Monetary Policy under Uncertainty about the Natural Unemployment Rate:

Brainard-Style Conservatism versus Experimental Activism *

Volker Wieland

Goethe University of Frankfurt, CFS and CEPR

This Version: March 2006

Abstract

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 tradeoff and the degree of inflation persistence in a simple macroeconomic model that incorporates rational learning by the central bank as well as market participants. Two conflicting motives drive the optimal policy. In the static version of the model, uncertainty provides a motive for the policymaker to move more cautiously than she would if she knew the true parameters. In the dynamic version, uncertainty also motivates an element of experimentation in policy. The optimal policy, which balances the cautionary and activist motives, is computed using empirical estimates of Phillips curve uncertainty as a benchmark. The effect due to experimentation is of quantitative relevance for moderate to high degrees of uncertainty. However, gradual inflation stabilization typically remains optimal, that is, the optimal policy response to inflation is still less aggressive than a policy that disregards parameter uncertainty. Exceptions occur when uncertainty is very high and inflation close to target.

JEL Classification System: E52, E24, D8, C61.

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

---

* Correspondence: Professur für Geldtheorie und -politik, Johann-Wolfgang-Goethe Universität, Mertonstrasse 17, D-60325 Frankfurt am Main, Germany, tel.: +49 69 798-25288, e-mail: wieland@wiwi.uni-frankfurt.de, homepage: http://www.volkerwieland.com.

This paper has benefitted substantially from comments by B.E. Press Journals in Macroeconomics editor David Romer and two anonymous referees. Also, I would like to thank Jeffrey Fuhrer and Mark Watson for providing details regarding their estimation results. Any remaining errors are the sole responsibility of the author. Two earlier working paper versions of this article appeared as CEPR Discussion Paper No. 3811 in 2003 and FEDS Working Paper No. 1998-22 in 1998 under the title 'Monetary Policy and Uncertainty about the Natural Unemployment Rate'. Helpful comments on those versions by John Leahy, Andrew Levin, Athanasios Orphanides, Richard Porter, Matthew Shapiro, Peter von zur Muehlen, Carl Walsh and David Wilcox are gratefully acknowledged. This paper is dedicated to Professor Reinhard Hujer of Johann-Wolfgang-Goethe University of Frankfurt at the occasion of his retirement. This dedication is meant to honor his exemplary achievements as a serious and enthusiastic researcher in empirical labor economics. He stood out in German academia setting an example for highest-quality econometric analysis and achieving international recognition in his field.

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, that is, the Phillips curve. 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. Fuhrer (1995), for example, states that *"the Phillips curve is alive and well"* as an empirical relationship in the United States economy, while Staiger, Stock and Watson (1997a, 1997b, 2002) emphasize that a typical 95% confidence interval for the natural rate is about 2.5 percentage points wide.[3] Of course, 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 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 degree of inflation persistence in the

[1] As 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.

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

[3] For the recent debate on the U.S. economy, see also Gordon (1997), Blanchard and Katz (1997), Akerlof, Dickens and Perry (1996) and 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.

Phillips curve. 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 sector 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-style 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 policy[5] 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

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

[5] See for example Sargent (1999a) and 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 include Bertocci and Spagat (1993), Balvers and Cosimano (1994), Wieland (2000b), Ellison and Valla (2001) and Yetman (2003). Analytical results concerning optimal policy are largely absent from the literature and numerical results are rare, because of the nonlinear nature of the dynamic learning problem. Among these studies, Wieland (2000b) analyzes the most general learning problem—a linear regression with two unknown parameters—and computes the optimal policy numerically.[7] Using this approximation Wieland (2000b) then concentrates on quantifying the likelihood of incomplete learning.[8] The problem studied in the present paper is further complicated by the presence of a lag as well as a forward-looking expectation of the dependent variable. Contrary to the model in Wieland (2000b), the model in the present paper exhibits complete learning of the true values of constant parameters, because the central bank's continued efforts to stabilize the economy following inflationary shocks generate sufficient information about the inflation-unemployment tradeoff in the long run. Therefore, this paper concentrates instead on quantifying the extent of experimentation and gradualism implied by alternative policies for an empirically plausible range of prior beliefs.

This paper makes the following contributions. First, extending Brainard's original analysis, I derive the Brainard-style cautionary policy rule in the presence of forward-looking

[6] One part of the literature focused 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 focused on characterizing optimal decision rules (cf. Prescott (1972), Kendrick (1981), Kendrick (1982), Mizrach (1991), Amman and Kendrick (1995), Keller and Rady (1999) and Wieland (2000a)).

[7] Ellison and Valla (2001) study discretionary monetary policy with uncertainty about the slope of the Phillips curve, but they only consider a very stylized learning problem with two possible parameter values similar to the illustrative example in Wieland (2000a). Yetman (2003) revisits the model of Wieland (1998) considering alternative degrees of inflation persistence, but he only studies a simplified one-period learning problem.

[8] Asymptotic properties of beliefs and policies in this framework have been studied by Easley and Kiefer (1988) and Kiefer and Nyarko (1989), who have shown that incomplete learning may occur. 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.

behavior by market participants. 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 parameters of the Phillips curve. The 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.

Second, the paper presents numerical results concerning the optimal policy in a dynamic model with rational learning by the central bank.[9] Using empirical estimates of the Phillips curve with adaptive expectations by Fuhrer (1995) and Staiger, Stock and Watson (2002) as a benchmark, I quantify the optimal degree of gradualism and experimentation. I find that the optimal policy incorporates a quantitatively significant degree of experimentation—meaning a more aggressive policy response than under cautionary Brainard-style 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.

Third, the paper investigates the influence of forward-looking expectations formation and rational learning by market participants on optimal central bank policy. The qualitative conclusions remain the same under rational learning as under adaptive expectations. However, the optimal extent of experimentation is smaller, because forward-looking behavior by market participants introduces an expectations channel of monetary policy transmission.

Fourth, 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

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

policy.[10] This literature has focused on evaluating the performance of monetary policy rules in different macroeconometric models under the assumption that all parameters are known with certainty. Here, I show 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.

The next section documents empirical estimates of the Phillips curve from Fuhrer (1995) and Staiger et al. (2002) and introduces a theoretical specification to be used in the subsequent policy analysis. Section 3 completes the macroeconomic model and derives the optimal policy rule of an inflation-targeting central bank when the relevant model parameters are known. In section 4, the cautionary Brainard-style policy rule is derived analytically in a static version of the model. Section 5 presents the dynamic framework with learning. A quantitative comparison of the optimal, cautionary and certainty-equivalent policy rules is provided in sections 6 to 8. Section 9 concludes and discusses avenues for future research. The numerical algorithm used in this paper is described in more detail in the appendix.

2 The Phillips curve: Two empirical examples and a theoretical specification

A standard empirical specification of the Phillips curve, which is estimated in many of the econometric studies mentioned in the introduction, takes the following linear form:

$$\pi_t = \alpha + \sum_{i=0}^{I} \beta_i u_{t-i} + \sum_{j=1}^{J} \gamma_j \pi_{t-j} + \kappa z_t + e_t. \tag{1}$$

This regression equation includes several lags of the unemployment rate, u_t, and the inflation rate, π_t, as well as a vector z_t containing proxy variables for supply shocks and various dummy variables, and a random noise term, e_t. Regression parameters are denoted by Greek letters. In the following analysis, I will abstract from the proxy and dummy variables, z_t, and focus on estimates of the intercept, α, the sum of slope coefficients, $\beta = \sum_{i=0}^{I} \beta_i$, and the sum of coefficients on lagged inflation, $\gamma = \sum_{j=1}^{J} \gamma_j$. The estimates will be denoted

[10] See for example Taylor (1993, 1999), Rotemberg and Woodford (1997), McCallum (1999), Clarida, Gali and Gertler (1998, 1999), Levin, Wieland and Williams (1999, 2003) and many others.

by Roman letters, i.e. a for the estimate of the intercept, $b = \sum_{i=1}^{I} b_i$ for the slope of the Phillips curve, and $c = \sum_{j=1}^{J} c_j$ for the degree of inflation persistence. An estimate of the natural unemployment rate can be obtained from the ratio of the negative of the regression constant and the sum of the coefficients on current and lagged unemployment rates, $-ab^{-1}$.[11] Thus, the degree of uncertainty regarding NAIRU estimates discussed in the literature depends on the precision of the slope estimate of the Phillips curve.

Table 1: Phillips Curve Estimates

Source	a	b	c	v^a	v^b	v^{ab}
Fuhrer (1995)	1.68	-0.28	1.0	0.58	0.015	-0.09
Staiger et al. (2002)		-0.28	1.0		0.01	

Notes: b and c refer to the sums of point estimates $\sum b_i$ and $\sum c_j$. v^a, v^b and v^{ab} denote the variances and covariances of the estimates where available.

To pick an example from the literature on the U.S. economy, I turn to Fuhrer (1995) who estimates equation (1) with a constraint that the coefficients on lagged inflation rates sum to one, i.e. $c = \sum_{j=1}^{J} c_j = 1$. Thus, the overall degree of inflation persistence is imposed rather than estimated. The specification satisfying this constraint is often referred to as the 'accelerationist' Phillips curve. Fuhrer's estimates are reported in the first row of **Table 1**.[12] This row also contains the variance of the intercept estimate, v^a, the variance of the sum of slope estimates, v^b, and their covariance, v^{ab}.[13] The NAIRU estimate obtained by

[11] An 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 regressors and normal errors.

[12] These estimates are taken from Table 1a on page 47 of Fuhrer (1995). They 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, 2 lags of unemployment, and, in addition, the price of oil as a proxy for supply shocks.

[13] I have computed the variances and covariance of the sums of coefficient estimates using the complete regression results sent to me by the author.

Fuhrer (1995) is 6% with a standard error of .56. For comparison, Table 1 also reports an estimate of the slope of the Phillips curve and its variance from Staiger, Stock and Watson (2002).[14]

In the remainder of this paper, I will study optimal monetary policy design under uncertainty about the Phillips curve treating it effectively as stable under a range of alternative policy rules. However, as emphasized by the Lucas critique, equation (1) need not necessarily be stable with respect to alternative policies in spite of its empirical success. For example, the sum of lags of inflation, $\sum_{j=1}^{J} \gamma_j \pi_{t-j}$, may at least partially reflect forward-looking expectations by market participants. Forward-looking market participants would take into account systematic changes in policy in forming their expectations. To allow for this possibility, I will extend the Phillips curve used in the subsequent policy analysis to the 'backward and forward-looking components' model of Buiter and Miller (1985) and Clark, Goodhart and Huang (1999):[15]

$$\begin{aligned} \pi_t &= \beta(u_t - u^*) + \gamma\pi_{t-1} + (1-\gamma)\pi_t^e + e_t, \quad \text{where } \beta < 0, 0 < \gamma \leq 1, \qquad (2) \\ &= \alpha + \beta u_t + \gamma\pi_{t-1} + (1-\gamma)\pi_t^e + e_t, \quad \text{where } \alpha = -\beta u^*. \end{aligned}$$

Here, the long-run equilibrium rate of unemployment is denoted by u^*. It is also the non-accelerating-inflation rate of unemployment (NAIRU). 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 and to a normally-distributed random shock $e_t \sim N(0, \sigma^2)$. Furthermore, current inflation depends on lagged inflation π_{t-1} as well as price setters' expectation of inflation π_t^e. The backward-looking component reflects inertia in inflation that may be derived from some types of overlapping wage contracts (cf. Fuhrer and Moore (1995)) or may be attributed to the presence of rule-of-thumb price setters. The degree of inflation inertia is determined by the index parameter γ. The parameter on the forward-looking component is set at $1-\gamma$ so

[14] This estimate is taken from the first column of Table 1.2 on page 18 of Staiger et al. (2002). They use the GDP deflator as measure of prices and the civilian unemployment rate. Contrary to Fuhrer's specification of the Phillips curve, their specification allows for a time-varying intercept and thus a time-varying NAIRU.

[15] Contrary to these authors, however, I use the unemployment gap rather than the output gap in the Phillips curve to match the empirical studies noted above.

that the two components sum to unity.

