{t-1} - \pi^*) + \frac{(v_\infty^{ab} - v_\infty^b \frac{a_\infty}{b_\infty})}{(b_\infty^2 + v_\infty^b)} \quad (30)$$

Because policy responds to the preceding period's inflation rate, the unemployment rate effectively is a function of the preceding period's price shock ϵ_{t-1} . In each time period, a new ϵ shock is realized. Thus, even if the policymaker's beliefs were to remain constant, the unemployment rate would keep changing over time. Consequently, u_t does not converge; and, according to theorem 4.2 in KN, the process of posterior beliefs converges to the point mass on the true parameter values, $(a_t, b_t, v_t^a, v_t^b, v_t^{ab}) \rightarrow (\alpha, \beta, 0, 0, 0)$ with probability 1 as $t \rightarrow \infty$.

To build further intuition concerning the asymptotic properties of beliefs and policies it is useful to consider the relationship between posterior beliefs and the sequence of unemployment rates more directly. The elements of the covariance matrix are related to the sequence $\{u_i\}_{i=0}^t$ as follows:

$$\begin{aligned} v_t^b &= \frac{\sigma_\epsilon^2}{\sum_{i=0}^t (u_i - \bar{u}_t)^2} \\ v_t^{ab} &= \bar{u}_t v_t^b \\ v_t^a &= \frac{\sigma_\epsilon^2}{t} + \bar{u}_t^2 v_t^b \end{aligned} \quad (31)$$

where \bar{u}_t is the sample average. Clearly, whether the covariance matrix converges to the zero matrix as $t \rightarrow \infty$, will depend on the behavior of the sum of squared deviations of

unemployment from its sample mean $\sum_{i=0}^t (u_i - \bar{u}_t)^2$. This is a non-decreasing series and as t increases it may either go towards infinity or towards a positive number K . If unemployment varies sufficiently so that $\sum_{i=0}^t (u_i - \bar{u}_t)^2 \rightarrow \infty$ as $t \rightarrow \infty$, then $\Sigma_t \rightarrow 0$. Then also the point estimates $(a_t, b_t) \rightarrow (\alpha, \beta)$ as a consequence of the martingale convergence theorem. Alternatively, if the sequence of unemployment rates u_t were to settle down to a fixed value fairly soon, $\sum_{i=0}^t (u_i - \bar{u}_t)^2 \rightarrow K$, then the deviation between u_t and its sample mean would go towards zero and uncertainty about the parameter estimates would remain even in the limit. As argued above, this case will not arise here, because under the hypothesis of an accelerationist Phillips curve, controlling inflation requires an active stabilization policy and thus continuing variations in unemployment.

8 Conclusion

The preceding analysis provides several new results concerning optimal monetary policy when there is uncertainty about the parameters of the Phillips curve. I consider both, a static and a dynamic versions of a simple macroeconomic model, which incorporates a short-run inflation-unemployment tradeoff and a time-varying natural unemployment rate. In the static model, parameter uncertainty induces a cautionary policy stance that responds gradually to inflation deviations due to the uncertainty about the slope of the Phillips curve. Because of the linear-quadratic nature of the baseline model, certainty-equivalence applies; and additive uncertainty—such as about the intercept of the Phillips curve and implicitly the natural rate—does not affect the optimal policy rule. However, the covariance of the intercept and slope estimate (and implicitly the covariance of the natural rate and the slope estimate) has an important effect on the optimal policy in the static model. This covariance effect also implies a form of gradualism, because it recommends a policy that gradually moves out of disinflationary (or reflationary) episodes even at the expense of temporarily under-shooting (or over-shooting) the inflation target.²⁹

²⁹A note of caution is in order when it comes to the relationship between caution and gradualism. In the framework with parameter uncertainty considered here, caution leads to gradualism because it reduces

In the dynamic model with learning, the policymaker recognizes that there is a tradeoff between current control and exploration for the sake of reducing uncertainty and improving future stabilization policy. She then needs to strike the optimal balance between caution and experimentation. The optimal policy is found to incorporate a quantitatively significant degree of experimentation, but typically remains less aggressive than a policy that disregards uncertainty. Thus, under a moderate degree of uncertainty the recommendation of gradualist policymaking survives in the dynamic model with learning. Only when uncertainty is very high and inflation relatively close to target, does optimal learning require a policy that responds more aggressively to inflation than the certainty-equivalent policy which disregards parameter uncertainty. In this situation optimal policy significantly increases current variability in both, unemployment and inflation.

Of course, any normative conclusions that may be drawn from these results depend on the maintained hypothesis that a short-run inflation-unemployment tradeoff and a natural unemployment rate exist. However these results are also useful in a positive sense for explaining the gradualist policy conducted by policymakers who believe that such a tradeoff exists but are uncertain about the relevant parameter values.

In terms of future research it would be helpful to allow for nonlinearities in the Phillips curve, for a non-quadratic loss function, and for interactions between the policymaker's and the public's learning processes. First, when the Phillips curve is nonlinear, then uncertainty about the natural rate has first-order effects.³⁰ For example, if the short-run Phillips curve is convex, then increased variability of the unemployment gap will imply a higher average rate of unemployment. The associated optimal policy will likely be asymmetric. Similarly,

the contribution of parameter uncertainty to the conditional variance of inflation. Caution however may also mean guarding against worst-case alternatives, which may require aggressive rather than gradualist policymaking. An earlier discussion of such cases is by von zur Muehlen (1982). More recently, Sargent (1998b), using a robust control framework with Knightian uncertainty, shows how caution can lead to more aggressive policy when this caution is exercised relative to a worst-case serial-correlation pattern of the shock process.

