

DISCUSSION PAPER SERIES

No. 3811

**MONETARY POLICY AND
UNCERTAINTY ABOUT THE NATURAL
UNEMPLOYMENT RATE**

Volker Wieland

INTERNATIONAL MACROECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP3811.asp

MONETARY POLICY AND UNCERTAINTY ABOUT THE NATURAL UNEMPLOYMENT RATE

Volker Wieland, Johann-Wolfgang-Goethe Universität, Frankfurt and CEPR

Discussion Paper No. 3811
March 2003

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Volker Wieland

CEPR Discussion Paper No. 3811

March 2003

ABSTRACT

Monetary Policy and Uncertainty about the Natural Unemployment Rate*

Inflation-targeting central banks have only imperfect knowledge about the effect of policy decisions on inflation. An important source of uncertainty is the relationship between inflation and unemployment. This Paper studies the optimal monetary policy in the presence of uncertainty about the natural unemployment rate, the short-run inflation-unemployment trade-off and the degree of inflation persistence in a simple macroeconomic model, which incorporates rational learning by the central bank as well as private sector agents. 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 they would if they 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, that is, it still remains less aggressive than a policy that disregards parameter uncertainty. Exceptions occur when uncertainty is very high and inflation close to target.

JEL Classification: C61, D80, E24 and E52

Keywords: inflation targeting, monetary policy, natural unemployment rate, optimal learning and parameter uncertainty

Volker Wieland
Chair, Monetary Theory and Policy
Goethe Universitaet Frankfurt, PF 94
Mertonstrasse 17
D-60054 Frankfurt am Main
GERMANY
Tel: (49 69) 7982 5288
Fax: (49 69) 7982 5272
Email: wieland@wiwi.uni-frankfurt.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=152242

*Helpful comments by John Leahy, Andrew Levin, Athanasios Orphanides, Richard Porter, Matthew Shapiro, Peter von zur Muehlen, Carl Walsh, David Wilcox and by seminar participants at the AEA Meetings 1997, the NBER Summer Institute 1997, the NBER URC Conference 1998, the Federal Reserve Board, Georgetown University, Goethe University Frankfurt, the University of Munich, the Institute of International Economic Studies of Stockholm University and Humboldt University Berlin as well as two anonymous referees are gratefully acknowledged. Also, I would like to thank Jeffrey Fuhrer for providing details on his estimation results. Any remaining errors are the sole responsibility of the author. This Paper is a substantially revised and extended version of an earlier working paper that appeared in April 1998 in the Federal Reserve Board Discussion Paper Series, FEDS-98-22.

Submitted 20th February 2003

1 Introduction

A number of central banks of industrialized countries have committed themselves to an explicit inflation targeting strategy and such a strategy has also been recommended for the European Central Bank and the U.S. Federal Reserve System. In implementing this strategy central banks are faced with considerable uncertainty concerning the exact effect of their principal instrument, the short-term nominal interest rate, on inflation.¹ A particularly important and much discussed source of uncertainty regarding the transmission of monetary policy to inflation is the relationship between unemployment and inflation. In implementing policy, central banks have to rely on empirical estimates of the natural unemployment rate (or NAIRU)², the slope of the short-run inflation-unemployment tradeoff and the degree of inflation persistence. Estimates of these parameters have changed over time and their precision is the subject of a continuing active debate.³ Indeed, Staiger, Stock and Watson (1997a, 1997b, 2002) investigate a variety of empirical specifications for the United States 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.

In general, a policy that would be optimal if the parameters of the inflation-unemployment relationship were known with certainty will be recognized as suboptimal once the uncertainty associated with these parameters is taken into account. In this paper, I characterize the optimal policy in the presence of uncertainty about the natural unemployment rate, the short-run inflation-unemployment tradeoff and the weight on forward-looking

¹As a result, inflation-targeting central banks such as the Bank of England and the Sveriges Riksbank have given the discussion of inflation uncertainty center stage in their inflation reports.

²An acronym for non-accelerating inflation rate of unemployment.

³For the recent empirical debate in the United States see Gordon (1997), Staiger, Stock, and Watson (1997a), (1997b) and (2002), Blanchard and Katz (1997), Fuhrer (1995), Akerlof, Dickens and Perry (1996), Phelps and Zoega (1997). There also exists a large literature on Phillips curves in other countries, see for example Debelle and Laxton (1997) and others.

expectations versus lagged inflation in the determination of current inflation. 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 by the central bank and private agents, 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*”, and proposed in his Marshall lectures that “*central banks should calculate the change in policy required to get it right and then do less*”.⁴ 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. Alternatively, proponents of robust control in monetary policy⁵ have argued that worst-case outcomes may best be prevented by following policy rules that are rather aggressive in responding to inflation deviations from target. A further reason, investigated in this paper, is that a more aggressive policy rule 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?*”

The tradeoff between current stabilization and exploration for the sake of better control

⁴See also Blinder (1998) for a discussion of this strategy.

⁵See for example Sargent (1999a), Hansen and Sargent (2001).

in the future has been the focus of a theoretical as well as a computational literature on optimal learning.⁶ Recent applications to monetary policy under uncertainty have been provided by Bertocci and Spagat (1993), Balvers and Cosimano (1994) and Wieland (2000b). Among these, Wieland (2000b) studies the most general learning problem—a linear regression with two unknown parameters—and numerically computes the optimal policy.⁷ Analytical results concerning optimal policy under parameter uncertainty are largely absent from the literature and numerical results are rare, because of 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 presence of a lag as well as a forward-looking expectation of the dependent variable. The numerical algorithm used in this paper is described in more detail in appendix A.

This paper makes the following contributions. First, extending Brainard’s analysis I derive a cautionary policy rule in a model that incorporates rational forward-looking behavior by private sector agents in labor and financial markets. I focus on the case of an inflation-targeting central bank that commits to a specific interest rate rule in the face of uncertainty about the NAIRU, the short-run slope of the Phillips curve and the weight on inflation persistence in terms of lagged inflation versus forward-looking inflation expectations. This cautionary rule represents the optimal policy under commitment in the static version of the model, where the central bank only cares about current performance and disregards dynamic learning effects. I find that the cautionary rule implies gradualism, that is, policy responds to inflationary or disinflationary shocks such that inflation gradually returns to target and policy remains tight or expansive for several periods.

⁶One part of the literature focussed primarily on the asymptotic properties of beliefs and actions. (cf. Taylor (1974), Anderson and Taylor (1976), Easley and Kiefer (1988), Kiefer and Nyarko (1989) and Aghion et al. (1991)), while the other part focussed on characterizing optimal decision rules (cf. Prescott (1972), Kendrick (1981), Kendrick (1982), Mizrach (1991), Amman and Kendrick (1995), Keller and Rady (1999), Wieland (2000a)).

⁷Asymptotic 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. Kasa (1999) also discusses the possibility of incomplete learning by a central bank. Wieland (2000b) 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.

Second, the paper presents numerical results concerning the optimal policy in a dynamic model with rational learning⁸ by the central bank as well as forward-looking agents in labor and financial markets. I find that 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 that completely disregards parameter uncertainty. Thus, in most cases the recommendation for gradualist policymaking under parameter uncertainty survives 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. I proceed to quantify the optimal degree of gradualism and experimentation using empirical estimates of Phillips curve parameter uncertainty by Fuhrer (1995) and Staiger, Stock and Watson (2002). In analyzing the optimal extent of experimentation I also investigate in detail how it is influenced by the presence of forward-looking rational inflation expectations in the Phillips curve and financial markets (i.e. the money market relevant for the definition of the short-term real interest rate). The qualitative properties of the optimal policy are the same under rational learning by private agents as under adaptive expectations. However, the optimal extent of experimentation is smaller, because forward-looking behavior by private agents introduces an expectations channel of monetary policy transmission.

Third, the policy rules derived in this paper are directly comparable to Taylor-style interest rate rules that have been studied extensively in the recent literature on monetary policy.⁹ This literature has focused on evaluating the performance of monetary policy rules in different macroeconomic models under the assumption that all parameters are known with certainty. The analysis in this paper shows how the response coefficients of such a policy rule need to be adjusted in the presence of uncertainty about the relationship between unemployment and inflation.

⁸Boundedly rational learning by central banks is studied by Sims (1988) and Sargent (1999b).

⁹See for example Taylor (1993, 1999), Rotemberg and Woodford (1997), McCallum (1999), Clarida, Gali and Gertler (1998, 1999), Levin, Wieland and Williams (1999, 2002) and many others.

The next section introduces the macroeconomic model that forms the basis of the subsequent analysis of optimal policy rules under uncertainty. In section 3 the cautionary Brainard-type policy rule is derived analytically. Section 4 presents the dynamic framework with learning. A quantitative comparison of the optimal, cautionary and certainty-equivalent policy rules is provided in section 5. Section 6 relates these findings to empirically documented Phillips curve uncertainty. Section 7 concludes and discusses several extensions and avenues for future research.

2 The model

To begin, I consider a central bank that pursues a strict inflation-targeting strategy as defined by Svensson (1997a). Such a central bank conducts monetary policy so as to minimize expected squared deviations from its inflation target π^* .¹⁰ This loss function can be decomposed in two terms indicating the possibility of a tradeoff between the conditional expectation of inflation deviations from target and the conditional variance of inflation:

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

In the following, I characterize optimal monetary policy under uncertainty for this central bank within a simple model of the macroeconomy that follows Clark, Goodhart and Huang (1999). The centerpiece of this model is a Phillips curve that explicitly introduces inflation persistence, in the form of lagged inflation, together with forward-looking inflation expectations:

$$\pi_t = \gamma\pi_{t-1} + (1 - \gamma)\pi_t^e + \beta(u_t - u_t^*) + \epsilon_t \quad \text{where } \beta < 0, 0 < \gamma \leq 1 \quad (2)$$

Current inflation π_t is related to the deviation of the unemployment rate u_t from the natural rate with a negative slope parameter β . Furthermore, it depends on lagged inflation π_{t-1} , price-setters' expectation of inflation π_t^e , and a normally-distributed random shock $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. This specification is sometimes called the 'backward and forward-looking

¹⁰An extension to flexible inflation targeting, which incorporates an output or unemployment stabilization objective, will be discussed in the final section of the paper.

components' model of the Phillips curve (cf. Buiter and Miller (1985)). The backward-looking component reflects inertia in inflation that may be derived from some types of overlapping wage contracts¹¹ or may be attributed to the presence of rule-of-thumb price setters. The coefficients of the two components sum to unity so that in the long-run equilibrium, $\pi_t = \pi_{t-1} = \pi_t^e$. As γ approaches zero, this specification simplifies to the standard Lucas surprise supply function. For $\gamma = 1$ it corresponds to the traditional accelerationist Phillips curve.¹² The long-run equilibrium rate of unemployment, i.e. the natural rate, is denoted by u_t^* . An often used assumption regarding the evolution of u_t^* is:

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

where $\eta_t \sim N(0, \sigma_\eta^2)$. For $\sigma_\eta^2 > 0$, the natural rate follows a random walk. For $\sigma_\eta^2 = 0$, the natural rate is constant, $u_t^* = u_0^* \forall t$.

