ABSTRACT

We examine the implications for the optimal interest rate rule that follow from relaxing the assumption that the policymaker's loss function is quadratic. We investigate deviations from quadratics for both symmetric and asymmetric preferences for a single target and find that (i) other characterisations of risk aversion than implied by the quadratic only affect dead-weight losses, unless there is multiplicative uncertainty; and (ii) asymmetries affect the optimal rule under both additive and multiplicative uncertainty but result in interest rate paths observationally similar, and in some cases equivalent, to those implied by a shifted quadratic. Our results suggest that in the context of monetary policymaking the convenient assumption of quadratic losses may not be that drastic after all.

Keywords: Loss functions, uncertainty, optimal monetary policy rules

JEL Classification: E42, E52, E61

MONETARY POLICY LOSS FUNCTIONS: TWO CHEERS FOR THE QUADRATIC[1]

Jagjit S Chadha[2] *University of Cambridge*

Philip Schellekens[3] ‡ *London School of Economics*

NOVEMBER 1999

"The assumption of a quadratic is, of course, subject to the objection that it treats positive and negative deviations from target as equally important. The use of a fancier utility function would provide additional reasons for departing from certainty equivalence." William Brainard (1967, p 413)

1. INTRODUCTION

Since the inception of the Tinbergen-Theil framework for analysing monetary policy, in the 1950s, there has been an uneasy acceptance of the quadratic loss function.[4] Brainard's "realistic"[5] extensions to the basic framework recognised this potential

[1] This paper is dedicated to the memory of Andrew Dumble. We are grateful for comments from seminar participants at the Bank of England, University of Southampton and the 1999 Royal Economic Society Annual Conference at the University of Nottingham. We acknowledge helpful comments from Matt Canzoneri, Alec Chrystal, John Driffill, Charles Goodhart, Nobu Kiyotaki, Jozef Plasmans, Chris Salmon, Paul Tucker, Danny Quah and two referees. Jenny Salvage provided excellent research assistance. Philip Schellekens acknowledges financial support from the Fund for Scientific Research – Flanders (Belgium). Any views expressed or errors which remain are in the paper are the solely the responsibility of the authors.

[2] Clare College, Cambridge and Department of Applied Economics, University of Cambridge.
Address: Department of Applied Economics, Austin Robinson Building, Sidgwick Avenue, Cambridge, CB3 9DE. Tel: +44 1223 335280. E-mail: jagjit.chadha@econ.cam.ac.uk

[3] Financial Markets Group, London School of Economics and University of Antwerp.
Address: FMG, LSE, Houghton Street, London WC2A 2AE. Tel: +44 171 955 7894.
E-mail: p.schellekens@lse.ac.uk

[4] Both Tinbergen (1954, pp 49-51) and Theil (1966) themselves were clearly aware of the potential limitations of quadratic losses both in terms of describing risk and possible prejudice to the robustness of results. For example, Theil writes (p 19) "...[T]here is no particular reason to assume that the loss function should always be quadratic...th[e] assumption [is a] convenient first approximation. When we try to generalise...it appears that the results become much more complicated...it turns out frequently that the results become completely unmanageable. This is undoubtedly why the quadratic loss function has such a prominent place in several fields... ".

[5] See Tobin's (1990) appreciation, *inter alia*, of the 1967 Brainard paper.

limitation and this makes the lack of attention paid in the subsequent literature surprising. One of the most respected of academics cum policymaker, Alan Blinder (1997), asks for similar consideration, and it is our intention to provide a response.[6]

This paper re-examines "the Brainard conservatism principle" (Blinder, 1997, p 11) with respect to the optimal policy rule with one instrument and one policy objective under non-quadratic preferences. Recall that the standard result, assuming a linear Phillips curve and quadratic losses in the presence of additive uncertainty, means that the policy instrument is set to offset any shock completely and immediately. It is only when uncertainty is represented in a multiplicative form - where imperfect control over the economy is represented by uncertainty over the impact of policy changes on the target variable – that the policy instrument is moved cautiously and gradually to offset a shock. The adoption of quadratic losses would seem to be an important part of this story as these suggest a particular, and possibly perverse, attitude to risk. One where, for example, the policymaker is indifferent between a one-period *under*shoot of the inflation target by 4% and a four-period *over*shoot by 2% (assuming, of course, that there is no discounting). Also, the use of the quadratic involves the implicit assumption of symmetry and it is worth examining how possible asymmetries would interact with the presence of additive and multiplicative uncertainty.

It seems quite plausible that if the characterisation of the policymaker's behaviour were made in a more appealing manner than quadratic utility then the specific generation of the "conservatism principle", in response to multiplicative uncertainty alone, may be overturned. In fact, much recent work in both consumption theory and applied finance has involved examining the integration of newer concepts of utility to older pricing puzzles.[7] Again, given the influence of this healthy literature it is surprising how little impact this has made on the analysis of optimal policy. And it is the examination of the robustness of the "conservatism principle" to the deviations (sic) from quadratic losses that will be the focus of this paper.

The rest of the paper is structured as follows. Section 2 examines the impact on the optimal interest rule when the loss function reflects constant absolute risk aversion (CARA) in the face of additive uncertainty, and by analogy other classes of risk aversion (such as CRRA). Section 3 examines the impact on the optimal interest rate rule of both additive and multiplicative uncertainty when preferences are asymmetric. Section 4 analyses simulations of the resulting optimal rules in four different cases and a graphical general solution to the time path of interest and inflation rates following an inflation

[6] Blinder (1997, p 6) writes "[A]cademic macroeconomists tend to use quadratic loss functions for reasons of mathematical convention, without thinking much about their substantive implications. The assumption is not innocuous...I believe that both practical central bankers and academics would benefit from more serious thinking about the functional form of the loss function."

[7] See Deaton (1992) for recent developments in consumption and Shiller (1998) for a signpost to the next generation of applied finance work in the non-expected utility paradigm.

shock. Section 5 offers concluding remarks, discusses some implications for the optimal delegation of monetary policy and suggests some possible further work.

2. DEVIATIONS FROM QUADRATICS: OTHER ATTITUDES TO RISK

The Brainard conservatism principle (leading to a cautious and gradualist setting[8] of the instrument) results from the interaction of multiplicative uncertainty with quadratic preferences. Alternatively, one could take the view that such smoothing is simply caused by a form of risk aversion (with respect to inflation volatility) other than the one implied by quadratics. For example, in terms of risk, two well-known properties of the quadratic are that the coefficient of relative risk aversion is 1 and that its third derivative is zero: the former implies that the elasticity of the policymaker's marginal loss with respect to inflation is always 1 and the latter implies that the variability of inflation does not affect marginal loss. The use of non-quadratics might be analogous to agents smoothing consumption in response to temporary income shocks. One might expect that the introduction of loss functions that deliver such smoothing in a consumption setting would also produce interest rate gradualism in a setting of monetary policymaking. This is, however, not necessarily true. It is shown below that caution and gradualism may not follow from non-quadratic preferences as long as losses are symmetric and uncertainty is additive.

2.1 The Framework

Consider the following simple control problem:

(1) $$\underset{\{i_t\}}{Min}\ \Omega \equiv E_0\left\{ \sum_{t=0}^{\infty} \delta^{t+1} L\left(\pi_{t+1};\pi^*\right)\right\}$$

subject to

(2) $$\pi_{t+1}-\bar{\pi}=a\left(\pi_t-\bar{\pi}\right)-b\left(i_t-\bar{i}\right)+e_{t+1},\quad e_{t+1}\sim N\left(0,\sigma_e^2\right),\ \text{i.i.d.}$$

---

[8] Caution has a particular meaning in this paper: it refers to a long-run policy stance that is closer to its natural level than if there were no multiplicative instrument uncertainty. Gradualism refers to the smoothing of an instantaneous policy adjustment into smaller adjustments over time such that the loss occurred through induced policy variability (due to the presence of multiplicative instrument uncertainty) is optimally reduced.

where $\pi_t, \bar{\pi}$ and π^* refer to the inflation rate at time t, the unconditional mean of inflation and the socially optimal rate of inflation; δ is the discount factor; a measures the persistence of the inflation process; b is the policy multiplier; $i_t - \bar{i}$ is the deviation of the policy instrument from its natural level; e_{t+1} is an additive shock at time t+1.