³⁰Clark, Laxton and Rose (1996) and DeBelle and Laxton (1996) provide some evidence for convexity. Bean (1996) studies the first-order implications of uncertainty for the case of such a convex short-run inflation-unemployment relationship. Eisner (1996) and Stiglitz (1996) argue that this relationship may be locally concave.

with non-quadratic preferences, higher-order moments will matter to the policymaker and affect the optimal policy rule. Finally, allowing the wage setters' and the policymaker's learning processes to depend on different information sets would make it possible to jointly address questions of credibility, time consistency and learning by the policymaker.

References

- Akerlof, G., Dickens, W., and G. Perry (1996), "The Macroeconomics of Low Inflation", *Brookings Papers on Economic Activity*, 27:1.
- Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991), "Optimal Learning by Experimentation," *Review of Economic Studies*, Vol. 58, pp. 621-654.
- Amman, H. and D. Kendrick (1995), "Nonconvexities in Stochastic Control Models," *International Economic Review*, Vol. 36, 2, pp. 455-475.
- Anderson, T., and J.B. Taylor (1976), "Some Experimental Results on the Statistical Properties of Least Squares Estimates in Control Problems" *Econometrica*, Vol. 44 (.) 1289-1302.
- Ball, L. (1997), "Disinflation and the NAIRU", in Romer, Christina and David Romer, eds., *Reducing Inflation: Motivation and Strategy*. 1997.
- Balvers, R. and T. Cosimano (1994), "Inflation Variability and Gradualist Monetary Policy," *Review of Economic Studies*, Vol. 61, pp. 721-738.
- Bean, C. (1996), "The Convex Phillips Curve and Macroeconomic Policymaking under Uncertainty," Manuscript, London School of Economics.
- Bertocci, G. and M. Spagat (1993), "Learning, Experimentation and Monetary Policy", *Journal of Monetary Economics*, Vol. 32, No.1, August, pp. 169-178.
- Blanchard, O. and L. Katz (1997), "What We Know and Do Not Know About the Natural Rate of Unemployment", *Journal of Economic Perspectives*, 11 (1), 51-72.
- Blinder, A. (1995), "Central Banking in Theory and Practice, Lecture I: Targets, Instruments, and Stabilization", Marshall Lecture presented at the University of Cambridge.
- Blinder, A. (1998), "Central Banking in Theory and Practice", MIT Press, January 1998.
- Brainard, W. (1967), "Uncertainty and the Effectiveness of Policy", *American Economic Review*, 57, 1967, pp. 411-25.
- Brayton, F., Levin A., Tryon R. and J.C. Williams, (1996), The Evolution of Macro Models at the Federal Reserve Board, presented at the Carnegie- Rochester Conference, November 1996.
- Buiter, W. and M. Miller (1985), "Costs and Benefits of an Anti- Inflationary Policy: Questions and Issues," in V. Argy and J. Neville, eds., *Inflation and Unemployment: Theory, Experience, and Policy-Making*.
- Caplin, A. and J. Leahy (1996), "Monetary Policy as a Process of Search," *American Economic Review*, 86 (4), 689-702.
- Caplin, A. and J. Leahy (1993), "Sectoral Shocks, Learning, and Aggregate Fluctuations," *Review of Economic Studies*, 60, 777-794.
- Chung, H. (1990), Did Policy Makers Really Believe in the Phillips Curve?: An Econometric Test, unpublished Ph.D. dissertation, University of Minnesota.
- Clarida, R., Gali, J. and M. Gertler (1997), "The Science of Monetary Policy", Manuscript, New York University, September.

- Clarida, R., Gali, J. and M. Gertler (1998), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory", NBER Working Paper, No. 6442, March.
- Clark, P., Goodhart, C. and H. Huang (1996), "Optimal Monetary Policy Rules in a Rational Expectations Model of the Phillips Curve", Discussion Paper No. 247, LSE Financial Markets Group.
- Clark, P., Laxton, D. and D. Rose (1996), "Asymmetry in the U.S. Output- Inflation Nexus", *IMF Staff Papers*, Vol. 43, pp. 216-51.
- Craine, Roger (1979), "Optimal Monetary Policy with Uncertainty", *Journal of Economic Dynamics Control*, Vol. 1, pp. 59-83.
- Debelle, G. and D. Laxton, (1996), "Is the Phillips Curve Really a Curve ? Some Evidence for Canada, the United Kingdom, and the United States", IMF Working Paper, WP/96/111.
- Easley, D. and N.M. Kiefer (1988), "Controlling a Stochastic Process with Unknown Parameters", *Econometrica* 56, 1045-1064.
- Eisner, R., (1995), "A New View of the NAIRU", Manuscript, Northwestern University, July 1996.
- Estrella, A., and F. Mishkin (1998), "The Role of NAIRU in Monetary Policy: Implications of Uncertainty and Model Selection", Manuscript, Federal Reserve Bank of New York, Research Department.
- Fuhrer, J., (1995), "The Phillips Curve is Alive and Well," *New England Economic Review*, (March/April), pp. 41-56.
- Fuhrer, J. and G. Moore (1995), "Monetary Policy Trade-offs and the Correlation Between Nominal Interest Rates and Real Output", *American Economic Review*, 85: 219-239.
- Gordon, R., (1996), "The Time-Varying NAIRU and its Implications for Economic Policy", NBER Working Paper, No. 5735, August 1996.
- Gordon, R., (1977), "Can the Inflation of the 1970s Be Explained?", *Brookings Papers on Economic Activity*, 8:1, pp 253-77.
- Gordon, R., (1982), "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment", in Baily, M.N. ed., *Workes, Jobs and Inflation* Washington: Brookings Institution, 1982, pp 88-152.
- Greene, W. (1993), *Econometric Analysis*, 2nd edition, Macmillan.
- Henderson, D. and W.J. McKibbin, (1993), "A comparison of some basic monetary policy regimes for open economies: implications of different degrees of instrument adjustment and wage persistence" *Carnegie-Rochester Conference Series on Public Policy*, 39, pp. 221-318.
- Henderson, D. and S. Turnovsky, (1972), "Optimal Macroeconomic Policy Adjustment under conditions of risk" *Journal of Economic Theory*, 4, pp. 58-72 .
- Horvath, B. (1991), "Are Policy Variables Exogenous? The Econometric Implications of Learning while Maximizing," *Lecture Notes in Economics and Mathematical Systems*, Vol. 364, Springer Verlag.

