

DISCUSSION PAPER SERIES

No. 4222

**PRUDENT MONETARY POLICY:
APPLICATIONS OF CAUTIOUS LQG
CONTROL AND PREDICTION**

Frederick van der Ploeg

INTERNATIONAL MACROECONOMICS



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP4222.asp

PRUDENT MONETARY POLICY: APPLICATIONS OF CAUTIOUS LQG CONTROL AND PREDICTION

Frederick van der Ploeg, European University Institute and CEPR

Discussion Paper No. 4222
February 2004

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

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

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

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

Copyright: Frederick van der Ploeg

ABSTRACT

Prudent Monetary Policy: Applications of Cautious LQG Control and Prediction*

Risk-adjusted linear-quadratic-Gaussian (LQG) optimal control with perfect and imperfect observation of the economy gives prudent feedback policy rules and cautious Kalman filters. The reaction coefficients depend on the variances of the modelling errors, since precaution reduces the power of its instruments by these variances. Also, prudence implies less weight to new uncertain observations that are less important for welfare and introduces a bias in prediction. Prudent predictions are thus neither efficient nor unbiased. A prudent central bank adjusts the nominal interest rate more aggressively to changes in the inflation gap, especially if the volatility of cost-push shocks is large. If the interest rate impacts the output gap after a lag, the interest also responds to the output gap, especially with strong persistence in aggregate demand. Prudence pushes up this reaction coefficient as well. If data are poor and appear with a lag, a prudent central bank responds less strongly to new measurements of the output gap. Prudence attenuates this policy reaction and biases the prediction of the output gap upwards, however, particularly if output targeting is important. Finally, prudence requires an extra upward (downward) bias in its estimate of the output gap before it feeds into the policy rule if inflation is above (below) target. This reinforces nominal interest rate reactions.

JEL Classification: C60, D80, E40, E50 and E60

Keywords: measurement errors, optimal monetary policy, prediction, prudence and Taylor rules

Frederick van der Ploeg
Economics Department
European University Institute
Villa San Paolo
Via della Piazzuola 43
50133 Florence
ITALY
Email: rick.vanderploeg@iue.it

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=100358

*I thank Mike Artis, Roel Beetsma and seminar participants at CentER, Tilburg University and the Catholic University of Milan for helpful comments.

Submitted 05 December 2003

1. Introduction

One of the most precious commodities of a modern capitalist economy is a stable price level or at least a low inflation rate. Central banks also attempt to avoid unemployment and steer towards near-zero output gaps. This form of flexible inflation targeting gives rise to Taylor (1993) rules, which indicate how much the nominal interest rate should react to the inflation gap and the output gap. One of the main features of a central bank is that it operates cautiously and prudently:

“A prudent man (or perhaps, I should say a prudent Bayesian) carries an umbrella even when the forecast says there is only a small chance of rain. If there is no rain, he suffers the small inconvenience of carrying the umbrella. But if he does not bring the umbrella and it does rain, he may suffer the much larger inconvenience of being caught in a downpour. The prudent central bank should behave similarly, accepting a high probability of a small adverse outcome in order to avoid the small risk of a very serious bad outcome” (Feldstein, 2003).

Feldstein argues that Greenspan’s policy of lowering the federal interest rate from 1.75 per cent to 1.25 per cent last year was prudent, because of the asymmetric nature of the risk faced at that time. The potential upturn could loose steam and there was a risk of deflation, while an unnecessarily strong stimulus could do little harm at the time. A large body of research on optimal monetary policy is based on the certainty equivalence principle, which says that uncertainty can be ignored. In calculating optimal interest rate rules future disturbances are set to their expected values. This approach is only valid under very special conditions (i.e., linear models, quadratic preferences, normally distributed errors). It abstracts from prudence and thus bears little relation to the practice of central banking.

It is surprising that there is so little research on the behaviour of prudent central banks. This requires departure from the certainty-equivalent linear-quadratic-Gaussian framework that is usually adopted - e.g., Svensson (1997), Rudebusch and Svensson (1999), Judd and Rotemberg (1998), Rotemberg and Woodford (1997, 1999), Woodford (2001). One approach is to explicitly recognise that statistical properties and order of the processes driving the modelling disturbances are not known and to derive robust (min-max) rules that perform well under different views of the world – e.g., Onatski and Stock (2000), Giannoni and Woodford (2002), and Onatski and Williams (2003). Another approach is to employ model averaging in a Bayesian context – see Brock, Durlauf and West (2003). Yet another approach is to allow room for judgement of central bankers in the derivation of optimal monetary policy rules – see Svensson (2002).

Our approach is complementary. We allow for precautionary behaviour of central banks and derive closed-loop monetary policy rules analytically. The objective is to minimise the expected value of an exponential transformation of the quadratic welfare loss criterion.

The coefficient of the exponential transformation corresponds to an Arrow-Pratt measure of absolute risk aversion, which we call the degree of prudence. This leads to linear policy rules with reaction coefficients that depend on the covariance matrices of the stochastic process driving the modelling disturbances - see Jacobson (1973, 1977), Speyer, Deyst and Jacobson (1974), and Whittle (1981, 1982 and 1990). Effectively, a prudent policy maker downplays the power of its instruments if the volatility of shocks hitting the economy is large.

The derivation of monetary policy rules must recognise that national accounts consist of poor quality data. They contain many measurement errors and observation lags. This problem is particularly severe for output data. Typically, various 'flash' estimates of GDP appear fairly quickly and are then subsequently substantially revised. Measurement errors show up, because the raw data do not satisfy the national accounting identities. Various studies have used subjective estimates of the reliability of data to adjust the data under the restriction that all the accounting identities must be satisfied – e.g., van der Ploeg (1982) and Barker, van der Ploeg and Weale (1984). The subjective variances of the raw data are provided by the national accountants and subsequently reduced by imposing the accounting restrictions. Unfortunately, they are seldom used in econometric analysis or in the derivation of optimal economic policy rules. Here we allow for measurement errors and lags in the observation of output data - see also Orphanides (2000). If the central bank adjusts its interest rate in reaction to changes in output gaps, it presumably does this less intensively if substantial measurement errors and lags in output data cause a deterioration of the signal-to-noise ratio - cf., Rudebusch (2001). Taylor rules also allow for reactions to changes in inflation. However, inflation data are more readily and accurately available than output data.

We investigate how measurement errors and lags in output data affect the Taylor rule. Pearlman (1986, 1992) demonstrated the usefulness of the Kalman filter for predicting the states of the economy in monetary models with forward-looking expectations. In backward-looking models and forward-looking models where policy makers and private agents have access to the same partial information sets, the Kalman filter calculations can be performed independently of deriving the optimal monetary policy rule. The separation of control and prediction is more tricky in forward-looking monetary models with commitment and asymmetric information – see also Svensson and Woodford (2003).

In general, optimal predictions of the output gap make use of wage and price data. We analyse how a prudent central banker takes account of incoming unreliable output data and uses this in the Taylor rule. A modified separation principle holds - see Whittle (1981, 1982 and 1990). The prudent Kalman filter depends on welfare preferences and yields biased predictions. In particular, a prudent policy maker gives less weight to new observations with large standard errors and that are less relevant for welfare. Conversely, to avoid costly mistakes prudence requires more weight to faulty data that are relevant for welfare. In the

umbrella example a prudent person assigns a larger subjective probability of rain than the objective probability of rain, especially if he or she dislikes rain a lot.

Section 2 presents the problem of risk-adjusted LQG control and prediction. Prudence implies that the policy maker plays a min-max game against nature. The policy maker hedges against undesirable outcomes by postulating that shocks damage its objectives even though, from a purely statistical point of view, they do not hurt on average. It still leads to linear feedback policy rules and a recursive scheme for the prediction of state variables. The Appendix gives explicit solutions. The main differences are that policy rules depend on the covariance matrices of state disturbances and the recursive prediction scheme depends on the penalty matrices of the welfare loss criterion. This scheme yields inefficient and biased predictions of the state variables. Section 3 shows how prudence affects the optimal inflation-output trade-off, given that the central bank faces an expectations-augmented accelerationist Phillips curve and no measurement errors and lags. We show that the optimal nominal interest rate of a prudent central bank reacts more aggressively to the inflation gap, especially if cost-push disturbances are volatile. Section 4 derives a prudent Taylor rule if the real interest rate impacts aggregate demand after one period. We demonstrate that the optimal interest rate again responds more aggressively to the output gap if prudence and volatility of cost-push shocks are large and if there is substantial persistence in aggregate demand. We also show that more weight to output targeting weakens policy responses of the central bank, particularly if triggered by changes in the inflation gap. Section 5 allows for measurement errors and lags in observing output data. We show that the reactions of the nominal interest rate to the measured output gap are less strong, especially if incoming data are relatively unreliable. We also show that a prudent central bank attenuates these policy reactions and furthermore biases its estimate of the output gap upwards. This makes the reactions of the central bank to the output gap more aggressive, particularly if cost-push shocks are volatile and output targeting is important. Finally, we show that a prudent central bank introduces an extra upward (downward) bias in its estimate of the output gap to be fed into the policy rule if inflation is above (below) target. This makes the nominal interest rate reactions more aggressive. Section 6 concludes and offers some suggestions for further research.