Two more equations are needed to account for the transmission of monetary policy from the central bank's principal policy instrument, that is, the short-term nominal interest rate, i_t , to the policy target, i.e. the inflation rate. First, the unemployment rate is related to real aggregate demand, y_t according to a version of "Okun's Law."¹³ Then, aggregate demand is related to the short-term real interest rate, that is the difference between the nominal rate i_t and expected inflation:

$$u_t = \phi y_t \quad \text{where } \lambda < 0 \quad (4)$$

$$y_t = \lambda(i_t - \pi_t^e) \quad \text{where } \phi < 0 \quad (5)$$

It remains to specify private sector expectations of inflation, π_t^e . The benchmark for the following analysis will be the case of rational expectations. For comparison I will also

¹¹See for example Fuhrer and Moore (1995a,b).

¹²Typical empirical estimates of this specification indicate a significant degree of inflation persistence ranging from 0.5 to near unity (cf. Fuhrer (1997) and Roberts (1997)). The New-Keynesian Phillips curve, that has received much interest in the recent literature (cf. Galí and Gertler (1999)), differs from the above specification only in the timing of the forward-looking inflation term which concerns period $t + 1$. The preferred empirical specification of the New-Keynesian Phillips curve also embodies a significant degree of inflation persistence.

¹³For a textbook discussion of this empirical regularity see Dornbusch and Fischer (1990). Clark et al. (1999) do not need this relationship, because they specify the Phillips curve in terms of the output gap rather than the unemployment gap.

consider the case of adaptive, random-walk expectations, i.e. $\pi_t^e = \pi_{t-1}$. In that case the Phillips curve specification (2) simplifies again to the traditional accelerationist Phillips curve and the ex-post real interest rate appears in the aggregate demand relationship. Under forward-looking behavior, expected inflation will be a function of the state of the economy, the parameters and importantly also the policy rule pursued by the central bank:

$$\pi_t^e = (\gamma + \beta\lambda\phi)^{-1}(\gamma\pi_{t-1} - \beta u_{t-1}^* + \beta\lambda\phi i_t^e) \quad (6)$$

Here the private sector's expectation regarding monetary policy is denoted by i_t^e . It indicates that we need to distinguish between discretionary policy and a possible commitment by the central bank to a specific policy rule. Under discretion, the central bank optimizes policy taking private sector expectations as given and unaffected by its choice of interest rate. Under commitment, the central bank internalizes the impact of its decision rule on private sector expectations and commits to delivering the state-contingent interest rate setting that is expected under this rule. The recent literature on monetary policy rules¹⁴ has emphasized the benefits of adhering to a rule rather than pursuing discretionary policy. Thus, in the following analysis I will focus on the optimal policy under commitment and only return to the case of discretion in the last section of the paper.

The central bank will set the nominal interest rate i_t so as to minimize $L(\cdot)$ based on its knowledge of the state of the economy (i.e. lagged inflation), the parameters (i.e. β, γ and the natural rate) but before the shocks ϵ_t and η_t are realized. It can predict and respond to impending changes in inflation only to the extent that they result from endogenous inflation persistence but not to the current-period random shocks. When the central bank is committed to a state-contingent rule such as

$$i_t = H(\pi_{t-1}, \beta, \gamma, u_{t-1}^*, \lambda, \phi), \quad (7)$$

it will implicitly take into account how its actions affect private sector expectations. Clarke et al. (1999) show that the optimal policy under commitment to such a rule can be obtained

¹⁴See for example the contributions in Taylor (1999).

by minimizing the loss function $L(\cdot)$ with respect to i_t and i_t^e under the explicit restriction that the ex-ante expected nominal interest rate, i_t^e , is equal to its rational expectation:¹⁵

$$i_t^e = E_{t-1}[i_t | \pi_{t-1}, \beta, \gamma, u_{t-1}^*, \phi, \lambda] \quad (8)$$

From (7) and (8) and the assumption made earlier that neither the central bank nor the private sector have prior information on the random shocks (ϵ_t, η_t) when choosing i_t and i_t^e respectively, it follows that the private sector's ex-ante rational expectation of the nominal interest rate will be equal to the interest rate prescribed by the state-contingent policy rule, $E_{t-1}[i_t | \pi_{t-1}, \beta, \gamma, u_{t-1}^*, \lambda, \phi] = H(\pi_{t-1}, \beta, \gamma, u_{t-1}^*, \lambda, \phi)$. As a result, the private sector's rational expectation of inflation is

$$\pi_t^e = E_{t-1}[\pi_t] = (\gamma + \beta\lambda\phi)^{-1}(\gamma\pi_{t-1} - \beta u_{t-1}^* + \beta\lambda\phi H(\pi_{t-1}, \beta, \gamma, u_{t-1}^*, \lambda, \phi)) \quad (9)$$

Furthermore, due to symmetric information between central bank and private sector, the expectation derived in (9) also corresponds to the central bank's rational expectation of inflation that enters the loss function (1).¹⁶ The second element of the loss function is the conditional variance of inflation,

$$VAR_{t-1}[\pi_t] = \sigma_\epsilon^2 + \beta^2 \sigma_\eta^2 \quad (10)$$

which turns out not to depend on the interest rate rule. Thus, the central bank will be able to minimize its loss $L(\cdot)$ simply by setting the interest rate to the value that induces an expected inflation rate equal to the inflation target $E_{t-1}[\pi_t] = \pi^*$. This corresponds to a strategy of "inflation forecast targeting" as defined by Svensson (1997a). As a result, the expected deviation from target will be equal to zero and the minimized loss will correspond to the exogenous conditional variance (10). The implied optimal interest rate rule is:

$$i_t = H(\pi_{t-1}, \beta, \gamma, u_{t-1}^*, \lambda, \phi) = (\lambda\phi)^{-1}u_{t-1}^* - (\beta\lambda\phi)^{-1}(\gamma\pi_{t-1} - (\gamma + \beta\lambda\phi)\pi^*) \quad (11)$$

¹⁵Note that in this notation, the private sector's expectations of inflation and the nominal interest rate, π_t^e and i_t^e are variables, while $E_{t-1}\pi_t$ and $E_{t-1}i_t$, the rational expectations at $t-1$ are functions of the policy rule, lagged inflation and the parameters. Committing to $E_{t-1}[i_t]$ has also been used as commitment strategy by Svensson (1997b) and many others.

¹⁶Possible extensions allowing for asymmetric information are discussed in the final section of the paper.

The first term of this rule essentially represents the equilibrium real interest rate, which is related to the natural unemployment rate. The second term represents the central bank's response to past inflation that is intended to return inflation to its target value in the next period. Note that the parameters β , λ and ϕ are all negative and the central bank responds to an increase in inflation by raising the nominal interest rate in the following period. The magnitude of the necessary policy response depends on the degree of inflation persistence and the slope of the Phillips curve as well as the slope parameter of Okun's law and the aggregate demand equation. As can be seen from (11) the neutral setting of the nominal interest rate when inflation is on target is: $i_t = (\lambda\phi)^{-1}u_{t-1}^* + \pi^*$.

To clarify the effect of rational expectations in this model I also derive the optimal policy rule under adaptive random-walk expectations (i.e. $\pi_t^e = \pi_{t-1}$):

$$i_t = (\lambda\phi)^{-1}u_{t-1}^* + \pi_{t-1} - (\beta\lambda\phi)^{-1}(\pi_{t-1} - \pi^*) \quad (12)$$

A comparison of the partial derivatives of (11) and (12) with respect to lagged inflation shows that the central bank needs to respond to an increase in inflation by raising the nominal interest rate in the subsequent period to a much greater extent if the private sector forms adaptive rather than rational expectations, $(\beta\phi\lambda - 1)(\beta\phi\lambda)^{-1} > -\gamma(\beta\phi\lambda)^{-1} > 0$. The reason is that under rational expectations the private sector expects the central bank to raise interest rates sufficiently to return inflation to target in the next period and forms its inflation forecast accordingly. Under commitment, the central bank in turn takes into account this beneficial effect of private sector expectations in the formulation of its monetary policy rule. In the literature this effect is typically referred to as the 'expectations channel' of monetary policy transmission.

So far we have focussed on the static problem under certainty, where the central bank minimizes current period losses with perfect knowledge regarding the parameters of the model.¹⁷ In the next section, we introduce uncertainty with respect to the parameters of the Phillips curve.

¹⁷Under strict inflation targeting with known parameters the optimal rule in the static model is in fact also dynamically optimal.

3 Parameter uncertainty and cautionary policy

The policy rule (11) cannot be implemented if the parameters of the Phillips curve, that is, the natural unemployment rate u_t^* , the slope β and the index of persistence γ are unknown. However the policymaker can obtain recursive estimates from the following equation:

$$\pi_t - E_{t-1}\pi_t = \alpha_t + \beta u_t + \gamma(\pi_{t-1} - E_{t-1}\pi_t) + \epsilon_t \quad (13)$$

The intercept of this equation corresponds to the product of the unknown slope and the natural rate and may therefore vary over time according to:

$$\alpha_t = \beta u_t^* = \beta u_{t-1}^* + \beta \eta_t = \alpha_{t-1} + \nu_t \text{ where } \nu_t \sim N(0, \sigma_\nu^2) \quad (14)$$

The means of the intercept, slope and persistence parameters as of $t - 1$ are denoted by

$$E_{t-1}[(\alpha_t, \beta, \gamma)] = (a_{t-1}, b_{t-1}, c_{t-1}), \quad (15)$$

while the degree of uncertainty about these parameters based on $t - 1$ information is characterized by the following covariance matrix:

$$\Sigma_{t|t-1} = \begin{pmatrix} v_{t-1}^a & v_{t-1}^{ab} & v_{t-1}^{ac} \\ v_{t-1}^{ab} & v_{t-1}^b & v_{t-1}^{bc} \\ v_{t-1}^{ac} & v_{t-1}^{bc} & v_{t-1}^c \end{pmatrix} = \begin{pmatrix} v_{t-1}^a + \sigma_\nu^2 & v_{t-1}^{ab} & v_{t-1}^{ac} \\ v_{t-1}^{ab} & v_{t-1}^b & v_{t-1}^{bc} \\ v_{t-1}^{ac} & v_{t-1}^{bc} & v_{t-1}^c \end{pmatrix} \quad (16)$$

The twelve variables $(a_{t-1}, b_{t-1}, \Sigma_{t|t-1})$ define a trivariate normal distribution which comprises all relevant information about the unknown parameters $(\alpha_t, \beta, \gamma)$ at time $t-1$. This distribution represents the policymaker's beliefs about the parameters of the Phillips curve.¹⁸

The NAIRU is not an explicit element of this distribution, but as in the empirical literature on Phillips curves the ratio of the means can be used as an estimator for the NAIRU:

$$\hat{u}_t^* = a_{t-1} b_{t-1}^{-1} \quad (17)$$

¹⁸Note 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. The normality assumption regarding the parameter estimates does not take into account the restrictions on the sign or the magnitude of the parameters, s.t. the non-negativity of the natural unemployment rate or the constraint of the persistence index to $(0,1]$. We maintain this assumption throughout this section for tractability of the theoretical analysis. The numerical analysis later on will recognize the explicit constraint on the index parameter γ in assessing the impact of policy on future estimates.

This illustrates that the uncertainty about NAIRU estimates that is emphasized in the empirical literature on Phillips curves cannot be discussed separately from the uncertainty that is associated with slope of the short-run inflation-unemployment tradeoff.