- Kasa, K. (1996), Will the Fed Ever Learn?, Manuscript, Federal Reserve Bank of San Francisco, Research Department.
- King, R. and M. Watson (1994), "The post-war U.S. Phillips curve: a revisionist econometric history," *Carnegie-Rochester Conference Series on Public Policy*, 41, 157-219.
- Kalchbrenner, J. and P. Tinsley (1977), On Filtering Auxiliary Information in Short-run Monetary Policy, in 'Optimal Policies, Control Theory and Technology Exports', *Carnegie-Rochester Conference Series on Public Policy*, Vol. 7, 1977.
- Kendrick, D. (1981), *Stochastic Control For Economic Models*, Economic Handbook Series, McGraw Hill, New York.
- Kendrick, D. (1982), "Caution and Probing in a Macroeconomic Model," *Journal of Economic Dynamics and Control*, 15:2, pp. 149-170
- Kiefer, N., and Y. Nyarko (1989), "Optimal Control of an Unknown Linear Process with Learning," *International Economic Review*, 30, pp. 571-586.
- McCallum, B., "Issues in the Design of Monetary Policy Rules," NBER Working Paper No 6016, April.
- Mizrach, B. (1991), Non-Convexities in a Stochastic Control Problem with Learning, *Journal of Economic Dynamics and Control*, Vol. 15, pp. 515-538.
- Orphanides, A., D. Small, V. Wieland and D. Wilcox (1997), "A Quantitative Exploration of The Opportunistic Approach to Disinflation," Finance and Economic Discussion Series, 97-36, Board of Governors of the Federal Reserve System, May 1997.
- Phelps, E.S, and G. Zoega (1997), "The Rise and Falling Trend of the Natural Rate," *American Economic Review*, 87 (2), 283-289.
- Prescott, E., (1972), "The Multi-Period Control Problem under Uncertainty", *Econometrica*, Vol.40, No.6, November 1972.
- Rotemberg, J. and M. Woodford (1997), "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," NBER Macroeconomics Annual 1997.
- Sack, B. (1997), "Uncertainty and Gradual Monetary Policy," Manuscript, Division of Monetary Affairs, Federal Reserve Board.
- Sargent, T. (1998a), "The Conquest of American Inflation," Manuscript, Hoover Institution, February 1998.
- Sargent, T. (1998b), "Discussion of 'Policy Rules for Open Economies' by Laurence Ball", forthcoming in, Taylor, J. (ed.), "Monetary Policy Rules", National Bureau of Economic Research.
- Sims, C. A., (1988), "Projecting Policy Effects with Statistical Models", *Revista de Analisis Economico*, 3, 3-20.
- Staiger, D., J. Stock, and M. Watson, (1997a), "How Precise are Estimates of the Natural Rate of Unemployment?," in Romer, Christina and David Romer, eds., *Reducing Inflation: Motivation and Strategy*. 1997.
- Staiger, D., J. Stock, and M. Watson, (1997b), "The NAIRU, Unemployment and Monetary Policy," *Journal of Economic Perspectives*, 11 (1), 33-49.

- Stiglitz, J. (1997), "Reflections on the Natural Rate Hypothesis," *Journal of Economic Perspectives*, 11 (1), 33-49.
- Svensson, L. (1997a), "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *European Economic Review* 41 (6), 1111-1146.
- Svensson, L. (1997b), "Inflation Targeting: Some Extensions," NBER Working Paper No. 5962, March.
- Taylor, J.B., (1974), "Asymptotic Properties of Multi-period Control Rules in the Linear Regression Model", *International Economic Review*, 15, pp. 472-484.
- Taylor, John. B., (1993), "Discretion versus Policy Rules in Practice", *Carnegie-Rochester Conference Series on Public Policy*, 39, pp. 195-214.
- von zur Muehlen, P., (1982), "Activist vs. Non-Activist Monetary Policy: Optimal Policy Rules under Extreme Uncertainty, Manuscript, Federal Reserve Board, Washington, DC, April 1982.
- Wieland, V. (1996a), "Learning by Doing and the Value of Optimal Experimentation, Finance and Economics Discussion Series," 96-5, Federal Reserve Board, Washington, DC.
- Wieland, V. (1996b), "Monetary Policy, Parameter Uncertainty and Optimal Learning," Manuscript, Division of Monetary Affairs, Federal Reserve Board, August 1996.
- Wieland, V. (1995), "Optimal Control with Unknown Parameters: A Study of Optimal Learning Strategies with an Application to Monetary Policy," Ph.D. Dissertation, July 1995.
- Williams, J.C. (1997), "Simple Rules for Monetary Policy", Manuscript, Division of Research and Statistics, Federal Reserve Board.