2. Risk-adjusted LQG control and prediction

Jacobson (1973, 1977), Speyer, Deyst and Jacobson (1974), and Whittle (1981, 1982 and 1990) have extended the standard theory of optimal linear-quadratic-Gaussian (LQG) control, discussed by Chow (1975), Aoki (1976), Kendrick (1981), Preston and Pagan (1982) and others, to allow for risk-sensitivity. Van der Ploeg (1984a,b) discusses risk-sensitive LQG control with perfect observation of the states of the economy and applies it to the derivation of the optimal inflation-output trade-off, Keynesian demand management, monopoly

behaviour and the management of exhaustible resources. Hansen and Sargent (1995) extend the analysis to allow for discounting. Van der Ploeg (1993) applied these ideas to the analysis of precautionary saving - see also Hansen, Sargent and Tallarani (1999). There is a brief mention of risk-sensitive control in the textbook of Ljungqvist and Sargent (2000).

Here we allow for imperfect observation of states of the economy. We base our presentation on Whittle (1981, 1982). Consider the quadratic welfare loss criterion:

$$(1) \quad \Gamma \equiv \sum_{t=0}^T \gamma_t \quad \text{with} \quad \gamma_t \equiv (x_t - x^*)' Q (x_t - x^*) + (u_t - u^*)' S (u_t - u^*), \quad 0 \leq t < T$$

and $\gamma_T \equiv (x_T - x_T^*)' Q_T (x_T - x_T^*),$

where the vectors x_t and x^* contain the actual and desired (or bliss) values for the state variables, respectively, the vectors u_t and u^* contain the actual and desired values for the policy instruments, respectively, and Q and S denote symmetric and positive-definite penalty matrices. It is straightforward to allow for cross products of x_t and u_t in the welfare loss function. The state variables follow recursively from the state-space model:

$$(2) \quad x_t = A x_{t-1} + B u_{t-1} + a_t + \varepsilon_t, \quad \varepsilon_t \sim \text{IN}(0, \Sigma), \quad x_0 \sim \text{N}(x_{p0}, V_0),$$

where A denotes the state-transition matrix, B stands for the matrix of policy impact multipliers and a_t is the vector of deterministic forcing terms. The vector of serially uncorrelated, normally distributed modelling errors ε_t has zero mean and covariance matrix Σ . These error terms do not become known to the policy maker until period t , so he can only react to *past* realisations of state variables. Higher-order lag structures in state variables and in policy instruments and/or more general ARIMA-error structures can be dealt with by suitably augmenting the vector x_t . The initial state variables conditional on the information available at time zero may not be known perfectly.

We define θ as the Arrow-Pratt measure of constant absolute risk aversion with respect to welfare $-\frac{1}{2}\Gamma$. The policy maker thus minimises the following criterion:

$$(3) \quad \Phi(\theta) \equiv (2/\theta) \log E[\exp(\frac{1}{2} \theta \Gamma)] \quad \text{with} \quad \theta > 0.$$

The risk-neutral case corresponds to $\Phi(\theta) \rightarrow E[\Gamma]$ as $\theta \rightarrow 0$. (The case of a risk-loving policy maker corresponds to $\theta < 0$.) The policy maker minimises the expected risk-adjusted welfare loss to go. A prudent policy maker also penalises variability in welfare. In fact, for small $\theta \text{var}(\Gamma)$ the decision maker approximately minimises $E(\Gamma) + \frac{1}{4}\theta \text{var}(\Gamma)$. Two policy makers

may share the same welfare loss criterion under certainty, even though their aversion to risk may differ. The coefficient θ captures this specific risk aversion. It makes the policy maker particularly sensitive to occasional large deviations of Γ from $E(\Gamma)$. A policy maker with a positive θ is called pessimistic, cautious or prudent.

With no observation lags and no measurement errors the policy instruments u_t can be conditioned on x_t . Often not all the state variables are observable and other variables are measured imperfectly after a lag. We thus assume that the measured variables at time t , z_t , are:

$$(4) \quad z_t = C x_{t-1} + \omega_t, \quad \omega_t \sim \text{IN}(0, \Omega),$$

where ω_t is the vector of serially uncorrelated, normally distributed measurement errors with zero mean and covariance matrix Ω . We assume that modelling errors ε_t and measurement errors ω_t are uncorrelated. The policy instruments at time t react to measurements at time t , z_t , of the imperfectly measured past states of the economy, x_{t-1} . The case of no measurements of certain states of the economy is also captured by (4). To allow for observation lags longer than one period, suitably augment the vector of state variables.

Define the residual sum of squares (the quadratic function in the exponent of the joint density of state and measured variables) as:

$$(5) \quad \Delta \equiv (x_0 - x_{p0})' V_0^{-1} (x_0 - x_{p0}) + \sum_{t=0}^{T-1} (\delta_{xt} + \delta_{zt})$$

with $\delta_{xt} \equiv (x_{t+1} - A x_t - B u_t - a_{t+1})' \Sigma^{-1} (x_{t+1} - A x_t - B u_t - a_{t+1})$
and $\delta_{zt} \equiv (z_{t+1} - C x_t)' \Omega^{-1} (z_{t+1} - C x_t)$.

Stress occurs if states and policy instruments deviate from their desired values as measured by Γ . Total stress is defined as:

$$(6) \quad \Psi \equiv \Gamma - \Delta/\theta.$$

Hence, if there is no uncertainty, total stress simply equals the quadratic welfare loss. If there is uncertainty in the state dynamics or in the measurements of the state variables, $\Delta > 0$ and total stress is reduced, particularly if the degree of prudence θ is small.

Theorem 1 (Saddlepoint solution and stress):

Suppose total stress (6), given (1) and (5), is freely minimised with respect to the instruments u_t, \dots, u_{T-1} and freely maximised with respect to the state variables x_0, \dots, x_T and the measured

variables z_{t+1}, \dots, z_T conditional on all information available at time t , say I_t . The order in which these operations are carried out is irrelevant. The value of u_t thus obtained is the optimal value of the vector of policy instruments at time t . The calculations only have meaning if θ is not too large, that is if total stress (6) is negative definite in the maximising variables.

Proof: Whittle (1981, Chapter 19, Theorem 3.1).

Standard certainty equivalence and separation hold for the risk-neutral case. The optimal values for the instruments at time t , u_t , can thus be obtained by minimising the welfare cost to go Γ from time t onwards, where the unknown disturbance terms from time $t+1$ onwards are set to their expected values conditional on the information available at time t , I_t . These expected values minimise the residual sum of squares Δ given I_t and yield BLUE-predictions of the future state variables.

With risk aversion the expected *objective* values for the disturbances at time t are replaced by prudent *subjective* values that maximise total stress Ψ given I_t . These prudent values are estimates of current observable variables that maximise stress. The values of u_{t+1}, \dots, u_{T-1} in Theorem 1 are chosen to minimise stress; they are current estimates of optimal future decisions. The policy maker effectively plays a min-max non-cooperative game against nature. He treats nature as a belligerent player producing damaging shocks and designs a min-max strategy for the ‘worst possible state of the world’. He hedges against unanticipated increases in the welfare loss caused by unfavourable shocks, particularly if prudence is large. (A love for risk, $\theta < 0$, implies that the policy maker believes that nature helps to achieve its objectives, so a min-min strategy for a cooperative game with nature is called for.)

To gain more insight, define past stress at time t as

$$(7) \quad \Psi_{Pt}(x_t, I_t) = -\theta^{-1} (x_0 - x_{P0})' V_0^{-1} (x_0 - x_{P0}) + \sum_{s=0}^{t-1} [\gamma_s - \theta^{-1} (\delta_{xs} + \delta_{zs})]$$

and future stress at time t as

$$(8) \quad \Psi_{Ft}(x_t) = \underset{u_t, \dots, u_{T-1}}{\text{Min}} \underset{x_{t+1}, \dots, x_T}{\text{Max}} \underset{z_{t+1}, \dots, z_T}{\text{Max}} \sum_{s=t}^{T-1} [\gamma_s - \theta^{-1} (\delta_{xs} + \delta_{zs})] + \gamma_T$$

$$= \underset{u_t, \dots, u_{T-1}}{\text{Min}} \underset{x_{t+1}, \dots, x_T}{\text{Max}} \sum_{s=t}^{T-1} (\gamma_s - \theta^{-1} \delta_{xs}) + \gamma_T.$$

Expression (7) and (8) obey, respectively, a forward and a backward recursion:

$$(7') \quad \Psi_{P_{t+1}}(x_{t+1}, I_{t+1}) = \underset{x_t}{Max} [\Psi_{P_t}(x_t, I_t) + \gamma_t - \theta^{-1} (\delta_{xt} + \delta_{zt})]$$

$$\text{with } \Psi_{P_0}(x_0, I_0) = -\theta^{-1} (x_0 - x_{P_0})' V_0^{-1} (x_0 - x_{P_0})$$

$$(8') \quad \Psi_{F_t}(x_t) = \underset{u_t}{Min} \underset{x_{t+1}}{Max} [\gamma_t - \theta^{-1} \delta_{xt} + \Psi_{F_{t+1}}(x_{t+1})]$$

$$\text{with } \Psi_T(x_T) = \gamma_T = (x_T - x_T^*)' Q_T (x_T - x_T^*).$$

The forward recursion (7') corresponds to a risk-adjusted Kalman filter. In contrast to the standard risk-neutral case, the filter (7') depends on preference parameters and gives prudent estimates of states of the economy as new information comes in. The backward recursion (8') corresponds to the risk-adjusted Bellman equation and revises the current policy rule as the horizon recedes by a single period. In contrast to the risk-neutral case, the optimal policy rules depend on the covariance matrices of the modelling disturbances. Hence, uncertainty independence (certainty equivalence) of the optimal policy rules and preference indifference of the optimal state predictions no longer holds. Theorem 1 can be rephrased to show that a modified separation principle holds for risk-adjusted LQG control and prediction.