Given our starting assumption of symmetric information between central bank and private sector this distribution also represents the private sector beliefs about the unknown parameters. Furthermore, for the remainder of this analysis we will continue to treat the other parameters, (λ, ϕ) , as known. Then, the central bank's and private sector's expectation of inflation as of $t - 1$ that appeared in equation (13) above can be expressed as a function of lagged inflation and last period's beliefs as well as the central bank's policy rule $H(\pi_{t-1}, a_{t-1}, b_{t-1}, c_{t-1}, \Sigma_{t|t-1}, \lambda, \phi)$ conditional on those same beliefs.

$$E_{t-1}\pi_t = (c_{t-1} + b_{t-1}\lambda\phi)^{-1}(a_{t-1} + b_{t-1}\lambda\phi H(\cdot) + c_{t-1}\pi_{t-1}) \quad (18)$$

This brings us back to the central question of the paper, namely what is an appropriate implementable policy rule $H(\cdot)$ under parameter uncertainty. A first potential candidate is the optimal rule under certainty in (11), which can be rendered implementable by replacing actual parameter values with available estimates:

$$i_t = H^{ceq}(\cdot) = (\lambda\phi)^{-1}a_{t-1}b_{t-1}^{-1} - (b_{t-1}\lambda\phi)^{-1}(c_{t-1}\pi_{t-1} - (c_{t-1} + b_{t-1}\lambda\phi)\pi^*) \quad (19)$$

This “certainty-equivalent” policy rule is useful as a benchmark for comparison but it is clearly not optimal in the presence of parameter uncertainty. The one-period optimal policy rule can be derived analytically by minimizing the current expected loss, $L(\cdot)$, conditional on all available information, including the degree of uncertainty associated with the parameter estimates. As noted previously, the loss $L(\cdot)$ consists of two components. The first component, the square of the expected deviation of inflation from target, uses the inflation expectation defined by (18). The second component is the conditional variance of inflation:

$$\begin{aligned} VAR_{t-1}[\pi_t] &= \sigma_\epsilon^2 + v_{t|t-1}^a + u_t^2 v_{t-1}^b + (\pi_{t-1} - E_{t-1}\pi_t)^2 v_{t-1}^c \\ &+ 2u_t v_{t-1}^{ab} + 2v_{t-1}^{ac}(\pi_{t-1} - E_{t-1}\pi_t) + 2v_{t-1}^{bc}u_t(\pi_{t-1} - E_{t-1}\pi_t) \\ &\text{where } u_t = \phi\lambda(H(\cdot) - E_{t-1}\pi_t) \text{ and } E_{t-1}\pi_t \text{ defined as in (18)} \end{aligned} \quad (20)$$

This variance depends on the degree of parameter uncertainty and on the chosen policy rule. As a consequence, the central bank faces a trade-off between the expected deviation of inflation from target and the conditional variance of inflation. Thus, the optimal rule does not simply imply inflation-forecast targeting but importantly takes into account inflation uncertainty. Since its response to uncertainty reflects caution as a motive in monetary policymaking, I will refer to it as the 'cautionary' rule. It takes the form:

$$i_t = H^{cau}(\cdot) = -((b_{t-1}\lambda\phi)^2 + R_1))^{-1} ((c_{t-1}b_{t-1}\lambda\phi - R_1)\pi_{t-1} - (c_{t-1} + b_{t-1}\lambda\phi)b_{t-1}\lambda\phi\pi^* + a_{t-1}(b_{t-1}\lambda\phi + R_2) + R_3) \quad (21)$$

where the effect of the parameter variances and covariances is summarized by (R_1, R_2, R_3) :

$$\begin{aligned} R_1 &= v_{t-1}^b c_{t-1}^2 + v_{t-1}^c (b_{t-1}\lambda\phi)^2 - 2v_{t-1}^{bc} (c_{t-1}b_{t-1}\lambda\phi) \\ R_2 &= (v_{t-1}^c + v_{t-1}^{bc})(b_{t-1}\lambda\phi) - (v_{t-1}^b + v_{t-1}^{bc})c_{t-1} \\ R_3 &= v_{t-1}^{ab} c_{t-1} (c_{t-1} + b_{t-1}\lambda\phi) - v_{t-1}^{ac} b_{t-1}\lambda\phi (c_{t-1} + b_{t-1}\lambda\phi) \end{aligned} \quad (22)$$

Each of these three coefficients would be zero in the absence of uncertainty and the policy rule would simplify to (19). Under uncertainty, however, optimal policy depends on the parameter variances (v^b, v^c) and covariances (v^{ab}, v^{ac}, v^{bc}) . It does not depend on the variance of the intercept in the estimated equation, v^a , the reason being that for a linear model and quadratic objective function certainty-equivalence applies with respect to additive uncertainty.

In his seminal paper Brainard showed that multiplicative parameter uncertainty such as the uncertainty captured by v^b provides a motive for cautious, gradualist policymaking.¹⁹ A comparison between the rules (21) and (19) shows that this result extends to the model with rational expectations and inflation persistence considered in this paper. Here, gradualism arises in two ways. First, the response of the cautionary rule (21) subsequent to an increase in inflation is more muted than under the certainty-equivalent rule (19).

¹⁹Other papers that have looked at this effect recently are Clarida, Gali and Gertler (1999), Estrella and Mishkin (1998) and Svensson (1999). Sack (1999) shows how parameter uncertainty can explain the high degree of serial correlation in interest rates.

The partial derivative $\delta H^{cau}/\delta\pi_{t-1}$ is a function of the variance of the slope estimate, v^b , the variance of the index of persistence, v^c and their covariance, v^{bc} . Once one recognizes that the parameter R_1 in (22) is equivalent to the variance, $VAR(\beta c_{t-1} - \gamma b_{t-1}\lambda\phi)$, and consequently must be non-negative, it is straightforward to show that:

$$\frac{\delta H^{cau}}{\delta\pi_{t-1}} = \frac{c_{t-1}b_{t-1}\lambda\phi - R_1}{(b_{t-1}\lambda\phi)^2 + R_1} < \delta H^{ceq}\delta\pi_{t-1} = \frac{c_{t-1}}{b_{t-1}\lambda\phi} \quad (23)$$

Thus, under parameter uncertainty the central bank will increase the nominal interest rate subsequent an inflationary shock by less than in the absence of uncertainty. This increase in the interest rate will not be sufficient to return next period inflation to target in expectation. Thus, under parameter uncertainty even a strict inflation- targeting central bank will not pursue a pure inflation-forecast targeting strategy defined as keeping expected inflation always on target. Instead, inflation will remain elevated and only return to target gradually over the next few periods. Similarly, the interest rate will exhibit gradualism in that it will be expected to remain elevated and only return to its neutral level after some time.

The cautionary policy rule (21) also exhibits a second element of caution or gradualism. This becomes apparent when we consider its implications in a situation where last period's inflation rate is equal to the central bank's target, $\pi_{t-1} = \pi^*$. It may seem surprising at first that the cautionary rule does not prescribe the same neutral setting as the certainty-equivalent rule in this case. The certainty-equivalent rule would set the nominal interest rate equal to its estimated neutral rate, $(\lambda\phi)^{-1}a_{t-1}b_{t-1}^{-1} + \pi^*$. Why does the cautionary rule not adopt the same neutral setting when $\pi_{t-1} = \pi^*$? The reason again is related to the second component of the central bank's loss function, that is the conditional variance of inflation, i.e. inflation uncertainty. The setting of the interest rate that minimizes the conditional variance of inflation need not be equivalent to its estimated neutral level. Rather the cautionary rule sets the nominal interest rate according to a simple weighted average of its neutral level and the variance-minimizing level:

$$H^{cau}(\pi_{t-1} = \pi^*) = \frac{(b_{t-1}\lambda\phi)^2}{(b_{t-1}\lambda\phi)^2 + R_1} \left(\pi^* - \frac{1}{\lambda\phi} \frac{a_{t-1}}{b_{t-1}} \right) \quad (24)$$

$$+\frac{R_1}{(b_{t-1}\lambda\phi)^2 + R_1} \left(\pi^* - \frac{a_{t-1}R_2 + R_3}{R_1} \right)$$

Here the first term in large parentheses corresponds to the neutral or natural level of the nominal interest rate, which ensures that unemployment is equal to the NAIRU, while the second term corresponds to the variance-minimizing level of the nominal interest rate, given $\pi_{t-1} = \pi^*$.

By definition the variance-minimizing level of the nominal interest rate is that level where the central bank will be able to assess the impact of the nominal interest rate on unemployment and inflation with the highest possible precision. Thus, the tendency to set the nominal interest rate near that level whenever inflation is on target clearly reflects a cautionary motive. Again, this motive implies a gradualist pursuit of policy. This becomes particularly clear when considering the case of constant natural unemployment rate $u_t^* = u_0^*$, i.e. a constant intercept, α , in the estimated inflation equation (13) discussed at the beginning of this section. In this case, the parameters (α, β, γ) may be estimated by recursive least squares. By definition, least squares estimates imply that the variance of the dependent variable is minimized at the means of the explanatory variables. Since the explanatory variables in (13), u_t and $\pi_{t-1} - E_{t-1}\pi_t$, are both linear functions of the ex-post real interest rate $i_t - \pi_{t-1}$, their historical mean values will coincide with the average ex-post real interest rate. Thus, when the optimal rule leans towards the variance-minimizing level of the interest rate it effectively keeps the ex-post real interest rate closer to its past average. As a result, the policy stance changes more gradually over time as would be prescribed by a policy rule that disregards inflation uncertainty arising from imprecise parameter estimates.

Clearly, the cautionary policy rule (21) that minimizes expected one-period loss is not necessarily optimal in a dynamic context. It is a “myopic” policy, because it disregards the effect of the current interest rate setting on future parameter estimates and policy performance. In the next section, I show how the estimates of the parameters of the inflation equation (13) may be updated over time and how such learning introduces an important dynamic link between current policy decisions and future parameter uncertainty

and stabilization performance.

4 Rational learning and the optimal policy rule

As new observations on inflation and unemployment become available the central bank and private sector agents can update their estimates of the unknown parameters (α, β, γ) in the inflation equation (13). As long as they share the same information and start off with the same prior belief about the unknown parameters their estimates and updating equations will coincide. The relevant updating equations for their beliefs, $\theta_{t-1} = (a_{t-1}, b_{t-1}, c_{t-1}, \Sigma_{t|t-1})$, can be cast in form of the Kalman filter. To be able to present the updating equations in a compact manner it is helpful to define vector of explanatory variables $X_t = (1 \quad u_t \quad (\pi_{t-1} - E_{t-1}\pi_t))'$ where $E_{t-1}\pi_t$ is as defined in equation (18). Then the updating equations for the elements of θ , that is the means and the covariance matrix, can be expressed as:

$$\begin{pmatrix} a_t \\ b_t \\ c_t \end{pmatrix} = \begin{pmatrix} a_{t-1} \\ b_{t-1} \\ c_{t-1} \end{pmatrix} + \Sigma_{t|t-1} X_t F^{-1} (\pi_t - E_{t-1}\pi_t - a_{t-1} - b_{t-1}u_t - c_{t-1}(\pi_{t-1} - E_{t-1}\pi_t)) \quad (25)$$

$$\Sigma_{t|t} = \Sigma_{t|t-1} - \Sigma_{t|t-1} X_t F^{-1} X_t' \Sigma_{t|t-1}$$

$$\text{where } F = X_t \Sigma_{t|t-1} X_t' + \sigma_\epsilon^2$$

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 updating²⁰ of the trivariate normal distribution that represents the central bank's and private sector's beliefs about α_t, β and γ .²¹ Under the assumption of a constant natural unemployment rate, these updating equations are also equivalent to recursive least squares. In that case the intercept parameter α will be constant

²⁰The asymptotic behavior of these beliefs is discussed further in appendix B.

