‘You Might as Well be Hung for a Sheep as a Lamb’: The Loss Function of an Agent.

Margaret Bray and Charles Goodhart
Financial Markets Group,
London School of Economics,
Houghton Street,
London WC2A 2AE
M.M.Bray@lse.ac.uk

July 18, 2002

Abstract

Most of those who take macro and monetary policy decisions are agents. The worst penalty which can be applied to these agents is to sack them if they are perceived to have failed. To be publicly sacked as a failure is painful, often severely so, but the pain is finite. Agents thus have loss functions which are bounded above, in contrast to the unbounded quadratic loss functions which are usually assumed for policy analysis. We find a convenient mathematical form for such a loss function, which we call a bell loss function. We contrast the different behaviour of agents with quadratic and bell loss functions in three settings. Firstly we consider an agent seeking to reach multiple targets subject to linear constraints. Secondly we analyse a simple dynamic model of inflation with additive uncertainty. In both these settings certainty equivalence holds for the quadratic, but not the bell loss function. Thirdly we consider a very simple model with one target and multiplicative (Brainard) uncertainty. Here certainty equivalence breaks down for both loss functions. Policy is more conservative than in the absence of multiplicative uncertainty, but less so with the bell than the quadratic loss function.

1 Introduction

Most of those who take macro and monetary policy decisions are agents, not themselves principals. The government is an agent of the electorate; the Central Bank is an agent of the government, and through them of the public more widely. By the same token many, perhaps most, financial decisions are similarly taken by agents. Bank and fund managers are agents of those that have committed funds to them.

The thesis of this paper is that insufficient attention has been given to this fact in analysing the likely behaviour of such decision-making agents, in particular to the implications for the shape of such an agent's loss function.

What is the worst penalty, or sanction, that principals can normally apply to their agent? The standard answer is to sack them, if they are perceived to have failed. It is feasible to think of applying more severe penalties, as the scale of failure rises, but this leads to greater difficulties in attracting high-quality people to act as agents. Be that as it may, we shall assume that in the present state of affairs the main sanction for failure is dismissal. To be (publicly) sacked as a failure is painful, often severely so, especially for agents with previously established reputations, but the pain is finite. As the likelihood of being sacked approaches unity, with the outcomes deviating increasingly from the objective agreed with the principal, so the loss function will become asymptotically equal to this finite loss.

This contrasts sharply with the implications of the standard quadratic loss function, where the loss increases towards infinity as the outcome differs from that desired. This has some natural justification in certain physical cases (e.g. heat, fluid intake) where deviation from the optimum (in either direction) at some point leads to death. Being removed from office is only rarely perceived as being on the same plane!

Indeed the main justification usually given for employing a quadratic loss function, apart from the fact that everyone else does so, is that it is mathematical tractable, and also that, within limits, it may be a reasonably robust model of reality, (Chadha and Schellekens 1999; for some recent variants, see Schellekens, 2002 and al-Nowaihi and Stracca, 2001; for a more generalised critique of quadratic loss functions on behavioural grounds, see Kahneman and Tversky, 2000). Our purpose here is to suggest an alternative, and more realistic, loss function for an agent, which has a reasonably simple mathematical formulation, and to examine how agents' behaviour, with such a loss function, will differ from that of someone (e.g. a principal) with a quadratic loss function.

Principals will (should) normally be able to specify relatively clearly to their agents what their objectives may be, though even here a multiplicity of objectives and horizons may complicate matters. But in a world of uncertainty the ‘best’ results in any time period may occur because the agent is luckier, less risk-averse, cuts legal corners, or for a variety of other reasons not directly connected with either ‘effort’ or ability. So how do principals decide when to abandon (sack) their initial choice of agent, and move their custom (e.g. money or vote) to another, especially given that frictions (e.g. information linkages; ‘the devil you know is better than the devil that you do not know’) cause any such moves to be expensive to the principal?

The standard answer is to apply some form of ‘bench-marking’. (See for example Basak, Shapiro and Teplá, 2002; Basak, Pavlova and Shapiro, 2002; Jorion, 2000; Teplá, 2001; Chan, Karceski and Lakonishok, 1999; Chevalier and Ellison, 1997; Fung and Hsieh, 1997, Grossman and Zhou, 1996). That is the principal compares the results obtained by the agent either to some absolute, or to some relative, measure of achievement. So long as the agent remains on the right side of the benchmark, she is regarded as ‘successful’, and would as a generality expect to be continued in position as agent. Indeed, an agent who was summarily sacked without proper cause while still being on the right side of the agreed benchmark would often be able to sue for unfair dismissal. In contrast the agent that failed to meet the pre-arranged benchmark by a large margin might not only be sacked, but even face a legal suit for negligence; the case in the UK in 2001 of Unilever against Mercury Asset Management was in point. In the monetary field, the establishment of publicly-announced ranges for the maintenance of inflation is another example; the requirement for the UK’s Monetary Policy Committee to write a letter to the Chancellor when inflation diverges by more than 1% from its current objective is again a case.