Appendix: The Numerical Dynamic Programming Algorithm

The algorithm used in this paper computes the value function and stationary optimal policy by iterating over the functional equation of this dynamic programming problem. It takes advantage of the contraction mapping defined by the following functional operator T

$$TW = \underset{u}{Min} \left[L(\pi, u, \theta) + \delta \int W(\pi', u, \theta') f(\pi' | \pi, u, \theta) d\pi' \right] \quad (32)$$

where π and θ are last period's values of the inflation rate and the policymaker's beliefs, which are the state variables of this DP problem, u is the current unemployment rate and effectively the control variable. π' is the inflation rate to be realized subsequent the policy action and θ' refers to the policymaker's beliefs and the end of the period based on the new inflation and unemployment observations. The relevant updating equations for these state variables are shown in (25) and (26). $f(\pi' | \pi, u, \theta)$ is the predictive distribution of the inflation rate and is a normal distribution, because both the error terms and the policymaker's beliefs are normal distributions. $L(\cdot)$ denotes the expected current loss while $W(\cdot)$ refers to a continuous function defined on the state space.

Successive application of the operator T will generate a sequence of functions W_n that will converge to the value function V , if T is a contraction mapping. Note that the space of continuous bounded functions is a complete and separable metric space in the sup metric defined:

$$\rho(W_n, W_{n+1}) = \underset{(\theta, \pi)}{Sup} |W_n(\theta, \pi) - W_{n+1}(\theta, \pi)| \quad (33)$$

Standard arguments can be used to show that Blackwell's sufficiency conditions are satisfied and T is a contraction mapping in the space of continuous and bounded functions (see for example Kiefer and Nyarko (1989)) such that:

$$\rho(TW_{n+1}, TW_n) \leq \delta \rho(W_{n+1}, W_n) \quad (34)$$

Thus, T has a unique fixed point V , which is the value function and a stationary optimal policy $H(\pi, \theta)$ exists. This optimal policy corresponds to the set of u 's which minimize the right-hand side of (32) based on the current state (π, θ) .

V can be computed by value iteration, meaning successive application of the operator T , since $T_n W \rightarrow V$ uniformly for any continuous bounded function W . A convenient starting value W_0 is the single period loss function $L(\cdot)$ or alternatively a constant. If $W_{n+1} = TW_n$, then $\rho(W_{n+1}, W_n) \leq \rho(W_n, W_{n-1})$ and after iterating $\rho(W_{n+1+i}, W_{n+i}) \leq \delta^{1+i} \rho(W_n, W_{n-1})$. This implies an upper bound on the error in approximating V by W_n :

$$\rho(V, W^n) \leq \sum \rho(W^{n+1+i}, W^{n+i}) \leq \frac{\delta}{1-\delta} \rho(W^n, W^{n-1}) \quad (35)$$

This upper bound can easily be calculated since it only depends on the discount factor and the distance between the approximations obtained from the last and the preceding iteration. The time needed for convergence within a maximal error bound can be reduced significantly by introducing policy iterations in between every value iteration. A policy iteration implies the application of the following operator:

$$T^P W_n = L(\pi, H_n(\pi, \theta), \theta) + \delta \int W(\pi', H_n(\pi, \theta), \theta') f(\pi' | \pi, H_n(\pi, \theta), \theta) d\pi' \quad (36)$$

where $H_n(\pi, \theta)$ is the approximation of the policy function obtained from the preceding value iteration n .

The computational algorithm then proceeds as follows: first, compute starting values W_0 for a grid of points in the state space (π, θ) and save them in a table; secondly, calculate W_1 by applying the operator T to W_0 and update said table. This second step requires calculating the minimum in u for each of the grid values of the state variables (π, θ) . For this purpose next period's expected value is calculated by evaluating the following integral:

$$\int W_0(\pi', u, \theta') f(\pi' | \pi, u, \theta) d\pi' \quad (37)$$

The functions $W(\cdot)$ and the updating equations to obtain π' and θ' are known functions and the conditional density of π' is normal. Thus the integral can be calculated using Gaussian quadrature and values of W_0 from the table, where $W(\cdot)$ is evaluated in between grid points by linear interpolation.

Given an approximation for this integral the minimization problem on the right-hand side of the functional equation can be solved by standard numerical optimization procedures. However the search for the minimum turns out to be difficult because there may exist multiple local minima. As a consequence there may be kinks in the value function and discontinuities in the optimal policy. Therefore I use a slow but secure optimization procedure such as golden section search supplemented by a rough initial grid search. For each value of (θ, π) , the minimum in u gives the value of $W_1(\cdot)$ used to update the table. The maximum of $|W_1(\theta, \pi) - W_0(\theta, \pi)|$ is used to calculate the upper bound of the approximation error. Finally, the whole procedure is repeated to obtain W_2 and so on until the difference between two successive approximations is sufficiently small ($< 0.5\%$).

Computation Costs

The numerical dynamic programming problems dealt with in this paper require an immense computational effort largely because of the so-called curse of dimensionality. Since each of the six state variables is approximated based on N grid points, the integration and optimization procedures described above have to be carried out N^6 times to complete one value iteration. The optimization step is especially time-consuming because of the existence of multiple local optima.

Two steps have been taken to reduce computation time: (i) the introduction of policy iterations, which reduce the number of value iterations needed for convergence, and thus the number of times that the optimization procedure has to be executed; (ii) as shown below, I use a trick to reduce the number of state variables by one, which means that the integration and optimization steps only have to be carried out N^5 times per value iteration.

The policy functions presented in the paper have been approximated for a grid of 10 to 20 grid points for each of the five remaining state variables. Convergence as defined by a 0.5% maximal difference between the two final approximations was achieved after almost 2 days on a SPARC 20 (multi-processor) work station. Depending on the discount factor such an approximation required between 5 and 8 value iterations and a declining number of policy iterations (50 or less) in between every value iteration.