²¹This dynamic learning model extends earlier analysis of optimal learning by Wieland (2000a) and (2000b) in a simple regression framework. It allows for a lagged dependent variable with an additional unknown parameter and includes the rational expectation of the dependent variable. These extensions raise the number of state variables and increase computational complexity, but their benefit is to allow application of the optimal learning framework to a simple but completely specified model of the monetary policy transmission mechanism.

and the only change in the updating equations will be that $v_{t|t-1}^a = v_{t-1}^a$.

As a result of rational learning by the central bank and the private sector the current choice of the interest rate will affect the precision of the point estimates as well as the estimates themselves through its impact on current unemployment, inflation expectations and inflation. 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. Thus, the optimal policy rule $H(\pi_{t-1}, \theta_{t-1}, \pi^*, \lambda, \phi)$ in this dynamic model with learning solves the following optimization problem:

$$\text{Min}_{H(\cdot)} E \left[\sum_{t=0}^{\infty} \delta^t (\pi_t - \pi^*)^2 \mid (\pi_0, \theta_0) \right] \quad (26)$$

$$\text{s.t. } i_t = H(\pi_{t-1}, \theta_{t-1}, \pi^*, \lambda, \phi) \text{ for } t = 1, \infty$$

$$\text{and s.t. equations (4), (5), (13), (18) and (25)}$$

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). The state variables of this dynamic programming problem are lagged inflation π_{t-1} and last period's beliefs θ_{t-1} . Denoting the value function for this dynamic program by $V(\pi, \theta)$ the associated Bellman equation corresponds to:

$$\begin{aligned} V(\pi_{t-1}, \theta_{t-1}) &= \text{Min}_{H(\cdot)} L(\pi_{t-1}, \theta_{t-1}, H(\cdot)) \\ &\quad + \delta \int V(\pi_t(H(\cdot), \cdot), \theta_t(H(\cdot), \cdot)) f(\pi_t \mid \pi_{t-1}, \theta_{t-1}, H(\cdot)) d\pi \\ &= \text{Min}_{H(\cdot)} L(\pi_{t-1}, \theta_{t-1}, H(\cdot)) + \delta \int V(\alpha_t, \beta, \gamma, \epsilon_t, H(\cdot), \pi_{t-1}, \theta_{t-1}) \\ &\quad p(\alpha_t, \beta, \gamma \mid \pi_{t-1}, \theta_{t-1}, H(\cdot)) q(\epsilon) d\alpha d\beta d\gamma d\epsilon \end{aligned} \quad (27)$$

Two terms on the right-hand side of the upper equation in (27) characterize the tradeoff between current control and estimation. $L(\cdot)$ is the expected current loss, while the second term denotes the expectation of next period's value function, which summarizes all future losses and is multiplied with the discount factor δ . This second 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. Inflation, unemployment and next period's beliefs all depend on the central bank's choice of interest rate i_t and thus on its policy rule $H(\pi_{t-1}, \theta_{t-1}, \pi^*, \lambda, \phi, \sigma_\epsilon, \sigma_\eta)$ that feeds back on all currently available information.

In the lower equation in (27), time t values of inflation and beliefs have been substituted out using equations (4), (5), (13), (18) and (25). They are functions of the previous period's inflation rate 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. $p(\alpha_t, \beta, \gamma|\cdot)$ is the trivariate normal distribution that describes the policymaker's beliefs about α_t , β and γ . $q(\epsilon)$ refers to the normal density function of the shocks in the Phillips curve.

Associated with this Bellman equation is a stationary optimal policy function which maps the state variables $(\pi_{t-1}, \theta_{t-1})$ into a value for the nominal interest rate:

$$i_t = H^{opt}(\pi_{t-1}, \theta_{t-1}, \pi^*, \gamma, \phi, \delta, \sigma_\epsilon, \sigma_\eta) \quad (28)$$

It is the dynamically optimal counterpart of the certainty-equivalent and cautionary policy rules (19) and (21) that were derived analytically in the preceding section. Unfortunately analytical solutions for $H^{opt}(\cdot)$ 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 rule.

5 The optimal balance of caution and experimentation

The Bellman 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. However, 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, the optimal learning problem with three unknown parameters in (26), which has a total of 10 continuous state variables is too large to be solved numerically with reasonable precision.²² Instead I provide numerical results for three simpler versions of this learning problem. First, I discuss a learning problem with one unknown parameter, the slope of the Phillips curve, β . Then, I report results on two generalizations of this problem, each with learning about two unknown parameters. In one case the central bank and private sector agents are learning about the intercept and slope of the Phillips curve (α, β) and in the other case about the slope and the index of persistence (β, γ) . In presenting these results I will discuss the differences in the optimal policy under adaptive expectations versus rational learning of private sector agents in labor markets and financial markets in detail. Inflation expectations relevant to labor market decisions appear in the Phillips curve with their importance depending on γ , while inflation expectations concerning financial market decisions appear in the definition of the ex-ante real interest rate.

To conduct numerical analysis of the optimal learning problem it is necessary to specify numerical values for several parameters. The chosen values are reported in **Table 1**.

Policy with uncertain β

²²The numerical algorithm and associated computation costs are discussed in more detail in appendix A.

Table 1: Calibration

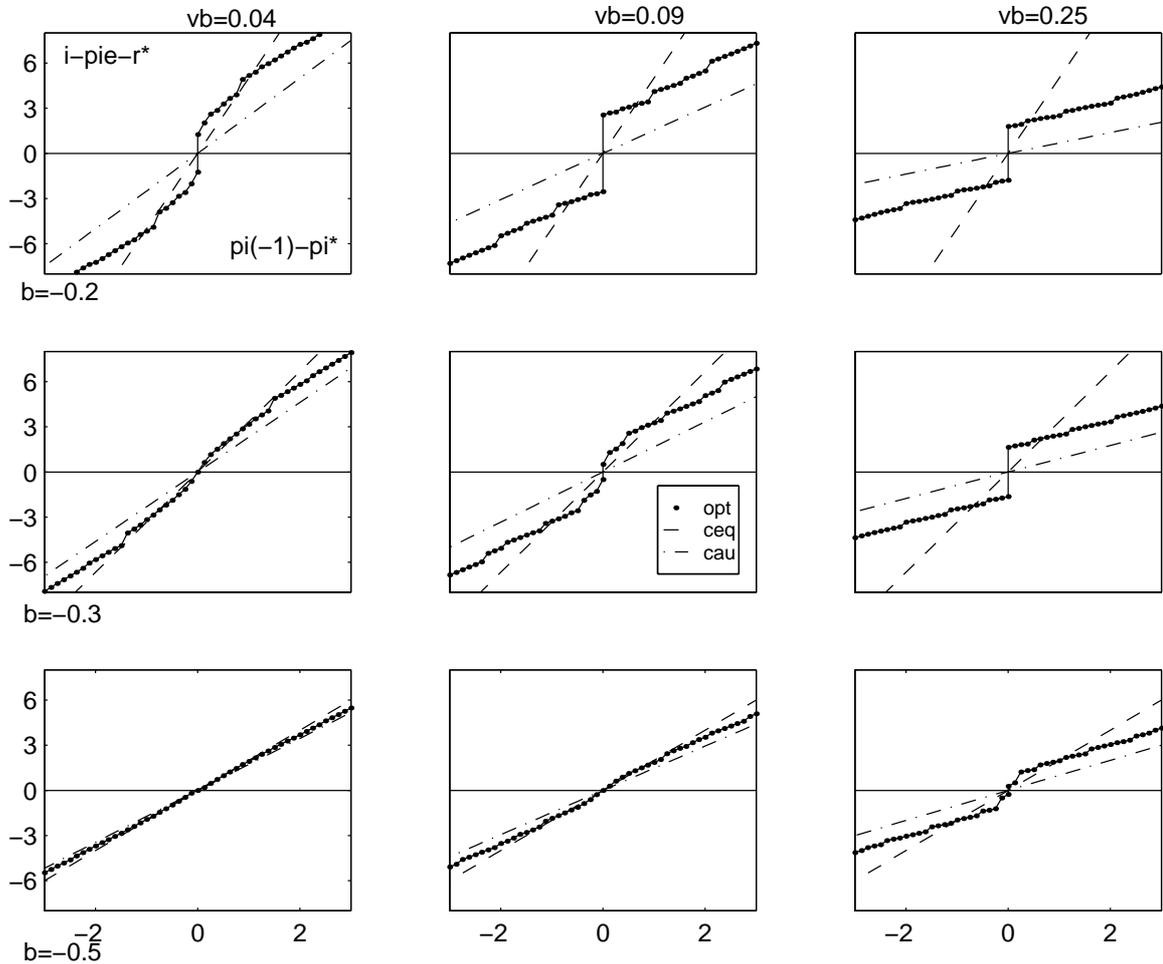
Economy	$\phi = -1$	$\lambda = -1$
Central Bank	$\pi^* = 0$	$\delta = 0.95$
Noise	$\sigma_\epsilon^2 = 1$	$\sigma_\eta^2 = 0$

First, I consider optimal learning by the central bank when private sector agents form adaptive expectations, i.e. $\pi_t^e = \pi_{t-1}$. **Figure 1** compares the response to lagged inflation implied by the dynamically optimal policy rule with that of the cautionary and certainty-equivalent rules. The figure contains nine panels. In each panel the horizontal axis measures the deviation of lagged inflation π_{t-1} from the zero inflation target. The vertical axis corresponds to the expected deviation of unemployment from its natural rate that would occur given the nominal interest rate set by the central bank in response to the inflation deviation from target. Given the parameter settings ($\lambda = -1, \phi = -1$), the vertical axis is also equal to the expected deviation of the real interest rate $i_t - \pi_t^e$ from its equilibrium value, $(\phi\lambda)^{-1}u^*$. Each panel compares the optimal policy rule (solid line with thick dots)²³ to the certainty-equivalent rule (dashed line) and the cautionary rule (dashed-dotted line) for a given combination of the slope estimate b and its variance v^b . The first row of panels corresponds to an estimate of $b = -0.2$, the second row to $b = -0.3$ and the third row to $b = -0.5$. The columns are ordered as follows: the first column corresponds to a variance of $v^b = 0.04$ the second column to $v^b = 0.09$ and the third column to $v^b = 0.25$. Thus, the lower-left panels show the policy response to inflation when the parameter β is estimated rather precisely, while the upper-right panels consider policy responses under high to extreme uncertainty.

The following five findings are directly apparent from **Figure 1**. First, not surprisingly the certainty-equivalent and cautionary policy rules respond linearly to inflation for a given degree of uncertainty, while the optimal rule responds in a nonlinear fashion. Second, the

²³The thick dots correspond to the grid points used in the numerical approximation.