Such benchmarks are inherently somewhat arbitrary. Why, for example, was the trigger for the Monetary Policy Committee to write a letter set at 1%, rather than say $1\frac{1}{2}$%? Historical experience of (absolute and relative) deviations from the (optimal) target is likely to play a large role, and benchmarks may well be adjusted in the light of such developing experience. But even when such benchmarks have been set, it is usually well understood that they may be broken for reasons that are no ‘fault’ of the agent. In the case

ment, the manager may have taken a rational, strategic view that the rise in price of some asset class, e.g. Japanese equities, TMT shares, etc., etc., was overdone, and hence be short of that class of assets that was driving the index up. For whatever reason, an initial, and/or minor, infringement of a benchmark is usually taken as a trigger for a formal explanation, and discussion, rather than leading to an immediate dismissal. This is certainly the case with the letter-writing requirement fo-2.5(4.7(h31(tthe)-294.2(M)73.8(o)0.7nu)6.2(e)-0.7tha)5
k93atkngbcom98.9(e73.8st)279494(l)-(i(k9.3(lly892.8.,)-4209(Sa)5.7of)0.2,y)-210.6(e)en tho bnaam98.9ha
nlv,theyecisel(y)-69017de

At some point the pay-off to the principal is so far from the optimum that sacking is certain. At that point the loss to the agent (asymptotically) reaches a maximum, which will usually be finite. Since the agent is concerned with keeping her (his) job inter-temporally, the starting point for the next play will usually be the outcome of the current game. So, even though being within the range of 'success' between the benchmark triggers may guarantee with certainty the agents' job on the next play, the closer that the outcome comes to one of the symmetric benchmarks, the greater the likelihood of (stochastically) triggering the benchmark on the very next play. For example, if the current rate of (RPIX) inflation in the UK is currently 1.6%, the likelihood of falling below the trigger of 1.5% next month is much greater than if the rate of inflation was $2\frac{1}{2}\%$ in the middle of the band.

If we assume a reasonably risk averse agent, then the loss function between the benchmark triggers, within the region of 'success', is convex. But, since the maximum loss to the agent is, we have argued finite, i.e. loss of position, then the loss function must eventually flatten out. As the outcomes move from the region of 'success' to the region of 'failure', so the loss function must, logically, pass through a point of inflexion, and the curvature of the function change from convex to concave.

The mathematical functional form that most closely and simply meets these desiderata has the functional form $1 - \exp\left(-k(x-a)^2\right)$. Figure 1 illustrates the curve for $a = 0$, and $k = \frac{1}{2}$. Mathematically this is closely related to the bell shaped normal density function $\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}\left(x-\mu\right)^2\right)$, so we call it a bell loss function. Figure 1 plots the quadratic loss function $\left(1 - \exp\left(-\frac{1}{2}\right)\right) x^2$ and the bell loss function $1 - \exp\left(-\frac{1}{2}x^2\right)$, which has points of inflection at 1 and -1. The two functions coincide at $x = -1, 0, 1$. In the

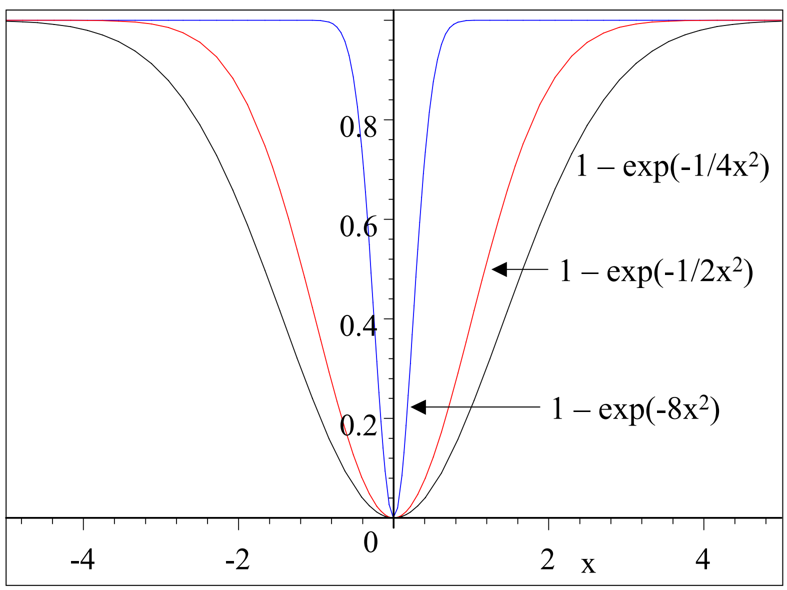

Figure 2: The bell loss function $1 - \exp(-kx^2)$ for $k = 1/4, 1/2$ and 8.