Transforming the DP Problem to Reduce the Number of State Variables

Any simplification which makes it possible to reduce the number of relevant state variables substantially reduces the computation time required for convergence of the DP algorithm. As shown here, one can ignore one state variable when dealing with a simple loss function such as $(\pi - \pi^*)^2$. In this case the value function $V(\pi, a, b, v^a, v^b, v^{ab})$ is given by the supremum of

$$E \left[\sum_{i=t}^{\infty} \delta^i (\pi_i - \pi^*)^2 \mid a_t, b_t, \Sigma_t \right] \quad \text{where } \pi_i = \pi_{i-1} + \alpha + \beta u_i + \epsilon_i \quad (38)$$

Proposition: The value function $V(\pi_t, a_t, b_t, v_t^a, v_t^b, v_t^{ab})$ has the following property (for any k):

$$V(\pi_t, a_t, kb_t, v_t^a, k^2 v_t^b, kv_t^{ab}) = V(\pi_t, a_t, b_t, v_t^a, v_t^b, v_t^{ab}) \quad (39)$$

Proof: Consider the transformed problem

$$\pi_t = \pi_{t-1} + \alpha + \check{\beta} \check{u}_t + \epsilon_t \quad \text{where } \check{\beta} = k\beta \quad \text{and} \quad \check{u} = \frac{u}{k} \quad (40)$$

If the prior on (α, β) is the bivariate normal distribution $N(a_t, b_t, v_t^a, v_t^b, v_t^{ab})$, then the prior on $(\check{\alpha}, \check{\beta})$ is $N(a_t, kb_t, v_t^a, k^2 v_t^b, kv_t^{ab})$. Since these two problems are equivalent, the proposition follows.

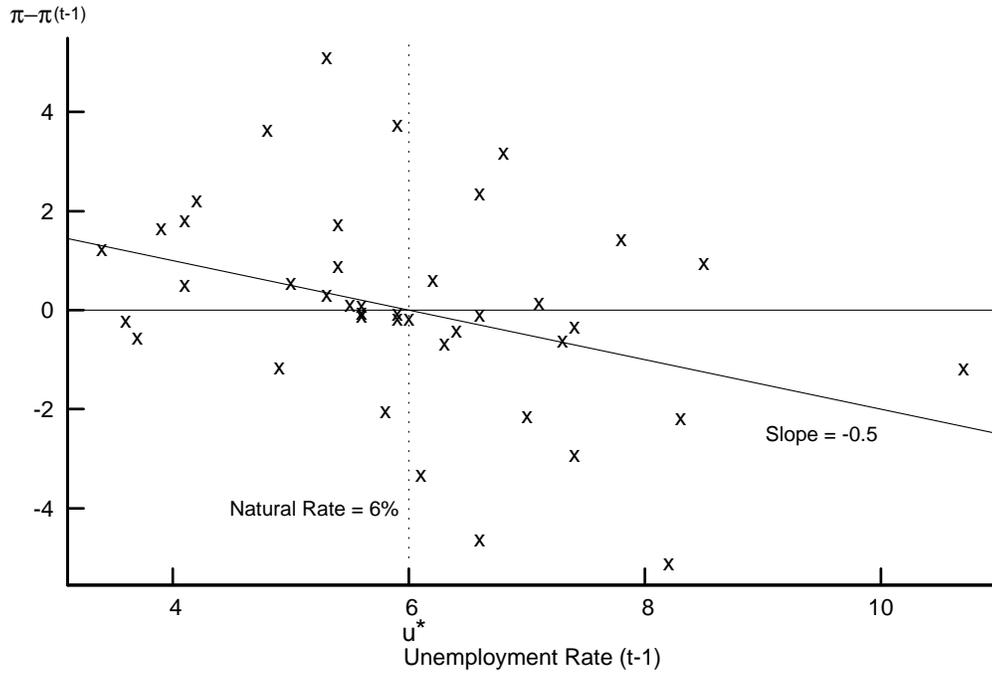
For example, by setting $k = (v_t^b)^{-\frac{1}{2}}$ and replacing it in the above equation one obtains:

$$V(\pi_t, a_t, b_t (v_t^b)^{-\frac{1}{2}}, v_t^a, 1, v_t^{ab} (v_t^b)^{-\frac{1}{2}}) = V(\pi_t, a_t, b_t, v_t^a, v_t^b, v_t^{ab}) \quad (41)$$

Thus it is sufficient to approximate the value function for a grid of five state variables. Such a transformation was used in computing the optimal policies shown in this paper.

Figure 1

The Change in Inflation and the Unemployment Rate
1955-1995: Annual Observations



Note: Inflation is measured from the fourth quarter of the previous year to the fourth quarter of the current year. The unemployment rate is from the fourth quarter of the previous year.