Theorem 2 (Min-max solution and separation of control and prediction):

Let $u(x_t, t)$ be the minimising value of u_t in equation (8') and let x_{Mt} be the value of x_t that maximises $\Psi_{P_t}(x_t, I_t) + \Psi_{F_t}(x_t)$, a function of I_t . Then the optimal value of u_t is $u(x_{Mt}, t)$, a function of I_t .

With imperfect observation of states, the optimal instruments at time t depend on the maximum-stress estimate x_{Mt} . If states are perfectly observed, the optimal control rule is simply $u_t = u(x_t, t)$. For many problems in economics it is more convenient to exploit the explicit model structure and apply Theorem 2 directly. The Appendix gives the recursive solutions to the general problem of risk-adjusted LQG control and prediction.

3. Prudence and the optimal inflation-output trade-off

To illustrate prudence in the optimal inflation-output trade-off, consider the accelerationist Phillips-curve, e.g., Phelps (1967):

$$(9) \quad \pi_{t+1} = \pi_t + \alpha y_t + \varepsilon_{\pi t+1} \quad \text{with } \alpha > 0, \quad \varepsilon_{\pi t} \sim \text{IN}(0, \sigma_\pi^2),$$

where π_t and y_t denote the inflation rate and the output gap at time t , respectively, and $\varepsilon_{\pi t}$ is a serially uncorrelated, normally distributed cost-push disturbance with zero mean and variance σ_π^2 . The long-run Phillips curve is vertical, so that systematic deviations from the natural rate of unemployment lead to an ever-ending spiral of inflation or deflation.

The intra-temporal welfare loss function allows for flexible inflation targeting. It is given by $\gamma_t = (\pi_t - \pi^*)^2 + \kappa y_t^2$, where π^* indicates desired inflation (say, 2% per year) and $\kappa > 0$ is the weight given to full employment (output) targeting. Strict inflation targeting corresponds to $\kappa = 0$. The equilibrium level of employment and output is efficient, so the desired output gap is zero. We thus abstract from extra inflation bias induced by time inconsistency problems as in Kydland and Prescott (1977) and Barro and Gordon (1983) and the consequent need for an ultraconservative central banker as in Rogoff (1985). The government can by appropriate use of fiscal and monetary policy immediately control the output gap, so y_t is the intermediate policy instrument to attain price stability and full employment. Later we allow the nominal interest rate to be the monetary policy instrument.

Although the infinite-horizon limits of the prudent policy rules may be derived from Theorem A.2, it is instructive to apply Theorem 2 directly. The postulated cost-push disturbance and the optimal output gap at time t follow from the min-max problem:

$$(10) \quad \Psi_{Ft}(\pi_t) = \underset{y_t}{Min} \underset{\varepsilon_{\pi t+1}}{Max} [(\pi_t - \pi^*)^2 + \kappa y_t^2 + \Psi_{Ft+1}(\pi_t + \alpha y_t + \varepsilon_{\pi t+1}) - (\varepsilon_{\pi t+1}^2 / \theta \sigma_\pi^2)].$$

If we postulate $\Psi_{Ft}(\pi_t) = R(\pi_t - \pi^*)^2$, the postulated ‘worst case’ inflation shock is given by:

$$(11) \quad \varepsilon_{M\pi t+1} = \theta \sigma_\pi^2 R_A (\pi_t + \alpha y_t - \pi^*) \quad \text{with} \quad R_A \equiv R / (1 - \theta \sigma_\pi^2 R) \geq R > 0,$$

where $\theta \sigma_\pi^2 R < 1$ must hold for (11) to be a maximum. R and R_A indicate the infinite-horizon limits of R_t and R_{At} . A prudent central bank assumes the ‘worst’ by postulating that cost-push disturbances hit the economy adversely even though expected values of these shocks are zero. When the *objective* expectation of future inflation, $\pi_{t+1|t} \equiv E(\pi_{t+1} \text{ given } I_t)$, is above (below) target, the central bank biases its *subjective* estimate of future inflation, $(\pi_{t+1|t} + \varepsilon_{M\pi t+1})$, upwards (downwards), particularly so if prudence is important and the volatility of cost-push shocks is large. The optimal value of the output gap is:

$$(12) \quad y_t = -\eta (\pi_t - \pi^*) \quad \text{with} \quad \eta \equiv \alpha R / [\kappa + (\alpha^2 - \theta \sigma_\pi^2 \kappa) R] > 0,$$

so that the central bank depresses aggregate demand if inflation is above target. Upon substitution of (11) and (12) into (10) and equating coefficients on π_t^2 , we obtain the quadratic $(\alpha^2 - \theta \sigma_\pi^2 \kappa) (R^2 - R) - \kappa = 0$. The positive root yields the relevant value of R:¹

$$(13) \quad R = \frac{1}{2} + \frac{1}{2} [1 + 4 \kappa (\alpha^2 - \theta \sigma_\pi^2 \kappa)^{-1}]^{1/2} \geq 1.$$

With strict inflation targeting ($\kappa=0$), we have $R=1$ and $\eta=1/\alpha$. Output is adjusted to achieve target inflation without any systematic error. Inflation is thus, apart from white noise, equal to desired inflation, $\pi_t=\pi^*+\varepsilon_{\pi t}$, regardless if there is prudent behaviour or not. Hence, with strict inflation targeting the postulated 'worst case' inflation shock is zero.

With flexible inflation targeting ($\kappa>0$), the postulated closed-loop dynamics for the inflation rate is given by:

$$(14) \quad \pi_{M t+1} - \pi^* = A_P (\pi_t - \pi^*) \quad \text{with} \quad 0 < A_P \equiv \kappa / [\kappa + (\alpha^2 - \theta \sigma_\pi^2 \kappa) R] < 1.$$

The minimisation of criterion (3) or (6) only converges and makes sense if $A_P < 1$ and $R > 1$, so $\theta \sigma_\pi^2 < \alpha^2 / \kappa$. In fact, this inequality holds automatically if the second-order condition for the prevention of breakdown, i.e., $R_A > 0$ and thus $\theta \sigma_\pi^2 R < 1$ and using (13) $\theta \sigma_\pi^2 < \alpha^2 / (\alpha^2 + \kappa)$ hold. The actual closed-loop dynamics for the inflation rate is then stable as well:

$$(15) \quad \pi_{t+1} - \pi^* = A_C \pi_t + \varepsilon_{\pi t+1} \quad \text{with} \quad 0 < A_C \equiv 1 - \alpha \eta = \kappa (1 - \theta \sigma_\pi^2 R) / [\kappa + (\alpha^2 - \theta \sigma_\pi^2 \kappa) R] < 1.$$

The reaction coefficient is positive and smaller than with strict inflation targeting, i.e., $0 < \eta < 1/\alpha$, so that the central bank plays safe and offsets systematic shocks less than fully. Monetary policy 'leans against the wind', because the central bank responds to excess inflation by deflating the economy (e.g., by raising the nominal interest rate). Since both R and R_A are increasing functions of $\theta \sigma_\pi^2$, the reaction coefficient η is an increasing function of $\theta \sigma_\pi^2$. Hence, the policy reactions are particularly severe if the central bank is very prudent and cost-push shocks are highly unpredictable.

Prudence raises the effective shadow penalty on inflation ($R_A > R$). A prudent central bank suffers a greater loss in utility from an unexpected rise in inflation than it gains from an unexpected fall in inflation of equal magnitude. The central bank thus puts more effort into cutting inflation to counteract possible stagflationary shocks. Hence, as prudence θ increases, the shadow penalty on inflation R_A rises and the reaction coefficient η rises.

¹ The economy is controllable (if $\beta \neq 0$) and the penalty matrix Q is positive definite (if $\kappa \neq 0$), so the Riccati equation converges (see Appendix). A steady-state Riccati coefficient R thus exists.

It can be shown that the adjustment speed for the actual closed-loop economy declines with the degree of prudence and the volatility of cost-push shocks ($\partial A_C / \partial \theta \sigma_\pi^2 < 0$). In contrast, the adjustment speed for the postulated closed-loop economy rises with prudence and volatility ($\partial A_P / \partial \theta \sigma_\pi^2 > 0$). A very prudent government reacts ‘nervously’ and is controlling the economy very precisely. We assume that the degree of prudence is not too large in order to avoid breakdown. This occurs as $\theta \sigma_\pi^2 R \rightarrow 1$, that is as $\theta \sigma_\pi^2 \rightarrow \alpha^2 / (\alpha^2 + \kappa)$. In that case, it can be shown that $R \rightarrow 1 + (\kappa / \alpha^2)$, $R_A \rightarrow \infty$, $\eta \rightarrow \alpha^{-1}$ and $\pi_t \rightarrow \varepsilon_{\pi t}$. For too high degrees of prudence the central bank has become too neurotic, since it is unable to offset the shocks that it fears and the risk-adjusted policy rule breaks down. The central bank's pessimistic assessment of shocks to the inflation rate so outweighs its knowledge of its statistical distribution that his expected cost criterion becomes infinite. Hence, optimisation becomes pointless and the central bank suffers a breakdown. As prudence increases and breakdown is approached, the central bank raises the reaction coefficient to the level that corresponds to strict inflation targeting ($\eta \rightarrow 1/\alpha$). In this sense, prudent central bankers attach greater weight to inflation than to output targeting. It is a very different rationale for a conservative central bank than the credibility argument provided by Rogoff (1985).