Figure 1: Optimal vs Cautionary and Certainty Equivalent Policy Rules
 Private sector agents in labor and financial markets form adaptive expectations



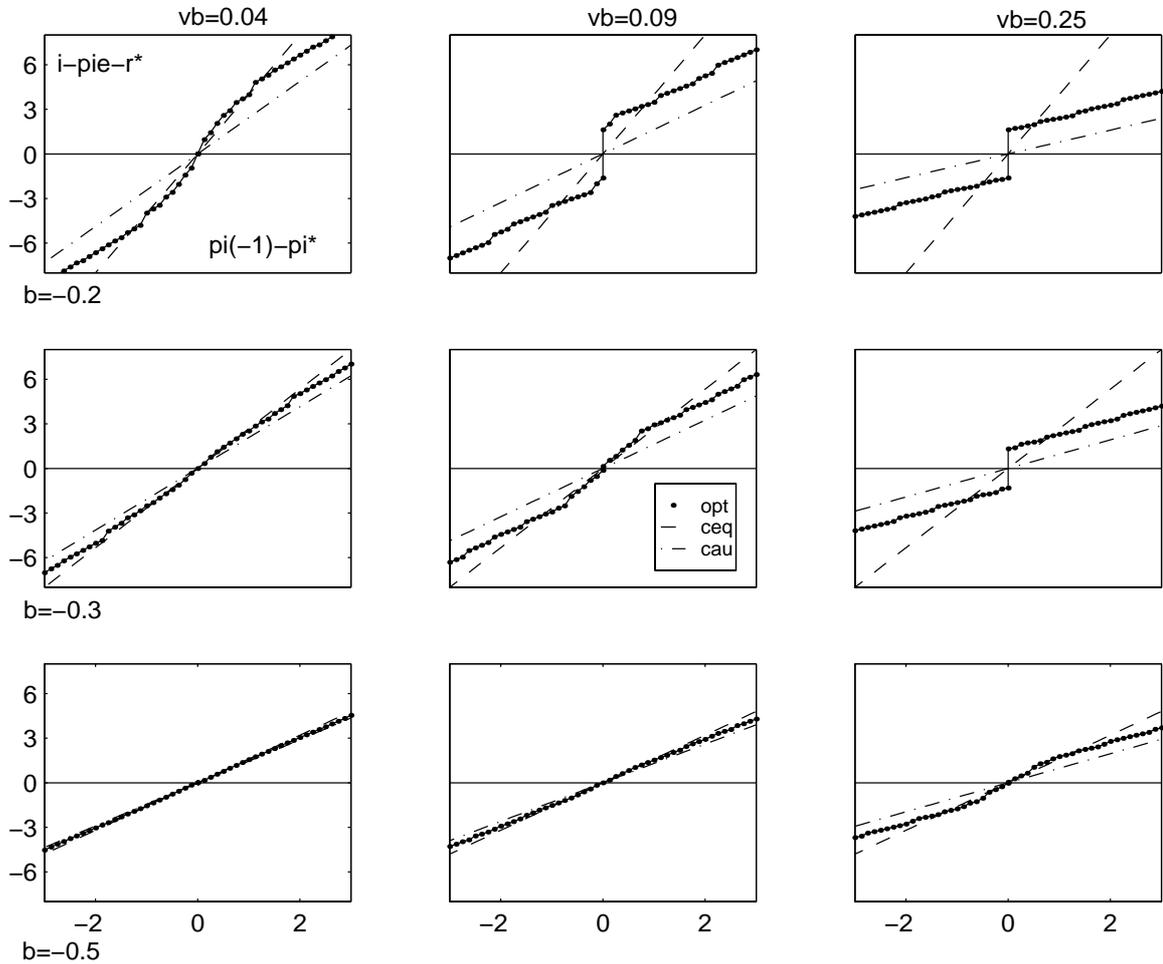
optimal rule always requires a more aggressive policy response to inflation than the cautionary rule. In absolute terms this difference, which represents the extent of experimentation incorporated in the optimal rule, changes little for moderate to high inflation deviations from target.²⁴ Third, the optimal rule typically implies a less aggressive policy stance than the certainty-equivalent rule that disregards parameter uncertainty. Thus, in spite of the

²⁴In other words, the relative importance of experimentation declines with the size of the inflation deviation from target. If inflation is substantially above target even the cautionary policy will result in a substantial policy response that will be expected to generate quite a bit of information about the inflation-unemployment tradeoff and the location of the natural rate.

incentive to experiment, the optimal policy exhibits gradualism. The policy tightening (or easing) following a shock to inflation is expected to persist for more than one period and implies a gradual return of inflation towards the target. Fourth, when lagged inflation is near the target (typically within less than 1 percentage point), the optimal policy response is somewhat more aggressive than the certainty-equivalent rule. Fifth, when uncertainty is extremely high (i.e. t -statistic < 1), the optimal policy rule exhibits a clear discontinuity at zero inflation. In other words, the optimal policy response at zero inflation differs from the neutral setting that would ensure an expected inflation rate equal to zero. It implies that the central bank accepts higher or lower expected inflation in order to obtain more precise parameter estimates and improve inflation stabilization in the future.

Next, we allow for rational learning by private sector agents in the labor market. Given symmetric information and policy commitment to a rule, rational forward-looking agents form the same beliefs about the unknown parameter β as the central bank. The index of persistence γ is set to 0.8 implying a weight of 0.2 on forward-looking expectations in the Phillips curve. The presence of forward-looking expectations adds an expectation channel of monetary policy transmission. As a result, the central bank can be less aggressive in responding to inflation deviations from target because private sector expectations take into account future policy action and move towards the inflation target. **Figure 2** provides a comparison of optimal, certainty-equivalent and cautionary rules for this case. Again the figure contains nine panels, with each panel corresponding to same combination of point estimate b and uncertainty v^b as in **Figure 1**. Three findings are directly apparent. First, all three policies are less aggressive in their response to lagged inflation due to the presence of an expectations channel for monetary policy. Second, the qualitative properties of the optimal rule are the same as under the case of adaptive expectations. The optimal extent of experimentation, that is, the difference between the cautionary and the optimal rule is slightly smaller. The intuitive reason is that with γ assumed known to the central bank, a change in inflation expectations has a direct, known effect on inflation. Thus, the central bank does not need to rely as much on the nominal interest rate, whose effect on inflation

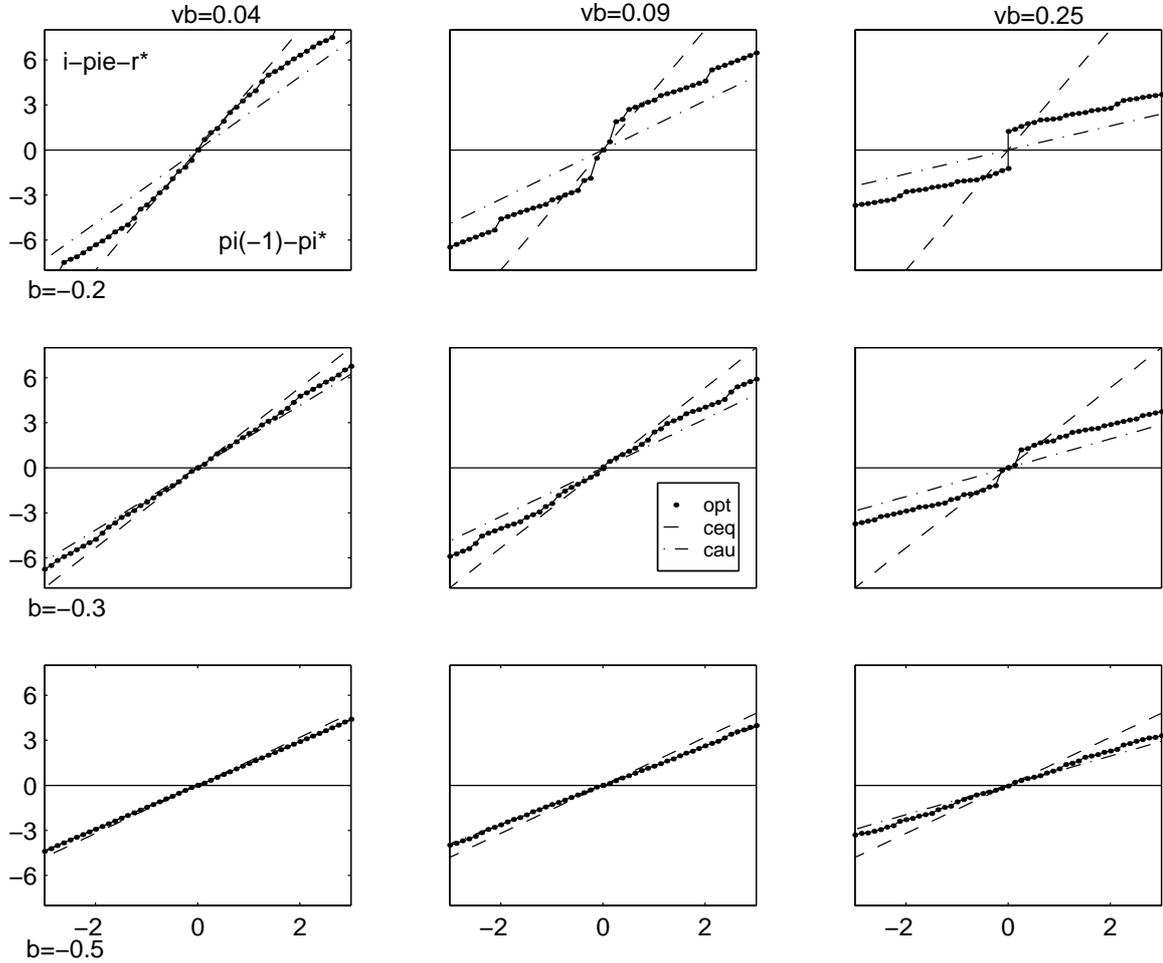
Figure 2: Optimal vs Cautionary and Certainty Equivalent Policy Rules
Rational learning in labor markets ($\gamma = 0.8$)



is imprecise due to uncertainty about β .

Finally, I also add rational learning by private sector agents in financial markets, that is, I incorporate forward-looking expectations of inflation in the definition of the real interest rate. The policy response to lagged inflation under the alternative rules is reported in **Figure 3**. Again, the optimal rule exhibits the same qualitative properties as previously. However, the presence of forward-looking expectations in financial markets further increases the power of the expectations channel of policy transmission. As a result,

Figure 3: Optimal vs Cautionary and Certainty Equivalent Policy Rules
Rational learning in labor markets ($\gamma = 0.8$) and financial markets