With $\gamma = 1$ equation (2) simplifies to the accelerationist Phillips curve estimated by Fuhrer (1995). Alternatively, with the assumption of *adaptive (random-walk) expectations formation*, $\pi_t^e = \pi_{t-1}$, equation (2) simplifies to the accelerationist Phillips curve for any value of γ. Fuhrer's (1995) estimate of this specification forms the starting point for the comparative quantitative policy analysis in the remainder of this paper. Building on this benchmark an assessment of the consequences of forward-looking, rational expectations formation is obtained for the intermediate case with $0 < \gamma < 1$.[16]

3 Optimal monetary policy when the parameters of the Phillips curve are known

A central bank that pursues a strict inflation-targeting strategy as defined by Svensson (1997a) chooses its main policy instrument, the short-term nominal interest rate i_t, so as to minimize the discounted sum of expected inflation deviations from its inflation target π^*:[17]

$$\underset{[i_t]_{t=0}^{\infty}}{\text{Min}} \; E\left[\sum_{t=1}^{\infty} \delta^{t-1}(\pi_t - \pi^*)^2 | \pi_0\right]. \tag{3}$$

δ refers to the central bank's discount factor. The expected *per-period* loss denoted by $L(\pi_t)$ in the following can be decomposed into two terms indicating the possibility of a tradeoff between the conditional expectation of inflation deviations from target and the conditional variance of inflation:

$$L(\pi_t) = E_{t-1}\left[(\pi_t - \pi^*)^2\right] = \left(E_{t-1}\left[\pi_t - \pi^*\right]\right)^2 + VAR_{t-1}\left[\pi_t\right]. \tag{4}$$

[16] Empirical estimates of this structural Phillips curve specification tend to indicate a significant degree of inflation persistence with γ ranging from 0.5 to near unity (cf. Fuhrer (1997) and Roberts (1997)). The New-Keynesian Phillips curve, which has received much interest in the recent literature (cf. Gali and Gertler (1999)), differs from the above specification in the timing of the forward-looking inflation term, which concerns period $t+1$, and in the use of the output gap rather than the unemployment gap. The preferred empirical specification of the New-Keynesian Phillips curve also embodies a significant degree of inflation persistence.

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

In addition to the Phillips curve equation (2) two more equations are needed to complete a simple model of the transmission of monetary policy from the nominal interest rate to inflation. Using a version of "Okun's Law,"[18] the unemployment rate is related to aggregate demand, q_t. Aggregate demand in turn depends on the short-term real interest rate, i.e. on the difference between the nominal rate and expected inflation and thereby on the policy instrument:

$$u_t = \phi q_t, \tag{5}$$

$$q_t = \lambda(i_t - \pi_t^e), \tag{6}$$

$$\text{where } \phi = \lambda = -1 \text{ in the following.}$$

In the remainder of the paper the notation is simplified by setting the Okun's law parameter, ϕ, and the interest-rate sensitivity of aggregate demand, λ, equal to negative unity. For these values the short-term real interest rate can simply be substituted for the unemployment rate in the Phillips curve equation (1).[19]

In order to render the exposition as straightforward as possible, optimal policy is first derived in a static version of the model where the central bank is only concerned with current expected loss (i.e. $\delta = 0$), and then in the dynamic version (i.e. $0 < \delta < 1$). In addition, two specifications of market participants' expectations of inflation, π_t^e, will be considered separately: first, the case of adaptive, random walk expectations corresponding to the benchmark specification of Fuhrer (1995), $\pi_t^e = \pi_{t-1}$, and then the case of forward-looking, rational expectations, $\pi_t^e = E_{t-1}[\pi_t]$.

Optimal policy in the static model with adaptive inflation expectations

In the static version of the model the central bank sets the nominal interest rate i_t so as to minimize current expected loss $L(.)$ based on its knowledge of the current state and the

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

[19] This notation seems appropriate given that the paper focuses on the policy impact of uncertainty regarding the parameters of the Phillips curve only. The reader interested in the impact of alternative values for ϕ and λ is referred to the earlier working paper version Wieland (2003). In addition, an extension considering uncertainty with respect to these parameters is briefly discussed in the concluding section.

parameters of the economy. As indicated by equation (4) current expected loss contains two elements, the central bank's conditional expectation of inflation as well as the conditional variance of inflation. With market participants' expectations equal to lagged inflation the central bank's conditional expectation of inflation corresponds to:

$$E_{t-1}[\pi_t] = \alpha + \beta(i_t - \pi_{t-1}) + \pi_{t-1} = \beta(i_t - \pi_{t-1} - u^*) + \pi_{t-1}. \tag{7}$$

The conditional variance turns out to be independent of policy i_t and equal to the exogenous variance of the random shock e_t:

$$VAR_{t-1}[\pi_t] = \sigma^2. \tag{8}$$

Thus, the central bank will be able to minimize $L(.)$ 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 approach has been termed "inflation forecast targeting" 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 (8). The implied optimal interest rate rule is:

$$i_t = u^* + \pi_{t-1} - \frac{1}{\beta}(\pi_{t-1} - \pi^*). \tag{9}$$

The first term essentially represents the equilibrium real interest rate, which is related to equilibrium unemployment. The second term stands for market participants' adaptive inflation expectations. It ensures that policy moves sufficiently to achieve the desired change in the real interest rate, $i_t - \pi_{t-1}$. The third term represents the central bank's response to past inflation deviations that is intended to return inflation to its target value in the next period. The neutral setting of the nominal interest rate when inflation is on target corresponds to: $i_t = u^* + \pi^*$.

Optimal policy in the static model with rational inflation expectations

If market participants form expectations rationally, they will take into account the state and the parameters of the economy including the policy rule pursued by the central bank.

Thus, substituting equations (5) and (6) in (2), market participants' inflation expectations correspond to:

$$\begin{aligned}
\pi_t^e &= E_{t-1}[\pi_t] \\
&= \beta\ (i_t^e - E_{t-1}[\pi_t]) \ \ - \beta u^* + \gamma \pi_{t-1} + (1-\gamma) E_{t-1}[\pi_t] \\
&= (\gamma + \beta)^{-1} (\beta\ i_t^e - \beta u^* + \gamma \pi_{t-1}\,). \qquad (10)
\end{aligned}$$

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 rules[20] 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.

As shown in the case of adaptive expectations, the central bank will set the nominal interest rate i_t so as to minimize $L(.)$ based on its knowledge of the state of the economy (i.e. lagged inflation) and the parameters (i.e. β, γ and u^*) but before the shock e_t is 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

$$i_t = H(\pi_{t-1}, \pi^*, \beta, \gamma, u^*), \qquad (11)$$

it implicitly takes 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 by minimizing the loss function $L(.)$ with respect to i_t and i_t^e under the explicit restriction

[20] See for example the contributions in Taylor (1999).

that the ex-ante expected nominal interest rate, i_t^e, is equal to its rational expectation:[21]

$$i_t^e = E_{t-1}[i_t \,|\, \pi_{t-1}, \pi^*, \beta, \gamma, u^*]. \tag{12}$$

From (11) and (12) and the assumption that neither the central bank nor the private sector have prior information on the random shock (e_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, $E_{t-1}[i_t \,|\pi_{t-1}, \pi^*, \beta, \gamma, u^*] = H(\pi_{t-1}, \pi^*, \beta, \gamma, u^*)$. Consequently, the private sector's rational expectation of inflation is

$$\pi_t^e = E_{t-1}[\pi_t] = (\gamma + \beta)^{-1}(\beta\, H(\pi_{t-1}, \pi^*, \beta, \gamma, u^*) - \beta u^* \; + \gamma\pi_{t-1}). \tag{13}$$

Symmetric information between the central bank and the private sector implies that the expectation derived in (13) also constitutes the central bank's rational expectation of inflation entering the current expected loss $L(.)$.[22] The conditional variance of inflation is independent of policy and equal to (8). Thus, the central bank is again able to minimize $L(.)$ by means of inflation forecast targeting. The optimal interest rate rule that ensures that expected inflation always equals the target rate corresponds to:

$$i_t = H(\pi_{t-1}, \beta, \gamma, u^*) = u^* + \pi^* - \frac{\gamma}{\beta}(\pi_{t-1} - \pi^*). \tag{14}$$

Contrary to the optimal policy under adaptive expectations in equation (9), the optimal policy under rational expectations depends on the index of inflation persistence γ that measures the relative importance of the backward- and forward-looking components in the Phillips curve. A comparison of the partial derivatives of (14) and (9) 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 greater extent if the private

[21] 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 the rational expectations at time $t-1$, $E_{t-1}\pi_t$ and $E_{t-1}i_t$, are functions of the policy rule, lagged inflation and the parameters. Committing to $E_{t-1}[i_t]$ has also been used as a commitment strategy by Svensson (1997b) and many others.

[22] Possible extensions allowing for asymmetric information are discussed in the final section of the paper.

sector forms adaptive rather than rational expectations, $(\beta-1)(\beta)^{-1} > -\gamma(\beta)^{-1} > 0$ where $\beta < 0$. The reason is that forward-looking market participants expect the central bank to raise interest rates sufficiently to return inflation to target in the next period. Under commitment, the central bank in turn takes into account this beneficial effect of market participants' expectations in the formulation of the policy rule. In the literature this effect is typically referred to as the 'expectations channel' of monetary policy transmission.

Optimal policy in the dynamic model

The question remains whether optimal policy in the dynamic model with positive discount factor δ differs from the optimal rules (9) and (14) derived in the static version of the model, respectively with adaptive and rational expectations. Recognizing that these rules implement inflation forecast targeting, the answer to this question is simple. Following these rules in every period will ensure that expected inflation remains equal to the target in every period. With expected inflation deviations from target pinned down at zero, the infinite horizon loss will be equal to the discounted sum of conditional variances of inflation, $\sum_{t=0}^{\infty} \delta^t(\sigma^2) = (1-\delta)^{-1}\sigma^2$. Since the variance of shocks, σ^2, is exogenous the expected discounted sum of losses cannot be reduced further and the rules (9) and (14) also achieve the global minimum in the dynamic version of the model, respectively with adaptive and rational expectations.

4 Parameter uncertainty and Brainard-style conservatism

We now turn to the central question of the paper, namely what is an appropriate policy rule $H(.)$ when the parameters of the Phillips curve are unknown. A possible solution would be to use the rules derived in the preceding section but to replace actual parameters with the estimates available from section 2:

$$E\left[\begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix}\right] = \begin{pmatrix} a \\ b \\ c \end{pmatrix}. \tag{15}$$

Using these estimates certainty-equivalent policy rules, $i_t = H^{ceq}(.)$, can be defined as follows:

$$H^{ceq}(.) = \frac{a}{b} + \pi_{t-1} - \frac{1}{b}(\pi_{t-1} - \pi^*) \text{ with adaptive expectations,} \quad (16)$$

$$H^{ceq}(.) = \frac{a}{b} + \pi^* - \frac{c}{b}(\pi_{t-1} - \pi^*) \quad \text{with rational expectations.} \quad (17)$$

These rules may be useful as benchmarks for comparison but they are not optimal, because they disregard the uncertainty associated with the parameter estimates and summarized by their covariance matrix:

$$Var\left[\begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix}\right] = \Sigma = \begin{pmatrix} v^a & v^{ab} & v^{ac} \\ v^{ab} & v^b & v^{bc} \\ v^{ac} & v^{bc} & v^c \end{pmatrix}. \quad (18)$$

Optimal policy in the static model with adaptive inflation expectations

With $\pi_t^e = \pi_{t-1}$, the central bank's conditional expectation of inflation corresponds to:

$$E_{t-1}[\pi_t] = a + b(i_t - \pi_{t-1}) + \pi_{t-1}. \quad (19)$$

The conditional variance of inflation, which forms the second component of the central bank's current expected loss $L(.)$, now depends on the degree of parameter uncertainty and on the policy instrument i_t:

$$VAR_{t-1}[\pi_t] = \sigma^2 + v^a + v^b(i_t - \pi_{t-1})^2 + 2v^{ab}(i_t - \pi_{t-1}) \quad (20)$$

As a consequence, the central bank faces a trade-off between the expected deviation of inflation from target and the conditional variance of inflation. The optimal rule does not imply inflation-forecast targeting and instead takes into account inflation uncertainty:

$$i_t = H^{cau}(.) = -\frac{ab + v^{ab}}{b^2 + v^b} + \pi_{t-1} - \frac{b}{b^2 + v^b}(\pi_{t-1} - \pi^*) \quad (21)$$

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.[23] Brainard's finding of cautionary policy is confirmed by showing that the response coefficient on the inflation gap in equation (21) is smaller in absolute value than in the certainty-equivalent rule in equation (16), $|-b(b^2+v^b)^{-1}| < |-b^{-1}|$. This conclusion follows directly from the fact that the variance v^b is positive in the presence of parameter uncertainty. For this reason, I will refer to this rule as the 'cautionary' rule. In the presence of 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 expected inflation to target by the next period. Thus, even a strict inflation-targeting central bank will not attempt to keep expected inflation always on target. Instead, inflation will remain elevated and return to target gradually over the next few periods. As the quote by Blinder in the introduction to this paper suggested parameter uncertainty leads to gradualist policy-making. Interestingly, however, policy does not depend on the variance of the intercept, v^a, the reason being that for a linear model and quadratic objective function certainty-equivalence applies with respect to additive uncertainty.