2 Attitudes to risk

The key mathematical difference between the quadratic loss function $\mathcal{L}^q(x) = (x-a)^2$ and the bell loss function $\mathcal{L}^b(x) = 1 - \exp\left(-k(x-a)^2\right)$ is that the quadratic loss function is convex for all values of x, whereas the bell loss function is convex for values of x close to a and concave for values of x far from a (see Figures 1 and 2). The derivative of the quadratic loss function $\frac{d\mathcal{L}^q(x)}{dx} = 2(x-a)$ increases in size as x moves away from its target a. Thus with a quadratic loss function the further x is from its target value a, the larger the gain from moving x towards a. Compare the bell loss function with first derivative

$$\frac{d\mathcal{L}^b}{dx} = 2k(x-a)\exp\left(-k(x-a)^2\right)$$

and second derivative

$$\frac{d^2\mathcal{L}^b}{dx^2} = \left(2k - 4k^2(x-a)^2\right)\exp\left(-k(x-a)^2\right).$$

The second derivative is positive, so the loss function is convex when the distance between x and a is less than $\frac{1}{\sqrt{2k}}$; within this region the further x is from a the greater the gain from moving x towards a. However the second derivative is negative and the bell loss function is concave, when the distance between x and a is greater than $\frac{1}{\sqrt{2k}}$; in this case the gain from moving x towards a becomes smaller as the distance between x and a increases.

This can have major implications for policy. Suppose that there is a very slight possibility of a major disruption of oil supply, which would push both inflation and output way outside their target ranges, i.e. x would deviate far from a. A government with a quadratic loss function would respond to that slight possibility, perhaps by using taxes to increase the domestic price of oil in order to maintain oil stocks and encourage the search for substitutes. With a quadratic loss function the prospective gain in the unlikely event of a future shock is worth the cost in terms of current output and inflation. But the resulting self-administered supply side shock might well cause the government to be regarded as a failure, and voted out of office. With a bell loss function the prospective gain in the unlikely event of an oil price shock would be too small to outweigh the current losses from policy measures that anticipate the shock.

The concavity of the bell loss function once beyond some distance from the target implies risk-accepting behaviour in certain circumstances. An agent with a bell loss function will be willing to take a gamble giving some probability of hitting the target, and some of missing it by a long way, which an agent with a quadratic loss function would reject.

3 Targets and Certainty Equivalence

We look at a simple example to gain further insight into the differences between the behaviour of an agent with a bell loss function, and the behaviour of an agent with a quadratic loss function. The most important feature of this example is that the certainty equivalence which holds with a quadratic loss function breaks down with a bell loss function. In this simple model there is a manager with a finite amount of resources M trying to hit n different targets $\{a_1, a_2, ...a_n\}$. If the manager puts μ_i resources into meeting the target i the gap between target and outcome is $x_i - a_i$ where $x_i = \mu_i + e_i$, and $\{e_1, e_2, ...e_n\}$ are independent normal random variables with mean $Ee_i = 0_i$ and var $e_i = \sigma_i^2$. The manager's problem is to choose resources $\{\mu_1, \mu_2, ...\mu_n\}$, subject to a constraint

$$\sum_{i=1}^{n} \mu_i \leq M,$$

in order to minimize the expectation of either the quadratic loss function

$$\mathcal{L}^q = \sum_{i=1}^{n} k_i (x_i - a_i)^2$$

or the bell loss function

$$\mathcal{L}^b = 1 - \exp\left(-\sum_{i=1}^{n} k_i (x_i - a_i)^2\right).$$

The loss functions are symmetric; overshooting a target is as bad as undershooting. We could tell stories about why managers might be penalized in this way, e.g. for wasting resources. We are not completely convinced by these stories; even if there are penalties for overshooting we think loss functions may well not be symmetric. However assuming symmetry makes the mathematics much simpler; which is why symmetric loss functions are so widely used in the literature.