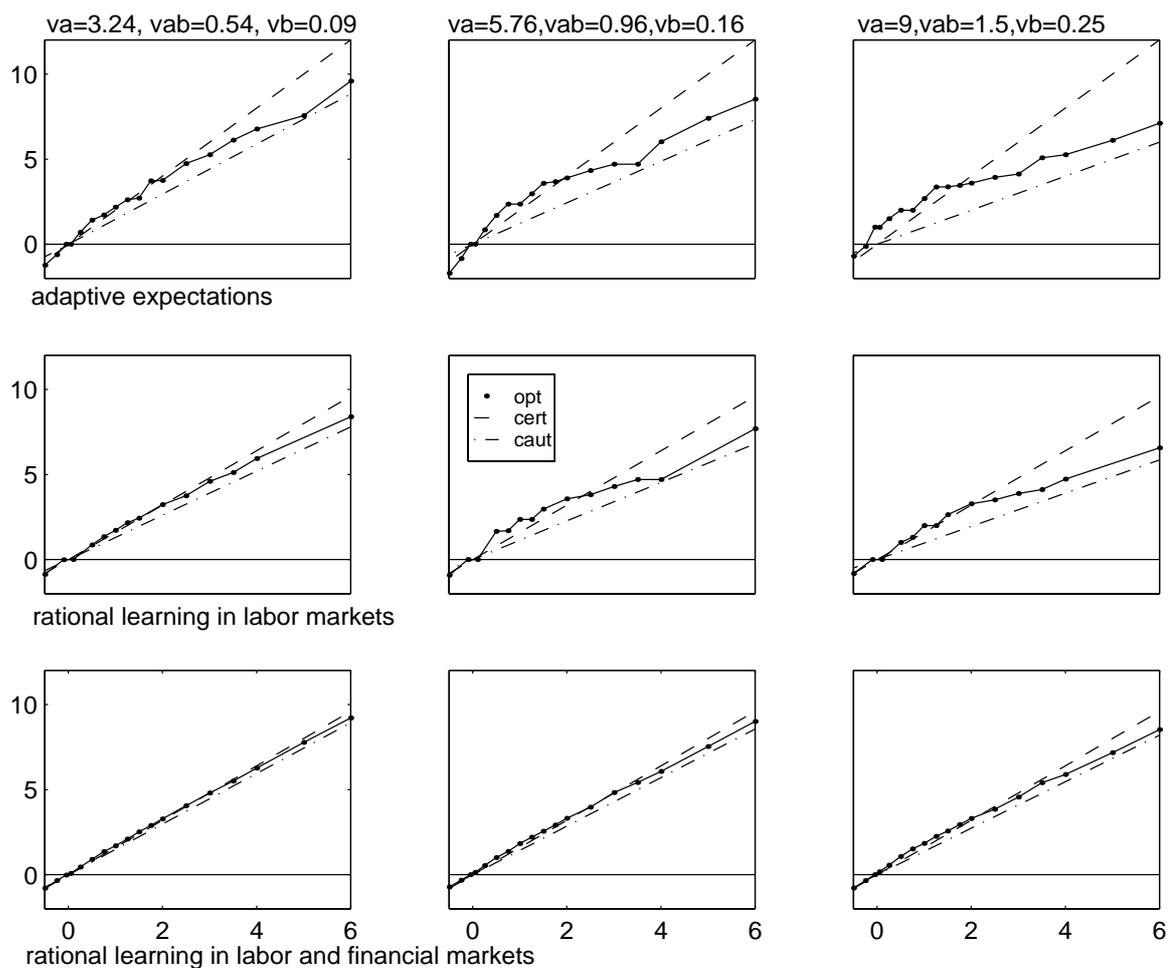
the policy rules are less aggressive and the optimal extent of experimentation is reduced.

Policy with uncertain α and β

Next, I consider the learning problem with unknown slope and intercept. This problem has six state variables. The dynamic programming algorithm described in the appendix provides numerical approximations for the value and policy functions over a wide range of the state space. To keep the number of charts in this section manageable I restrict

attention to a comparison of different policies for intercept and slope estimates, $a_{t-1} = 3.0$ and $b_{t-1} = -0.5$, respectively. The implied estimate of the natural unemployment rate is 6%.²⁵ Furthermore, the index of persistence will be set equal to 0.8.

Figure 4: Optimal vs Cautionary and Certainty Equivalent Policy Rules
Two Unknown Parameters: Intercept and Slope ($a = 3, b = -0.5$)



As discussed in section 3, even when lagged inflation is on target, the variance-minimizing setting of the nominal interest rate with two unknown parameters need not coincide with the neutral setting, which ensures future expected inflation to be on target.

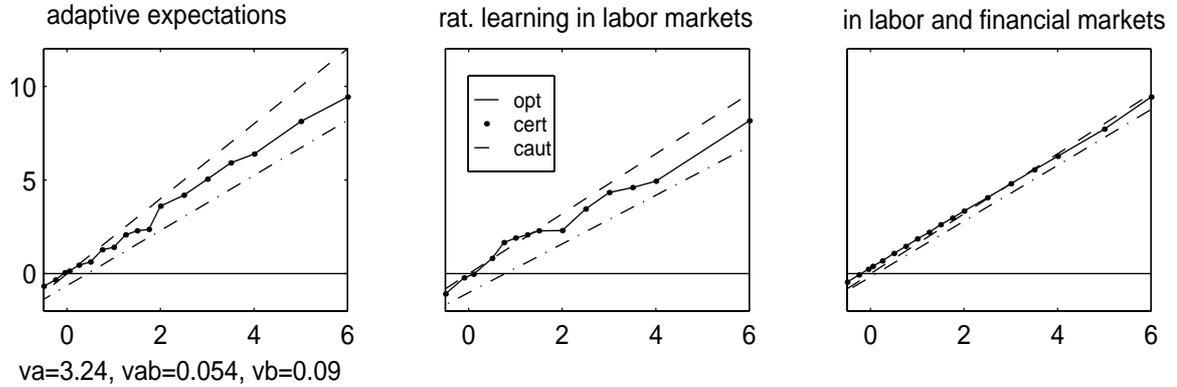
²⁵Results for alternative values of the slope and intercept estimates can be provided upon request.

In the following I present two sets of results. For the first set, displayed in **Figure 4**, the covariance v^{ab} is chosen exactly so that the variance-minimizing setting of the nominal interest rate coincides with the neutral setting when inflation is on target. Thus, the expected real interest rate and unemployment rate will coincide with their natural levels when lagged inflation equals the target of zero inflation. The second set of results, displayed in **Figure 5** considers values of the covariance which imply a variance-minimizing level of the interest rate below the natural level.

Figure 4 also contains nine panels. The first row of three panels refers to the case of adaptive expectations (as in Figure 1), the second row to the case of rational learning in labor markets, and the third row to the case of rational learning in labor and financial markets. Each column shows the same policy under alternative degrees of uncertainty. The variance of the slope estimate in each column corresponds to that in the respective columns of Figures 1, 2 and 3. For illustrative purposes, all panels focus on the policy response when inflation is above target since the response to negative inflation deviations is symmetric.

The first key result is that the findings from the learning problem with unknown slope concerning the optimal policy response to inflation carry over to the learning problem with unknown intercept and slope. The optimal response typically falls inside the wedge created by the aggressive certainty-equivalent rule and the cautionary rule. Optimal policy always incorporates a small extent of experimentation but mostly remains gradualist, that is, less aggressive than a certainty-equivalent rule that disregards parameter uncertainty. Exceptions to the principle of gradualism only occur near the inflation target or under extreme uncertainty. Even then, these exceptions are small in magnitude. The second result is that the stronger the expectations channel of monetary policy, the smaller the required policy response to inflation and thus the smaller the differences between the optimal, certainty-equivalent and cautionary rules. As discussed above a second source of gradualism may arise from the covariance of intercept and slope. This effect is illustrated by **Figure 5**. The three panels in this figure correspond to the first column of panels in **Figure 4**. The only difference is the value of the covariance v^{ab} . It is now set so that

Figure 5: Optimal vs Cautionary and Certainty Equivalent Policy Rules
 Two unknown parameters: Intercept and slope, covariance effect



the variance-minimizing level of the interest rate lies below its natural rate even when lagged inflation is on target. This occurs whenever the policy stance has been on average below the natural rate for the sample with which the parameters have been estimated. As a result, the cautionary rule implies a substantially easier policy stance than the certainty-equivalent rule, that is unaffected by the parameter covariance. However, this covariance effect is substantially reduced under the optimal rule indicating again a degree of experimentation.

Policy with uncertain β and γ

One key result emerging from the preceding comparisons is that rational learning by private sector agents tends to reduce the wedge between the certainty-equivalent and cautionary rules and consequently also the optimal extent of gradualism (i.e. the difference between the certainty-equivalent and optimal rules) and of experimentation (i.e. the difference between optimal and cautionary rules). To some degree this finding may depend on the fact that the parameter γ which determines the impact of inflation expectations on inflation and governs the strength of this expectations channel of monetary transmission has so far been treated as known with certainty.

Figure 6: Optimal vs Cautionary and Certainty Equivalent Policy Rules
 Two Unknown Parameters: Index of Persistence and Slope ($c = 0.8, b = -0.5$)

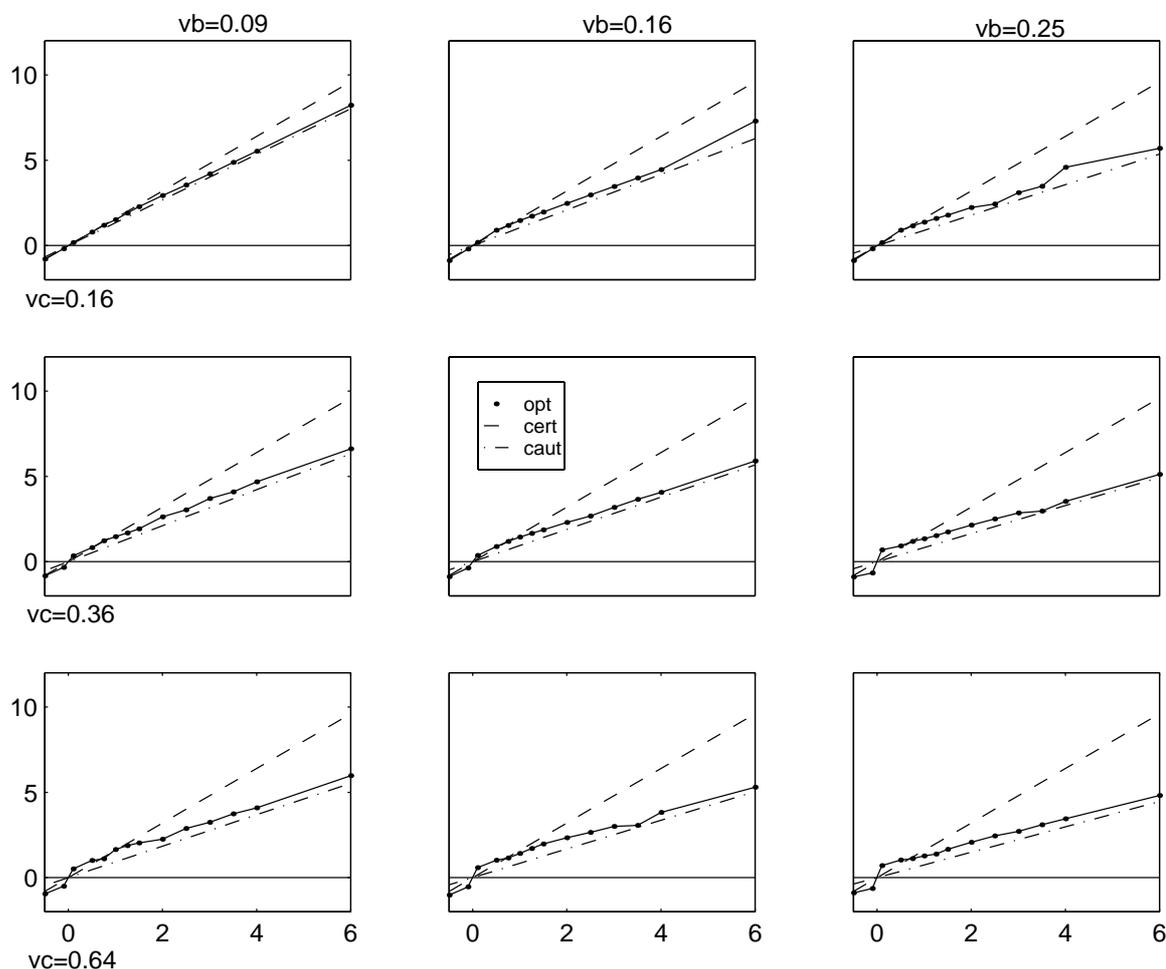


Figure 6 shows the alternative policies when both β and γ are uncertain. Again, there are 9 panels showing the policy response to inflation deviations from target under the three rules. Parameter estimates are $b_{t-1} = -0.5$ and $c_{t-1} = 0.8$ respectively. The three columns reflect the same values of the variance of the slope considered before, $v^b = (0.04, 0.09, 0.25)$, while the rows now refer to alternative values of $v^c = (0.16, 0.36, 0.64)$. The covariances range from -0.012 to -0.04 . The results shown are obtained for the model specification with rational learning in labor markets.