Another interesting finding concerns the neutral setting of the interest rate that turns out to depend on the variance v^b and the covariance v^{ab}. Due to this 'covariance effect' it may differ from the neutral setting in the certainty-equivalent rule, which corresponds to the sum of the natural rate estimate and the inflation target, $-ab^{-1}+\pi^*$. Why does the cautionary rule not adopt the same neutral setting when inherited inflation is on target? The reason is related to 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:

$$H^{cau}(\pi_{t-1}=\pi^*) = \frac{b^2}{b^2+v_b}\left(-\frac{a}{b}+\pi^*\right) + \frac{v^b}{b^2+v_b}\left(-\frac{v^{ab}}{v^b}+\pi^*\right). \tag{22}$$

[23]Other papers that have studied this effect are Clarida, Gali and Gertler (1999), Estrella and Mishkin (1998) and Svensson (1999). Sack (2000) shows how parameter uncertainty can explain the high degree of serial correlation in interest rates.

Here, the first term in large parentheses corresponds to the level of the nominal interest rate, which ensures that unemployment is equal to the natural rate, while the second term corresponds to the variance-minimizing level of the nominal interest rate. This tendency to set the nominal interest rate near the variance-minimizing level clearly reflects a cautionary motive.

By definition, the variance-minimizing level corresponds to that level where the central bank can assess the impact of policy on inflation (via unemployment) with the highest possible precision. This is the case when unemployment and interest rates are at their average level throughout the sample used to obtain the estimates a and b. Ordinary least squares imply that the covariance of the Phillips curve estimates is simply the negative of the product of the sample mean of unemployment denoted by $\bar{u}$ and the variance of the slope, i.e. $v^{ab} = -\bar{u}v^{b}$.[24] Furthermore, ordinary least squares implies that $a = \overline{\Delta\pi} - b\bar{u}$. Thus, if the first difference of inflation is zero over the sample, it follows that the sample mean of unemployment equals the natural rate $\bar{u} = ab^{-1}$, and the neutral level of the cautionary rule collapses to the neutral level of the certainty-equivalent rule. This turns out to be the case for the estimates obtained by Fuhrer (1995) reported in Table 1: $\bar{u} = -v^{ab}(v^{b})^{-1} = 0.09(0.015)^{-1} = 6 = -ab^{-1} = 1.68(-0.28)^{-1}$.

Optimal policy in the static model with rational inflation expectations

Given symmetric information between the central bank and market participants regarding the data used to estimate the Phillips curve, forward-looking market participants will utilize those same estimates (a, b, c) and covariance matrix Σ in forming expectations. Thus, the central bank's and market participants' expectation of inflation as of $t-1$ can be expressed as a function of lagged inflation, the parameter estimates and the central bank's policy rule $H(\pi_{t-1}, a, b, c, \Sigma)$, which in turn depends on the parameters and covariance

[24] For a derivation see Greene (1993), pp 155-157.

matrix defined in (18):

$$\begin{aligned} E_{t-1}[\pi_t] &= a + b(H(.) - E_{t-1}[\pi_t]) + c\pi_{t-1} + (1-c)E_{t-1}[\pi_t] \\ &= (c+b)^{-1}(a + bH(.) + c\pi_{t-1}). \end{aligned} \tag{23}$$

The second component of current expected loss L(.), the conditional variance of inflation, then corresponds to:

$$\begin{aligned} VAR_{t-1}[\pi_t] \quad &= \sigma^2 + v^a + v^b u_t^2 + v^c(\pi_{t-1} - E_{t-1}\pi_t)^2 \\ &+2v^{ab}u_t + 2v^{ac}(\pi_{t-1} - E_{t-1}\pi_t) + 2v^{bc}u_t(\pi_{t-1} - E_{t-1}\pi_t)), \\ &\text{where } u_t = (H(.) - E_{t-1}\pi_t) \text{ and } E_{t-1}\pi_t \text{ is defined as in (23).} \end{aligned} \tag{24}$$

The optimal policy rule trades off the impact of policy on the conditional expectation and the conditional variance on inflation. It takes the form:

$$\begin{aligned} i_t = H^{cau}(.) &= \frac{b^2}{b^2+R_1}\left(-\frac{a}{b} + \pi^*\right) + \frac{R_1}{b^2+R_1}\left(-\frac{aR_2+R_3}{R_1} + \pi^*\right) \\ &\quad - \frac{cb - R_1}{b^2+R_1}(\pi_{t-1} - \pi^*), \end{aligned} \tag{25}$$

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

$$\begin{aligned} R_1 &= v^b c^2 + v^c b^2 - 2v^{bc}bc \\ R_2 &= (v^c + v^{bc})b - (v^b + v^{bc})c \\ R_3 &= v^{ab}c(c+b) - v^{ac}b(c+b). \end{aligned} \tag{26}$$

Each of these three coefficients would be zero in the absence of uncertainty and the policy rule would simplify to (17). Under uncertainty, however, optimal policy depends on the parameter variances (v^b, v^c) and covariances (v^{ab}, v^{ac}, v^{bc}).

Similarly to the case of adaptive expectations two Brainard-style effects of uncertainty can be identified under rational expectations. The first effect is reflected by the two terms in the first row of (25). These two terms constitute the weighted average of the natural level of the interest rate and its variance-minimizing level. Again, the cautionary policy leans towards the variance-minimizing level and may therefore differ from the natural level even when inflation is on target.

The second effect concerns the policy response to inflation deviations from target. Again, the response of the cautionary rule (25) subsequent to an increase in inflation is more muted than under the certainty-equivalent rule (17). The partial derivative $\partial H^{cau}/\partial \pi_{t-1}$ 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}. R_1 in (26) corresponds to the variance of $(\beta c - \gamma b)$. Thus, $R_1 > 0$ in the presence of parameter uncertainty. With $b < 0$ and $-b < c$ it is then straightforward to show that:

$$\frac{\partial H^{cau}}{\partial \pi_{t-1}} = -\frac{cb - R_1}{b^2 + R_1} \;<\; \frac{\partial H^{ceq}}{\partial \pi_{t-1}} = -\frac{c}{b}. \tag{27}$$

Optimal policy in the dynamic model with constant parameter uncertainty

In the case of known parameters it is straightforward to show that the optimal policy in the static version of the model also achieves the global minimum in the dynamic version where the central bank aims to minimize expected current and discounted future losses and the discount factor is positive, $(0 < \delta < 1)$. The cautionary policy rule (25), however, is not necessarily optimal in the dynamic model. There exist several dynamic links in the model with unknown parameters. Most interestingly, the central bank may re-estimate the unknown parameters every time new data arrives and thereby learn over time. In the next section, I show how the estimates of the parameters of the inflation equation (2) 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.

Following Brainard (1967) in assuming that parameter estimates and parameter uncertainty remained fixed over time, one can instead solve for the dynamically optimal policy conditional on the sole remaining dynamic link that arises from lagged inflation in the

Phillips curve:

$$\underset{H(.)}{\text{Min}} E\left[\sum_{t=1}^{\infty} \delta^{t-1}(\pi_t - \pi^*)^2 \mid \pi_0\right] \tag{28}$$

$$\begin{aligned}
&\text{s.t. } \; i_t = H(\pi_{t-1}, \pi^*, a, b, c, \Sigma) \\
&\text{and } \; \pi_t = \alpha + \beta(H(.) - E_{t-1}[\pi_t]) + \gamma\pi_{t-1} + (1-\gamma)E_{t-1}[\pi_t] + e_t \\
&\text{where } \; E_{t-1}[\pi_t] \; \text{ is defined by } \; (23) \\
&\text{and } \; (a, b, c, \Sigma) \; \text{ is constant for } \; t = 1, \infty.
\end{aligned}$$

Here, the Phillips curve with rational expectations (defined as in equation (23)) is considered. The dynamic problem under adaptive expectations is formulated analogously by making the optimization instead s.t. $\pi_t = \alpha + \beta(H(.) - \pi_{t-1}) + \pi_{t-1} + e_t$.

Contrary to the case with known parameters, the presence of a trade-off between expected inflation and inflation variance could be the source of differences between the optimal rule in the static model derived previously and the solution to this dynamic problem. To investigate this possibility I rewrite the dynamic optimization problem defined above as a standard dynamic programming problem with lagged inflation π_{t-1} as its single state variable. Denoting the value function for this dynamic program by $V(\pi)$ the associated Bellman equation corresponds to:

$$\begin{aligned}
V(\pi_{t-1}) \quad = \quad & \underset{H(.)}{\text{Min}} \; L(\pi_{t-1}, H(.), a, b, c, \Sigma\,) \\
& \qquad + \delta \int V(\pi_t(\,H(.)\,, ..)\,)\, f(\,\pi_t |\, \pi_{t-1}, a, b, c, \Sigma\,)\, d\pi,
\end{aligned} \tag{29}$$

subject to constraints summarized in (28). $f(\pi_t|.)$ denotes the predictive distribution of inflation.

In the absence of a simple analytical solution, I solve for the optimal policy numerically by iterating over the Bellman equation. The numerical methods are described in more detail in the appendix. Numerical solutions are obtaining using the parameter estimates and covariances also employed in section 6 and subsequent sections. Interestingly, the numerical solution to the dynamic problem is indistinguishable (up to the accuracy of the numerical approximation) from the solution to the static problem, (i.e. equation (25)). Absent a formal proof, my conjecture is that optimal policy is identical in the dynamic and static cases.

5 Rational learning and the optimal policy rule

As new observations on inflation and unemployment become available, the central bank and forward-looking market participants may update their estimates of the unknown Phillips curve parameters. In order to analyze *rational learning* by the central bank and market participants, I will draw on the well-developed framework of the Bayesian learning literature cited in the introduction. In particular, Easley and Kiefer (1988), Kiefer and Nyarko (1989) and Wieland (2000a) study Bayesian learning in controlled regressions. These authors consider the following regression, where X_t denotes the vector of explanatory variables, $\tilde{\beta}$ the vector of unknown parameters, y_t the dependent variable and e_t a random shock:

$$y_t = \tilde{\beta}' X_t + e_t. \tag{30}$$

At least one of the explanatory variables in X_t is chosen by a decision maker with market power, such as the central bank considered in this paper. This choice is made conditional on its prior belief regarding the unknown parameters that is modelled with a probability distribution $p(\tilde{\beta}|\theta_{t-1})$. The vector θ_{t-1} contains all state variables which are required to describe this distribution based on the information available in period $t-1$. Then a random shock e_t with distribution $N(0, \sigma^2)$ occurs and a new realization of the dependent variable y_t can be observed. Before choosing next period's control, the decision maker updates his beliefs using the new information (y_t, X_t). The posterior distribution is obtained using Bayes rule,

$$p(\tilde{\beta}|\theta_t) = \frac{p(y_t|\tilde{\beta}, X_t, \theta_{t-1})p(\tilde{\beta}|\theta_{t-1})}{p(y_t|x_t, \theta_{t-1})}, \tag{31}$$

which implies a set of nonlinear updating equations for the state variables θ_t:

$$\theta_t = B(\theta_{t-1}, X_t, y_t). \tag{32}$$

In the following, I specify the elements of the Bayesian learning framework when applied to the model in this paper, respectively with adaptive and rational expectations in the Phillips curve.