Equations **(1)** and **(2)** assume that the policymaker sets a path of interest rates such that future deviations of inflation from its target are minimised subject to an inflation relationship and a particular specification of preferences. The reduced-form process for inflation in **(2)** is kept deliberately simple as the emphasis of this framework is on the specification of the preferences of the policymaker.[9] The minimal features we require are persistence in inflation (ensuring a role for policy) and uncertainty (of the additive and later the multiplicative form). Inflation is described as an autoregressive process with a long-run mean equal to $\bar{\pi}$. Inflationary persistence is captured by parameter a $(0 \le a < 1)$. As well as the additive shock, e_{t+1}, inflation can be influenced by deviations of the policy instrument i_t from its neutral level $\bar{i}$ - which will be set to zero for the rest of the paper. The policy multiplier, b $(b > 0)$, translates policy actions into inflation outcomes and is assumed to be non-stochastic in this section. The only source of uncertainty is the additive shock, e_{t+1}, which is normal and i.i.d. with mean 0 and variance σ_e^2. Note that the instrument is set at the beginning of each period, whereas the shock occurs at the end of each period. As a result, a shock has one-for-one first-round effects on inflation during the current period, but stabilisation policy can offset its second-round effects in subsequent periods. Finally, the intertemporal loss function in (1) consists of the infinitely discounted sum of per-period losses $L(\pi_{t+1};\pi^*)$. Discount factor δ takes some value between 0 and 1.

Let us now turn to the specification of the per-period loss function. Natural candidates for a richer description of the policymaker's behaviour towards risk would be the exponential (or CARA) and the isoelastic (or CRRA) loss functions

(3) $$L^{cara}\left(\pi_{t+1};\pi^*\right) = \exp\left[\beta\left(\pi_{t+1} - \pi^*\right)\right] - 1 \qquad \text{with } \beta > 0$$

and

(4) $$L^{crra}\left(\pi_{t+1};\pi^*\right) = \tfrac{1}{1-\rho}\left(1 + \pi_{t+1} - \pi^*\right)^{1-\rho} - 1 \qquad \text{for } \rho \neq 1 \text{ and } \rho > 0$$

(4') $$L^{crra}\left(\pi_{t+1};\pi^*\right) = \ln\left(1 + \pi_{t+1} - \pi^*\right) - 1 \qquad \text{for } \rho = 1$$

[9] In a model that does incorporate a private sector with forward-looking expectations, the sluggishness in equation **(2)** could be derived from the existence of nominal rigidities such as menu costs or overlapping nominal contracts.

As the name suggests, the CARA loss function is characterised by constant *absolute* risk aversion (equal to β), whereas CRRA implies constant *relative* risk aversion (equal to ρ).

Recall that quadratic losses,

$$L^{q}\left(\pi_{t+1};\pi^{*}\right)=\tfrac{1}{2}\left(\pi_{t+1}-\pi^{*}\right)^{2},$$

imply *increasing absolute* risk aversion.

In order to substantiate our claim that non-quadratic preferences (and thus other descriptions of risk aversion than implied by the simple quadratic) do not deliver policy-caution or policy-gradualism, we will focus subsequently on the CARA loss function.[10]

There is however one important caveat before proceeding. In the consumption literature, smoothing occurs due to the interaction between risk aversion and an inter-temporal budget constraint ensuring an intertemporal trade-off between consumption today and consumption in the future. In the Tinbergen-Theil setting, however, there is no natural constraint on the inter-temporal behaviour of inflation - higher inflation does not necessarily imply lower inflation tomorrow. Due to the absence of a properly defined resource constraint, optimisation under CARA preferences will yield unrealistic solutions for the setting of interest rates. If the inflation target were for example equal to zero, optimality would require the interest rate to be set at plus infinity because the resulting negative rates of inflation imply policy gains. While the introduction of an output term in the loss function could certainly offset some of this perverse tendency, we have opted for a modification of the CARA function such that it incorporates the concept of a target. This will also allow for a more natural and direct comparison with the quadratic paradigm.

2.2 Risk Aversion Only Affects Dead-Weight Losses

Consider the following symmetric two-part CARA loss function,[11]

(5) $$L\left(\pi_{t+1};\pi^{*}\right)=\begin{cases}\exp\left[-\beta_{1}\left(\pi_{t+1}-\pi^{*}\right)\right]-1, & \text{for}\,\pi_{t+1}<\pi^{*}\\ \exp\left[\beta_{2}\left(\pi_{t+1}-\pi^{*}\right)\right]-1, & \text{for}\,\pi_{t+1}\geq\pi^{*}\end{cases}$$

with $\beta_1=\beta_2=\beta$ for symmetry. This two-part function is displayed in Figure 1. Using the indicator function, the loss function can be re-written as follows

(6) $$L\left(\pi_{t+1};\pi^{*}\right)=I_{t+1}\exp\left[-\beta_{1}\left(\pi_{t+1}-\pi^{*}\right)\right]+\left(1-I_{t+1}\right)\exp\left[\beta_{2}\left(\pi_{t+1}-\pi^{*}\right)\right]$$

[10] The CRRA case is entirely analogous.

[11] Horowitz (1987) adopts a similar framework.

where I_{t+1} takes the value 1 for inflation draws below target and zero for draws above. Equation **(1)** can therefore be re-written as

FIGURE 1
A Symmetric Two-Part CARA Loss Function

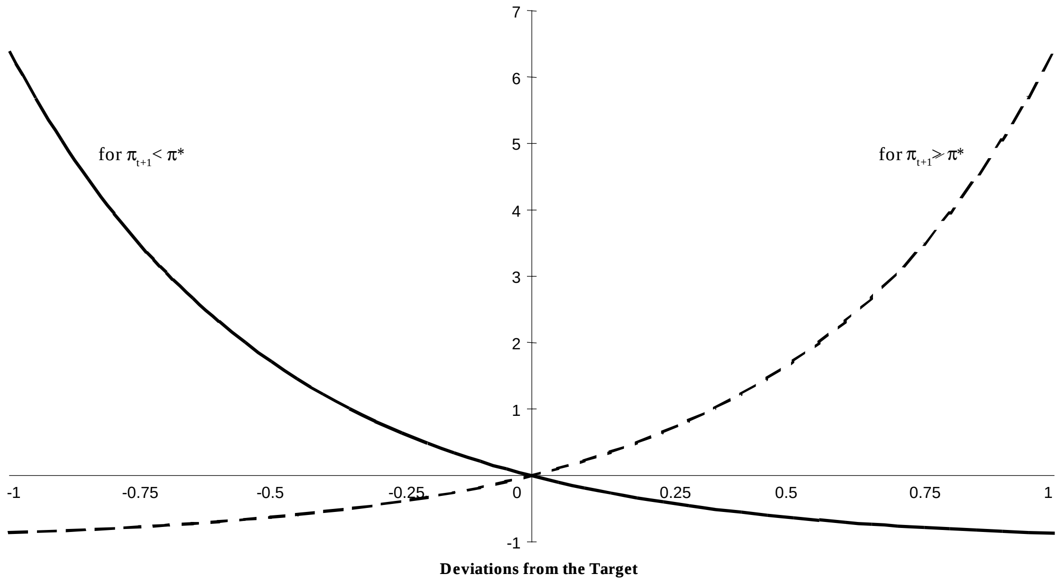

(7) $$\underset{\{i_t\}}{Min}\, E_0 \sum_{t=0}^{\infty} \delta^{t+1} \left\{ I_{t+1} \exp\left[-\beta_1\left(\pi_{t+1} - \pi^*\right)\right] + \left(1 - I_{t+1}\right)\exp\left[\beta_2\left(\pi_{t+1} - \pi^*\right)\right] \right\}$$

As a result of the one-period control problem suggested by equation (2), the control problem in **(7)** can be reduced to

(8) $$\underset{i_t}{Min}\, E_t \left\{ \begin{array}{l} I_{t+1} \exp\left[-\beta_1\left(\bar{\pi} + a\left(\pi_t - \bar{\pi}\right) - bi_t + e_{t+1} - \pi^*\right)\right] + \\ \left(1 - I_{t+1}\right) \exp\left[\beta_2\left(\bar{\pi} + a\left(\pi_t - \bar{\pi}\right) - bi_t + e_{t+1} - \pi^*\right)\right] \end{array} \right\}.$$

For notational convenience, set X equal to $\pi^* - \bar{\pi} - a\left(\pi_t - \bar{\pi}\right) + bi_t$. In order to evaluate the probability of the additive shock being on one side of the split-distribution or the other, we need to examine the probability of next period's inflation rate being larger or smaller than the target. The expectations over the indicator functions are given by:

(9) $$\begin{aligned} E_t\left(I_{t+1}\right) &= \text{Prob}\left\{\text{e}_{t+1} < X\right\} = \Phi\left(X / \sigma_e\right) \\ E_t\left(1 - I_{t+1}\right) &= \text{Prob}\left\{\text{e}_{t+1} \geq X\right\} = 1 - \Phi\left(X / \sigma_e\right), \end{aligned}$$

where $\Phi(\cdot)$ is the cumulative density of the standard normal. The argument in **(8)** then becomes:

(10) $$\Phi(X/\sigma_e)\underset{-\infty}{\overset{X}{E_t}}\exp[-\beta(-X+e_{t+1})] + (1-\Phi(X/\sigma_e))\underset{X}{\overset{+\infty}{E_t}}\exp[\beta(-X+e_{t+1})]$$

where expectations are taken over the intervals $[-\infty, X]$ and $[X,\infty]$ respectively. Equation **(10)** can be evaluated to give the following expression