The output rule (12) can be implemented as a monetary policy rule. Let the output gap depend on the current real interest rate and ignore shocks to aggregate demand:

$$(16) \quad y_t = -\beta (i_t - \pi_{t+1|t} - r) = -\beta (i_t - \pi_t - \alpha y_t - r) = -\beta' (i_t - \pi_t - r)$$

with $\beta' \equiv \beta / (1 - \alpha\beta) > \beta > 0$,

where i_t is the (short) nominal interest rate at time t and r indicates the average Wicksellian real interest rate. Stability of the IS-curve requires $\alpha\beta < 1$, so $\beta' > \beta > 0$. The one-period-ahead rational-expectations private-sector forecast of inflation equals $\pi_{t+1|t} = \pi_t + \alpha y_t$, so the real interest rate is given by $i_t - \pi_{t+1|t} = i_t - \pi_t - \alpha y_t$. Substitution of (16) into the Phillips curve (9) yields $\pi_{t+1} = (1 + \alpha\beta')\pi_t - \alpha\beta'(i_t - r) + \varepsilon_{\pi t+1}$, so the uncontrolled inflation dynamics is unstable. The output rule (12) gives the following reaction function for the nominal interest rate:

$$(17) \quad i_t = r + \pi^* + [1 + (\eta/\beta')] (\pi_t - \pi^*)$$

so if inflation is above target the central bank tightens monetary policy. This interest rate rule can also be expressed as $i_t = r + \pi^* - (\eta^{-1} + \beta'^{-1})y_t$. But then the central bank depresses the economy in periods of unemployment by raising the nominal interest rate. This is obviously unrealistic. Section 4 allows for a richer structure of the economy, so that the optimal interest rule corresponds to a Taylor rule that reacts to *both* the inflation gap and the output gap.

We conclude that the constant in the nominal interest rate rule equals the sum of the average real interest rate and desired inflation and that the reaction coefficient for the inflation gap exceeds unity. We also observe that policy reactions are more aggressive if the central bank is prudent and cost-push shocks are volatile. Sargent (1999) derives a similar result in a different context. If the degree of prudence increases, the central bank tends towards strict inflation targeting and pushes up the reaction coefficient on the inflation gap towards $1/\alpha\beta > 1$. Equations (12) and (13) show that $\partial\eta/\partial\kappa = -(1-\theta\sigma_\pi^2) < 0$ if $\kappa=0$, hence introducing output targeting ($\kappa > 0$) raises R , but reduces η . Hence, if the central bank attaches more weight to output targeting (higher κ), rule (17) indicates that the interest rate responds more vigorously to changes in the inflation gap.

Proposition 1: If inflation follows an accelerationist Phillips curve and the real interest rate impacts aggregate demand instantaneously, the optimal nominal interest rate is given by the sum of the average real interest rate and desired inflation plus a term that reacts to the inflation gap. The nominal interest rate reacts more strongly to the inflation gap if output targeting is important. A prudent central bank assigns less power to the nominal interest rate and thus reacts more aggressively to the inflation gap, particularly so if the volatility of cost-push shocks is large. More prudent behaviour leads in the limit to strict inflation targeting.

(In the unlikely case that the central bank loves risk, i.e., $\theta < 0$, it becomes optimistic rather than pessimistic and is gambling on unexpected shocks to reduce inflation. This complacency lowers the shadow penalty on inflation and thus reduces the feedback reaction coefficient η . Consequently, the economy tends more towards the unregulated economy, i.e., $R \rightarrow 1$, $R_A \rightarrow 0$, and $\eta \rightarrow 0$, as the degree of risk loving or imprudence increases. In that case, the dynamics of inflation is no longer stable.)

4. Cautious Taylor rules with perfect observation of the economy

Section 3 demonstrated the effects of prudence on the optimal inflation-output trade-off. To gain a better understanding of prudent monetary policy, we analyse how caution affects optimal Taylor rules – see Taylor (1993). For a discussion of the forward-looking New Keynesian Phillips curve based on Calvo (1983), $\pi_t = \delta\pi_{t+1} + \alpha y_t + \varepsilon_{\pi,t+1}$ with $\delta \leq 1$, see Roberts (1995), Clarida, Gáli and Gertler (1999) and Woodford (2003). However, Mankiw and Reiss (2002) criticise this literature. Its predictions that disinflation causes a boom and that monetary policy quickly impacts inflation are patently unrealistic. Instead, they suggest that sluggish price adjustment may occur if information disseminates slowly throughout the population. Optimal monetary policy is then described by an elastic price standard in which

there is no base drift in the price level and the price level deviates from target if output deviates from the natural rate – see Ball, Mankiw and Reiss (2003). Christiano et al. (2003) obtain inflation inertia as well by assuming that the aggregate price index becomes available with a lag and replacing the New Keynesian Phillips curve by $\pi_t - \gamma\pi_{t-1} = \delta(\pi_{t+1} - \gamma\pi_t) + \alpha y_t + \varepsilon_{\pi t+1}$ with $0 \leq \gamma \leq 1$. Since these issues are still debated, we show the effects of prudence with the familiar backward-looking Phillips curve. In fact, we use the backward-looking model of a closed economy developed by Svensson (1997). Rudebusch and Svensson (1999) show that this model seems to fit US data fairly well - see also Judd and Rudebusch (1998).

We use the accelerationist Phillips curve (9) and allow for output persistence and stochastic shocks to aggregate demand. The nominal interest rate affects output with a lag of one period and inflation with a lag of two periods. Aggregate demand is thus given by:

$$(16') \quad y_{t+1} = \lambda y_t - \beta (i_t - \pi_{t+1|t} - r) + \varepsilon_{y t+1} = \lambda y_t - \beta (i_t - \pi_t - \alpha y_t - r) + \varepsilon_{y t+1} \\ = \lambda' y_t - \beta (i_t - \pi_t - r) + \varepsilon_{y t+1} \\ \text{with } 0 < \lambda < 1, \lambda' \equiv \lambda + \alpha\beta > \lambda, \beta > 0 \text{ and } \varepsilon_{y t} \sim \text{IN}(0, \sigma_y^2),$$

where λ indicates persistence in the output gap and $\varepsilon_{y t}$ is a serially uncorrelated, normally distributed demand shock with zero mean and variance σ_y^2 . The one-period-ahead private-sector inflation forecast and the real interest rate are given by, respectively, $\pi_{t+1|t} = \pi_t + \alpha y_t$ and $i_t - \pi_{t+1|t} = i_t - \pi_t - \alpha y_t$. The state variables are π_t and y_t . The instrument is the (short) nominal interest rate i_t . Each period corresponds roughly to three quarters. We postpone the treatment of lags and errors in measuring the output gap to section 5.

We use Theorem 2 (or alternatively Theorem A.1) to derive prudent policy rules for the optimal nominal interest rate as a function of last period's inflation gap and output gaps. These policy rules introduce caution into the well-known Taylor rules for monetary policy. Postulating that future stress is given by $\Psi_{\text{Fit}}(\pi_t, y_t) = R_{\pi\pi}(\pi_t - \pi^*)^2 + R_{yy}y_t^2 + 2R_{\pi y}(\pi_t - \pi^*)y_t$ and solving the min-max problem (8') with respect to next period's disturbances and considering the infinite-horizon limits, we obtain $\varepsilon_{M y t+1} = 0$ and

$$(17) \quad \varepsilon_{M \pi t+1} = \theta \sigma_\pi^2 \xi_1 [1 - \theta \sigma_\pi^2 \xi_1]^{-1} (\pi_t + \alpha y_t - \pi^*) \quad \text{with } \xi_1 \equiv R_{\pi\pi} - R_{\pi y}^2 / R_{yy} > 0,$$

where we require $0 \leq \theta \sigma_\pi^2 \xi_1 < 1$ for a meaningful solution. Since the matrix R must be positive definite, the coefficient ξ_1 is positive. The central bank sets the value of the future aggregate demand disturbance equal to its expected value. In contrast, the central bank postulates that the value of the cost-push shock exceeds its expected value if the one-period-ahead inflation

forecast is above target. A prudent central bank thus assumes that cost-push shocks hurt objectives more than their statistical expectation suggests they would.

The min-max problem (8') also yields the prudent Taylor rule for the optimal nominal interest rate:

$$(18) \quad i_t = r + \pi^* + (1 + \mu_\pi) (\pi_t - \pi^*) + \mu_y y_t$$

with $\mu_\pi \equiv \xi_2 [\beta(1-\theta\sigma_\pi^2 \xi_1)]^{-1} > 0$, $\mu_y \equiv (\lambda'/\beta) + \alpha\mu_\pi > 0$ and $\xi_2 \equiv R_{\pi y}/R_{yy} > 0$.

Substitution of (17) and (18) together with (9) and (16') into the recursion (8') and equating coefficients on π_t^2 , y_t^2 and $\pi_t y_t$, we obtain the algebraic equations:

$$(19) \quad \xi_1 = R_{\pi\pi} - [\alpha^2 (R_{\pi\pi} - 1)^2]/[\kappa + \alpha^2 (R_{\pi\pi} - 1)] > 0$$

$$R_{\pi\pi} = 1 + \xi_1 (1-\theta\sigma_\pi^2 \xi_1)^{-2} > 1 \quad \text{and} \quad 0 < \xi_2 = \alpha/[\alpha^2 + \kappa/(R_{\pi\pi} - 1)] < 1/\alpha.$$

The first two equations of (19) can be solved simultaneously for $R_{\pi\pi}$ and ξ_1 and subsequently the third equation of (19) gives ξ_2 . The reaction coefficients then follow from (18). Note that $R_{\pi y} = \alpha(R_{\pi\pi} - 1) > 0$ and $R_{yy} = \kappa + \alpha^2(R_{\pi\pi} - 1) > 0$.