Updating the central bank's beliefs when market participants form adaptive expectations

In the case of the Phillips curve with adaptive expectations estimated by Fuhrer (1995) and Staiger et al. (2002) the relevant regression equation is:

$$\pi_t - \pi_{t-1} = \alpha + \beta u_t + e_t. \tag{33}$$

This regression fits the regression equation of the Bayesian learning model defined by (30) as follows: $X_t = (1\ u_t)'$, $y_t = \pi_t - \pi_{t-1}$ and $\tilde{\beta} = (\alpha\ \beta)'$. With e_t distributed normally with known variance σ^2 the central bank's prior beliefs are described by the following normal distribution:

$$p\left(\left(\begin{array}{c}\alpha\\ \beta\end{array}\right)|\theta_{t-1}\right) = N\left(\left(\begin{array}{c}a_{t-1}\\ b_{t-1}\end{array}\right), \Sigma_{t-1} = \left(\begin{array}{cc}v^a_{t-1} & v^{ab}_{t-1}\\ v^{ab}_{t-1} & v^b_{t-1}\end{array}\right)\right). \tag{34}$$

Thus, θ_{t-1}, the vector of state variables describing beliefs, has five elements: $(a_{t-1}, b_{t-1}, v^a_{t-1}, v^{ab}_{t-1}, v^b_{t-1})$. The subscript $t-1$ is appended to the point estimates and covariances to denote the estimation conditional on information up to and including period $t-1$. As new information (u_t, π_t) becomes available the point estimates and covariances representing the central bank's beliefs are updated. Bayes rule implies the following updating equations:

$$\begin{aligned}\left(\begin{array}{c}a_t\\ b_t\end{array}\right) &= \left(\begin{array}{c}a_{t-1}\\ b_{t-1}\end{array}\right) + \Sigma_{t-1}X_tF^{-1}(\pi_t - \pi_{t-1} - a_{t-1} - b_{t-1}u_t)\\ \Sigma_t &= \Sigma_{t-1} - \Sigma_{t-1}X_tF^{-1}X_t'\Sigma_{t-1} \text{ where } F = X_t\Sigma_{t-1}X_t' + \sigma^2.\end{aligned} \tag{35}$$

A derivation of the updating equations using Bayes rule can be found in Zellner (1971). In the case considered here, the updating equations correspond exactly to recursive least squares as can be seen from Harvey (1992). F refers to the conditional variance of the dependent variable. The means a and b are only updated if the observed change in inflation differs from the expected change in inflation. The covariance matrix, however, changes deterministically as a function of the square of the explanatory variables X. The updating is conditional on a known normal distribution of random shocks. This assumption is standard in the optimal learning literature. It guarantees that given a normal prior, the posterior belief will also be a normal distribution. The asymptotic behavior of posterior beliefs is discussed in section 7.

Updating beliefs when the central bank and market participants learn rationally

As new observations on inflation and unemployment become available, not only the central bank but also forward-looking market participants can update their estimates of the unknown parameters in the Phillips curve. 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.

In the case of the Phillips curve with partially forward-looking expectations the central bank and market participants will estimate the following equation recursively:

$$\pi_t - E_{t-1}\pi_t = \alpha + \beta u_t + \gamma(\pi_{t-1} - E_{t-1}\pi_t) + e_t. \tag{36}$$

Somewhat surprisingly perhaps, this regression can also be mapped into the regression equation (30) of the Bayesian learning model by defining explanatory and dependent variables as follows: $X_t = \begin{pmatrix} 1 & u_t & (\pi_{t-1} - E_{t-1}\pi_t) \end{pmatrix}'$, $y_t = \pi_t - E_{t-1}\pi_t$ and $\tilde{\beta} = (\alpha \;\; \beta \;\; \gamma)'$. The essential insight is to recall that beliefs are updated recursively. When updating estimates and covariance matrices from period $t-1$, the conditional expectation of inflation based on $t-1$ information is known and defined by equation (23). Thus, the conditional expectation may be subtracted from dependent and explanatory variables in period t as suggested above to match the simple regression of the Bayesian learning model.

Due to the addition of the index of inflation persistence, γ, the central bank's and forward-looking market participants' beliefs are now represented by a trivariate normal distribution:

$$p\left(\begin{pmatrix}\alpha\\ \beta\\ \gamma\end{pmatrix}|\theta_{t-1}\right)=N\left(\begin{pmatrix}a_{t-1}\\ b_{t-1}\\ c_{t-1}\end{pmatrix},\Sigma_{t-1}=\begin{pmatrix}v^{a}_{t-1} & v^{ab}_{t-1} & v^{ac}_{t-1}\\ v^{ab}_{t-1} & v^{b}_{t-1} & v^{bc}_{t-1}\\ v^{ac}_{t-1} & v^{bc}_{t-1} & v^{c}_{t-1}\end{pmatrix}\right). \tag{37}$$

The vector of state variables θ that characterize beliefs now contains nine variables, the three means $(a_{t-1}, b_{t-1}, c_{t-1})$, the three variances $(v^a_{t-1}, v^b_{t-1}, v^c_{t-1})$ and the three covariances $(v^{ab}_{t-1}, v^{ac}_{t-1}, v^{bc}_{t-1})$. The updating equations implied by Bayes rule again correspond to recursive least squares:

$$\begin{aligned}\begin{pmatrix}a_t\\ b_t\\ c_t\end{pmatrix} &= \begin{pmatrix}a_{t-1}\\ b_{t-1}\\ c_{t-1}\end{pmatrix}+\Sigma_{t-1}X_tF^{-1}(\pi_t-E_{t-1}\pi_t-a_{t-1}-b_{t-1}u_t-c_{t-1}(\pi_{t-1}-E_{t-1}\pi_t)) \\ \Sigma_t &= \Sigma_{t-1}-\Sigma_{t-1}X_tF^{-1}X_t'\Sigma_{t-1} \text{ where } F=X_t\Sigma_{t-1}X_t'+\sigma^2\end{aligned} \tag{38}$$

where $E_{t-1}\pi_t$ is defined by (23).

Optimal policy with rational learning

Due to learning by the central bank, the current choice of the interest rate will affect the precision of the point estimates and 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 $H(\pi_{t-1}, \theta_{t-1}, \pi^*)$ with learning solves the following optimization problem:

$$\underset{H(.)}{\text{Min}}\, E\left[\sum_{t=1}^{\infty}\delta^{t-1}(\pi_t-\pi^*)^2 \mid (\pi_0,\theta_0)\right] \tag{39}$$

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

and $\pi_t = \alpha + \beta(H(.) - E_{t-1}[\pi_t]) + \gamma\pi_{t-1} + (1-\gamma)E_{t-1}[\pi_t] + e_t$

where $E_{t-1}[\pi_t]$ is defined by (23)

and s.t. the belief updating equations (38).

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 operator E, which denotes the statistical expectation. However, one can still use a standard contraction mapping argument as in Kiefer and Nyarko (1989) to show that a unique value function exists that solves the dynamic program and corresponds to the infimum of the sum of expected current and discounted future losses in (39). 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:

$$
\begin{aligned}
V(\pi_{t-1}, \theta_{t-1}) &= \underset{H(.)}{\text{Min}} \; L(\pi_{t-1}, \theta_{t-1}, H(.)\,) \\
&\qquad + \delta \int V(\pi_t(\,H(.)\,,..)\,, \theta_t(\,H(.), ..)\,)\, f(\,\pi_t|\,\pi_{t-1}, \theta_{t-1}, H(.)\,)\, d\pi \\
&= \underset{H(.)}{\text{Min}} \; L(\pi_{t-1}, \theta_{t-1}, H(.)\,) \;+\; \delta \int V(\alpha, \beta, \gamma, e_t,\; H(.)\,, \pi_{t-1}, \theta_{t-1}) \\
&\qquad p(\alpha, \beta, \gamma \mid \pi_{t-1}, \theta_{t-1}, H(.)\,)\, f(e)\, d\alpha\, d\beta\, d\gamma\, de.
\end{aligned}
\tag{40}
$$

The two terms on the right-hand side of the upper equation in (40) characterize the tradeoff between current control and estimation. $L(.)$ 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 by 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. $f(\pi_t|.)$ 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 $H(\pi_{t-1}, \theta_{t-1}, \pi^*, \sigma^2)$ that feeds back on all currently available information.

In the lower equation in (40), time t values of inflation and beliefs have been substituted out using equations (5), (6), (2), (23) and (38). They are functions of the previous period's inflation rate and beliefs and also of the unknown parameters and random shock

e_t. Expectations are taken with respect to the unknown parameters and the random shock. $p(\alpha, \beta, \gamma|.)$ is the trivariate normal distribution that describes the policymaker's beliefs about the unknown parameters. $f(e)$ 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:

$$i_t = H^{opt}(\pi_{t-1}, \theta_{t-1}, \pi^*, \delta, \sigma^2). \tag{41}$$

Unfortunately analytical solutions for $H^{opt}(.)$ 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.

Numerical approximation

In order to obtain numerical approximations it is necessary to specify numerical values for the central bank's inflation target, its discount factor and the variance of shocks. These values are reported in **Table 2**.

Table 2: Calibrated Parameters

Inflation Target	Discount Factor	Shock Variance
$\pi^* = 0$	$\delta = 0.95$	$\sigma^2 = 1$

The Bellman equation (40) 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 (40). It is well known that 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 value function iterations with policy iterations to speed up convergence. Nevertheless, the optimal learning problem with three unknown parameters in (39) that has a total of 10 continuous state variables is too large to be solved numerically with reasonable precision.[25]

Instead I provide numerical results for alternative versions of the learning problem with up to two unknown parameters in the following three sections. Section 6 presents optimal policies for the benchmark specification with adaptive expectations as estimated by Fuhrer (1995) and Staiger et al. (2002). First, I only treat the slope of the Phillips curve as unknown and then the intercept and the slope together. To illustrate the resulting dynamic learning behavior I simulate time paths under alternative policies in section 7. Section 8 extends the analysis to include forward-looking, rational learning by market participants.

6 The optimal balance of caution and experimentation

Our benchmark case for comparing optimal learning by the central bank with Brainard-style conservatism is the estimated Phillips curve with adaptive expectations, i.e. $\pi_t^e = \pi_{t-1}$ in equation (2). Fuhrer's (1995) estimates of this equation, $(a = 1.68, b = -0.28)$, imply a natural rate u^* of 6%. The associated variances, $(v^a = 0.58, v^b = 0.015, v^{ab} = -0.09)$, are fairly small and the respective ratios of point estimate and standard deviation, $(2.21, -2.28)$, indicate a high degree of statistical significance. From the Brainard-style analysis with constant beliefs in section 4 we learned that uncertainty about the multiplicative slope parameter β importantly affects the central bank's policy choice, while the variance of the intercept has no such effect. Thus, as a starting point for the policy comparison exercise, I focus on the case with only β unknown.

Optimal policy with unknown slope

Figure 1 compares the policy response to lagged inflation implied by optimal, cautionary and certainty-equivalent policy making. The horizontal axis measures the deviation of lagged inflation π_{t-1} from target. The vertical axis corresponds to the deviation of unem-

[25] The numerical algorithm and associated computation costs are discussed in more detail in appendix A.

ployment from its natural rate to be expected given the central bank's choice of interest rate in response to the inherited inflation gap. The vertical axis also equals the deviation of the real interest rate, $(i_t - \pi_t^e = i_t - \pi_{t-1})$, from its equilibrium value, $r^* = u^*$.

Figure 1: Alternative Policies with Unknown Slope β
Prior Belief from Fuhrer (1995): $(b = -0.28, v^b = 0.015)$

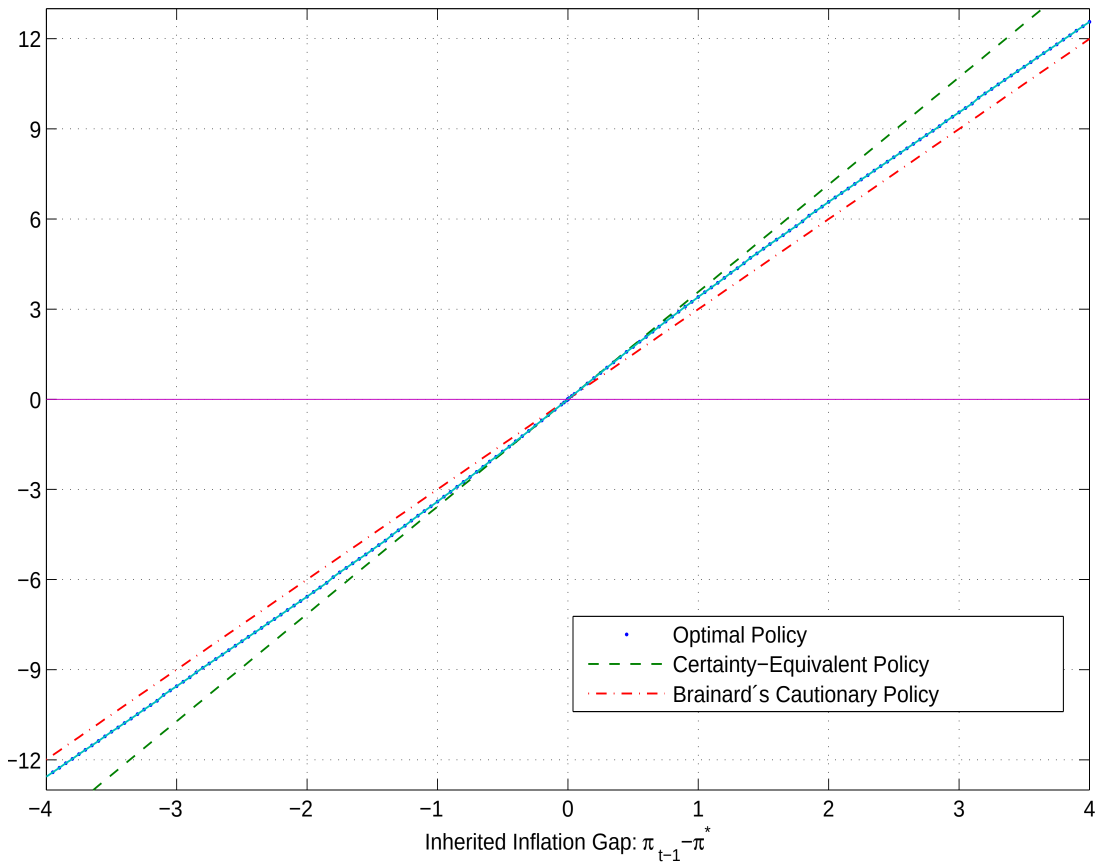

The optimal policy rule (solid line with bold dots)[26] is derived from the optimization problem with learning in (39). It is compared to the certainty-equivalent rule (dashed line) in equation (16) and to the cautionary policy rule (dashed-dotted line) derived from the Brainard-style dynamic optimization problem without learning in (28). While this latter

[26]The bold dots correspond to the grid used for numerical approximation.

rule is derived numerically with constant beliefs, it turns out to be indistinguishable from the cautionary rule for the static model shown previously in equation (21).