Figure 3
 Disinflation Path Expected by the Policymaker
 (Conditional on Prior Beliefs, the Policy Response and Updating of the Parameter Estimates)

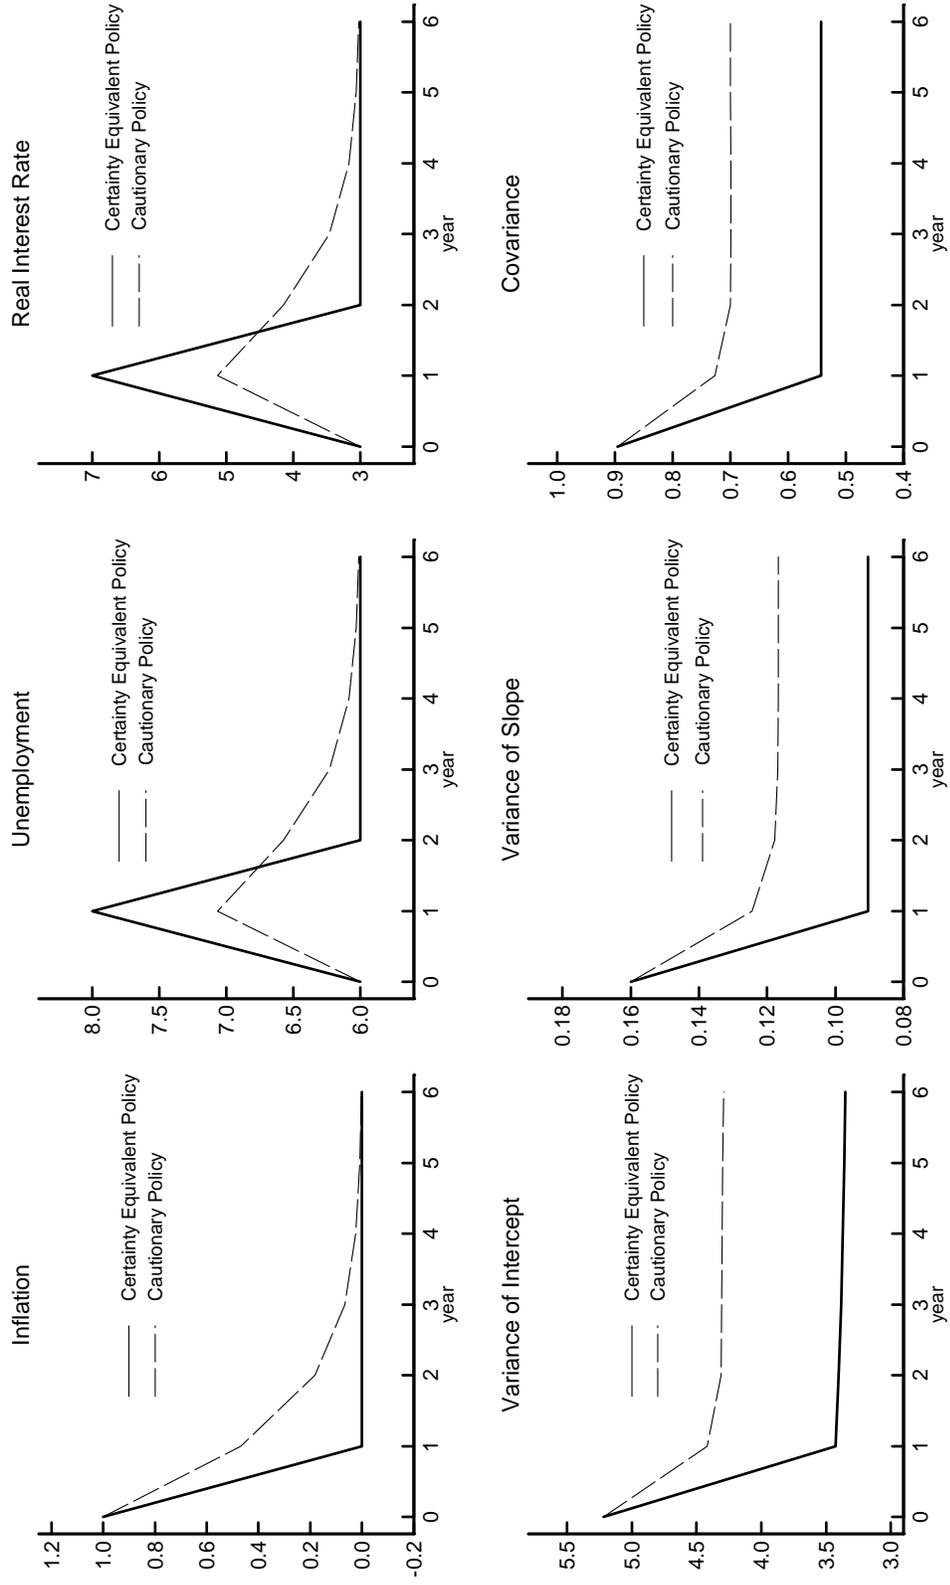


Figure 4

Policy Response to the Inflation Gap (Discount Factor = 0.95)