The constant in the prudent Taylor rule corresponds to the sum of the average real interest rate and target inflation. The optimal nominal interest rate 'leans against the wind', since the central bank deflates (reflates) the economy by raising (lowering) the nominal interest rate if inflation and output are above (below) target. The central bank reacts more vigorously to the output gap if there is a relatively large degree of persistence in aggregate demand. The private-sector one-period-ahead inflation forecast, $\pi_{t+1|t} = \pi_t + \alpha y_t$, is predetermined in period t and independent of the instrument i_t . Hence, we can rewrite the prudent Taylor rule (18) so that the nominal interest reacts to the one-period-ahead inflation gap forecast and the output gap:

$$(18') \quad i_t = r + \pi^* + (1 + \mu_\pi) (\pi_{t+1|t} - \pi^*) + (\lambda'/\beta) y_t.$$

The reaction coefficient with respect to the one-period-ahead inflation forecast also exceeds unity if there is prudence – cf., Woodford (2001) for the case without prudence. Clarida, Gáli and Gertler (2000) and Svensson and Woodford (2002) discuss the merits of the nominal interest rate reacting to future forecasts of inflation and the output gap if there is no prudence.

To better understand the prudent Taylor rule (18), we consider some special cases. If there is no prudence or there are no cost-push disturbances ($\theta\sigma_\pi^2=0$), we have:

$$(18') \quad \mu_\pi = \alpha \xi_1 / [\beta(\kappa + \alpha^2 \xi_1)] < 1/\alpha\beta \quad \text{and} \quad \xi_1 = \frac{1}{2} + \frac{1}{2} (1 + 4\kappa/\alpha^2)^{1/2} \geq 1.$$

Flexible inflation targeting without prudence gives the well-known Taylor (1993) rule, which was estimated as $i_t = 4\% + 1.5(\pi_t - \pi^*) + 0.5 y_t$. Rotemberg and Woodford (1997, 1999) provide other estimates of Taylor rules. Rotemberg and Woodford (1999) and Woodford (2001) argue that the optimal output-response coefficient is much smaller if de-trended output is used instead of the theoretically correct output gap measure. The special case of strict inflation targeting and no prudence yields $\xi_1=1$ and the reaction coefficients $\mu_\pi=1/\alpha\beta$, $\mu_y=(1+\lambda')/\beta$. It gives the maximum reaction to the inflation and output gaps. More emphasis on output targeting (higher κ) reduces both reaction coefficients.

If we allow for prudence in the face of cost-push shocks ($\theta\sigma_\pi^2 > 0$), strict inflation targeting ($\kappa=0$) implies $\xi_1=1$, $\xi_2=1/\alpha$, $\mu_\pi=1/[\alpha\beta(1-\theta\sigma_\pi^2)] > 1/(\alpha\beta)$ and $\mu_y > (1+\lambda')/\beta$. In contrast to the case where the real interest rate affects aggregate demand instantaneously (see section 3), prudence and volatile cost-push disturbances imply more aggressive reactions to changes in the inflation and output gaps with strict inflation targeting.

A case of special interest is what happens just before the degree of prudence or volatility of cost-push shocks reaches breakdown-level. This occurs as $\theta\sigma_\pi^2 \rightarrow 1$ and thus as $\xi_1 \rightarrow 1$, $\xi_2 \rightarrow 1/\alpha$ and $\mu_\pi, \mu_y \rightarrow \infty$. The policy makers have become so neurotic and their policy reactions so aggressive that optimisation becomes pointless. In fact, the central bank's *subjective* assessment of cost-push shocks has become so pessimistic that it outweighs its *objective* knowledge of its statistical distribution. Hence, the expected cost criterion becomes infinite and the central bank's policy rule breaks down.

For the general case of flexible inflation targeting ($\kappa > 0$) and prudence with stochastic cost-push shocks ($\theta\sigma_\pi^2 > 0$), total differentiation of (19) shows that $R_{\pi\pi}$, ξ_1 and ξ_2 are increasing functions of $\theta\sigma_\pi^2$ and thus that the reaction coefficients μ_π and μ_y are increasing function of $\theta\sigma_\pi^2$. Hence, the nominal interest rate reacts more aggressively to inflation and output gaps if the central bank is more prudent and cost-push disturbances are more volatile. If aggregate demand is not much affected by the real interest rate (small β), the policy reactions are more aggressive as well. More weight to output targeting (higher κ) raises both $R_{\pi\pi}$ and ξ_1 and, typically, reduces ξ_2 . Hence, the nominal interest rate rule becomes less aggressive (μ_π and μ_y fall) but less so for a prudent central bank. More weight to output targeting thus attenuates the central bank's policy reactions to changes in the output gap and, especially, the inflation gap.

Proposition 2: If the nominal interest rate impacts output after one period and inflation after two periods, the optimal nominal interest rate also reacts to the output gap, especially if there

is substantial persistence in aggregate demand. The central bank reacts more vigorously to both the inflation gap and the output gap if prudence is important and the volatility of shocks to the dynamics of inflation is large, even under strict inflation targeting. More weight to output targeting weakens policy responses of the central bank, particularly if caused by changes in the inflation gap. Volatility of shocks to aggregate demand does not affect the monetary policy rules.

5. Prudent monetary policy with errors and lags in measuring the output gap

In practice central bankers conduct monetary policy in an environment in which they observe the inflation rate pretty well, but measure the output gap with a lag and imperfectly. We thus assume that the central bank observes the inflation rate instantaneously and without error while it observes the output gap with a lag of one period and measurement error $\omega_{y,t}$:

$$(20) \quad z_t = y_{t-1} + \omega_{y,t}, \quad \omega_{y,t} \sim \text{IN}(0, \tau^2), \quad y_0 \sim \text{N}(y_{p0}, v_{y0}),$$

where $\omega_{y,t}$ is the serially uncorrelated, normally distributed measurement error with zero mean and variance τ^2 . The optimal nominal interest rate rule thus reacts to the imperfectly measured output gap of the previous period and this period's realised inflation rate. With the aid of Theorem 2 a prudent central bank can still separate the problems of control and prediction. First, it solves the min-max problem (8') to find the nominal interest rate rule (18). Then, it solves for the prediction of the output gap y_{Mt} to be used in the prudent Taylor rule,

$$(18'') \quad i_t = r + \pi^* + (1 + \mu_\pi) (\pi_t - \pi^*) + \mu_y y_{Mt},$$

by maximising total past and future stress $\Psi_{Pt}(y_t, I_t) + \Psi_{Ft}(\pi_t, y_t)$. Past stress depends on only one stochastic variable, i.e., the unknown output gap and not the perfectly and instantaneously observable inflation rate. We thus postulate that past stress equals $\Psi_{Pt}(y_t, I_t) = -(\theta v_{yt})^{-1} (y_t - y_{Pt})^2$. Maximising total stress with respect to the current output gap yields:

$$(21) \quad y_{Mt} = [y_{Pt} + \theta v_{yt} R_{\pi y} (\pi_t - \pi^*)] / [1 - \theta v_{yt} R_{yy}].$$

A risk-neutral central bank uses the output gap predictions coming from the conventional Kalman filter, y_{Pt} . A prudent central bank, however, biases its measurement of the output gap predictions upwards (downwards) if inflation is above (below) target. To fight too high (low) inflation, the central bank tightens (loosens) monetary policy by raising (lowering) the nominal interest rate. Consequently, the bias in the output gap reinforces this nominal interest

rate response, especially if the variance of the measurement error in the output gap is large and prudence is important. In fact, even if inflation is on target, the central bank amplifies the objective estimate of the output gap and thus effectively reacts more aggressively to deviations of the statistical Kalman filter estimate of the output gap from target.

The final step is to derive the risk-adjusted updating predictions for the output gap and their variances, y_{Pt} and v_{yt} , from the recursions:

$$(22) \quad \Psi_{Pt+1}(y_{t+1}, I_{t+1}) = \underset{y_t}{Max} \{ \Psi_{Pt}(y_t, I_t) + \kappa y_t^2 - (\theta \tau^2)^{-1} (z_{t+1} - y_t) - (\theta \sigma_y^2)^{-1} [y_{t+1} - \lambda' y_t + \beta(i_t - \pi_t - r)]^2 \}$$

with $\Psi_{P0}(y_0, I_0) \equiv -(\theta v_{y0})^{-1} (y_0 - y_{P0})^2$.

Solving for the maximising value of y_t , substituting back into the recursion (22), and equating coefficients on y_{t+1}^2 and y_{t+1} yields the following risk-adjusted updating formulae for the mean and variance of the output gap:

$$(23) \quad v_{yt+1} = \sigma_y^2 + \lambda'^2 (v_{yt}^{-1} + \tau^2 - \theta \kappa)^{-1}$$

$$(24) \quad y_{Pt+1} = \lambda' y_{Pt} - \beta (i_t - \pi_t - r) + \lambda' (v_{yt}^{-1} + \tau^2 - \theta \kappa)^{-1} [\tau^2 (z_{t+1} - y_{Pt}) + \theta \kappa y_{Pt}],$$

where prudence must be small for a meaningful solution (i.e., $\theta < (\kappa v_{yt})^{-1}$, $\forall t > 0$ must hold).