The following three findings are directly apparent from Figure 1. First, 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 optimal rule always requires a more aggressive policy response to inflation than the cautionary rule. This difference, which represents the *extent of experimentation* incorporated in the optimal rule, changes little in absolute terms for moderate to high inflation deviations from target.[27] 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 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.

Table 3 reports the optimal extent of gradualism and experimentation for inflation deviations from target of two and three percentage points. The values shown in the table represent the differences in the expected unemployment rate due to central bank policy.

Table 3: Optimal Degree of Gradualism and Experimentation
Differences in Expected Unemployment Rates when β is Unknown

	Gradualism ($H^{ceq} - H^{opt}$)		Experimentation ($H^{opt} - H^{cau}$)	
$b = -0.28$	$v^b = 0.01$	$v^b = 0.015$	$v^b = 0.01$	$v^b = 0.015$
$\pi_{t-1} - \pi^* = 2$	0.45	0.58	0.35	0.57
$\pi_{t-1} - \pi^* = 3$	0.85	1.17	0.37	0.55

Notes: ($b = -0.28, v^b = 0.015$) estimated in Fuhrer (1995). ($b = -0.28, v^b = 0.01$) estimated in Staiger et al. (2002).

[27] In 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.

Here, gradualism refers to the difference between the certainty-equivalent and the optimal policy rule: $H^{ceq}(\pi_{t-1} - \pi^*, b, v^b) - H^{opt}(\pi_{t-1} - \pi^*, b, v^b)$. It indicates that the certainty-equivalent rule, which disregards parameter uncertainty, is too aggressive in fighting inflation. Experimentation refers to the difference between the optimal and the cautionary policy rule: $H^{opt}(\pi_{t-1} - \pi^*, b, v^b) - H^{cau}(\pi_{t-1} - \pi^*, b, v^b)$. It indicates that the cautionary rule responds too litte to inflation deviations from target. The columns in Table 3 refer to the parameter variances estimated by Fuhrer (1995) and Staiger et al. (2002) since they obtain the same point estimate. The differences between the optimal policy and its two alternatives are economically significant, ranging from one half to a full percentage point of the unemployment rate. However, they are small relative to the overall policy response.

Given these differences the question arises to what extent the optimal policy improves expected performance compared to Brainard-style conservatism or compared to complete disregard of uncertainty. **Figure 2** provides a comparison of expected losses as a function of the lagged inflation gap given the beliefs underlying the policy comparison in Figure 1. Of course, expected losses are lowest under the optimal policy. The next higher curve measures the expected performance under the cautionary policy. The difference in performance is largest near the inflation target, i.e. near an inflation gap of zero. Expected losses under the certainty-equivalent policy that disregards uncertainty completely are even greater. In percentage terms, however, the differences in losses are relatively small. This should not be too surprising given that the degree of uncertainty about the parameter estimate b is rather small. Also, in calculating expected losses for the cautionary and certainty-equivalent policy rules it was assumed that the central bank would nevertheless update its beliefs over time. If one assumes instead that beliefs remain constant expected losses will be significantly larger. In the case of Brainard-style cautionary policy this is indicated by the highest curve in Figure 2.

So far, I have only considered a scenario with very low uncertainty as suggested by the estimates presented in section 2. However, uncertainty may be renewed occasionally as structural changes occur. Furthermore, the precision of these estimates likely overstates the

Figure 2: Value Functions under Alternative Policies with Unknown Slope β
Prior Belief from Fuhrer (1995): $(b = -0.28, v^b = 0.015)$

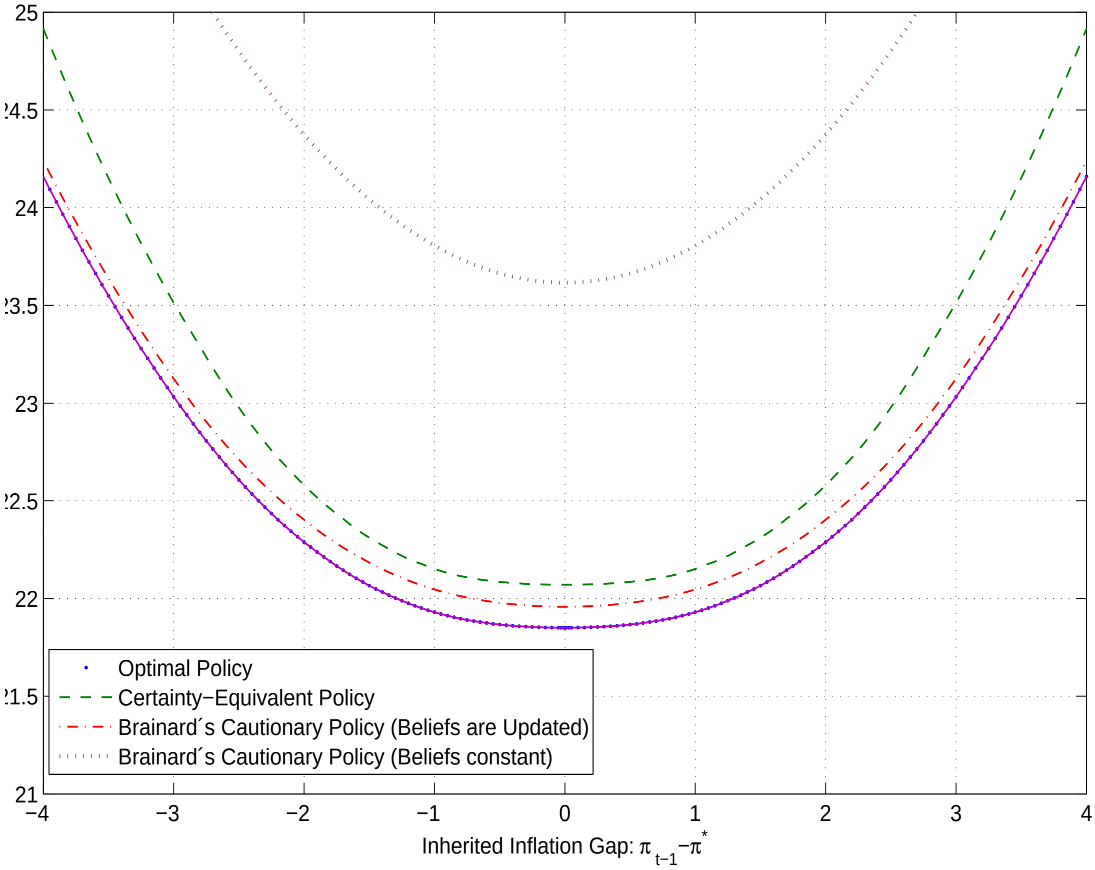

confidence with which policymakers (or the economics profession at large for that matter) would rely on them in actual policy practice. **Figure 3** extends the policy comparison to alternative point estimates and considers scenarios of greater uncertainty. There are nine panels. The first panel in the middle row corresponds to the benchmark case from Figure 1, shown only for positive inflation deviations from target. The upper and lower rows of panels are computed for point estimates of $b = (-0.22, -0.34)$, which are within one-half standard deviation from the empirical estimate of -0.28. The columns are ordered

Figure 3: Alternative Policies with Unknown Slope β
A Range of Point Estimates and Variances

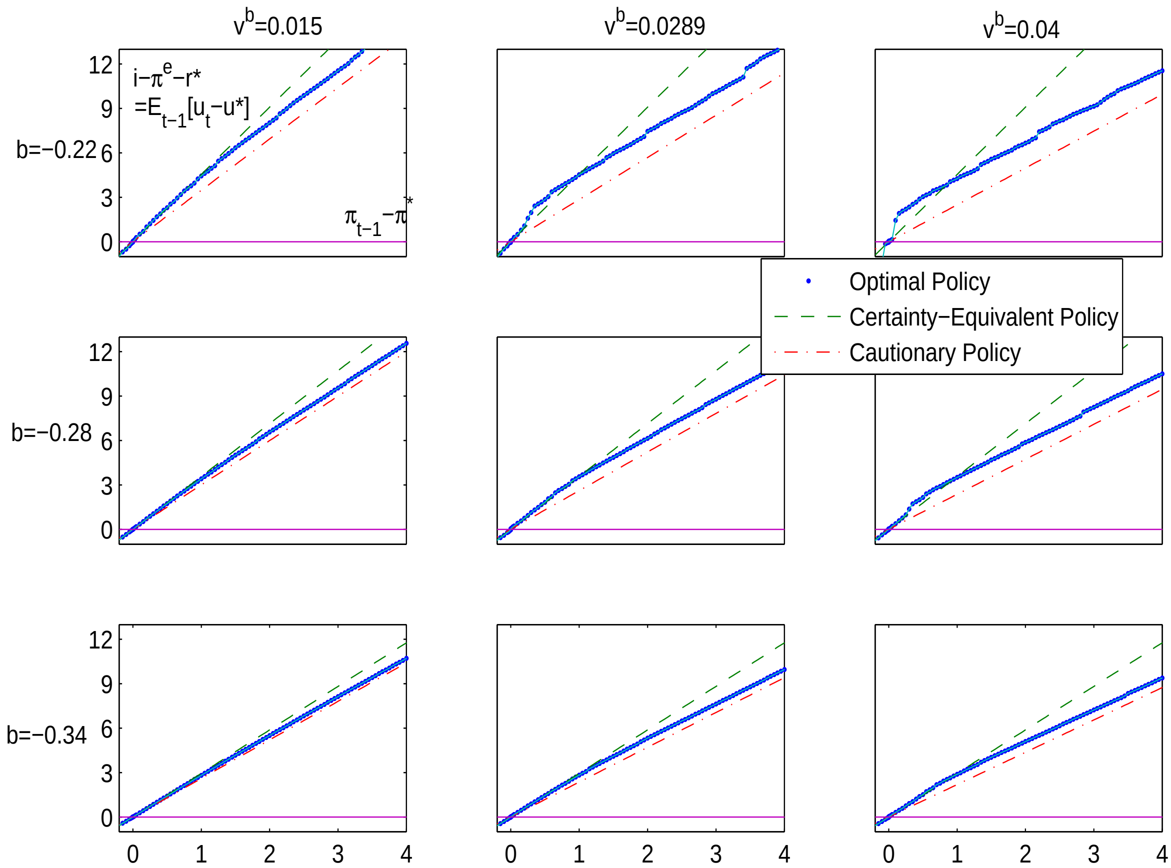

as follows: the first column corresponds to a variance of $v^b = 0.015$ as in Fuhrer (1995), the second column to $v^b = 0.0289$ and the third column to $v^b = 0.04$. For $b = -0.28$ the ratio of point estimate and standard deviation corresponds to 1.65 and 1.4, respectively, indicating scenarios of borderline statistical significance. As this ratio becomes smaller (and uncertainty greater) the wedge between the inflation responses of the certainty-equivalent and cautionary rules widens substantially. Thus, it is widest in the upper-right panel and smallest in the lower-left panel.

From the nine panels in Figure 3 it is directly apparent that the optimal policy tends to remain in between the certainty-equivalent and cautionary rules in terms of the policy response to inflation in most scenarios. Thus, the optimal policy still embodies some degree of gradualism, because it responds less actively than the certainty-equivalent rule, but also some degree of experimentation, because it responds more actively than the cautionary rule.