(11) $$\exp\left(\frac{\beta^2\sigma_e^2}{2}\right)\times\left\{\begin{array}{l}\Phi(X/\sigma_e)\ \Phi(\beta\sigma_e + X/\sigma_e)\ \exp(\beta X)\ + \\ (1-\Phi(X/\sigma_e))\ (1-\Phi(\beta\sigma_e + X/\sigma_e))\ \exp(-\beta X)\end{array}\right\}$$

and we can see that this function will be minimised when *X* equals zero, i.e. when interest rates are set to close the gap between current inflation and target completely and immediately.[12]

(12) $$X = 0 \iff i_t = \frac{a}{b}(\pi_t - \bar{\pi}) + \frac{1}{b}(\bar{\pi} - \pi^*)$$

This expression is exactly the same as the one obtained under quadratic losses. It says that the deviation of the optimal interest rate from its neutral level (recall that the neutral level has been set to zero) is a function of two components. The first component is essentially a simple feedback rule, implying that the interest rate response depends on how far last period's inflation was away from its long-run mean. The second component derives from the possibility that the inflation target does not necessarily correspond with the long-run mean of the autoregressive inflation process. If the inflation target is such that inflation will have to be sustained above (below) its long-run level, then interest rates need to be permanently lower (higher).[13]

From **(12)** an important conclusion can be derived: deviations from quadratics (in the form of Equation **(5)**) do not affect the optimal rule, as long as the Tinbergen-Theil loss function is symmetric and uncertainty is additive. Interest rates will still be set so as to offset completely any shock to inflation last period.

These results also imply that richer descriptions of risk aversion (to that implied by quadratic losses) are irrelevant if the maintained hypothesis of additive uncertainty and symmetric preferences is not violated. To put it differently, risk aversion merely affects dead-weight losses.

Recall that nothing can be done about the first-round effects of an additive shock to inflation. Only the second round effects to the next, and subsequent, period's inflation rate can be stabilised. We have shown that stabilisation will be complete and immediate: there is no element of policy-caution or policy-gradualism. This is what is meant with risk aversion, *per se*, being irrelevant for the *optimal rule*.

---

[12] Global convexity of **(5)** ensures that this is the global minimum.

[13] When the inflation target coincides with the unconditional mean of the inflation process, then the optimal rule is $i_t = (a/b)(\pi_t - \pi^*)$.

Of course, the extent of risk aversion is not irrelevant for the *value of the loss*, which equals:

$$\textbf{(13)} \qquad \frac{1}{2}\exp\left(\frac{\beta^2\sigma_e^2}{2}\right)\times\left\{1 + \underset{-\beta\sigma_e}{\overset{\beta\sigma_e}{N}}(0,1)\right\}$$

if the optimal policy is implemented. Note that the equilibrium value of the loss in **(13)** increases both with the level of additive variability and the extent of the policymaker's aversion to risk. Intuitively, an increase in risk aversion, for example, means that a particular level of additive variability becomes more costly as the first-round inflationary effects cannot be undone. As a result, at the time that the policymaker can act upon the shock (i.e. the next period), the loss has occurred and is dead weight.[14]

2.3 Multiplicative Uncertainty is Required for Caution and Gradualism

Introducing multiplicative uncertainty should affect the optimal rule. Multiplicative uncertainty is taken to mean that parameter b is a uncertain with some strictly positive and finite variance. The reason is simply that the actions of the policymaker bring about an additional source of variability into the loss function. Thus, the dead-weight loss argument no longer applies. As in the quadratic case, this will make optimal policy cautious and gradualist.

In a framework with symmetric preferences and both additive and multiplicative uncertainty there are now two interactions going on. First of all, there is the earlier result that, for a given level of additive variability, risk aversion increases the dead-weight losses due to first-round effects on inflation. But this does not affect the optimal rule. Secondly, risk aversion will amplify the costs of a given degree of multiplicative uncertainty when the policymaker tries to stabilise the second-round effects on inflation. The more risk averse the policymaker is, the more cautious and gradualist policy will be. In contrast to the interaction between risk aversion and additive variability, risk aversion will not affect the dead-weight losses through the multiplicative uncertainty channel because this channel operates when interest rates are moved. Since in this model interest rate actions tomorrow cannot offset the first-round effects on inflation today, risk aversion does not amplify the dead-weight losses through this mechanism.

Unfortunately, the current split-CARA framework becomes analytically intractable when multiplicative uncertainty is introduced, so that the above claims still require verification. In any case, the framework has served our purpose, in that we show formally that risk aversion is irrelevant in a setting of additive variability and symmetric preferences. If one wishes to examine issues of policy-caution and policy-gradualism, non-quadratic preferences (and their implications for risk aversion), *per se*, are not sufficient. Moreover, they are not necessary as one can easily examine these issues in a quadratic framework.

3. DEVIATIONS FROM QUADRATICS: ASYMMETRY

The quadratic paradigm is sometimes criticised because positive and negative deviations from the target are treated symmetrically. In this section, we explore the implications arising from the assumption of asymmetric losses in a setting of monetary policymaking. The analysis will show that a non-quadratic loss function around a particular target is observationally similar to a quadratic loss function around a different target, even if we allow for a rich description of the stochastic nature of the economy. Asymmetric losses may be an interesting way of characterising the policymaker's attitude to policy outcomes, *if* such attitudes reflect either (i) a view about the social welfare function[15] or (ii) an exogenous view of the policymaker about the embarrassment costs of positive as compared to negative deviations from target or both.

As the inflation remit of the Bank of England is symmetric, this invalidates the application of the second political economy line of thought. Nevertheless, the possibility remains that the government, when determining the level of the inflation target, has taken into account possible asymmetries to the social cost of inflation. As a consequence, the level of a symmetric inflation target may internalise possible asymmetries in the social cost of deviations of inflation from that target.

Varian (1975), in his discussion on the losses faced by property valuers, suggests an asymmetric loss function which rises linearly on one side of zero and rises exponentially on the other side.[16] It is this loss function, the so-called LINEX (Linear Exponential), that we employ to examine the impact of both additive and multiplicative uncertainty on the optimal path of interest rates.

3.1 Introducing the LINEX Function

Varian (1975) introduced the following convex loss function:

(14) $$L(\xi) = \alpha \exp(\gamma\xi) - \beta\xi - \alpha, \qquad \text{with } \gamma, \beta \neq 0, \alpha > 0$$

---

[14] In finance theory terms, choosing a risk aversion parameter may alter the price of risk but as additive uncertainty is uncorrelated with policy risk, there is no impact on the insurable quantity of risk: this means that the optimal plan does not alter.

[15] The social welfare function with respect to inflation may be asymmetric because of shoe-leather-type arguments on the costs of inflation or views on the probability of debt deflation.

[16] The argument used by Varian was that underassessment of property values led to approximately linear revenue losses whereas overassessment may result in appeals, litigation and other costs. Zellner (1986)

where ξ is the deviation of the policy objective from target.[17] We can see that $L(0) = 0$ and that for a minimum to exist at $\xi = 0$ we must have $\gamma\alpha = \beta$.[18] So **(14)** can be re-written as:

(15) $L(\xi) = \alpha[\exp(\gamma\xi) - \gamma\xi - 1]$, with $\gamma \neq 0, \alpha > 0$.

Note that γ determines that extent of the asymmetry in the LINEX function and α scales the losses. Figure 2 shows the LINEX function for $\alpha = 1$ and for γ = 0.5, 1.0, 1.5. For comparison the quadratic losses are also plotted; the x-axis plots the deviation, ξ. Note that for small losses the difference between the LINEX and quadratic appear small and, in fact, if we expand $\exp(\gamma\xi) = 1 + \gamma\xi + \frac{1}{2}\xi^2\gamma^2$ we find that $L(\xi) = \frac{1}{2}\xi^2\gamma^2$. But, of course, for larger values of ξ the differences in losses tend to become substantial. Appendix A discusses some related points on the LINEX function.

FIGURE 2
LINEX and Quadratic Losses Compared

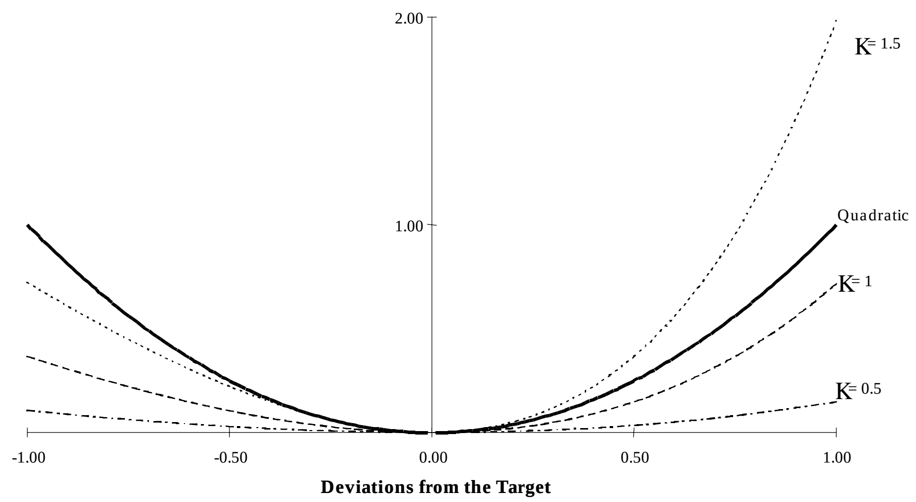

suggests an even clearer example by pointing out that in the construction of dams underestimate of peak flows is much more serious than an overestimate.