Firstly consider the case with no uncertainty. In this case it makes no difference whether the loss function is bell or quadratic, because $\mathcal{L}^b = 1 - \exp(-\mathcal{L}^q)$ which is a strictly increasing function of $\mathcal{L}^q$. If $\sum_{j=1}^n a_j \leq M$ and there is no uncertainty the manager is in the happy position of being able to meet all the targets simultaneously by setting $\mu_i = x_i = a_i$ for all i making both loss functions zero. If $\sum_{j=1}^n a_j > M$ it is impossible to meet all the targets. Minimising $\sum_{i=1}^n k_i (\mu_i - a_i)^2$ subject to $\sum_{j=1}^n \mu_j \leq M$ is easily solved using standard Lagrangian techniques. The Lagrangian is

$$L = -\sum_{i=1}^{n} k_i (\mu_i - a_i)^2 + \lambda \left(M - \sum_{i=1}^{n} \mu_i \right)$$

the first order conditions $-2k_i (\mu_i - a_i) - \lambda = 0$ imply that

$$\mu_i = a_i - \frac{1}{2} \lambda k_i^{-1}.$$

Assuming that the constraint is satisfied as an equality gives

$$\sum_{j=1}^{n} \mu_j = \sum_{j=1}^{n} a_j - \frac{1}{2} \lambda \sum_{j=1}^{n} k_j^{-1} = M,$$

so $\frac{1}{2}\lambda = \left(\sum_{j=1}^n a_j - M \right) / \left(\sum_{j=1}^n k_j^{-1} \right)$ and the solution is

$$\mu_i = a_i - \frac{k_i^{-1}}{\sum_{j=1}^n k_j^{-1}} \left(\sum_{j=1}^{n} a_j - M \right).$$

There is a shortfall on every target, the size of the shortfall $a_i - \mu_i$ is proportional to the gap $\sum_{j=1}^n a_j - M$ between the resources $\sum_{j=1}^n a_j$ needed to meet all the targets and the resources available M. The constant of proportionality for target i $\frac{k_i^{-1}}{\sum_{j=1}^n k_j^{-1}}$ depends inversely upon the weight k_i given to the target in the objective. The shortfall is largest for the targets with the lowest weight k_i.

We now consider the case with uncertainty. Our assumption that $x_i = \mu_i + e_i$ $Ee_i = 0_i$ and $\operatorname{var} e_i = \sigma_i^2$ implies that $E(x_i - a_i)^2 = (\mu_i - a_i)^2 + \sigma_i^2$.

With the quadratic loss function the objective becomes

$$E\mathcal{L}^q = E\left(\sum_{i=1}^{n} k_i (x_i - a_i)^2\right) = \sum_{i=1}^{n} k_i \left[(\mu_i - a_i)^2 + \sigma_i^2\right],$$

which is minimised by choosing $\{\mu_1, \mu_2, ...\mu_n\}$ to minimise $\sum_{i=1}^{n} k_i (\mu_i - a_i)^2$. This is mathematically the same problem as we solved for the certainty case. The solution is to set $\mu_i = a_i$ if $\sum_{j=1}^{n} a_j \leq M$, the total resources available and set

$$\mu_i = a_i - \frac{k_i^{-1}}{\sum_{j=1}^{n} k_j^{-1}} \left(\sum_{j=1}^{n} a_j - M\right) \tag{1}$$

if $\sum_{j=1}^{n} a_j > M$. This is an example of the well known phenomenon of certainty equivalence. The solutions to optimization problems with quadratic objective functions of random variables with linear constraints are the same as the solution to the same problem with the random variables replaced by their mean. However the bell loss function does not give certainty equivalence. The result which makes the bell loss function tractable is Proposition 1, which we prove in the appendix.

Proposition 1 *If x is normally distributed with mean μ and variance σ^2 and k and a are real numbers*

$$E\left[\exp\left(-k(x-a)^2\right)\right] = \frac{1}{\sqrt{1+2k\sigma^2}} \exp\left(-\frac{k(\mu - a)^2}{1+2k\sigma^2}\right).$$

Given our assumptions that $\{x_1, x_2, ...x_n\}$ are normally distributed and independent, but not identically distributed, $Ex_i = \mu_i$ and $\operatorname{var} x_i = \sigma_i^2$, this implies that

$$\begin{aligned} E\mathcal{L}^b &= 1 - E\exp\left(-\sum_{i=1}^{n} k_i (x_i - a_i)^2\right) \\ &= 1 - \Pi_{i=1}^{n}\left(\frac{1}{\sqrt{1+2k_i\sigma_i^2}} \exp\left(-\frac{k_i(\mu_i - a_i)^2}{1+2k_i\sigma_i^2}\right)\right) \\ &= 1 - \left(\Pi_{i=1}^{n}\frac{1}{\sqrt{1+2k_i\sigma_i^2}}\right) \exp\left(-\sum_{i=1}^{n}\left(k_i^{-1} + 2\sigma_i^2\right)^{-1}(\mu_i - a_i)^2\right). \end{aligned}$$