The infinite-horizon limit of the variance of the output gap, v_y , solves the quadratic:

$$(23') \quad (\tau^2 - \theta \kappa) v_y^2 - [\lambda'^2 + (\tau^2 - \theta \kappa) \sigma_y^2 - 1] v_y - \sigma_y^2 = 0.$$

It follows from total differentiation of (23') that the asymptotic variance of the output gap (v_y) is large if the measurements of the output gap are unreliable (τ^2 large), the shocks hitting aggregate demand are volatile (σ_y^2 large), and persistence in aggregate demand is substantial (λ' large). In addition, as a precautionary measure, a prudent central bank pushes up its subjective estimate of the variance of the output gap, especially if output targeting is important. This helps to avoid costly mistakes.

If there is no prudence, (24) gives the standard Bayesian updating formula for the BLUE-predictions of the imperfectly measured output gaps. New measurements of the output gap thus induce less substantial revisions if the signal-to-noise ratio is large, that is if new data are unreliable relative to the precision of the current estimate of the output gap (i.e., τ^2/v_{yt} large). The reactions of the central bank to the measured output gap are thus weakened, particularly if incoming data are unreliable. A prudent central bank committed to strict

inflation targeting ($\kappa=0$) uses these predictions of the output gap in the decoupling formulae (21) and subsequently in the prudent Taylor rule (18). A prudent central bank with no concern for output targeting thus simply biases the output gap predictions upwards (downwards) if inflation is above (below) target.

If the central bank is prudent and engages in flexible inflation targeting, equation (24) requires two adjustments to the predictions. First, the magnitude of the revisions of the output gap are increased if prudence and output targeting are important ($\theta\kappa$ large). Note that there is a direct effect of $\theta\kappa$ on the term multiplying $z_{t+1}-y_{pt}$ in (24) and an indirect effect through v_{yt} . A prudent central bank weakens its response to changes in the output gap rather less if it attaches a large weight to output targeting. Essentially, a prudent central bank makes more use of faulty data rather than risking the chance of big stochastic fluctuations in the output gap and thus adopts a more aggressive policy response. Second, a prudent central bank introduces as a precautionary measure an upward bias in its estimate of the output gap, especially if the weight given to output targeting is large. The resulting upward bias in the nominal interest rate depresses the economy, so that the central bank is cushioned against unexpected large positive output gaps. Clearly, the predictions that follow from (23) and (24) are neither efficient nor unbiased.

Proposition 3: If inflation is observed instantaneously and without error but the output gap is imperfectly observed with a period of one lag, the nominal interest rate reacts less strongly to the imperfectly measured output gap. The reaction is particularly weakened if the variance of the measurement error is large relative to the variance of the current estimate of the output gap, where the latter increases with the variance of the shocks to and the degree of persistence in aggregate demand. However, if prudence is important and the welfare weight given to unemployment is large, the central bank employs as a precautionary measure a larger variance of the output gap and attenuates its reactions to changes in the output gap rather less. Furthermore, a prudent central bank introduces as precaution an upward bias in its estimate of the output gap and thus in the nominal interest rate, especially if output targeting is important. Finally, a prudent central bank introduces an extra upward (downward) bias in its estimate of the output gap before it feeds into the policy rule if inflation is above (below) target.

6. Concluding remarks

We have introduced prudence into the standard LQG control and prediction framework. A prudent policy maker increases the shadow penalty on target variables, since he wants to hedge against stochastic shocks possibly hampering the achievement of desired values. This invalidates usual certainty equivalence, so that one cannot substitute future disturbances by

their expected values. Instead, a prudent policy maker uses subjective, cautious estimates of future disturbances that depend on preferences in order to avoid costly mistakes.

In practice a policy maker faces errors and lags in observing the state of the world. A risk-neutral policy maker uses the Kalman filter to revise BLUE-estimates of the mean and covariance matrix of the states of the world as new information comes in. These predictions are independent of preferences and are used in the optimal feedback policy rules. A prudent policy maker uses inefficient predictions of the states of the world, since he raises the variances of the states versus those of the incoming observations as he is afraid to incorporate possibly faulty information that may lead to significant welfare losses. Large penalties for the target variables and a high degree of prudence imply large reductions in the relative precision of the measurements. A prudent policy maker also introduces a bias in the prediction of the states. Hence, the prudent Kalman filter no longer yields BLUE-predictions of the states. With prudence special care must be taken to couple the derivation of the optimal control rules and the prediction of the states.

Since prudence is the essence of central banking, we investigated the effects of prudence on optimal monetary policy rules. In line with Sargent (1999), we found that allowing for prudence on the optimal inflation-output trade-off yields more aggressive reactions of the nominal interest rates to the inflation gap, particularly if the volatility of cost-push shocks is great. Craine (1979) and Söderström (2002) allow for parameter uncertainty in the dynamics of inflation and also find more vigorous policy responses. To be safe a prudent central bank assigns a lower effectiveness of its monetary instrument. More prudent behaviour is in the limit equivalent to strict inflation targeting. If the real interest rate affects aggregate demand with a lag, the nominal interest rate also reacts to the output gap, especially if there is substantial persistence in aggregate demand. Prudence and bigger volatility of cost-push disturbances imply stronger reactions to both the inflation and output gaps, even under strict inflation targeting. More weight to output targeting weakens policy responses of the central bank, particularly if caused by changes in the inflation gap.

Rudebusch (2001) shows that data uncertainty in the model of Rudebusch and Svensson (1999) leads to weaker monetary policy reactions. We also find that the reactions of the nominal interest rate to the measured output gap are less strong, especially if the incoming output gap data are relatively unreliable compared with the precision of the current output gap estimate. However, a prudent central bank attenuates its reactions to changes in the output gap much less, especially if output targeting is important. A prudent central bank also introduces as a precautionary measure an upward bias in its estimate of the output gap and thus in the nominal interest rate, again especially if output targeting is important. Both these elements make the reactions of the central bank to the output gap more aggressive, particularly if shocks to inflation are volatile and output targeting is important. Finally, a prudent central

bank introduces an extra upward (downward) bias in its estimate of the output gap before it feeds into the policy rule if inflation is above (below) target. This makes the nominal interest rate reactions more aggressive.

It is important to extend our methods for deriving prudent policy rules and prediction formulae to macroeconomic models with forward-looking expectations. This is crucial for a deeper understanding of monetary policy. However, the forward-looking New Keynesian Phillips curve reviewed in Clarida, Gáli and Gertler (1999) seems unrealistic, since it predicts that disinflation causes a boom and monetary policy impacts inflation quickly. Christiano et al. (2001) assume that individual prices are indexed to an aggregate price index and that the latter becomes available for indexation only after a lag. This modification gives rise to inflation inertia. Ball, Mankiw and Reiss (2003) assume that information disseminates slowly throughout the population and obtain more realistic impulse responses as well. They find that an elastic price standard is optimal. It is worthwhile to investigate how Taylor rules are affected by prudence in these more realistic models. It also seems worthwhile to investigate how prudence affects optimal fiscal and monetary policy when prices are sticky and government debt is used to smooth tax distortions. It may thus be interesting to introduce prudence in the analysis of Benigno and Woodford (2003). It seems realistic to suppose asymmetry, i.e., that the central bank is more prudent than the fiscal authority.

In future research, it is worthwhile to explore how other types of model uncertainty (e.g., multiplicative disturbances, uncertainty about model specification) affect prudent behaviour. For example, parameters such as the sensitivity of unanticipated inflation with respect to the output gap and the interest rate semi-elasticity of aggregate demand may not be known precisely. Brainard (1967) shows that it then pays to use more instruments than targets and, in particular, to use a mix of instruments in order to diversify risks and to not go the whole hog in reaching desired values – see also, e.g., Söderström (2002). An interesting question is how a prudent central banker adjusts its behaviour in response to multiplicative uncertainty about key parameters.

A related issue is robustness of policy rules - e.g., Onatski and Stock (2000), Giannoni and Woodford (2002, 2003), and Onatski and Williams (2003). Such rules should be simple, easy to understand, react to measured variables, stabilise the economy, and be determinate in the sense of supporting a rational expectations equilibrium. Robust rules should do well under a wide variety of assumptions about the precise statistical distribution and order of the additive disturbance processes affecting the economy and the processes driving the measurement errors. This amounts to a different type of prudence to that discussed here. It requires both robust Bayesian and min-max policy rules that do well under the 'worst' possible predictions. Central banks do not react very strongly to something that is likely to be estimated with considerable error, but prudence makes the reactions to poor data more

aggressive. Onatski and Williams (2003) also find that many of their robust policy rules are relatively aggressive, since central banks fear particularly very persistent increases in inflation arising from long-run deviations from a vertical Phillips curve. Such rules perform well at low frequencies, but consequently fare worse at business cycle frequencies. In general, robustness requires the central bank to react to a richer lag structure of output gaps and inflation rate and, perhaps, even to other indicators of monetary imbalance.

Svensson (2002) and Feldstein (2003) rightly argue that a prudent central bank will reserve the right to use informal judgement in addition to a formal policy rule. This allows for the use of other and more subjective information and for effects not captured by the formal model and rule. In addition, it may enhance its credibility. A prudent central bank does not mechanically apply a Taylor-type rule, but may use it as a rule of thumb for monetary policy. A useful area for further research is thus to develop a framework of prudent monetary policy formulation in which there is scope for judgement of central bankers. Svensson (2002) suggests a framework for doing just that, but the prudent Kalman filter predictions discussed here also provide a natural way of incorporating judgement of central bankers into more formal Taylor-type rules for the nominal interest rate.