[17] In a related vein, Christoffersen and Diebold (1997) study the optimal prediction problem under general asymmetric loss structures.

[18] This is simply found by differentiating **(14)** with respect to ξ and solving for β.

3.2 The Optimal Interest Rule Under Asymmetric Losses

The intertemporal maximisation problem with policy being subject to asymmetric preferences and multiplicative instrument uncertainty can be summarised as:

(16) $$\underset{\{i_t\}}{Min} E_0 \left\{ \sum_{t=0}^{\infty} \delta^{t+1} L(\pi_{t+1}) \right\},$$

with,

(17) $$L(\pi_{t+1}) = \exp[\gamma(\pi_{t+1} - \pi^*)] - \gamma(\pi_{t+1} - \pi^*) - 1 \qquad \text{with } \gamma > 0 \ .$$

subject to

$$\pi_{t+1} - \bar{\pi} = a(\pi_t - \bar{\pi}) - b_{t+1} i_t + \varepsilon_{t+1} \qquad \text{with} \quad t = 0,1,...;\ 0 \le a \le 1;\ b > 0$$

$$\begin{pmatrix} b_{t+1} \\ \varepsilon_{t+1} \end{pmatrix} \sim N\left[\begin{pmatrix} \bar{b} \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_b^2 & 0 \\ 0 & \sigma_e^2 \end{pmatrix}\right], \qquad \text{i.i.d.}$$

The preferences of the policymaker are described by a LINEX loss function. If for some reason overshooting the inflation target (π^*) is more costly than undershooting it, we can restrict γ to be strictly positive. This will imply that undershooting is penalised in an approximately linear fashion, whereas the marginal losses from overshooting are increasing in next period's inflation rate. Of course, the following analysis could also be completed for the case where negative deviations imply exponential losses and positive deviations imply linear losses. But for the remainder of this paper, we arbitrarily require γ to be strictly positive, without loss of generality.

The aim of the exercise is to find the optimal interest rate path which will minimise the intertemporal loss function subject to the relationships in the above equations. Note that control is imperfect due to both additive and multiplicative uncertainty.[19] Both sources of uncertainty are assumed to follow a normal distribution and to be independent of each other (i.e. σ_{be}). The parameter of inflation persistence is assumed to be a known constant.

[19] As to the precise nature of the uncertainty, the authorities may believe that the parameters of the model are random variables with a particular positive variance. Alternatively, they may regard the true (population) parameters values as being non-random quantities in the underlying model but put some margin of error on their estimated (sample) values. In what follows, we assume that the underlying additive shocks are genuine random variables (which will ensure a role for stabilisation policy) and that the multiplicative uncertainty mainly derives from imperfect inference (which will deliver a cautious setting of policy). For a discussion, see Brainard (1967, p 413-4).

The solution can be found by solving

$$\textbf{(18)} \quad \underset{i_t}{Min}\, E_t \left\{ \begin{array}{l} \exp\left[\gamma\left(\bar{\pi} + a(\pi_t - \bar{\pi}) - b_{t+1} i_t + \varepsilon_{t+1} - \pi^*\right)\right] - \\ \gamma\left(\bar{\pi} + a(\pi_t - \bar{\pi}) - b_{t+1} i_t + \varepsilon_{t+1} - \pi^*\right) - 1 \end{array} \right\}$$

subject to

$$\begin{pmatrix} b_{t+1} \\ \varepsilon_{t+1} \end{pmatrix} \sim N\left[\begin{pmatrix} \bar{b} \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_b^2 & 0 \\ 0 & \sigma_e^2 \end{pmatrix}\right], \qquad \text{i.i.d.}$$

The first-order condition of this optimisation problem implicitly defines the optimal interest rate setting (for mathematical details see Appendix B):

$$\textbf{(19)} \quad i_t = \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}\left(\bar{\pi} - \pi^*\right) + \frac{\gamma\sigma_e^2}{2\bar{b}} + \frac{\gamma\sigma_b^2}{2\bar{b}} i_t^2 + \frac{1}{\gamma\bar{b}} \ln\left(1 - \frac{\gamma\sigma_b^2}{\bar{b}} i_t\right).$$

It is analytically intractable to get a reduced form solution for the optimal rule in the general case. However, there are some interesting simple special cases.

<u>THE DEFAULT CASE</u>: No Asymmetry And Only Additive Uncertainty

If there is no multiplicative instrument uncertainty and the preferences of the policymaker tend to symmetry, then the optimal rule collapses to

$$\textbf{(20)} \quad i_t = \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}\left(\bar{\pi} - \pi^*\right).$$

In order to interpret this expression, assume that a one-off additive shock has occurred at time t, producing an overshooting of the inflation target by $x\%$ points. Nothing can be done about the initial boost in inflation, but as long as there is some persistence in inflation (i.e. $a > 0$), the second-round effects of the shock to inflation at time t+1 (another deviation from the inflation target by $a\,x\%$ points) will be fully neutralised.

<u>THE ASYMMETRY CASE</u>: Asymmetry and Only Additive Uncertainty

If multiplicative uncertainty is absent and preferences are asymmetric, the optimal rule becomes

$$\textbf{(21)} \quad i_t = \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) + \frac{\gamma\sigma_e^2}{2\bar{b}}.$$

The assumption of asymmetric risk aversion produces an upward bias in the optimal rule if overshooting is considered to be more costly than undershooting. It is clear that the interest rate 'premium' due to risk aversion increases with the extent of additive variability, as well as with the degree of asymmetry in the preferences of the policymaker.

<u>THE UNCERTAINTY CASE</u>: No Asymmetry and Both Types of Uncertainty

Letting preferences approach symmetry, L'Hôpital's rule delivers the optimal rule under multiplicative instrument uncertainty (see Appendix C):

$$\textbf{(22)} \quad i_t = \frac{1}{1 + \frac{\sigma_b^2}{\bar{b}^2}} \left(\frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) \right),$$

which is exactly the Brainard result (1967, p 414) that one would obtain under quadratic losses. If the coefficient of variation ($\sigma_b / \bar{b}$) exceeds zero, the optimal interest rate response will be such that the gap with the inflation target is not entirely closed.

Note that we can also solve for the long-run steady-state values of inflation and interest rates. This will be illustrated for the default and the uncertainty cases. Analytically, Equations **(20)** and **(22)** need to be matched with the steady-state condition

$$\textbf{(23)} \quad i = \frac{a-1}{b}(\pi - \bar{\pi})$$

which follows from Equation **(2)**.

For the default case, the steady-state values for the inflation and interest rate are respectively

$$\textbf{(24)} \quad \pi^{SS} = \pi^*$$

$$\textbf{(24')} \quad i^{ss} = \frac{a-1}{b}(\pi^* - \bar{\pi}).$$

This tells us that in the long run inflation will settle down at the inflation target. Unless the long-run mean is equal to the inflation target, this requires continual policy intervention ($i^{SS} \neq 0$).

For the uncertainty case, the steady state can be characterised by

$$\textbf{(25)} \quad \pi^{SS} = \lambda\pi^* + (1-\lambda)\overline{\pi} \qquad \text{where } \lambda = 1/(1+\{1-a\}\sigma_b^2/\overline{b}^2)$$

$$\textbf{(25')} \quad i^{ss} = \frac{1}{1+\sigma_b^2/b^2}\frac{a\lambda-1}{b}\left(\pi^* - \overline{\pi}\right)$$

Equation **(25)** shows that the long-run steady state in the uncertainty case can be represented as a weighted average of the long-run mean of the inflation process and the inflation target. If there is no multiplicative uncertainty, then λ equals 1 (as $\sigma_b^2 = 0$) and long-run inflation will hit the inflation target. The other extreme is the case of infinite multiplicative uncertainty which delivers λ equal to 0 (as $\sigma_b^2 \to \infty$) and a long-run inflation rate which reverts to the long run mean of the process. Similarly, Equation **(25')** shows that the degree of activism is inversely related to the degree of multiplicative uncertainty. This is what we mean by policy-caution in this particular setting: because of multiplicative uncertainty, the long-run response of the interest rate is biased towards its neutral level; as a result, inflation will settle down closer to its mean. Note also that if the long-run mean and the inflation target coincide, then the issue of caution entirely evaporates: inflation settles down at its target and interest rates at their neutral level.