Thus the solution minimises $\sum_{i=1}^{n}\left(k_i^{-1}+2\sigma_i^2\right)^{-1}(\mu_i-a_i)^2$ subject to the resource constraint $\sum_{i=1}^{n}\mu_i \leq M$. This is mathematically the same problem as before with k_i replaced by $\left(k_i^{-1}+2\sigma_i^2\right)^{-1}$. The solution is to set $\mu_i = a_i$ $i = 1, 2..n$ if $\sum_{j=1}^{n} a_j \leq M$, and if $\sum_{j=1}^{n} a_j > M$

$$\mu_i = a_i - \frac{k_i^{-1}+2\sigma_i^2}{\sum_{j=1}^{n}\left(k_j^{-1}+2\sigma_j^2\right)}\left(\sum_{j=1}^{n} a_j - M\right). \qquad (2)$$

Contrast the policy function for the quadratic loss function (equation 1) and the bell loss function (equation 2). In both cases the shortfall $a_i - \mu_i$ between the target a_i and the expected value μ_i of x_i is proportional to the overall shortfall $\sum_{j=1}^{n} a_j - M$. With the quadratic loss function there is certainty equivalence, the coefficient $\frac{k_i^{-1}}{\sum_{j=1}^{n} k_j^{-1}}$ depends only on the weights $\{k_i\}$ in the loss function. With the bell loss function there is no certainty equivalence, the coefficient $\frac{k_i^{-1}+2\sigma_i^2}{\sum_{j=1}^{n}\left(k_j^{-1}+2\sigma_j^2\right)}$ now depends both on the weights $\{k_i\}$ and the variances $\{\sigma_i^2\}$. There is a large shortfall for targets with small weights k_i and large variances σ_i^2, so considerable uncertainty about whether the target will be met even if resources are provided.

In the quadratic case only the relative sizes of $\{k_1, k_2, ...k_n\}$ matter; without loss of generality we can assume that $\sum_{i=1}^{n} k_i = 1$. In the bell case both the relative and absolute values of $\{k_1, k_2, ...k_n\}$ matter. Let $K = \sum_{i=1}^{n} k$

4 Inflation Targeting

We now turn to a very simple dynamic model of inflation,

$$x_t = -\beta (i_t - \pi_t) + u_t \tag{3}$$

$$\pi_{t+1} = \alpha x_t + \pi_t + e_{t+1}$$

where x_t is the output gap, i_t the interest rate, and π_t the inflation rate. We assume that the disturbance terms u_t and e_{t+1} are uncorrelated normal random variables with zero mean, and no serial correlation; $\operatorname{var} u_t = \sigma_u^2$ and $\operatorname{var} e_{t+1} = \sigma_e^2$. We consider a quadratic loss function

$$E_t \mathcal{L}_t^q = E_t \sum_{\tau=0}^{\infty} \delta^\tau \left[a (\pi_{t+\tau} - \pi^*)^2 + b x_{t+\tau}^2\right]$$

and a bell loss function

$$E_t \mathcal{L}_t^b = E_t \left\{1 - \exp\left[-\sum_{\tau=0}^{\infty} \delta^\tau \left(a (\pi_{t+\tau} - \pi^*)^2 + b x_{t+\tau}^2\right)\right]\right\}.$$

We prove propositions 2 and 3 in the appendix using dynamic programming.

Proposition 2 *The optimal policy interest rate policy at date t with the quadratic loss function minimises*

$$E_t \left[a (\pi_t - \pi^*)^2 + b x_t^2 + \delta c^q (\pi_{t+1} - \pi^*)^2\right]$$

and is implemented by setting

$$i_t = \pi_t + m^q (\pi_t - \pi^*) \tag{4}$$

where

$$m^q = \frac{\alpha \delta c^q}{\beta (b + \delta c^q \alpha^2)}$$

and c^q is the unique positive root of the equation

$$c = a + \frac{b\delta c}{b + \delta c\alpha^2}. \tag{5}$$

The optimal policy rule 4 depends upon the weights a and b given to output and inflation only through the ratio a/b. In the policy rule given by equation 4 the real interest rate $i_t - \pi_t$ is proportional to the deviation of inflation from target $\pi_t - \pi^*$.

There is no certainty equivalence with the bell loss function; the variance of the disturbance term in the inflation equation σ_e^2 affects the value the solution of 7 c^b, and the policy rule 6. As with the quadratic loss function the optimal policy rule with a bell loss function makes the real interest rate $i_t - \pi_t$ proportional to the gap between actual and target inflation $\pi_t - \pi^*$. The constant of proportionality m^b depends upon the parameters α and β of the inflation and output equations, the weights a and b given to output and inflation and the variance σ_e^2. Someone observing policy without knowing the policy weights would not be able to tell whether it stemmed from a quadratic or bell loss function. Agents, for example politicians in office and fund managers, are often accused of being myopic becuase of a wish to remain in office. In this section we have demopnstrated why this follows formally from the inherent nature of their loss function.