Figure 4: Alternative Policies with Unknown Slope β
Point Estimate from Fuhrer (1995) with Extreme Uncertainty: $(b = -0.28, v^b = 0.0784)$

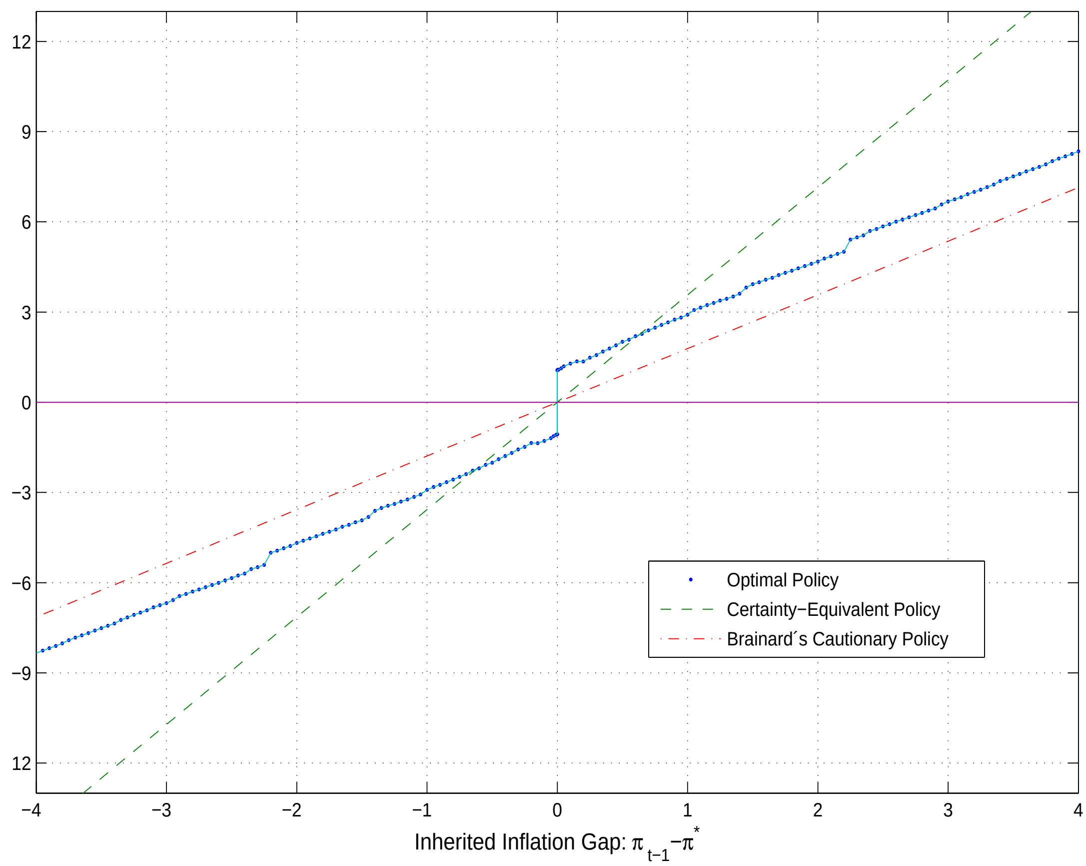

However, there are interesting exceptions where the finding of gradualism in optimal policy is overturned. In particular, when lagged inflation is near the target (typically within less than 1 percentage point) and uncertainty is moderate to high, the optimal policy response may be somewhat more aggressive than the certainty-equivalent rule. This is apparent, for example, from the second and third panel in the middle row of Figure 3. Furthermore, when uncertainty is extremely high, for example in the upper-right panel, the optimal policy rule exhibits a discontinuity near the inflation target. This discontinuity is illustrated more clearly in **Figure 4** which shows the optimal policy with Fuhrer's point estimate of -0.28 but extreme uncertainty $v^b = 0.0784$.

An implication of the discontinuity at zero is that the optimal policy response actively perturbs inflation a bit near the target, purely to generate information. In this case, the central bank accepts expected deviations from target relative to the certainty-equivalent policy solely in order to obtain more precise parameter estimates and improve inflation stabilization in the future. However, this case only arises in extreme scenarios where the ratio of point estimate and standard error is near 1 as in Figure 4.

Optimal policy with unknown slope and intercept

Next, I extend the analysis to the learning problem with two unknown parameters, that is, the intercept and slope of the Phillips curve. This problem has six state variables, namely the two means, two variances, the covariance and lagged inflation. The results are shown in **Figure 5** which contains three panels. The first panel provides a comparison of alternative policies for the full set of point estimates and covariances from Fuhrer (1995): $(a = 1.68, b = -0.28, v^a = 0.58, v^b = 0.015, v^{ab} = -0.09)$. The second and third panel contain scenarios with the same point estimates but greater variances. v^b is set to the values previously considered in Figure 3: (0.0289, 0.04). The variance of the intercept estimate v^a is increased in the same manner. The covariance is chose to keep the correlation coefficient constant at -0.965.

These results confirm that the main findings from the learning problem with unknown slope concerning the optimal policy response to inflation carry over to the learning problem

Figure 5: Alternative Policies with Unknown Intercept α and Slope β
Prior Belief: (a=1.68, b=-0.28) and Various Degrees of Uncertainty

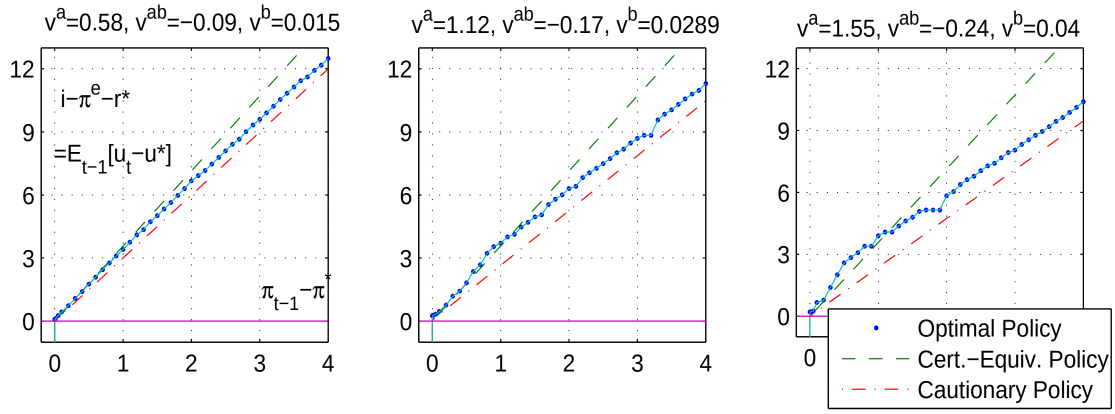

with unknown slope and intercept.[28] The optimal response typically falls inside the wedge created by the aggressive certainty-equivalent rule and the cautionary rule. Optimal policy always incorporates some experimentation compared to the cautionary rule. However, it still also remains gradualist, that is, less aggressive than the certainty-equivalent rule that disregards parameter uncertainty. Exceptions to the principle of gradualism only occur near the inflation target or under extreme uncertainty, and those exceptions tend to be small in magnitude.

Table 4 reports the optimal extent of gradualism and experimentation with a prior belief corresponding to the full set of Fuhrer (1995)'s estimates as in the first panel of Figure 4. Again, the values reported in the table represent the differences in the expected unemployment rate that arise under the alternative policy rules in response to inflation deviations from target of two and three percentage points as in Table 3. Interestingly, the optimal degree of experimentation is slightly greater if intercept and slope are unknown than if only the slope parameter is unknown as can bee seen by comparing the results in Table 4 to the columns in Table 3 for $v^b = .015$. Similarly, the optimal degree of gradualism

[28]The optimal policy rule in the second and third panels exhibits a certain "jumpiness", which reflects approximation error. Similar "jumpiness" in the policy rules for one unknown parameter shown in Figures 1 and 3 disappeared once the number of grid points was drastically increased.

is somewhat smaller than in the problem with one unknown parameter.

Table 4: Optimal Degree of Gradualism and Experimentation
Differences in Expected Unemployment Rates with α and β Unknown

	Gradualism ($H^{ceq} - H^{opt}$)	Experimentation ($H^{opt} - H^{cau}$)
$\pi_{t-1} - \pi^* = 2$	0.48	0.67
$\pi_{t-1} - \pi^* = 3$	1.12	0.60

Notes: Point estimates and covariance matrix as estimated in Fuhrer (1995).

As discussed in section 3 the variance-minimizing setting of the unemployment rate with a zero inherited inflation gap need not necessarily coincide with the estimated natural rate and therefore may not keep future inflation on target. Whether the variance-minimizing level of the unemployment rate deviates from the estimated natural rate depends on the ratio of covariance and variance of the slope estimate. Using the estimated covariance matrix from Fuhrer (1995) this ratio turns out to be equal to the natural rate estimate: $v^{ab}(v^b)^{-1} = a(b)^{-1} = 6$. Thus, the 'covariance effect' discussed in section 3 does not arise for our benchmark estimates. This property is therefore maintained for the two scenarios with greater uncertainty presented in the middle and right panel of Figure 5. For alternative scenarios studying this covariance effect the interested reader is instead referred to Wieland (2003).

7 The dynamics of learning and convergence

The preceding section has compared the stationary policy and value functions at different points in the state space, i.e. for given inherited inflation and given prior beliefs including point estimates *and* covariance matrix. Over time, of course, beliefs will change and the central bank will obtain more precise estimates of the unknown parameters. Thus, a given set of point estimates will not always have the same implications for policy. To illustrate the dynamic behavior of the economy under alternative policies I conduct a dynamic simulation taking into account that the central bank will be able to re-estimate the Phillips curve every time new inflation and unemployment observations become available. This simulation indicates the *expected disinflation path* when initial inflation is two percentage points above target. While the initial point estimates are set to the values from Fuhrer (1995), ($a_1 = 1.68, b_1 = -0.28$), the initial covariance matrix implies a much greater degree of uncertainty and is set to the same values as in the third panel in Figure 5, ($v_1^a = 1.55, v_1^b = 0.04, v_1^{ab} = -0.24$).[29]

The two panels in the top row of **Figure 6** show the expected paths of inflation and unemployment under the certainty-equivalent policy (solid line), the cautionary Brainard-style policy (dashed line) and the optimal policy (dotted line). All three policy rules require that the central bank raises interest rates in order to increase unemployment and bring inflation back to target. However, the speed of disinflation differs across rules. The certainty-equivalent rule ensures that inflation is expected to be back on target by the next period. The cautionary rule disinflates more gradually over the course of three periods. Thus, the peak of unemployment is smaller under the cautionary rule than under the certainty-equivalent rule. As one would expect based on the policy comparisons in the preceding section, the optimal speed of disinflation lies in between the cautionary and certainty-equivalent rules. In other words, the optimal policy exhibits gradualism compared to the certainty-equivalent policy, which disregards parameter uncertainty but disinflates more actively than the Brainard-style policy.

The lower four panels in Figure 6 report the expected dynamic paths of the point estimates, variances and covariance. The point estimates stay constant over time. By definition, the central bank cannot expect surprises that would change its point estimates. Therefore, I simulate the *expected* disinflation path by setting the random shocks e_1 to e_5 equal to their expected value of zero and the true parameter values α and β equal to the initial estimates a_1 and b_1.

[29] In other words, the initial situation can be thought of as a time period early in the sample of Fuhrer (1995) when estimates would still have been very uncertain. Alternatively, it could be thought of as a time period immediately following a structural change that would have renewed uncertainty.