As can be seen from all panels the qualitative findings of the one-unknown parameter case, also survive in this context. Quantitatively, there is still a noticeable extent of experimentation, but the optimal rule remains almost always less aggressive than the certainty-equivalent rule, except under very high uncertainty near the inflation target.

6 Empirical Examples

The question remains to what extent the differences between optimal, cautionary and certainty-equivalent policy rules discussed in the preceding section are of quantitative importance given the degree of uncertainty reflected in empirical estimates of the Phillips curve. Estimating the complete model with rational learning goes beyond the objective of this paper. Instead, I relate the numerical analysis conducted here to some of the empirical estimates available in the literature.

The standard framework for estimating the inflation-unemployment relationship takes the following linear form:

$$\pi_t = a + \sum_{i=1}^I b_i u_{t-i} + \sum_{j=1}^J c_j \pi_{t-j} + dz_t + \epsilon_t \quad (29)$$

This regression equation usually includes several lags of the inflation rate and the unemployment rate and a vector z_t that contains proxy variables for supply shocks and various dummy variables. 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 unemployment rates, $\hat{u}^* = a(\sum_{i=1}^I b_i)^{-1}$.²⁶ Thus, the degree of uncertainty regarding NAIRU estimates discussed in the literature is directly related to the precision of estimates of the slope of the Phillips curve. As shown in preceding section, this type of uncertainty

²⁶An 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. However, 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.

is more important to monetary policy decisions than the component of NAIRU uncertainty that derives from the variance of the intercept a in the above regression equation.

Table 2: Phillips Curve Estimates

	a	$\sum b$	$\sum c$
Fuhrer (1995)			
Estimates	1.68	-0.28	1.0
Variance	0.58	0.015	
Covariance: v^{ab}	-0.09		
Staiger et al. (2002)			
Estimates		-0.28	1.0
Variance		0.01	

I consider estimates from two contributions to the ongoing debate on Phillips curves and NAIRU uncertainty in the United States (cf. Fuhrer (1995) and Staiger, Stock and Watson (2002)). Both papers report estimates of a version of the above regression equation (29) with a constraint that the coefficients on lagged inflation rates sum to one. This corresponds to $\gamma = 1$ in the model of this paper, and thus to the Phillips curve without forward-looking component. The estimates I consider are summarized in **Table 2**. The first set of estimates is taken from Fuhrer (1995), page 47, Table 1a.²⁷ The intercept estimate, its variance and its covariance with the sum of slope coefficients reported here have been computed using the complete regression results that I received from the author. As to Staiger et al. (2002), I only use the slope estimate and its variance from the first column of Table 1.2 on page 18.²⁸

Table 3 reports the optimal extent of gradualism (the difference between the optimal and certainty-equivalent rule) and of experimentation (the difference between optimal and cautionary rule) given the empirically estimated degree of uncertainty. The two columns

²⁷These estimates were obtained using quarterly data on the CPI excluding food and energy and the civilian unemployment rate from 1960:2Q to 1993:4Q. The author uses 12 lags of inflation and 2 lags of the unemployment rate as well as the oil price as supply shock proxy.

²⁸This regression uses the GDP Deflator as measure of prices and the civilian unemployment rate. It differs importantly from Fuhrer's specification in that the intercept and thus the NAIRU is time-varying.

refer to inflation deviations from target of 2% and 3% respectively. The values reported in the table concern the differences in the expected unemployment rates that result from the alternative policy rules.

Table 3: Gradualism and Experimentation
Differences in Expected Unemployment Rates

	$\pi_{t-1} - \pi^* = 2$	$\pi_{t-1} - \pi^* = 3$
Fuhrer (1995)		
CEQ-OPT	0.51	0.99
OPT-CAU	0.62	0.72
Staiger et al. (2002)		
CEQ-OPT	0.32	0.67
OPT-CAU	0.35	0.33

In both cases I restrict attention to the optimal policy under adaptive expectations in the Phillips curves, which fits the specifications chosen by those authors. This also implies that the unemployment differences reported above are independent of the Okun's law parameter λ and the interest rate sensitivity of aggregate demand ϕ . For Fuhrer's estimates I consider the optimal learning problem with unknown intercept and slope and take into account the covariance effect. For Staiger et al.'s estimates I consider the optimal learning problem with unknown slope. The optimal policy always lies in between the certainty-equivalent and cautionary rules. Thus, the certainty-equivalent rule is too aggressive in fighting inflation while the cautionary rule implies too much gradualism. The differences between the optimal policy and these alternatives are economically significant, ranging from one half to a full percentage point of the unemployment rate.

7 Conclusions and extensions

The bottomline of the preceding analysis is that an inflation-targeting central bank should not disregard uncertainty about the relationship between unemployment and inflation. Typically, it will be optimal to respond more gradually to inflationary shocks than a central

bank that disregards such uncertainty. However, gradualism can be overdone. In particular, a central bank that implements Brainard’s recommendation of gradualist policymaking in a myopic manner and disregards dynamic learning effects will respond too cautiously to inflationary shocks. A central bank that recognizes the tradeoff between current control and experimentation for the sake of reducing uncertainty and improving future policy performance will be a more aggressive inflation fighter than the central bank that implements Brainard’s recommendation myopically. However, this central bank will still act more gradually than one that disregards parameter uncertainty. Exceptions to this rule arise only when uncertainty is very high and at the same time inflation close to target.

The preceding analysis can be extended along several dimensions. Some of these extensions are straightforward while others are interesting avenues for future research. The remainder of this section discusses four such extensions.

Flexible inflation targeting

So far, the paper studied policy rules for a strictly inflation targeting central bank that focuses exclusively on inflation. However, the framework developed in this paper carries over to flexible inflation targeting with a loss function that includes deviations of unemployment from its natural rate:

$$L(\pi_t) = E_{t-1} \left[(\pi_t - \pi^*)^2 + \omega(u_t - u_t^*)^2 \right] \quad (30)$$

As is well known, the unemployment stabilization objective introduces an alternative motive for the central bank to respond gradually to inflationary shocks. This is directly apparent from a comparison of the optimal interest rate rules under certainty for strict inflation targeting, equation (11), and the one for flexible inflation targeting:

$$i_t == -((\beta\lambda\phi)^2 + \omega\gamma^2)^{-1} \left((\beta\lambda\phi\gamma - \omega\gamma^2)\pi_{t-1} - (\beta\lambda\phi - \omega\gamma)\beta u_{t-1}^* \right) \quad (31)$$

A central bank that assigns a positive weight ω to unemployment deviations will respond less aggressively to inflation deviations from target. As a result, inflation will be expected

to return more gradually to the target subsequent a shock. Policy rules under uncertainty can be computed in the same manner as for strict inflation targeting. However, due to the non-normal distribution of u^* one needs to resort to numerical methods even in the case of the cautionary rule that is optimal in the static version of the model without learning. The qualitative properties of the optimal rule under strict inflation targeting will survive under flexible inflation targeting, but of course, quantitative results will differ.

Optimal policy under discretion

Having considered policy choices for a central bank that is able to commit to a specific rule, it is of interest to explore optimal policy under discretion. Under discretion, the central bank will optimize policy taking private sector expectations as given. The private sector will try to minimize expectational errors taking the central bank's response to private sector expectations as given. The main purpose of Clarke et al. (1999) is to compare optimal policy under discretion and commitment. However, they assume that the parameters of the economy are known with certainty. The optimal policy under discretion may be derived as follows. First, one determines the interest rate that minimizes the central bank's loss function (1) treating the private sector agents expectation of the interest rate, i_t^e , and thus their inflation expectation π^e as constant and independent of monetary policy. With known parameters this corresponds to:

$$i_t = -(\beta\lambda\phi)^{-1}(\gamma\pi_{t-1} + (1 - \gamma\beta\lambda\phi)\pi_t^e - \beta u_{t-1}^* - \pi^*) \quad (32)$$

Private sector agents set i_t^e and thus π_t^e to minimize forecasts errors taking the central bank's response as given. The rational expectation of inflation taking (32) as given corresponds to the inflation target π^* . The nominal interest rate in this Nash equilibrium is equivalent to the interest rate rule under commitment derived in (11). Differences between optimal policy under discretion and commitment arise once a policy tradeoff is introduced, such as the tradeoff between inflation and unemployment in the case of flexible inflation targeting under certainty, or the tradeoff between the expected inflation deviation from

target and its conditional variance under parameter uncertainty as in section 3 of this paper. The cautionary policy under discretion can be derived analytically following the procedure suggested here. The computation of the optimal policy under discretion in the dynamic model with learning poses an additional complication of the numerical analysis that would be an interesting problem to address in future research.²⁹

Demand uncertainty

While studying the implications of uncertainty about the relationship between unemployment and inflation in much detail, the other key relationships of the model have been treated as certain in the preceding analysis. An extension of the model would allow for uncertainty due to random shocks $(\epsilon_t^u, \epsilon_t^y)$ and imprecisely estimated parameters in the Okun's law and aggregate demand relationships such as

$$u_t = \phi y_t + \epsilon_t^u \quad (33)$$

$$y_t = \lambda(i_t - \pi_t^e) + \epsilon_t^y \quad (34)$$

with beliefs regarding ϕ and λ characterized by normal distributions $N(p, v^p)$ and $N(l, v^l)$. The presence of the random shocks ϵ_t^u and ϵ_t^y renders estimation of ϕ and λ nontrivial. The resulting imprecision of the parameter estimates will further increase the component of inflation uncertainty that is influenced by monetary policy. Thus, it will enhance the motive for caution and widen the wedge between certainty-equivalent and cautionary policy rules. Increased parameter uncertainty will also tend to strengthen the incentive for experimentation and consequently the difference between optimal and cautionary rules. The demand-side shocks, however, imply some random variation in output and unemployment that will improve the estimates of the parameters of the inflation equation and will tend to reduce the incentive for experimentation. It is possible to derive the cautionary policy rule for the case when the Phillips curve parameters as well as ϕ and λ are imprecisely estimated, but the curse of dimensionality prevents the numerical analysis of optimal learning treating

²⁹Such an analysis would be related to the theoretical framework developed by Nyarko (1998).

all these parameters as jointly unknown. Nevertheless, the techniques presented in this paper can be used to derive optimal policy rules under uncertainty about ϕ and λ separately.

Asymmetric information and heterogenous beliefs

A key assumption maintained throughout this paper is that rational forward-looking agents and the central bank have the same information set available when making decisions. As result of this assumption, agents and central bank update their beliefs regarding the parameters of the inflation equation (α, β, γ) in the same manner. In practice, it is reasonable to assume that the central bank has an informational advantage compared to the public. This is more likely with regard to current estimates of the state of the economy and short-horizon forecasts than with regard to fundamental issues concerning the structure of the economy. Heterogenous beliefs about the parameters of the economy are more likely to arise to differences between agents and central banks in terms of their priors on reasonable parameter values or their view regarding the appropriate structural model.