5 Multiplicative (Brainard) Uncertainty

Up to now we have looked at situations where uncertainty is additive, policy choices affect the mean, but not the variances of random variables. However if β is a random variable in equation 3, uncertainty is multiplicative, both mean and variance being affected by the interest rate (Brainard 1967). This makes the problem much less tractable; we cannot solve for the value function in the dynamic programming problem and characterise the optimal policy as we can with solely additive uncertainty. We can however say something about a simpler problem. Consider the simplest case of multiplicative uncertainty, minimising either the expected quadratic loss $\mathcal{L}^q = (\beta x + u - \alpha)^2$ or the expected bell loss $\mathcal{L}^b = 1 - \exp\left(-k(\beta x + u - \alpha)^2\right)$ where x is a policy variable, and β and u are independent normal variables, $E\beta = \beta_o$, $\operatorname{var}\beta = \sigma_\beta^2$, $Eu = 0$, $\operatorname{var} u = \sigma_u^2$. Firstly consider the quadratic loss function.

$$E\mathcal{L}^q = x^2\left(\beta_0^2 + \sigma_\beta^2\right) - 2\alpha\beta_0 x + \alpha^2 + \sigma_u^2,$$

which is minimised by setting $x = \frac{\alpha\beta_0}{\left(\beta_0^2 + \sigma_\beta^2\right)}$. If there is no uncertainty, $\sigma_\beta^2 = 0$, and $x = \frac{\alpha}{\beta_0}$, as σ_β^2 tends to infinity x tends to 0. The policy is more conservative than in the absence of multiplicative uncertainty, that is closer to the value of x (in this case 0) which minimises uncertainty.

Now consider the bell loss function. From proposition 1 in the appendix

$$E\mathcal{L}^b = 1 - \frac{1}{\sqrt{1+2k\,\mathrm{var}\left(\beta x + u\right)}} \exp\left(-\frac{k\left(E\left(\beta x+u\right)-\alpha\right)^2}{1+2k\,\mathrm{var}\left(\beta x+u\right)}\right)$$
$$= 1 - \frac{1}{\sqrt{1+2k\left(\sigma_\beta^2 x^2 + \sigma_u^2\right)}} \exp\left(-\frac{k\left(\beta_0 x-\alpha\right)^2}{1+2k\left(\sigma_\beta^2 x^2+\sigma_u^2\right)}\right), \quad (8)$$

we have been able to prove

Proposition 4 *If* $\alpha/\beta_0 > 0$ *the optimal policy* x^* *lies in the interval* $\left(\frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}, \frac{\alpha}{\beta_0}\right)$, *is decreasing in* σ_u^2 *and increasing in* k*;* $0 < \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2} < \lim_{k\to\infty} x^* < \frac{\alpha}{\beta_0}$ *If* $\alpha/\beta_0 < 0$, x^* *lies in the interval* $\left(\frac{\alpha}{\beta_0}, \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}\right)$, *is increasing in* σ_u^2 *and decreasing in* k*;* $0 > \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2} > \lim_{k\to\infty} x^* > \frac{\alpha}{\beta_0}$*. In either case*

$$\lim_{k\to 0} x^* = \lim_{\sigma_u^2\to\infty} x^* = \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}$$

and

$$\lim_{\sigma_\beta^2\to 0} x^* = \frac{\alpha}{\beta_0}.$$

The optimal policy with multiplicative uncertainty and a quadratic loss function is $x^* = \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}$. This is a compromise between the most conservative policy $x^* = 0$, which eliminates the effects of multiplicative uncertainty, and the policy $x^* = \frac{\alpha}{\beta_0}$ which would be optimal in the absence of multiplicative uncertainty when $\sigma_\beta^2 = 0$. Proposition 4 establishes that the optimal policy with a bell loss function lies between the quadratic loss policy $x^* = \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}$ which is optimal with multiplicative uncertainty, and the policy $x^* = \frac{\alpha}{\beta_0}$ which would be optimal in the absence of of multiplicative uncertainty. As k increases from zero towards infinity the optimal policy moves from $\frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}$ towards $\frac{\alpha}{\beta_0}$; but contrary to our original intuition it is bounded away from $\frac{\alpha}{\beta_0}$. Policy with a bell loss function and multiplicative uncertainty is to some degree conservative, but less so than with a quadratic loss function.