Returning to the most general case (i.e. multiplicative instrument uncertainty and asymmetric preferences), note that the last two terms in **(19)** result from the introduction of multiplicative uncertainty into the asymmetry case. An interesting issue is the extent to which these terms lead to qualitatively different results compared to the introduction of such uncertainty in the default case. In order to answer this question let us turn to some simulation results.

4. RESULTS

Section 3 derived expressions for the general form of the optimal rule, Equation **(19)** and for three relevant special cases: that of asymmetry (*A*), uncertainty (*U*) and the default case. Via simulation techniques, this section explores the implication of the optimal rule for the setting of interest rates both as a static first period choice and as a dynamic path.[20] We are also able to derive graphically the solution to the choice of optimal interest rates for all four cases. The final result allows us to answer the question of whether asymmetric preferences are sufficient to deliver gradualist interest rate responses and whether cautious interest rate responses are delivered.

For expositional purposes, we have made one modification to the general expression in Equation **(19)**: the long-run mean of the inflation process has been set to zero (i.e. $\bar{\pi} = 0$). This will give further insights on the interaction between a long-run mean and an inflation target, which are not necessarily equal.

4.1 The Initial Interest Rate Response

Figure 3 examines the initial interest rate response to inflation shocks under the four different cases. The size of inflation shocks (on the x-axis) is allowed to vary from -10% to +10% and the choice of optimal interest rates in the first period is shown on the y-axis. We chose the following parameter values for the simulations $\gamma = 1.5$, $a = 0.5$, $b = 1$, $\pi^* = 2.5$, $\sigma_e^2 = 0.05$ and $\sigma_b^2 = 0.5$. The parameter choice is explained as follows. As γ is the extent of asymmetry in the loss function, which tends to symmetry as $\gamma \to 0$, 1.5 from Figure 2 would seem a fair degree of asymmetry. Parameter a is the extent of non-policy related inflation persistence in the economy and is set to be something below the observed persistence - which includes policy reaction - typically found for modern industrialised economies.[21] Parameter b is the impact on inflation of an interest movement and set to allow for full pass through. Further, π^* is the inflation target and σ_e^2 is the variance of additive shocks, which is set to a small number with a value less than σ_b^2 (the variance of multiplicative uncertainty). This will ensure that the state-independent bias in interest rates, the third term on the r.h.s. of Equation **(19)**, does not swamp the interaction terms between asymmetry and multiplicative uncertainty, the last two terms in Equation (19).

Figure 3 shows that, for this range of single inflation draws, the optimal initial interest rate response rises linearly in the value of the inflation draw. In both the default case and in the *A* case the initial interest rate response rises at the rate a/b, at the rate $(a/b)(1+\sigma_b^2/b)$ in the *U* case, and approximately the same in the general case. We find that in both the default and *U* case the optimal initial interest rate response from a five percent inflation draw is zero but that in both the *A* and the general case, reflecting the asymmetry bias, an optimal initial interest rate response of zero occurs when the inflation draw is $\pi^*/a - \gamma\sigma_e^2/(2a)$. This means that in comparison to the symmetric cases the optimal initial interest rate response with an inflation draw equal to target is biased up by an intercept amount of $\gamma\sigma_e^2/(2b)$ in the *A* case but something less in the general case because of the interaction between risk aversion and uncertainty i.e. the last two terms in Equation **(19)**. [22]

[20] The optimal rule in the general case is solved using Gauss-Newton iterative procedures.

[21] The estimated sample persistence post-Bretton Woods has been in the order of 0.5-1.0 for OECD countries. Of course, the non-policy related persistence parameter will in the real world include some expectation of likely policy accommodation.

[22] Note that, for convenience, we assume that the inflation target exceeds the long-run mean of inflation. Of course, the opposite case may apply as well.

FIGURE 3
Initial Interest Rate Responses to Inflation Shocks

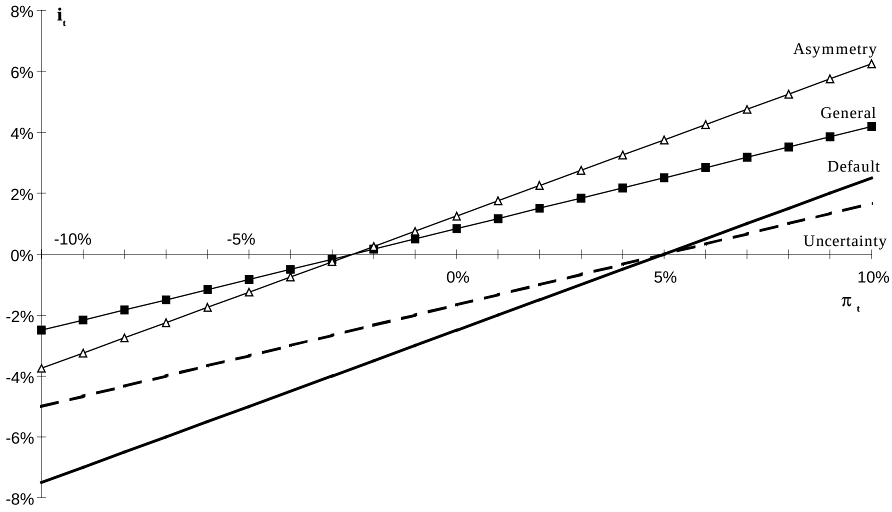

FIGURE 4
Initial Interest Rate Responses under Increasing Uncertainty

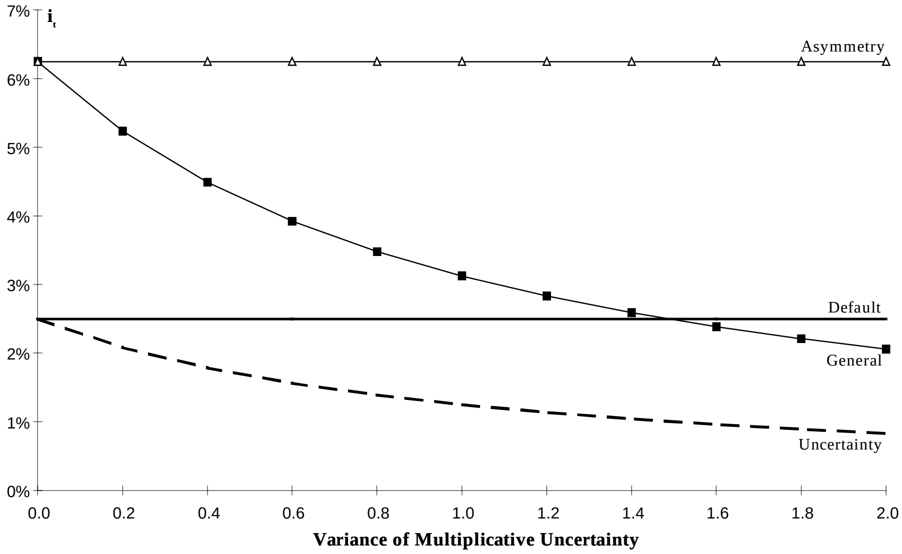

What happens to the initial interest rate choice under increasing parameter uncertainty? Figure 4 examines the initial response of interest rates to a given 10% inflation shock when the variance of *b* - the extent of multiplicative uncertainty - is allowed to increase from 0 to the implausible level of 2.[23] The chart shows that the initial interest response is state-independent of the level of multiplicative uncertainty in both the default case and *A*. But that the introduction of multiplicative uncertainty makes the

[23] For a value of *b* equal to 1, a variance of 2 may be considered to be implausible because there would be an approximately 20% chance of an increase in interest rates leading to a perverse response in inflation.

initial response of the optimal interest rule in the general case similar to the default case when the variance of b is set at around 1.5. One way of thinking about this result (if there is agreement on the other parameters) is to argue that if the authorities think that the optimal initial step in interest rates following a 10% inflation shock is 2.5%, they either live in a default world or a general world with relatively large multiplicative uncertainty. Also note that if multiplicative uncertainty rises to implausibly large levels (i.e. greater than 2) then the initial interest rate response looks similar for the U and for the general case. Or if the policymaker considers that the economic structure is chronically uncertain then, with other factors tending to be outweighed, the initial interest rate response will tend to zero.[24]

4.2 The Dynamic Interest Rate Path

We are now able to show the dynamic path of interest rates (and simultaneously for inflation) following the calculation of the initial response. We assume that the inflation rate is the beginning of period rate and the interest rate is the end of period rate. Following the initial inflation draw (π_t) and optimal interest rate response (i_t), the economy's inflation relationship, Equation **(2)** with $\bar{\pi} = \bar{i} = 0$, delivers a new inflation level (π_{t+1}), and this leads to second optimal interest rate response (i_{t+1}) and so on until the steady-state values are reached.