It is straightforward to introduce an informational advantage of the central bank in terms of an advance signal e_t about the inflation shock ϵ_t into the model of this paper. The optimal policy rule will then include a policy response to this signal quite similar in qualitative terms to the policy response to lagged inflation. The private sector's rational expectation of this component of the interest rate rule will be equal zero. Thus, private sector agents will only be able to predict the component of the policy rule that responds to lagged inflation. Since the signal e_t does not help in estimating the Phillips curve parameters, belief updating equations will remain the same for the central bank and the private sector.

The possibility of heterogenous beliefs due to differences in priors or in the reference model represents a particularly interesting area for future research. One difficulty in this regard is that one will need to keep track of the central bank's and private agents' beliefs separately. This will substantially increase the state space of the optimal learning problem. Nevertheless, it should be feasible to study a problem with one unknown parameter and two sets of alternative beliefs using the techniques developed in this paper.

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.
- Balvers, R. and T. Cosimano (1994), "Inflation Variability and Gradualist Monetary Policy," *Review of Economic Studies*, Vol. 61, pp. 721-738.
- 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.
- 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.
- Clarida, R., Gali, J. and M. Gertler (1999), "The Science of Monetary Policy", *Journal of Economic Literature*, 37, 1661-1707.
- Clarida, R., Gali, J. and M. Gertler (1998), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory", *Quarterly Journal of Economics*, 42, 1003-1067.
- Clark, P., Goodhart, C. and H. Huang (1999), "Optimal Monetary Policy Rules in a Rational Expectations Model of the Phillips Curve", *Journal of Monetary Economics*, 43, 497-520.
- Debelle, G. and D. Laxton (1997), "Is the Phillips Curve Really a Curve ? Some Evidence for Canada, the United Kingdom, and the United States", *IMF Staff Papers*, 44, 249-282.

- Dornbusch, R. and S. Fischer, (1990), Intermediate Macroeconomics.
- Easley, D. and N.M. Kiefer (1988), "Controlling a Stochastic Process with Unknown Parameters", *Econometrica* 56, 1045-1064.
- Estrella, A., and F. Mishkin (1999), "The Role of NAIRU in Monetary Policy: Implications of Uncertainty and Model Selection" in Taylor, J.B., ed., *Monetary Policy Rules*, NBER and Chicago Press.
- Fuhrer, J. (1997), "The (Un) Importance of Forward-Looking Behavior in Price", *Journal of Money, Credit and Banking*, 29, (3), 338-50.
- Fuhrer, J. (1995), "The Phillips Curve is Alive and Well," *New England Economic Review*, (March/April), pp. 41-56.
- Fuhrer, J. and G. Moore (1995a), "Monetary Policy Trade-offs and the Correlation Between Nominal Interest Rates and Real Output", *American Economic Review*, 85: 219-239.
- Fuhrer, J. C. and G. R. Moore (1995b), "Monetary policy trade-offs and the correlation between nominal interest rates and real output," *American Economic Review*, 85, 219-239.
- Gali, J. and M. Gertler (1999), "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 44 (2), 195-222.
- Gordon, R. (1997), "The Time-Varying NAIRU and its Implications for Economic Policy", *Journal of Economic Perspectives*, 11(1):11-32.
- Hansen, L. and T. Sargent (2001), Robust control and model uncertainty, *American Economic Review*, 91 (2).
- Kasa, K. (1999), Will the Fed Ever Learn?, *Journal of Macroeconomics*, 21, 279-292.
- Keller, G. and S. Rady (1999), Optimal Experimentation in a Changing Environment, *Review of Economic Studies*, 66 (3).
- 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.
- Levin, A., Williams, J.C., and V. Wieland (2002), "The performance of forecast-based monetary policy rules under model uncertainty," accepted for publication, *American Economic Review*.
- Levin, A., Williams, J.C., and V. Wieland (1999) Robustness of simple monetary policy rules under model uncertainty, in John B. Taylor (ed.), *Monetary Policy Rules*, NBER and

Chicago Press, June 1999.

McCallum, B. (1999) , “Issues in the Design of Monetary Policy Rules,” Handbook of Macroeconomics, ed. by J. B. Taylor and M. Woodford, North-Holland.

Mizrach, B. (1991), Non-Convexities in a Stochastic Control Problem with Learning, Journal of Economic Dynamics and Control, Vol. 15, pp. 515-538.

Nyarko, Y. (1998), “The truth is in the eye of the beholder: or Equilibrium in beliefs and rational learning in games, manuscript, New York University.

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.

Roberts, J. (1997), “Is inflation sticky?,” *Journal of Monetary Economics*, 39, 173-196.

Rotemberg, J. and M. Woodford (1997), “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” NBER Macroeconomics Annual 1997.

Sack, B. (1999), “Uncertainty and Gradual Monetary Policy,” *Journal of Monetary Economics*, 46 .

Sargent, T. (1999a), “Discussion of ‘Policy Rules for Open Economies’ by Laurence Ball”; in Taylor, J. (ed.), “Monetary Policy Rules”, NBER and Chicago Press.

Sargent, T. (1999b), “The Conquest of American Inflation,” Princeton University Press.

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.

Staiger, D., J. Stock, and M. Watson (2002), “Prices, Wages and the U.S. NAIRU in the 1990s, in, A. Krueger, and R. Solow (eds.), *The Roaring Nineties*, Russell Sage Foundation/The Century Fund: New York.

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), “Optimal inflation targets, ‘conservative’ central banks, and linear inflation contracts,” *American Economic Review* 87, 98-114.

Svensson, L. (1999), “Inflation targeting: Some extensions,” *Scandinavian Journal of Economics*, 101, 3, 337-361.

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.

Taylor, J.B. (1999), ed., “Monetary Policy Rules”, NBER and Chicago Press.

Wieland, V. (2000a), “Learning by Doing and the Value of Optimal Experimentation, Finance and Economics Discussion Series,” *Journal of Economic Dynamics and Control*, 24, 501-534.

Wieland, V. (2000b), “Monetary Policy, Parameter Uncertainty and Optimal Learning,” *Journal of Monetary Economics*, 46, 199-228.

Appendix A: The Numerical Dynamic Programming Algorithm

The algorithm used in this paper computes the value function and stationary optimal policy by iterating over the Bellman equation, which defines the following contraction mapping:

$$TW = \underset{i}{Min} \left[L(\pi, i, \theta) + \delta \int W(\pi', i, \theta') f(\pi' | \pi, i, \theta) d\pi' \right] \quad (35)$$

where T stands for the functional operator and π and θ are last period’s values of the inflation rate and the beliefs about the unknown parameters, that is the state variables of the problem. $W(\cdot)$ is a continuous function defined on the state space. $L(\cdot)$ denotes the expected current loss. The control variable i corresponds to the central bank’s policy instrument. π' is the inflation rate to be realized subsequent the policy action and θ' refers to the beliefs at the end of the period based on new inflation and unemployment observations. The relevant updating equations for these state variables are (13) and (25). Inflation expectations are solved out according to (18) and expressed in terms of the state variables. $f(\pi' | \pi, \cdot, \theta)$ is the predictive distribution of the inflation rate. It is a normal distribution, because both the error terms and the beliefs are normal distributions.

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:

$$\rho(W_n, W_{n+1}) = \underset{(\theta, \pi)}{Sup} |W_n(\theta, \pi) - W_{n+1}(\theta, \pi)| \quad (36)$$

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:

$$\rho(TW_{n+1}, TW_n) \leq \delta \rho(W_{n+1}, W_n) \quad (37)$$

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 (35) 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 :

$$\rho(V, W^n) \leq \sum \rho(W^{n+1+i}, W^{n+i}) \leq \frac{\delta}{1-\delta} \rho(W^n, W^{n-1}) \quad (38)$$

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:

$$T^P W_n = L(\pi, H_n(\pi, \theta), \theta) + \delta \int W(\pi', H_n(\pi, \theta), \theta') f(\pi' | \pi, H_n(\pi, \theta), \theta) d\pi' \quad (39)$$

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:

$$\int W_0(\pi', u, \theta') f(\pi' | \pi, i, \theta) d\pi' \quad (40)$$

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 substantial computational effort largely because of the so-called curse of dimensionality. The largest problem considered had six state variables. If each of the six state variables is approximated with a grid of N gridpoints, 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.

Several 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) a convenient reformulation of the problem allows the reduction of the state space by one state variable, which means that the integration and optimization steps only have to be carried out N^5 times per value iteration;³⁰ (iii) the algorithm is written in FORTRAN so as to reduce computation time relative to higher-level languages such as MATLAB.

The most time-consuming problems computed in this paper are those with two unknown parameters. The largest grid used in this case consisted of $13 * 15 * 19 * 16 * 29$ gridpoints. In this case, I also used 60-point Gaussian quadrature with respect to the shock ϵ . Convergence as defined by a 0.5% maximal difference between the two final value function approximations for these problems was achieved after about 60 hours on a 2.4 GHz Intel Pentium 4 Chip with 1 MB RAM. Typically convergence required 6 to 8 value iterations with a declining number of policy iterations (50 or less) in between every value iteration.

Appendix B: Convergence of Beliefs and Policies

Although the paper so far discussed in detail optimal policies with dynamic learning, it avoided the question whether the central bank and private agents 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 several ways: (i) the intercept may be time-varying, (ii) the regression includes a lagged dependent variable, and (iii) the regression includes an expectation of the 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 the following I discuss convergence for the formulation with adaptive expectations of the private sector $\pi_t^e = \pi_{t-1}$ but the argument can be generalized to the case with rational learning on the part of private agents. Under adaptive expectations the coefficient on lagged inflation in the Phillips curve corresponds to 1, and α and β can be estimated by means of this simple regression

$$\Delta\pi_t = \alpha - \beta u_t + \epsilon_t \tag{41}$$

with the change rather than the level of inflation as dependent variable.

The parameter estimates and covariance matrix are updated according to an appropriately simplified version of (25). This 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

³⁰For a discussion of this reformulation see the appendix of Wieland (2000a).

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 $E_{t-1}[\Delta\pi_t - a_{t-1} + b_{t-1}u_t] = 0$,³¹ it follows that $E_{t-1}[a_t] = a_{t-1}$ and $E_{t-1}[b_t] = b_{t-1}$.

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 (2000a) and (2000b).³²

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 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 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 (41) converges with probability one to a limit belief $(a_\infty, b_\infty, \Sigma_\infty)$ as $t \rightarrow \infty$. For any given belief, the unemployment rate 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:

$$u_t = \frac{a_\infty}{b_\infty} + \frac{b_\infty}{b_\infty^2 + v_\infty^b}(\alpha + \beta u_{t-1} + \epsilon_{t-1} - \pi^*) + \frac{(v_\infty^{ab} - v_\infty^b \frac{a_\infty}{b_\infty})}{(b_\infty^2 + v_\infty^b)} \quad (42)$$

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 $t \rightarrow \infty$.

³¹This is true because u_t is a deterministic function of i_t and effectively part of the central bank's information set at $t-1$.

³²Wieland (2000a) 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 (2000b) 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.

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:

$$\begin{aligned} v_t^b &= \frac{\sigma_\epsilon^2}{\sum_{i=0}^t (u_i - \bar{u}_t)^2} \\ v_t^{ab} &= \bar{u}_t v_t^b \\ v_t^a &= \frac{\sigma_\epsilon^2}{t} + \bar{u}_t^2 v_t^b \end{aligned} \tag{43}$$

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.