Figure 6: Disinflation and Learning with Alternative Policies
Point Estimates from Fuhrer (1995) with Initial Prior of High Uncertainty

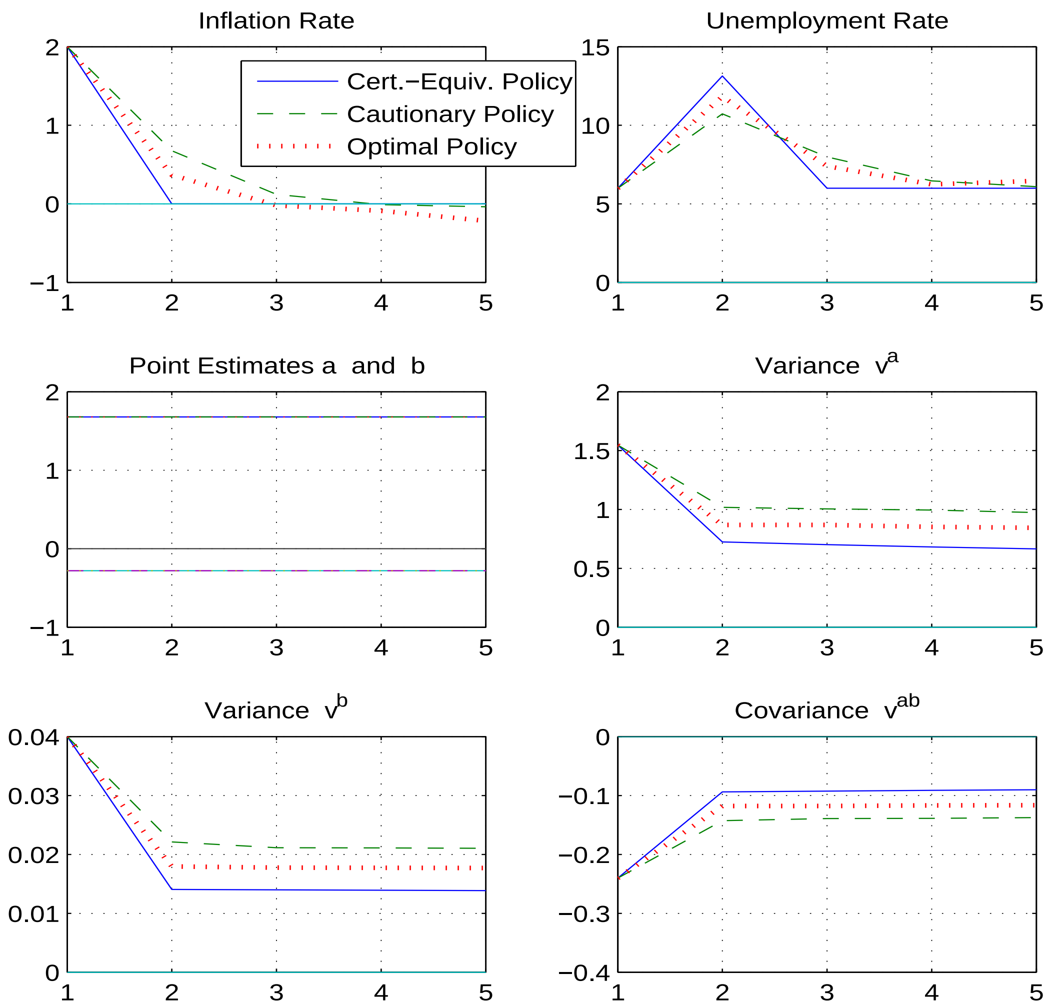

Turning from the point estimates to the dynamic path of the covariance matrix, the simulated paths in Figure 6 clearly indicate that uncertainty decreases over time. The variances and covariance are deterministic functions of the sum of squared deviations of the explanatory variable, i.e. the unemployment rate, from its sample average $\bar{u}_t$:[30]

$$\begin{aligned} v_t^b &= \frac{\sigma^2}{\sum_{j=1}^{t}(u_j - \bar{u}_t)^2} \\ v_t^{ab} &= -\bar{u}_t v_t^b \\ v_t^a &= \frac{\sigma^2}{t} + \bar{u}_t^2 v_t^b \end{aligned} \tag{42}$$

Thus, the central bank will expect the degree of uncertainty about Phillips curve parameters to decline over time as long as stabilization policy leads to variations in the unemployment rate. Since all three policy rules imply continued stabilization policy in the event of inflationary shocks the covariance matrix must converge to zero in the long-run. How quickly the reduction in uncertainty occurs depends on the magnitude of the policy response to inflation. Consequently, the simulation in Figure 6 shows that uncertainty is reduced most quickly under the certainty-equivalent rule and most slowly under the cautionary rule. The optimal policy indicates the optimal speed of learning.

The point estimates a_t and b_t follow a martingale relative to the central bank's information. Since the expected observation $E_{t-1}[\pi_t - \pi_{t-1} - a_{t-1} - b_{t-1}u_t] = 0$ it follows that $E_{t-1}[a_t] = a_{t-1}$ and $E_{t-1}[b_t] = b_{t-1}$ as already alluded to above. Whether the process of posterior beliefs converges to the true values α and β has been studied by Kiefer and Nyarko (1989). In particular, they show that if the explanatory variable u_t does not converge to a constant, then the process of posterior point estimates converges to the true values and the covariance matrix to zero.[31] As noted previously, u_t will never settle down to a constant value as long as inflationary shocks e_t lead to variations in the inflation rate and monetary policy responds by moving the interest rate and consequently unemployment in order to stabilize inflation. Thus, the central bank will eventually learn the true parameter values as long as it follows any of the three alternative policies considered in this paper. Incomplete learning about the parameters of the Phillips curve would only be a possibility if the central bank would focus exclusively on stabilizing unemployment and would achieve a constant

[30]For a derivation see Greene (1993).

[31]For a proof see theorem 4.2. on page 577 in Kiefer and Nyarko (1989).

unemployment rate, but then inflation would not be controlled.[32]

The speed of learning apparent from the reduction of uncertainty in Figure 6 seems quite fast. Over two to four periods, depending on the policy rule followed, the central bank transitions from a situation of very high uncertainty to the precision of parameter estimates obtained by Fuhrer (1995). This finding sheds some doubt on the regression specification chosen by Fuhrer and others in the empirical literature. An alternative specification would be one that allows for shifts in Phillips curve parameters. In this case, uncertainty would occasionally be renewed and the incentive to experiment would remain rather than die out over time. This extension is left for future research. For an illustrative example of optimal learning with time-varying parameters the reader is referred to Beck and Wieland (2002).

8 Optimal experimentation with forward-looking market participants

The question remains how the optimal policy for a central bank that is learning about Phillips curve parameters would change if forward-looking market participants also learn in the same manner. Given symmetric information and policy commitment to a rule, rational market participants form the same beliefs about the unknown parameters α and β as the central bank. The presence of forward-looking expectations adds an expectation channel of monetary policy transmission. Consequently, the central bank can afford to be less activist in responding to inflation deviations from target because private sector expectations take into account future policy action and move towards the inflation target.

Figure 7 provides a comparison of optimal, certainty-equivalent and cautionary rules when α and β are unknown and forward-looking market participants learn just like the central bank. The results are presented in nine panels. In the first column of panels prior beliefs correspond to the benchmark case with Fuhrer's (1995) estimates and covariance matrix, while the second and third panel consider scenarios with increased uncertainty as

[32]For a central bank learning problem that allows the possibility of incomplete learning the reader is referred to Wieland (2000b).

Figure 7: Alternative Policies with Unknown Intercept α and Slope β
Prior Belief (a=1.68, b=-0.28) - Different Extent of Forward-Looking Expectations

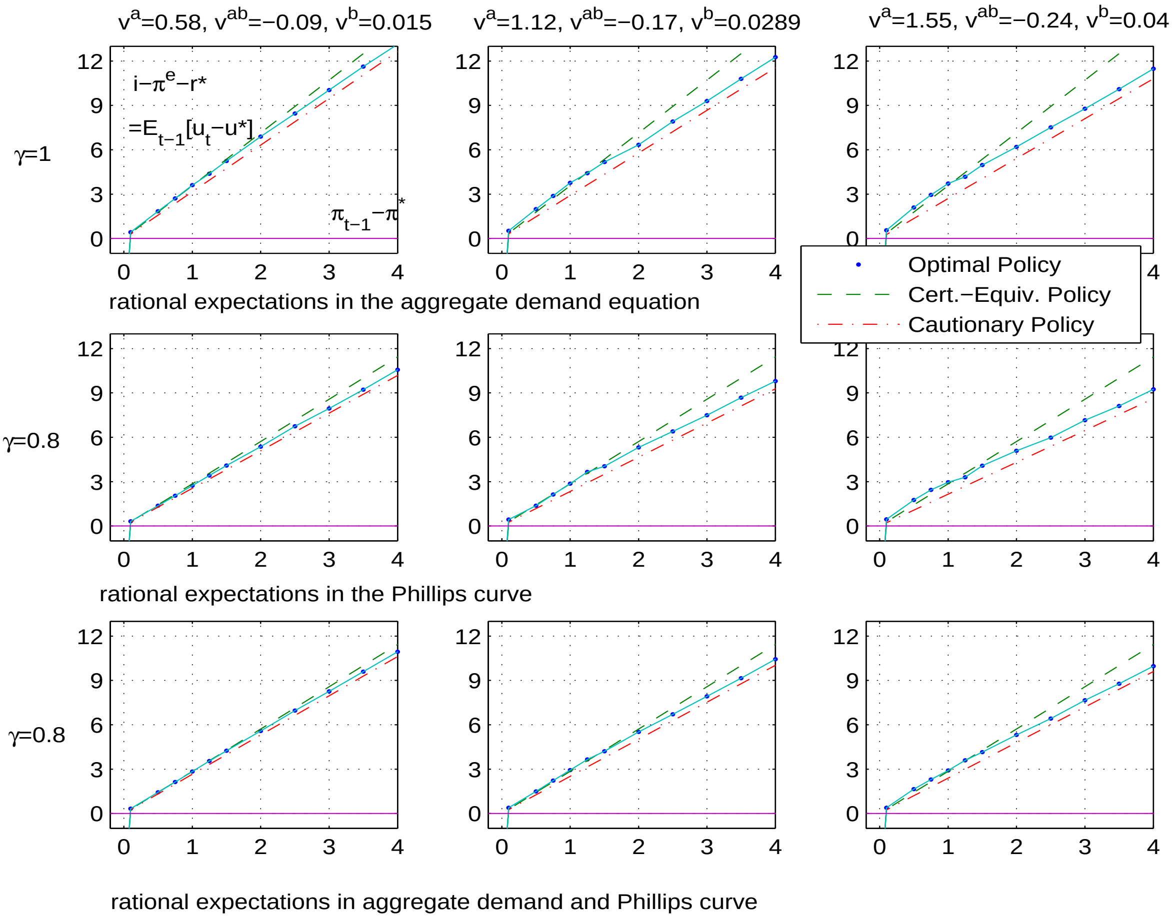

in Figure 5. I introduce forward-looking behavior into the analysis in three steps. First, I allow for forward-looking expectations in the aggregate demand equation, (6), i.e. in terms of the definition of the real interest rate, while maintaining the assumption of adaptive expectations in the Phillips curve (first row of panels). Secondly, I consider forward-looking expectations in the Phillips curve, while maintaining adaptive expectations in the aggregate demand equation (middle row of panels). The index of inflation persistence γ is set to 0.8 implying a weight of 0.2 on forward-looking expectations. In a third step I consider forward-looking expectations, both, in the Phillips curve and in the aggregate demand equation (bottom row of panels).

Three findings are directly apparent. First, all three policies are less aggressive in their response to lagged inflation due to the presence of the expectations channel of monetary policy transmission. Second, the qualitative properties of the optimal rule are the same as under the case of adaptive expectations depicted in Figure 5. Third, the optimal extent of experimentation, that is, the difference between the cautionary and the optimal rule, is somewhat smaller. The intuitive reason is that a change in inflation expectations has a direct effect on inflation. Thus, the central bank does not need to rely as much on the nominal interest rate and its uncertain impact on inflation. A further extension would be to consider learning about the index of inflation persistence γ, which governs the importance of forward-looking expectations in the Phillips curve. It turns out that the qualitative properties of optimal policy remain the same. For policy comparisons with learning about γ and β the interested reader is referred to the earlier working paper version (Wieland (2003)).