FIGURE 5a
Response Paths of Interest Rates

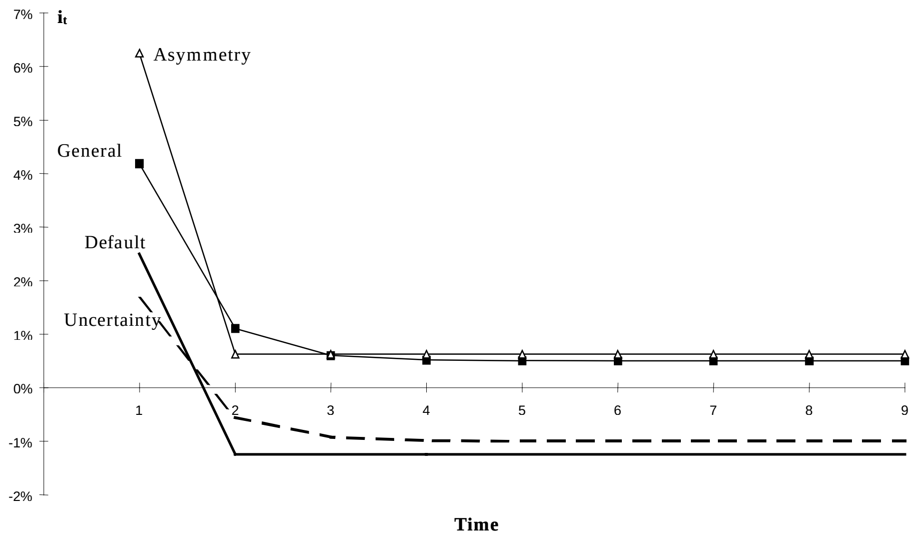

[24] This result is analogous to the Friedman (1951) argument for what might be termed "policy passivism".

Figure 5a plots the response of interest rates over time to an inflation shock of 10% under the four cases. The first point to note is that the level of steady state interest rates is different in the four cases and so then is the steady-state inflation rate, or implied target.[25] Second, note that in the default and *A* cases interest rates return to their steady-state path at the end of the second period - there is no gradualism. In the two cases involving uncertainty the return to the steady state occurs by the end of the fourth period - i.e. it is gradualist. Finally, as long as the long-run mean of the inflation process is not equal to the inflation target, the gradualist response also delivers one which is cautious, in the sense that the long-run steady-state value of the interest rate will be closer to its neutral level.

FIGURE 5b
Impulse Response Paths of Interest and Inflation Rates

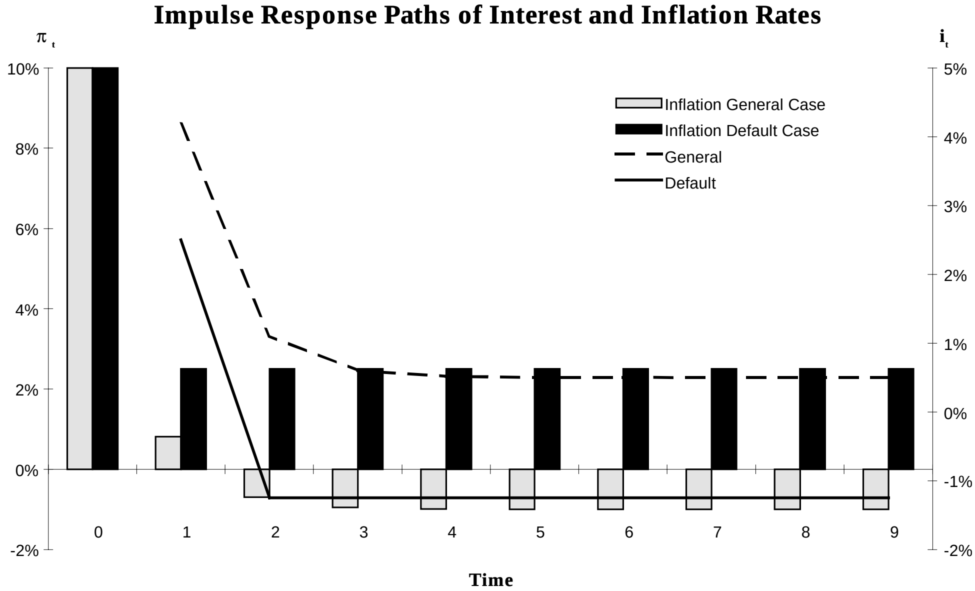

Figure 5b plots the dynamic response of interest rates and inflation to a 10% inflation shock in the default and general cases where the explicit inflation target has been set to 2.5%. For the former case, in the absence of gradualism, interest rates and inflation arrive at their steady-state values after one period.[26] In the general case, the economy is close to its steady-state at the end of the fourth period. In the general case, because of both uncertainty and risk aversion, the LINEX loss function forces the optimal policymaker to drive the economy towards a lower inflation target than explicitly stated. It is this implicit modification to the explicit target, down to some -1% in this case, and to the long run mean projected by the economy's inflation relationship (Equation **(2)**) which leads to the negative bias in the long-run inflation rate and the analogous positive bias to interest rates.

---

[25] Note from the discussion in Section 3 that the implicit inflation target is identical in the default and *U* case when the explicit inflation target, π^*, is the same as the long run mean, $\bar{\pi}$.

4.3 The Graphical Solution

Figure 6 plots a graphical solution to the simulations presented in Figures 3-5, namely the steady-state locus and the initial interest rate response. From the steady-state solution to Equation **(2)**, we find that the steady-state locus passes through the origin with slope *(a-1)/b* and cuts the initial interest rate responses at the steady state locus of inflation and interest rates. This means for the four cases shown that the default, uncertainty, general and asymmetry cases imply successively lower inflation targets and higher steady-state interest rate. Just as the dynamic paths in Figure 4 showed different steady-state interest rates, Figure 6 shows the same steady state in inflation/interest rate space. The rankings of the implied inflation targets in terms of their deviations from default case are parameter dependent but from Figure 6 we are able to say that, for positive asymmetry: (i) the non-default cases have implied inflation targets <u>lower</u> than for the default case and (ii) the implied inflation target for the general case will always be lower than that for the *U* case.

Figure 7 plots the dynamic response of interest rates in the inflation/interest rate space. To find the dynamic response to an initial inflation draw, a vertical line is displayed from the inflation draw to the initial interest rate response for each of the four cases. This vertical line shows the jump variable property of interest rates in the first period. In the cases without multiplicative uncertainty the next step (indicated by the little arrow) is the final one and represents the move back to the steady-state locus. In the cases involving uncertainty, the next steps involve exponentially decaying movements along the initial response line back to the steady-state locus (this is visualised by the periodic intersection of the arrow with the relevant response path). From the graph we can also see that the gradualist response is also cautious: the intersection of the response path and the steady-state locus is closer to the origin when there is multiplicative uncertainty.

So why is the inflation target lower in the non-default case? There are three separate reasons. The easiest way is to first examine the move from the default case, with no multiplicative uncertainty or from the asymmetry case to *U*, we can see that arithmetically the bias follows from the (square) of coefficient of variation in the denominator. Intuitively, this means that the lack of perfect control over the economy makes the optimal policymaker choose to base interest rate decisions around an implicit inflation target somewhat lower than the explicit inflation target. Because the long-run mean of inflation embedded in the Phillips curve relationship is not identical to that of the explicit inflation target, then the achievement of the inflation target actually requires some policy initiative towards which, in the presence of control uncertainty, the policymaker minimises expected losses by aiming too low.

---

[26] Interest rates are a negative deviation from base in this example because the inflation target is set above the long run mean (i.e. zero).