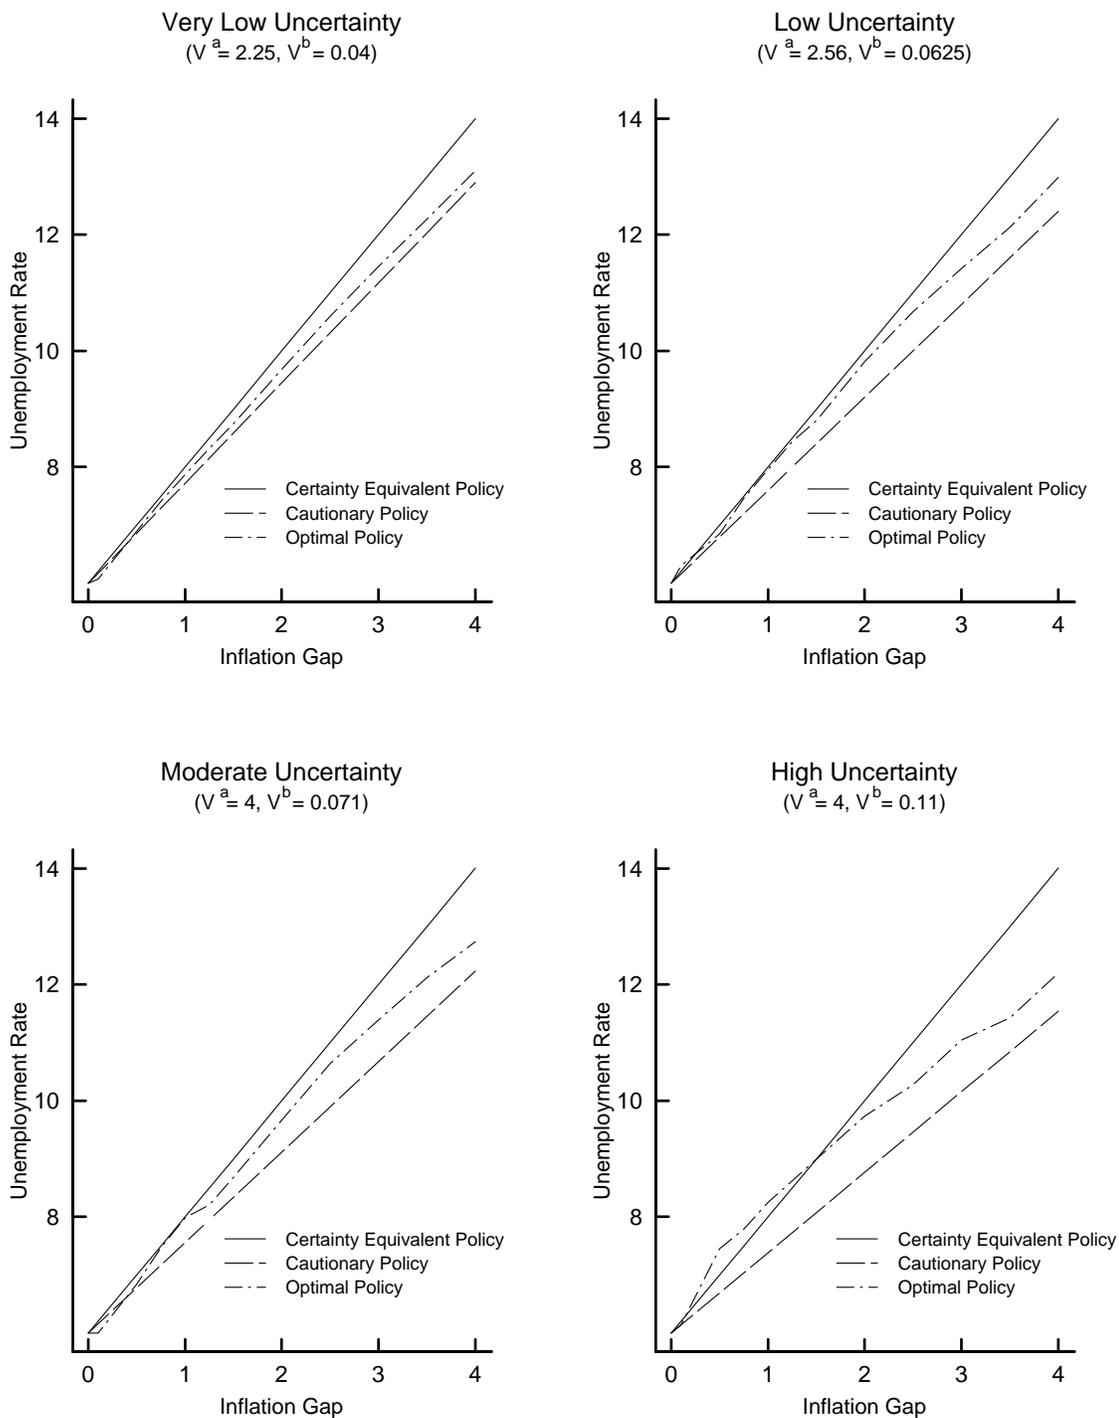


Figure 5
Alternative Discount Factors

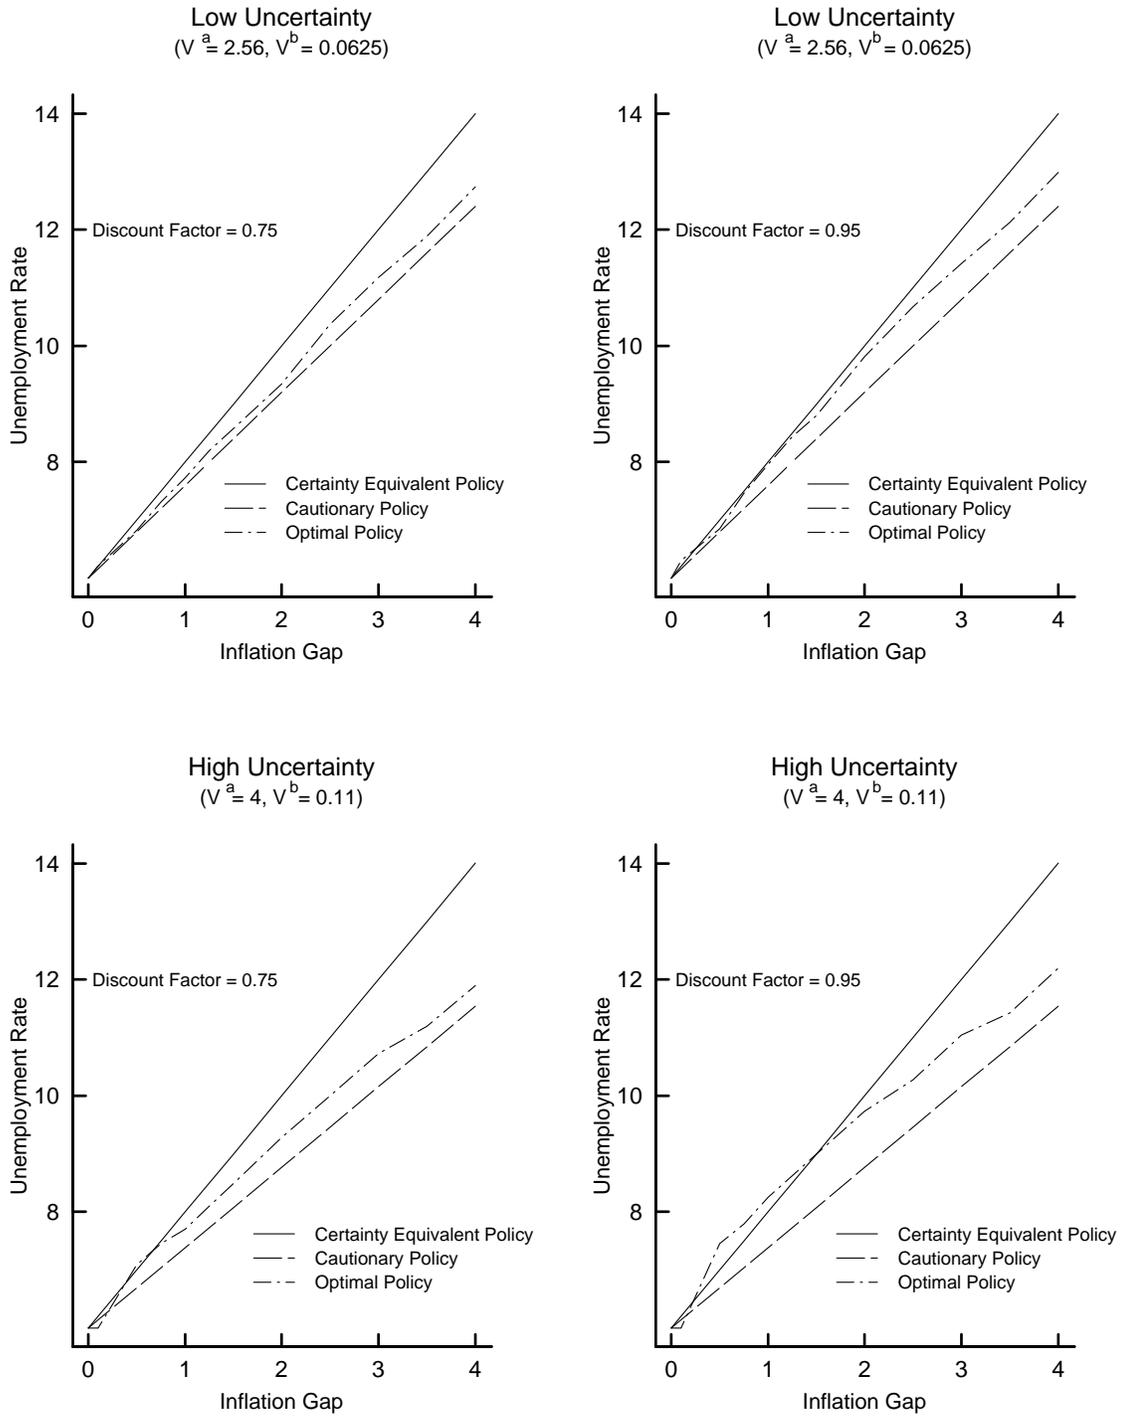


Figure 6