9 Conclusions and extensions

The main conclusion from 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 policy-making 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 has studied policy rules for a strictly inflation targeting central bank. 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:

$$L(\pi_t) = E_{t-1}\left[(\pi_t - \pi^*)^2 + \omega(u_t - u^*)^2)\right] \tag{43}$$

The unemployment stabilization objective introduces an alternative motive for the central bank to respond gradually to inflationary shocks. The optimal interest rate rule in the static model under certainty with rational expectations then corresponds to:

$$i_t = u^* + \pi^* - (\beta^2 + \omega\gamma^2)^{-1}(\beta\gamma - \omega\gamma^2)(\pi_{t-1} - \pi^*) \tag{44}$$

For $\omega = 0$ this rule simplifies to the optimal rule under strict inflation targeting, equation (14). 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 following a shock. Policy rules under uncertainty can be computed in the same manner as for strict inflation targeting. However, the cautionary rule, which is optimal in the static version of the model without learning, can only be computed by numerical methods due to the non-normal distribution of u^*. One can show that 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 (4) treating the private sector agents' expectation of the interest rate, i_t^e, and thus their inflation expectation π_t^e, as constant and independent of monetary policy. With known parameters this corresponds to:

$$i_t = -(\beta)^{-1}(\gamma\pi_{t-1} + (1-\gamma\beta)\pi_t^e - \beta u^* - \pi^*) \tag{45}$$

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 (45) 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 (14). 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.[33]

Demand uncertainty

While studying the implications of uncertainty about the relationship between unemployment and inflation in much detail, the remainder of the model has been treated as

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

known with certainty. An extension would allow for uncertainty due to random shocks and imprecisely estimated parameters in the Okun's law and aggregate demand relationships as follows:

$$u_t = \phi q_t + e_t^u \tag{46}$$

$$q_t = \lambda(i_t - \pi_t^e) + e_t^q. \tag{47}$$

Beliefs concerning ϕ and λ may then be characterized by normal distributions $N(p, v^p)$ and $N(l, v^l)$. The presence of the random shocks e_t^u and e_t^q would render estimation of ϕ and λ nontrivial. Uncertainty regarding these parameter estimates would further increase the component of inflation uncertainty that is influenced by monetary policy. Thus, it would 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 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 the central bank and forward-looking market participants have the same information set available when making decisions. As a consequence of this assumption, agents and the 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 from differences between agents and central banks in 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 s_t about the inflation shock e_t into the model of this paper. The optimal policy rule will then include a policy response to this signal very much like the policy response to lagged inflation. The private sector's rational expectation of this component of the interest rate rule will be equal to zero. Thus, market participants will only be able to predict the component of the policy rule that responds to lagged inflation. Since the signal s_t does not help in estimating the Phillips curve parameters, the updating equations for beliefs 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 market participants' 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.

Appendix: 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:

$$TW = \underset{i}{Min} \left[L(\pi, i, \theta) + \delta \int W(\pi', i, \theta') f(\pi' | \pi, i, \theta) d\pi' \right] \tag{48}$$

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. $W(.)$ is a continuous function defined on the state space. $L(.)$ 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 transition equations for these state variables are summarized in (39) and include those for beliefs in (38). $f(\pi' | \pi, i, \theta)$ is the predictive distribution of the inflation rate. It is a normal distribution, because the beliefs are normal distributions and the random shocks are also normally distributed.

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 as follows:

$$\rho(W_n, W_{n+1}) = \underset{(\theta, \pi)}{\text{Sup}} |W_n(\theta, \pi) - W_{n+1}(\theta, \pi)| \tag{49}$$

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:

$$\rho(TW_{n+1}, TW_n) \leq \delta \rho(W_{n+1}, W_n) \tag{50}$$

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 controls i, which minimize the right-hand side of (48) based on the current state (π, θ).

V can be computed by value iteration, meaning successive application of the operator T, since $T_n W \to V$ uniformly for any continuous bounded function W. A convenient starting value W_0 is the single period loss function $L(.)$ but alternatively a constant also suffices. If $W_{n+1} = TW_n$, then $\rho(W_{n+1}, W_n) \leq (W_n, W_{n-1})$ 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:

$$\rho(V, W_n) \leq \sum \rho(W_{n+1+i}, W_{n+i}) \leq \frac{\delta}{1-\delta} \rho(W_n, W_{n-1}) \tag{51}$$

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:

$$T^P W_n = L(\pi, H_n(\pi, \theta), \theta) + \delta \int W(\pi', H_n(\pi, \theta), \theta') f(\pi' | \pi, H_n(\pi, \theta), \theta) d\pi' \tag{52}$$

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 i 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:

$$\int W_0(\pi', u, \theta') f(\pi'|\pi, i, \theta) d\pi' \tag{53}$$

The functions $W(.)$ 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 $W(.)$ 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 Bellman 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 i gives the value of $W_1()$ 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 has 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;[34] (iii) the algorithm is written in FORTRAN in order 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 $10 * 14 * 10 * 36 * 60$ gridpoints. In this case, I also used 60-point Gaussian quadrature with respect to the shock e. Convergence as defined by a 0.5% maximal difference between the two final value function approximations for these problems was achieved after about 72 hours on a 2.21 GHz AMD Processor with 1.48 GB RAM. Typically convergence required 6 to 8 value iterations with a declining number of policy iterations (250 or fewer) in between every value iteration.

[34]For a discussion of this reformulation see the appendix of Wieland (2000a).

References

Akerlof, G., W. Dickens and G. Perry, "The Macroeconomics of Low Inflation", *Brookings Papers on Economic Activity*, 1996, 1, pp. 1-59.

Aghion, P., P. Bolton, C. Harris and B. Jullien, "Optimal Learning by Experimentation", *Review of Economic Studies*, 1991, 58, pp. 621-654.

Amman, H. and D. Kendrick, "Nonconvexities in Stochastic Control Models", *International Economic Review*, 1995, 36, pp. 455-475.

Anderson, T. and J.B. Taylor, "Some Experimental Results on the Statistical Properties of Least Squares Estimates in Control Problems", *Econometrica*, 1976, 44, pp. 1289-1302.

Balvers, R. and T. Cosimano, "Inflation Variability and Gradualist Monetary Policy", *Review of Economic Studies*, 1994, 61, pp. 721-738.

Bertocci, G. and M. Spagat, "Learning, Experimentation and Monetary Policy", *Journal of Monetary Economics*, 1993, 32, pp. 169-178.

Blanchard, O. and L. Katz, "What We Know and Do Not Know About the Natural Rate of Unemployment", *Journal of Economic Perspectives*, 1997, 11, pp. 51-72.

Blinder, A., "Central Banking in Theory and Practice, Lecture I: Targets, Instruments, and Stabilization", Marshall Lecture presented at the University of Cambridge, 1995.

Blinder, A., *Central Banking in Theory and Practice*, January 1998, MIT Press.

Brainard, W., "Uncertainty and the Effectiveness of Policy", *American Economic Review*, 1967, 57, pp. 411-25.

Buiter, W. and M. Miller, "Costs and Benefits of an Anti-Inflationary Policy: Questions and Issues", 1985, in V. Argy and J. Neville, eds., *Inflation and Unemployment: Theory, Experience, and Policy-Making.*

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

Clarida, R., J. Gali and M. Gertler, "The Science of Monetary Policy", *Journal of Economic Literature*, 1999, 37, pp. 1661-1707.

Clarida, R., J. Gali and M. Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory", *Quarterly Journal of Economics*, 1998, 42, pp. 1003-1067.

Clark, P., C. Goodhart and H. Huang, "Optimal Monetary Policy Rules in a Rational Expectations Model of the Phillips Curve", *Journal of Monetary Economics*, 1999, 43, pp. 497-520.

Debelle, G. and D. Laxton, "Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom, and the United States", *IMF Staff Papers*, 1997, 44, pp.

249-282.

Dornbusch, R. and S. Fischer, *Intermediate Macroeconomics*, 1990, McGraw-Hill.

Easley, D. and N.M. Kiefer, “Controlling a Stochastic Process with Unknown Parameters”, *Econometrica*, 1988, 56, pp. 1045-1064.

Ellison, M. and N. Valla, “Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions”, *Journal of Monetary Economics*, 2001, 48, pp. 153-171.

Estrella, A. and F. Mishkin,“The Role of NAIRU in Monetary Policy: Implications of Uncertainty and Model Selection”, 1999, in Taylor, J.B., ed., *Monetary Policy Rules*, NBER and Chicago Press.

Fuhrer, J., “The (Un) Importance of Forward-Looking Behavior in Price Specifications”, *Journal of Money, Credit and Banking*, 1997, 29, pp. 338-50.

Fuhrer, J., “The Phillips Curve is Alive and Well”, *New England Economic Review*, March/April 1995 , pp. 41-56.

Fuhrer, J. C. and G. R. Moore, “Inflation Persistence”, *The Quarterly Journal of Economics*, 1995, 110, pp. 127-159.

Gali, J. and M. Gertler, “Inflation Dynamics: A Structural Econometric Analysis”, *Journal of Monetary Economics*, 1999, 44, pp. 195-222.

Gordon, R., “The Time-Varying NAIRU and its Implications for Economic Policy”, *Journal of Economic Perspectives*, 1997, 11, pp. 11-32.

Greene, W., *Econometric Analysis*, 1993, 2nd edition, Macmillan.

Hansen, L. and T. Sargent, ”Robust Control and Model Uncertainty”, *American Economic Review*, 2001, 91, pp. 60-66.

Harvey, A., *Time Series Models*, 1992, 2nd edition, MIT Press.

Kasa, K., ”Will the Fed Ever Learn?”, *Journal of Macroeconomics*, 1999, 21, pp. 279-292.

Keller, G. and S. Rady, ”Optimal Experimentation in a Changing Environment”, *Review of Economic Studies*, 1999, 66, pp. 475-507.

Kendrick, D., Stochastic Control For Economic Models, 1981, *Economic Handbook Series*, McGraw-Hill.

Kendrick, D., “Caution and Probing in a Macroeconomic Model”, *Journal of Economic Dynamics and Control*, 1982, 15 (2), pp. 149-170

Kiefer, N. and Y. Nyarko, “Optimal Control of an Unknown Linear Process with Learning”, *International Economic Review*, 1989, 30, pp. 571-586.

Levin, A., J.C. Williams and V. Wieland, “The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty”, *American Economic Review*, 2003, 93, pp. 622-645.

Levin, A., J.C. Williams and V. Wieland, ”Robustness of Simple Monetary Policy Rules under Model Uncertainty”, June 1999, in John B. Taylor (ed.), *Monetary Policy Rules*, NBER and Chicago Press.

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

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

Nyarko, Y., “The Truth is in the Eye of the Beholder: or Equilibrium in Beliefs and Rational Learning in Games”, manuscript, 1998, New York University.

Phelps, E.S and G. Zoega, “The Rise and Falling Trend of the Natural Rate”, *American Economic Review*, 1997, 87, pp. 283-289.

Prescott, E., “The Multi-Period Control Problem under Uncertainty”, *Econometrica*, 1972, 40, pp. 1043-1058.

Roberts, J., “Is Inflation Sticky?”, *Journal of Monetary Economics*, 1997, 39, pp. 173-196.

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

Sack, B., “Does the Fed Act Gradually? A VAR Analysis”, *Journal of Monetary Economics*, 2000, 46, pp. 229-256.

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

Sargent, T., *The Conquest of American Inflation*, 1999b, Princeton University Press.

Sims, C. A., “Projecting Policy Effects with Statistical Models”, *Revista de Analisis Economico*, 1988, 3, pp. 3-20.

Staiger, D., J. Stock and M. Watson, “How Precise are Estimates of the Natural Rate of Unemployment?”, 1997a, in Romer, Christina and David Romer, eds., *Reducing Inflation: Motivation and Strategy*, NBER and Chicago Press.

Staiger, D., J. Stock, and M. Watson, “The NAIRU, Unemployment and Monetary Policy”, *Journal of Economic Perspectives*, 1997b, 11, pp. 33-49.

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

Stiglitz, J., “Reflections on the Natural Rate Hypothesis”,*Journal of Economic Perspectives*, 1997, 11, pp. 33-49.

Svensson, L.,“Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets”, *European Economic Review*, 1997a, 41, pp. 1111-1146.

Svensson, L.,“Optimal Inflation Targets, ’Conservative’ Central Banks, and Linear Inflation Contracts”, *American Economic Review*, 1997b, 87, pp. 98-114.

Svensson, L.,“Inflation Targeting: Some Extensions”,*Scandinavian Journal of Economics*, 1999, 101, 3, pp. 337-361.

Taylor, J.B., “Asymptotic Properties of Multi-Period Control Rules in the Linear Regression Model”, *International Economic Review*, 1974, 15, pp. 472-484.

Taylor, J.B., “Discretion versus Policy Rules in Practice”, *Carnegie-Rochester Conference Series on Public Policy*, 1993, 39, pp. 195-214.

Taylor, J.B., ed., *Monetary Policy Rules*, 1999, NBER and Chicago Press.

Wieland, V., “Learning by Doing and the Value of Optimal Experimentation”, *Journal of Economic Dynamics and Control*, 2000a, 24, pp. 501-534.

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

Wieland, V., “Monetary Policy and Uncertainty about the Natural Unemployment Rate”, Finance and Economics Discussion Series, 1998, 1998-22.

Wieland, V., “Monetary Policy and Uncertainty about the Natural Unemployment Rate”, 2003, CEPR Discussion Paper, 3811.

Yetman, J., “Probing Potential Output: Monetary Policy, Credibility, and Optimal Learning under Uncertainty”, *Journal of Macroeconomics*, 2003, 25, pp. 311-330.

Zellner, A., *Introduction to Bayesian Inference in Econometrics*, 1971, Wiley.