FIGURE 6
The Steady State Values of Inflation and Interest Rates

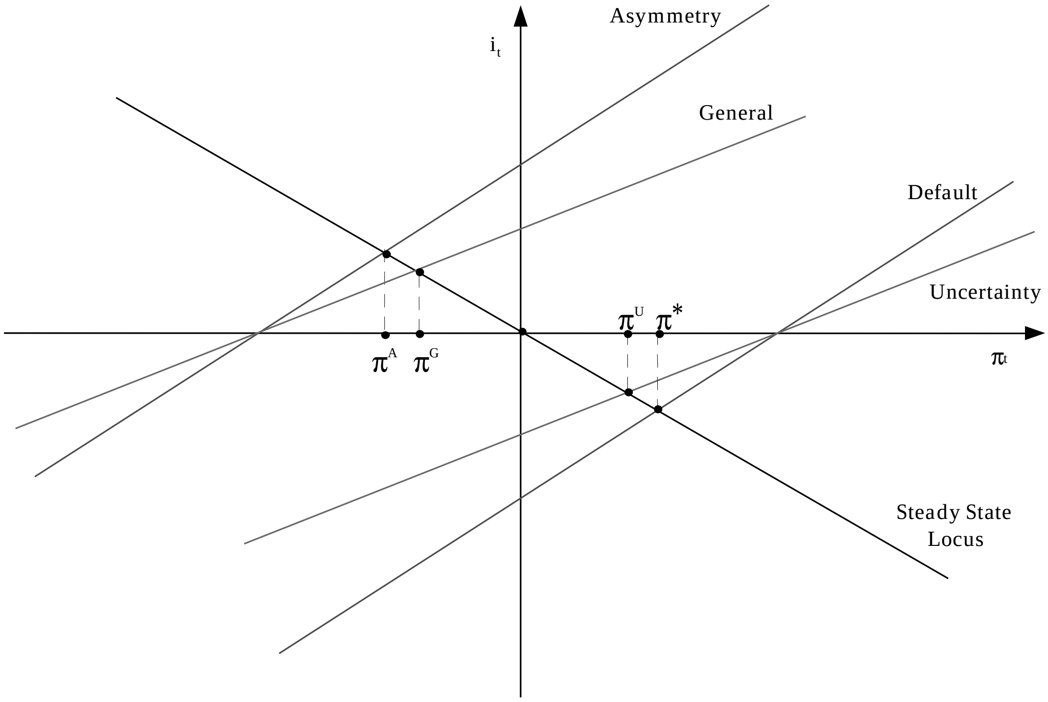

FIGURE 7
Response Paths in The Interest Rate Space

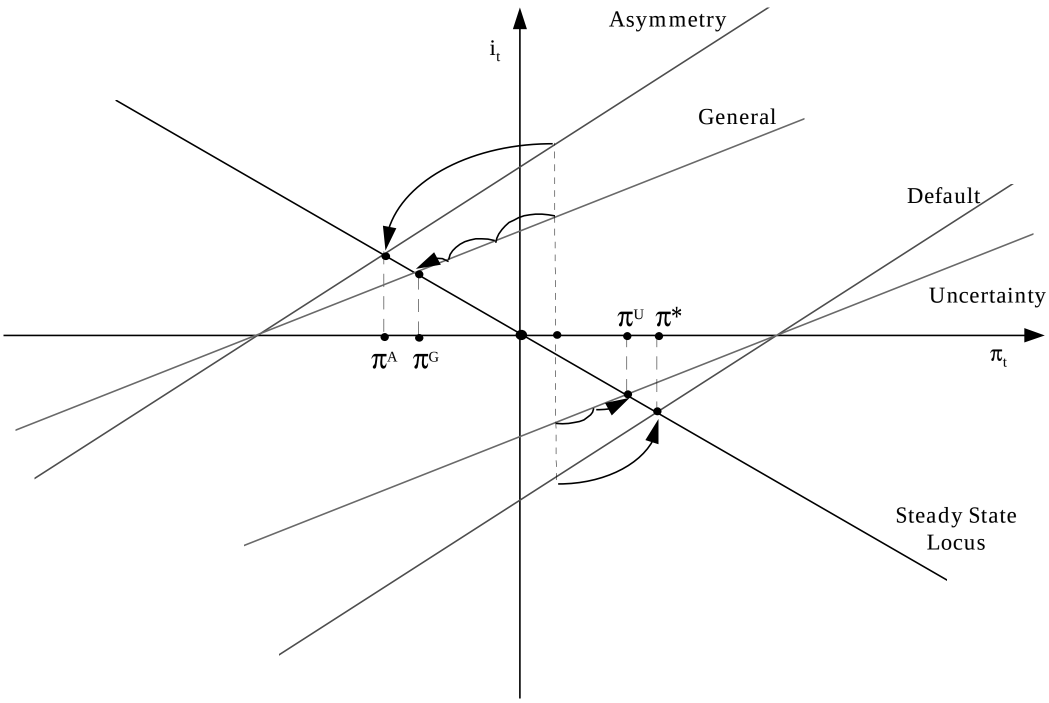

Secondly, if we introduce risk aversion alone to the default case, this introduces a bias to the explicit inflation target because losses due to positive deviations are simply magnified with asymmetric preferences and the rational policymaker is able to mitigate them by aiming for a point to the left of the minimum of the LINEX function.

Finally, when we combine the risk aversion and multiplicative uncertainty in the general case, for the parameters chosen, the interaction between uncertainty and risk aversion acts to increase the inflation target. This is because of two separate effects: (i) the dichotomy between the explicit inflation target and the economy's long-run inflation mean mitigating the policy action we noted in the A case and (ii) that the policy maker will, in any case, tend towards choosing a zero interest rate draw in the presence of uncertain control over the economy. For the cases involving uncertainty the lower inflation target results in smoothing of interest rates which with cautious responses delivers lower inflation. This is so because the interest rate response is in two periods given the persistence of inflation. In the A case the lower inflation target simply implies activist first period interest rate setting.

The main impact of asymmetric preferences is not to over-turn the "Brainard conservatism principle". It is still parameter (or what we might to think of as control) uncertainty that leads to gradual responses. With either symmetric or asymmetric preferences, risk aversion in itself does not deliver smoothing when shocks are additive. We can find biases in interest rate setting for additive, as well as multiplicative uncertainty, in the case of asymmetric preferences. But these seem to occur in a very Brainard way - the implicit quadratic is simply shifted to the left.

5. CONCLUSION

The paper examines the implications of non-quadratic loss functions for policy-gradualism and policy-caution within the context of monetary policymaking. We deviate from the quadratic framework in two respects: first, while retaining the assumption of symmetry, we allow the curvature of the loss function to change; secondly, we introduce an asymmetry into the loss function.

Changing the curvature of a symmetric loss function - for example, by introducing constant absolute (or relative) risk aversion - is shown not to matter for the optimal rule as long as uncertainty is additive. As a result, certainty equivalence also applies to non-quadratic loss functions provided that these are symmetric. So if the source of the uncertainty is about the type of the shock, deviating from quadratics does not buy us anything new: the optimal rule remains the same, and only the policymaker's dead-weight losses are different.

As with the quadratic case under additive uncertainty, welfare losses will be minimised at an inflation rate set equal to target. And so it continues to make sense to hit this target as soon and as closely as possible: there will thus be no case for gradualism or caution. In order to generate gradualism and caution, non-quadratic preferences, *per se,* are not sufficient, as one needs to introduce multiplicative uncertainty. Moreover, non-quadratic preferences are not necessary as one can easily examine these issues in a quadratic framework. This brings us to the conclusion that the analytically convenient assumption of quadratic losses may not be that unreasonable after all.

When we introduce an asymmetry in the loss function (with a LINEX function) we find that the optimal interest rate rule is biased in a state-independent way, if uncertainty is merely additive. Asymmetric preferences then result in an interest rate path that is equivalent to that implied by a shifted quadratic loss function. If upward risks are considered more (less) costly than downward risks, then the minimum of the quadratic loss function is smaller (larger) than that of the non-quadratic. In our framework, this means that the implied inflation target (which internalises the asymmetry) is smaller (larger) than the stated target.

With multiplicative uncertainty, the asymmetry does not yield qualitatively different conclusions from changing the curvature of a symmetric loss function: gradualism and caution only obtain when uncertainty is multiplicative. Moreover, simulations of the optimal rule under asymmetry and multiplicative uncertainty show that the interest rate paths are very similar to those implied by a shifted quadratic. As for caution, we have also established that multiplicative uncertainty is not sufficient. An additional requirement is that long-run policy interventions are necessary. This latter feature is illustrated in our rather simple model by letting the inflation target and the long-run mean of the inflation process differ.[27]

With reference to the delegation of monetary policy, the use of asymmetric loss functions leads to a number of important insights. First of all, if the government requires the central bank to be goal dependent, then the central bank should also be required to pursue the delegated goal in a symmetric way. This result is consistent with the inflation remits of many central banks operating in inflation targeting regimes, including the Bank of England. Secondly, if there is an asymmetry in the loss function of the government for some social welfare or political economy reason, this need not require the loss function of the central bank to be asymmetric as well. The asymmetry in the government's loss function would simply shift the increase or decrease the level of the mandated target (depending on the nature of the asymmetry) without necessarily altering the symmetry of the central bank's objectives.

Perhaps Alan Blinder (1998) had himself come to a conclusion similar to the one suggested by this paper because in the year following his plea quoted in the Introduction he wrote "Sceptics often object to certainty equivalence on the grounds that...there is no particular reason to think that the objective function is quadratic...[but] policymakers almost always will be contemplating changes in policy instruments that can be expected to lead to small changes in macroeconomic variables, For such changes...any convex objective function is approximately quadratic".

---

[27] Some permanent bias to policy intervention might, however, result from a variety of factors, for example optimal taxation, which have not been modelled in this set-up. It may therefore be possible to derive such a bias without maintaining that the long-run mean of inflation and its target are unequal.

As to future research, there may be considerable interest in exploring the implications of the results when the long-run mean and the inflation target are allowed to coincide gradually under some process of learning. To do so, the next step is to incorporate our results in a more realistic setting with agents whose expectations about inflation influence actual inflation outcomes. In addition, we might suggest at least three other possible uses of asymmetric loss functions: in the field of examining non-quadratic adjustment costs, for example, in models of investment; applications in explaining the excess returns in financial markets (i.e. that prices of assets may be biased); and finally, with respect to the maintenance of fixed exchange rate zones, where there is large asymmetry in the policymaker's preferences at either limit of the exchange rate band.