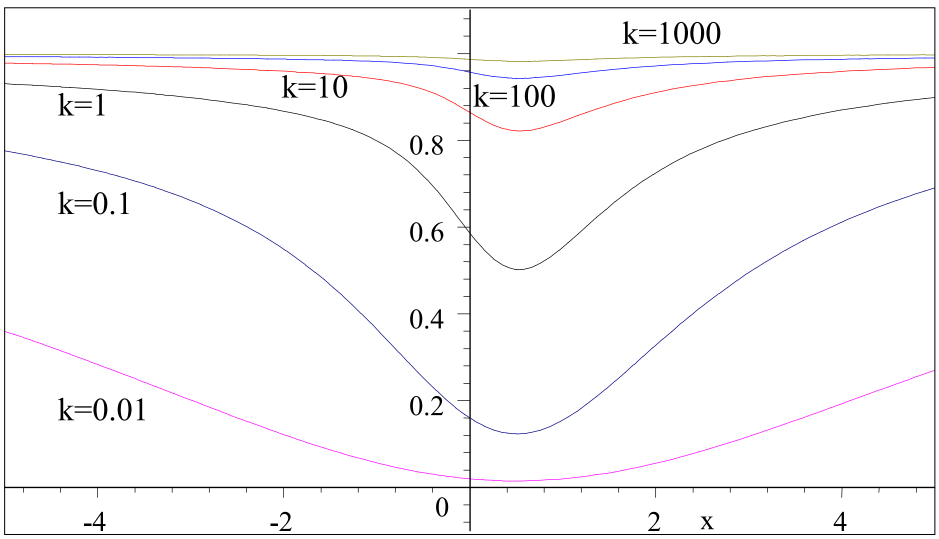

Figure 3: The expected loss function $E\mathcal{L}^b = E\left\{1 - \exp\left(-k\left(\beta x + u - \alpha\right)^2\right)\right\} = 1 - \frac{1}{\sqrt{1+2k\left(\sigma_\beta^2 x^2 + \sigma_u^2\right)}} \exp\left(-\frac{k(\beta_0 x - \alpha)^2}{1+2k\left(\sigma_\beta^2 x^2 + \sigma_u^2\right)}\right)$ with $\sigma_\beta^2 = \sigma_u^2 = \alpha = \beta_0 = 1$, plotted as a function of the policy variable x for different values of k.

Figure 3 illustrates this point, showing the value of the expected loss $1 - \frac{1}{\sqrt{1+2k\left(\sigma_\beta^2 x^2+\sigma_u^2\right)}} \exp\left(-\frac{k(\beta_0 x-\alpha)^2}{1+2k\left(\sigma_\beta^2 x^2+\sigma_u^2\right)}\right)$, for different values of x^* and k, with $\sigma_\beta^2 = \sigma_u^2 = \alpha = \beta_0 = 1$. The higher the value of k the greater the loss for all values of x, and the higher the loss minimising policy, which is, however, always less than than $\frac{\alpha}{\beta_0} = 1$.

In summary, a bell loss function can, at one extreme, when $k = 0$, mimic a quadratic loss function; not surprisingly, therefore, under these circumstances the optimal policy under multiplicative uncertainty remains the same in both cases. As k increases, the width of the convex range (of success) narrows. Since, outside that range, one might as well be hung for a sheep as a lamb, so policy in conditions of multiplicative uncertainty becomes more aggressive, less conservative, than under a quadratic loss function. We had, at one stage, thought that as k became infinitely large, i.e. that the region of success became restricted to a point, that policy would just aim to hit that one point, ignoring multiplicative uncertainty altogether. In practice, however, there is always sufficient curvature in the relationships to trade-off some variance against some chance of hitting the mean, i.e. the optimal policy is bounded away from $\alpha/\beta_{0.}$

6 Conclusion

We set out to compare policy making behaviour with a bell loss function and a conventional quadratic loss function. The most important difference is that certainty equivalence no longer holds with a bell loss function, even with additive uncertainty. In the two linear examples we studied behaviour with the policy rules have the same linear functional form with both loss functions, but the weights differ with the bell loss function and depend on variances. Brainard uncertainty is much less tractable with a bell loss function; in even the simplest case the optimal policy is characterised by a cubic equation. However we were able to show that in this case a bell loss function and Brainard uncertainty makes for conservative policy, although not as much so as with a quadratic loss function.

References

[1] al-Nowaihi, A. and L. Stracca, (2001), 'Non-standard central bank loss functions, skewed risks, and the certainty equivalence principle', European Central Bank Research Paper, (July).