Policy Under Extreme Uncertainty
($V^a = 16$, $V^b = 0.66$)

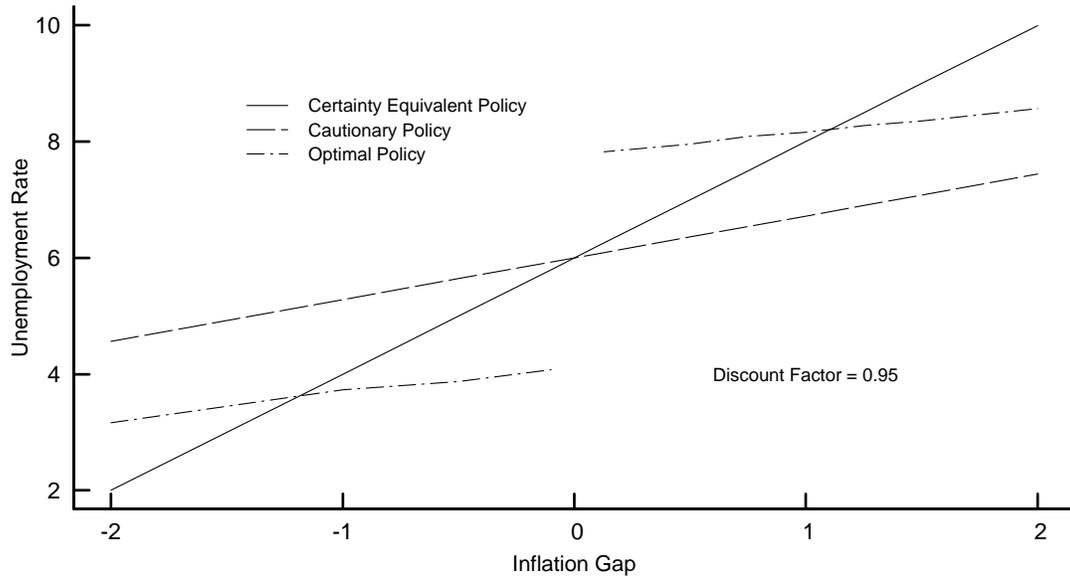


Figure 7
 Covariance Effect
 (Discount Factor = 0.95)