To conclude it may be useful to qualify our results in two important respects. Common sense of many practitioners dictates that prudence implies that the nominal interest rate should respond less strongly to changes in inflation and output gaps, while our analysis suggests more aggressive policy responses. There is no reason why prudence should imply passive behaviour of the policy maker. In fact, it is more likely to lead to overzealous, even uptight control of the economy. Just as a prudent driver may react strongly to every bend in the road in order to avoid driving into the curb, a prudent central bank changes the interest rate more frequently. Practitioners may be more concerned with nominal interest rate smoothing than output targeting - e.g., Lippi and Neri (2003). However, it helps to offer some kind of welfare-theoretic rationale and explain why interest rate volatility may damage welfare and financial markets - e.g., Woodford (2002). Perhaps, central bankers delay interest rate adjustments and prefer some unpredictability in order not to be seen to follow the pressure of market players and commentators. The second qualification is that neither the traditional nor the New Keynesian Phillips curves fully capture real world features such as credit constraints, equity constraints, bankruptcies and other market failures arising from imperfect information. Stiglitz and Greenwald (2003) point out that then the nominal interest rate affects aggregate demand and that monetary policy is associated with big allocative distortions and is as much about supervision and regulation as the interest rate.

Appendix: Recursive solutions

First, we present a theorem with the optimal risk-adjusted feedback policy rules for the case that states are perfectly observable without measurement lags and that instruments do not enter the intra-temporal welfare loss function ($S=0$). Subsequently, we present the recursive solution if state variables are not or imperfectly observable with a one-period lag.

Theorem A.1 (Risk-adjusted LQG control with perfect state observation):

Minimisation of the risk-adjusted welfare loss criterion (3) subject to the welfare loss (1) with $S=0$ and the state-space model (2) yields the following optimal policy feedback rules:

$$(A1) \quad u_t = G_{t+1} x_t + g_{t+1},$$

where the feedback gain matrix and the vector of policy constants are given by

$$(A2) \quad G_t \equiv - (B'R_{At}B)^{-1} B'R_{At} A \quad \text{and} \quad g_t \equiv - (B'R_{At}B)^{-1} B'R_{At} (a_t - r_t) \\ \text{with} \quad R_{At} \equiv (R_t^{-1} - \theta \Sigma)^{-1}$$

and the Riccatti and auxiliary equations are given by

$$(A3) \quad R_t = Q + (A+BG_{t+1})' R_{A,t+1} (A+BG_{t+1}) \quad \text{and} \quad R_t r_t = Q x^* + (A+BG_{t+1})' R_{A,t+1} (r_{t+1} - a_{t+1}) \\ \text{starting with} \quad R_T = Q_T \quad \text{and} \quad r_T = x_T^*,$$

where R_t and R_{At} denote symmetric and positive-definite Riccatti matrices. The postulated min-max value of the vector of disturbances at time t is given by

$$(A4) \quad \varepsilon_{p,t} = \theta \Sigma R_{At} (A x_{t-1} + B u_{t-1} + a_t - r_t).$$

The matrix $(I - \theta \Sigma R_t)$ must be positive definite for all t . Otherwise, there will be an infinite expected loss of utility. The more general case of non-zero S is treated in Theorem A.2.

Proof: Use the Bellman principle of dynamic programming – see van der Ploeg (1984).

The risk-neutral case ($\theta=0$) does not depend on the covariance matrix Σ and corresponds exactly to Chow (1975, Chapters 7 and 8). The principle of certainty equivalence no longer applies exactly if $\theta \neq 0$, since the optimal policy feedback rules (A1)-(A3) can no longer be obtained by simply setting the disturbance terms at their expected values (i.e. $\varepsilon_t=0$). The risk-

averse policy rules depend on the covariance matrix Σ , since in each period the policy maker replaces the (symmetric) shadow penalty matrix R_t by the risk-adjusted shadow penalty matrix R_{At} . A risk-averse policy maker thus increases the effective shadow penalty on uncertain target variables. The policy maker wants to be safe and avoid shocks frustrating the achievement of desired targets. The size of the increase in the risk-adjusted shadow penalty increases with the variance of the target variable. The prudent policy maker when deciding on the instruments for this period does not set next period's disturbances to their expected value of zero, but instead sets them equal to (A4).

Theorem A.2 (Risk-adjusted LQG control and prediction - homogenous case):

Consider the homogenous case with $a_t = x^* = 0$ and $u^* = 0$. Assume that the matrix $[R_s - (\theta \Sigma)^{-1}]$ is negative definite for $s > t$. Future stress is quadratic in the states, i.e., $\Psi_{F_t}(x_t) = x_t' R_t x_t$, where the symmetric and positive-definite matrix R obeys the risk-adjusted Riccati recursion:

$$(A5) \quad R_t = Q + A' (BS^{-1}B' - \theta \Sigma + R_{t+1}^{-1})^{-1} A \quad \text{with} \quad R_T = Q_T.$$

The optimal control is linear in x_{Mt} , the maximum-stress estimate, or x_{Pt} , the predicted states:

$$(A6) \quad u(x_{Pt}, t) = G_t x_{Mt} \quad \text{with} \quad x_{Mt} \equiv (I - \theta \Sigma R_t)^{-1} x_{Pt} \quad \text{and} \quad G_t \equiv -S^{-1}B'(BS^{-1}B' - \theta \Sigma + R_{t+1}^{-1})^{-1} A$$

Past stress equals (up to a constant) $\Psi_{Pt}(x_t, I_t) = -\theta^{-1} (x_t - x_{Pt})' V_t^{-1} (x_t - x_{Pt})$ where the matrix V_t and the predicted states x_{Pt} satisfy the prudent Kalman filter recursions provided that the matrix $[V_s - (\theta Q)^{-1}]$ is negative definite for $s \leq t$:

$$(A7) \quad V_t = \Sigma + A (V_{t-1}^{-1} + C' \Omega^{-1} C - \theta Q)^{-1} A', \quad t \geq 1$$

$$(A8) \quad x_{Pt} = A x_{Pt-1} + B u_{t-1} + A (V_{t-1}^{-1} + C' \Omega^{-1} C - \theta Q)^{-1} [C' \Omega^{-1} (z_{t+1} - C x_{Pt}) + \theta Q x_{Pt}], \quad t \geq 1.$$

If there are no measurement errors or lags in observation, (A6) becomes $u(x_t, t) = G_t x_t$.

Proof: Application of Theorem 2 - see Whittle (1981, Chapter 19).

A prudent policy maker has the pessimistic view that state disturbances push the states in an undesired direction. Hence, the policy maker reduces the effective power of its instruments $(BS^{-1}B')$ in (A5)-(A6) by $(\theta \Sigma)$. The extent by which it reduces the power of its instruments is thus large if prudence θ is big and shocks are volatile. A prudent policy maker also has a

pessimistic view on the use of new information and thus reduces the observation information matrix ($C'\Omega^{-1}C$) in (A7)-(A8) by θQ . Use of new information is particularly discounted if prudence θ is big and the state penalties are large. In the case of no prudence, $\theta=0$, the Kalman filter gives x_{pt} as the Bayesian estimate of the states, $E[x_t | I_t]$, and V_t is its covariance matrix. With prudence x_{pt} takes account of the 'worst state of the world' and thus maximises past stress. The term $(\theta Q x_{pt})$ in the prudent Kalman filter (A8) is the bias of incoming information due to prudence of the policy maker. Hence, the prudent Kalman filter (A7)-(A8) not only produces inefficient estimates, due to the reduction of the observation information matrix by θQ , but it also produces biased estimates.

The stationary infinite-horizon limit of R_t as $t \rightarrow \infty$, R , exists if, say, the matrices Q and $(BS^{-1}B' - \theta\Sigma)$ are positive definite and the system is controllable. If G denotes the limit of G_t as $t \rightarrow \infty$, the postulated closed-loop state-space system is:

$$(A9) \quad x_{t+1} = (A+BG) x_t + \varepsilon_{p,t+1} = A_p x_t \quad \text{with} \quad A_p \equiv [I + (BS^{-1}B' - \theta\Sigma) R] A,$$

where A_p has all eigenvalues inside the unit circle. The actual closed-loop system satisfies:

$$(A10) \quad x_{t+1} = A_C x_t + \varepsilon_{t+1} \quad \text{with} \quad A_C \equiv A + B G = (I - \theta \Sigma R) A_p .$$

Stability of (A10) is not guaranteed if $\text{abs}(\theta)$ is too large.

The possibility of breakdown occurs if θ is so large that the matrix $[R_s - (\theta\Sigma)^{-1}]$ is no longer negative definite for $s > t$. In that case, the policy maker is faced with an inability to control and becomes paralysed. The policy maker's pessimistic assessment of disturbances so outweighs its knowledge of its statistical distribution that his expected cost criterion becomes infinite. Hence, optimisation becomes pointless and the policy maker suffers a breakdown.

(With extreme risk loving, $\theta \rightarrow -\infty$, the attitude that disturbances work to one's advantage may lead to a policy rule which is too lax for stability. In fact, the controlled system approaches the uncontrolled system, i.e., $R \rightarrow Q$, $G \rightarrow 0$ and $A_C \rightarrow A$, as $\theta \rightarrow -\infty$. Hence, if the uncontrolled system is unstable, it is impossible to stabilise the controlled system for large degrees of risk loving.)