REFERENCES

BLINDER, A S (1997), "What central bankers can learn from academia - and vice versa", *Journal of Economic Perspectives*, **11(2)**, pp 3-19.

BLINDER, A S (1998), *Central Banking in Theory and Practice: The 1996 Lionel Robbins Lectures*, Cambridge: MIT Press.

BRAINARD, W C (1967), "Uncertainty and the Effectiveness of Policy", *American Economic Review*, **57**, pp 411-425.

CHRISTOFFERSEN, P F AND F X DIEBOLD (1997), "Optimal Prediction under Asymmetric Loss", *Econometric Theory*, **13(6)**, pp 808-17.

DEATON, A (1992), *Understanding Consumption*, Oxford: Oxford University Press.

FRIEDMAN, M (1951), "The Effects of a Full-Employment Policy on Economic Stability: A Formal Analysis", *Economie Appliquée*, **4**, pp 441-56.

HOROWITZ, A R (1987), "Loss functions and public policy", *Journal of Macroeconomics*, **9(4)**, pp 489-504.

SHILLER, R J (1998), "Human behaviour and the financial system", NBER Working Paper no 6375.

SVENSSON, L (1997), "Inflation targeting: implementing and monitoring inflation targets", *European Economic Review*, **41**, pp 1111-1146.

THEIL, H (1966), *Economic Forecasts and Policy*, Amsterdam: North Holland Press.

TINBERGEN, J (1954), *Centralisation and Decentralisation in Economic Policy*, Amsterdam: North Holland Publishing.

TOBIN, J (1990), "On the theory of macroeconomic policy", *De Economist*, **138(1)**, pp 1-14.

VARIAN, H R (1975), "A Bayesian Approach to Real Estate Assessment", in *Studies in Bayesian Econometrics and Statistics in Honour of Leonard J. Savage*, eds. S E Fienberg and Arnold Zellner, Amsterdam: North-Holland, pp 195-208.

ZELLNER, A (1986), "Bayesian estimation and prediction using the asymmetric loss functions", *Journal of the American Statistical Association*, **81**, pp 446-451.

APPENDIX A
SOME POINTS ON THE LINEX

The Arrow-Pratt coefficient of absolute risk aversion for this LINEX loss function is:[28]

(A.1) $$-\frac{\gamma \exp(\gamma\xi)}{\exp(\gamma\xi)-1}.$$

This coefficient has the property that there is risk neutrality at $\xi = 0$ and that $r_A' > 0$ and $r_A'' < 0$ for $\xi \neq 0$. The expectation of the LINEX function is given by the following expression:

$$E_\pi L(\xi) = \exp(\gamma\pi^{imp}) E_\pi \exp(\gamma\pi) - \gamma(\pi^{imp} - E_\pi \pi) - 1,$$

which is minimised by differentiating **(A.1)** and solving the first order conditions for π^{imp}, the implicit inflation target. This gives:

$$\pi^{imp} = -\frac{1}{\gamma} \ln\left(E_\pi \exp(-\gamma\pi)\right),$$

[28] We take the positive value of the second over the first derivative for $\zeta < 0$ and the negative value for $\zeta > 0$.

which can be evaluated analytically when π has a normal probability density function with mean μ and variance σ_e^2. In this case the moment generating function gives:

$$E_{\pi} \exp(-\gamma\pi) = \exp\left(-\gamma\mu + \tfrac{1}{2}\gamma^2\sigma_e^2\right)$$

which in turn gives:

(A.2) $\pi^{imp} = \mu - \gamma\sigma_e^2 / 2$.

Equation **(A.2)** tells us that the expectation of the loss function tends to move away from the quadratic as γ and v move away from 0.

APPENDIX B

DERIVATION OF THE NON-LINEAR EXPECTATION

$$\begin{aligned}
&E_t\{b_{t+1}\exp[-\gamma b_{t+1}i_t]\} \\
&= \int_{-\infty}^{+\infty} b_{t+1}\exp[-\gamma b_{t+1}i_t]\ \frac{1}{\sqrt{2\Pi}\sigma_b}\exp\left[-\frac{(b_{t+1}-\bar{b})^2}{2\sigma_b^2}\right]db_{t+1} \\
&= \int_{-\infty}^{+\infty} b_{t+1}\frac{1}{\sqrt{2\Pi}\sigma_b}\exp\left[\frac{-b_{t+1}+2(\bar{b}-\sigma_b^2\gamma i_t)b_{t+1}-\bar{b}^2}{2\sigma_b^2}\right]db_{t+1} \\
&= \int_{-\infty}^{+\infty} b_{t+1}\frac{1}{\sqrt{2\Pi}\sigma_b}\exp\left[-\frac{(b_{t+1}-(\bar{b}-\sigma_b^2\gamma i_t))^2}{2\sigma_b^2}\right]db_{t+1}\times\exp\left[-\bar{b}i_t+\frac{\gamma^2\sigma_b^2}{2}i_t^2\right] \\
&= (\bar{b}-\sigma_b^2\gamma i_t)\exp\left[-\bar{b}\gamma i_t+\frac{\gamma^2\sigma_b^2}{2}i_t^2\right].
\end{aligned}$$

DERIVATION OF THE OPTIMAL RULE IN THE GENERAL CASE

Consider the following optimisation problem

$$\underset{i_t}{Min}\ E_t\left\{\begin{array}{l}\exp[\gamma(\bar{\pi}+a(\pi_t-\bar{\pi})-b_{t+1}i_t+\varepsilon_{t+1}-\pi^*)]- \\ \gamma(\bar{\pi}+a(\pi_t-\bar{\pi})-b_{t+1}i_t+\varepsilon_{t+1}-\pi^*)-1\end{array}\right\}$$

subject to

$$\begin{pmatrix}b_{t+1}\\ \varepsilon_{t+1}\end{pmatrix}\sim N\left[\begin{pmatrix}\bar{b}\\ 0\end{pmatrix},\begin{pmatrix}\sigma_b^2 & 0\\ 0 & \sigma_e^2\end{pmatrix}\right].$$

The first-order condition to this problem is

$$E_t\{\gamma b_{t+1}\exp[\gamma(\bar{\pi}+a(\pi_t-\bar{\pi})-b_{t+1}i_t+\varepsilon_{t+1}-\pi^*)]\}=\gamma\bar{b},$$

And, since b_{t+1} and ε_{t+1} are independent,

$$\gamma \exp\left[\gamma(\bar{\pi} + a(\pi_t - \bar{\pi}))\right] E_t\{b_{t+1}\exp[-\gamma b_{t+1} i_t]\} E_t\{\exp[\gamma\varepsilon_{t+1}]\} = \gamma\bar{b} .$$

Using the earlier result for the first expectation and a similar line of reasoning for the second one, we have

$$\gamma \exp[\gamma(\bar{\pi} + a(\pi_t - \bar{\pi}))] \left(\bar{b} - \sigma_b^2 \gamma i_t\right) \exp\left[-\bar{b}\gamma i_t + \frac{\gamma^2\sigma_b^2}{2} i_t^2\right] \exp\left[\frac{\gamma^2\sigma_e^2}{2}\right] = \gamma\bar{b}$$

which after taking logarithms yields

$$i_t = \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) + \frac{\gamma\sigma_e^2}{2\bar{b}} + \frac{\gamma\sigma_b^2}{2\bar{b}} i_t^2 + \frac{1}{\gamma\bar{b}} \ln\left(1 - \frac{\gamma\sigma_b^2}{\bar{b}} i_t\right).$$

APPENDIX C

<u>DERIVATION OF THE OPTIMAL RULE IN THE PURE UNCERTAINTY CASE</u>

$$\begin{aligned}
i_t &= \lim_{\gamma\to 0}\left[\frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) + \frac{\gamma\sigma_e^2}{2\bar{b}} + \frac{\gamma\sigma_b^2}{2\bar{b}} i_t^2 + \frac{1}{\gamma\bar{b}}\ln\left(1 - \frac{\gamma\sigma_b^2}{\bar{b}} i_t\right)\right] \\
&= \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) + 0 + 0 + \lim_{\gamma\to 0}\left[\frac{1}{\gamma\bar{b}}\ln\left(1 - \frac{\gamma\sigma_b^2}{\bar{b}} i_t\right)\right] \\
&= \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) - \lim_{\gamma\to 0}\left[\frac{\sigma_b^2 i_t}{\left(1 - \frac{\gamma\sigma_b^2}{\bar{b}} i_t\right)\bar{b}^2}\right] \\
&= \frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*) - \frac{\sigma_b^2}{\bar{b}^2} i_t .
\end{aligned}$$

Consequently, the optimal rule in the pure uncertainty case is

$$i_t = \frac{1}{1 + \frac{\sigma_b^2}{\bar{b}^2}}\left(\frac{a}{\bar{b}}(\pi_t - \bar{\pi}) + \frac{1}{\bar{b}}(\bar{\pi} - \pi^*)\right).$$