To allow for non-zero desired values for states and instruments, redefine states and instruments as $x_t - x_t^*$ and $u_t - u_t^*$. This changes a_t to $a_t - (x_t^* - A x_{t-1}^* - B u_{t-1}^*)$. Tracking non-zero desired values for states and instruments can thus be dealt with by non-zero a_t . Whittle (1982, page 279) shows that with non-zero a_t (A6) becomes:

$$(A6') \quad u(x_{pt}, t) = G_t (I - \theta \Sigma R_t)^{-1} (x_{pt} + \theta V_t \chi_t),$$

where χ_t satisfies the recursion

$$(A11) \quad \chi_t = A' [I + (BS^{-1}B' - \theta \Sigma) R_{t+1}]^{-1} (R_{t+1} a_{t+1} + \chi_{t+1}) \quad \text{with} \quad \chi_T \equiv -R_T Q_T^*.$$

References

- Aoki, M. (1976). *Optimal Control and System Theory in Dynamic Economic Analysis*, North-Holland, Amsterdam.
- Ball, L., N.G. Mankiw and R.Reis (2003). *Monetary policy for inattentive economies*, NBER Working Paper No. 9491, Cambridge, Mass.
- Barker, T.S., F. van der Ploeg and M.R. Weale (1984). *A balanced system of national accounts for the U.K.*, *Review of Income and Wealth*, 30, 4, 459-483.
- Barro, R. and D. Gordon (1983). *A positive theory of monetary policy in a natural rate model*, *Journal of Political Economy*, 91, 589-610.
- Benigno, P. and M. Woodford (2003). *Optimal monetary and fiscal policy: A linear-quadratic approach*, NBER Working Paper No. 9905, Cambridge, Mass.
- Brainard, W. (1967). *Uncertainty and the effectiveness of monetary policy*, *American Economic Review, Papers and Proceedings*, 57, 2, 411-425.
- Brock, W.A., S.N. Durlauf and K.D. West (2003). *Policy evaluation in uncertain economic environments*, NBER Working Paper No. 10025, Cambridge, Mass.
- Calvo, G.A. (1983). *Staggered prices in a utility maximizing framework*, *Journal of Monetary Economics*, 12, 383-398.
- Chow, G.C. (1975). *Analysis and Control of Dynamic Economic Systems*, John Wiley & Sons, New York.
- Christiano, L.J., M.S. Eichenbaum and C.L. Evans (2001). *Nominal rigidities and the dynamic effects of a shock to monetary policy*, NBER Working Paper No. 8403, Cambridge, Mass.
- Clarida, R., J. Galí and M. Gertler (1999). *The science of monetary policy: a New Keynesian perspective*, *Journal of Economic Literature*, 37, 1661-1707.
- Clarida, R., J. Galí and M. Gertler (2000). *Monetary rules and macroeconomic stability: evidence and some theory*, *Quarterly Journal of Economics*, 115, 147-180.
- Craine, R. (1979). *Optimal monetary policy with uncertainty*, *Journal of Economic Dynamics and Control*, 1, 59-83.
- Feldstein, M. (2003). *Monetary policy in an uncertain environment*, NBER Working Paper No. 9969, Cambridge, Mass.

- Giannoni, M.P. and M. Woodford (2002). Optimal interest rate rules: I. General theory & II. Applications, NBER Working Papers 9419 & 9420, Cambridge, Mass.
- Giannoni, M.P. and M. Woodford (2003). Optimal inflation targeting rules, NBER Working Paper No. 9939, Cambridge, Mass.
- Hansen, L.P. and T.J. Sargent (1995). Discounted linear exponential quadratic Gaussian control, *IEEE Transactions Automatic Control*, 40, 968-971.
- Hansen, L.P., T.J. Sargent and T.D. Tallarini, Jr. (1999). Robust permanent income and pricing, *Review of Economic Studies*, 66, 4, 873-907.
- Jacobson, D.H. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games, *IEEE Transactions on Automatic Control*, 18, 124-131.
- Jacobson, D.H. (1977). *Extensions of Linear-Quadratic Control, Optimization and Matrix Theory*, Academic Press, New York.
- Judd, J.F. and G.D. Rudebusch (1998). Taylor's rule and the Fed: 1970-1997, *Federal Reserve Bank of San Francisco Economic Review*, 3, 3-16.
- Kendrick, D. (1981). Control theory with applications to economics, 111-158 in K.J. Arrow and M.D. Intrilligator (eds.), *Handbook of Mathematical Economics*, Vol. 1, Elsevier/North-Holland, Amsterdam.
- Kydland, F. and E. Prescott (1977). Rules rather than discretion: the inconsistency of optimal plans, *Journal of Political Economy*, 85, 473-490.
- Lippi, F. and S. Neri (2003). Information variables for monetary policy in a small structural model of the Euro area, CEPR Discussion Paper No. 4125, London.
- Ljungqvist, L. and T.J. Sargent (2000). *Recursive Macroeconomic Theory*, MIT Press, Cambridge, Mass.
- Mankiw, N.G. and R. Reiss (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve, *Quarterly Journal of Economics*, 1295-1328.
- Onatski, A. and J.H. Stock (2000). Robust monetary policy rules under model uncertainty in a small model of the US economy, NBER Working Paper No. 7490, Cambridge, Mass.
- Onatski, A. and N. Williams (2003). Modeling model uncertainty, NBER Working Paper No. 9566, Cambridge, Mass.
- Orphanides, A. (2000). The quest for prosperity without inflation, Working Paper No. 15, European Central Bank, Frankfurt.
- Pearlman, J.G. (1986). Diverse information and rational expectations models, *Journal of Economic Dynamics and Control*, 10, 333-338.
- Pearlman, J.G. (1992). Reputational and nonreputational policies under partial information, *Journal of Economic Dynamics and Control*, 16, 339-357.

- Phelps, E.S. (1967). Phillips curves, expectations of inflation and optimal unemployment over time, *Economica*, 34, 254-281.
- Ploeg, F. van der (1982). Reliability and the adjustment of sequences of large economic accounting matrices (with discussion), *Journal of the Royal Statistical Society, Series A*, 145, 2, 169-194.
- Ploeg, F. van der (1984a). Risk and linear-quadratic stabilization, *Economics Letters*, 73-78.
- Ploeg, F. van der (1984b). Economic policy rules for risk-sensitive decision making, *Zeitschrift für Nationalökonomie/Journal of Economics*, 44, 3, 207-235.
- Ploeg, F. van der (1993). A closed-loop solution for a model of precautionary saving, *Review of Economic Studies*, 60, 2, 385-396.
- Preston, A.J. and A.R. Pagain (1982). *The Theory of Economic Policy, Statics and Dynamics*, Cambridge University Press, Cambridge, UK.
- Roberts, K. (1995). New-Keynesian economics and the Phillips curve, *Journal of Money, Credit and Banking*, 27, 4, 975-984.
- Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target, *Quarterly Journal of Economics*, 100, 1169-1189.
- Rotemberg, J. and M. Woodford (1997). An optimization-based econometric framework for the evaluation of monetary policy, *NBER Macroeconomics Annual*, 297-346.
- Rotemberg, J. and M. Woodford (1999). Interest-rate rules in an estimated sticky-price model, in J.B. Taylor (ed.), *Monetary Policy Rules*, Chicago University Press, Chicago.
- Rudebusch, G.D. (2001). Is the Fed too timid? Monetary policy in an uncertain world, *Review of Economics and Statistics*, 83, 203-217.
- Rudebusch, G.D. and L.E.O. Svensson (1999). Policy rules for inflation targeting, 203-246 in J.B. Taylor (ed.), *Monetary Policy Rules*, Chicago University Press, Chicago.
- Sargent, T.J. (1999). Discussion of L. Ball, Policy rules for open economies, in J. Taylor (ed.), *Monetary Policy Rules*, University of Chicago Press, Chicago.
- Söderström, U. (2002). Monetary policy with uncertain parameters, *Scandinavian Journal of Economics*, 104, 125-145.
- Speyer, J.L., J. Deyst and D.H. Jacobson (1974). Optimisation of stochastic linear systems with additive measurement and process noise using exponential performance criteria, *IEEE Transactions on Automatic Control*, 19, 358-366.
- Stiglitz, J.E. and B. Greenwald (2003). *Towards a New Paradigm in Monetary Economics*, Raffaele Mattioli Lectures, Cambridge University Press, Cambridge, UK.
- Svensson, L.E.O. (1997). Inflation forecast targeting: implementing and monitoring inflation targets, *European Economic Review*, 41, 1111-1146.
- Svensson, L.E.O. (2002). What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules, NBER Working Paper No. 9421, Cambridge, Mass.

- Svensson, L.E.O. and M. Woodford (2001). Indicator variables for optimal policy under asymmetric information, *Journal of Economic Dynamics and Control*, to appear.
- Svensson, L.E.O. and M. Woodford (2002). Implementing optimal policy through inflation-forecast targeting, working paper, Princeton University.
- Taylor, J.B. (1993). Discretion versus policy rules in practice, *Carnegie-Rochester Conference Series in Public Policy*, 39, 195-214.
- Whittle, P. (1981). Risk-sensitive linear/quadratic/Gaussian control, *Advances in Applied Probability*, 13, 764-777.
- Whittle, P. (1982). *Optimization over Time – Dynamic Programming and Stochastic Control*, Volume 1, John Wiley & Sons, Chichester.
- Whittle, P. (1990). *Risk-Sensitive Optimal Control*, John Wiley & Sons, Chichester.
- Woodford, M. (2001). The Taylor rule and optimal monetary policy, *American Economic Review*, *Papers and Proceedings*, 91, 232-237.
- Woodford, M. (2002). Optimal interest-rate smoothing, *Review of Economic Studies*, forthcoming.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, Princeton.