[2] Basak, S., Pavlova, A. and A. Shapiro, (2002), 'Optimality of Observed Risk Management Restrictions', work in progress, LBS, MIT and NYU.

[3] Basak, S., Shapiro, A. and L. Teplá, (2002), 'Risk Management with Benchmarking', mimeo, February.

[4] Brainard, W.C.(1967) 'Uncertainty and the Effectiveness of Policy' American Economic Review Papers and Proceedings 57, p411-425.

[12] Schellekens, P., (2002), 'Caution and Conservatism in the Making of Monetary Policy', Journal of Money, Credit and Banking, Vol. 34, No. 1, (February).

[13] Teplá, L., (2001), 'Optimal Investment with Minimum Performance Constraints', Journal of Economic Dynamics and Control, 25, pp 1629-1645.

A Appendix.

A.1 Proof of Proposition 1

We are evaluating

$$E\exp\left(-k\left(x-a\right)^2\right)=\frac{1}{\sqrt{2\pi\sigma^2}}\int_{-\infty}^{+\infty}\exp\left(-k\left(x-a\right)^2-\frac{\left(x-\mu\right)^2}{2\sigma^2}\right)dx.$$

Expanding and then completing the square implies that

$$\begin{aligned}
& k\left(x-a\right)^2+\frac{\left(x-\mu\right)^2}{2\sigma^2} \\
= & \frac{1}{2\sigma^2}\left[\left(1+2k\sigma^2\right)x^2-2\left(\mu+2ka\sigma^2\right)x+\mu^2+2k\sigma^2a^2\right] \\
= & \frac{1}{2\sigma^2}\left[\left(1+2k\sigma^2\right)\left(x-\frac{\mu+2ka\sigma^2}{1+2k\sigma^2}\right)^2+\mu^2+2k\sigma^2a^2-\frac{\left(\mu+2ka\sigma^2\right)^2}{1+2k\sigma^2}\right] \\
= & \left[\frac{\left(1+2k\sigma^2\right)}{2\sigma^2}\left(x-\frac{\mu+2ka\sigma^2}{1+2k\sigma^2}\right)^2+\frac{k\left(\mu-a\right)^2}{1+2k\sigma^2}\right]
\end{aligned}$$

so

$$E\exp\left(-k\left(x-a\right)^2\right)=\frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{k\left(\mu-a\right)^2}{1+2k\sigma^2}\right)\int_{-\infty}^{+\infty}\exp\left(-\frac{\left(x-\widehat{\mu}\right)^2}{2\widehat{\sigma}^2}\right)dx$$

where $\widehat{\mu}=\frac{\mu+2ka\sigma^2}{1+2k\sigma^2}$ and $\widehat{\sigma}^2=\frac{\sigma^2}{1+2k\sigma^2}$. But as $\frac{1}{\sqrt{2\pi\widehat{\sigma}^2}}\exp\left(-\frac{(x-\widehat{\mu})^2}{2\widehat{\sigma}^2}\right)$ is a normal density function

$$\frac{1}{\sqrt{2\pi\widehat{\sigma}^2}}\int_{-\infty}^{+\infty}\exp\left(-\frac{\left(x-\widehat{\mu}\right)^2}{2\widehat{\sigma}^2}\right)dx=1.$$

Hence

$$
\begin{aligned}
E\exp\left(-k\left(x-a\right)^{2}\right) &= \frac{\sqrt{2\pi\widehat{\sigma}^{2}}}{\sqrt{2\pi\sigma^{2}}}\exp\left(-\frac{k\left(\mu-a\right)^{2}}{1+2k\sigma^{2}}\right) \\
&= \frac{1}{\sqrt{1+2k\sigma^{2}}}\exp\left(-\frac{k\left(\mu-a\right)^{2}}{1+2k\sigma^{2}}\right).
\end{aligned}
$$

A.2 Proof of Proposition 2

We are seeking to solve the dynamic programming problem of minimising the expectation of

$$
\mathcal{L}_{t}^{q}=E_{t}\sum_{\tau=0}^{\infty}\delta^{\tau}\left[a\left(\pi_{t+\tau}-\pi^{*}\right)^{2}+bx_{t+\tau}^{2}\right]
$$

by choosing the interest rate $i_{t+\tau}$ at dates $t+\tau$ when

$$
x_{t}=-\beta\left(i_{t}-\pi_{t}\right)+u_{t}
$$

and

$$
\pi_{t+1}=\alpha x_{t}+\pi_{t}+e_{t+1}.
$$

We conjecture that the value function is of the form $v\left(\pi_{t}\right)=c\left(\pi_{t}-\pi^{*}\right)^{2}+d$ in which case the Bellman equation is satisfied if there are numbers c and d with $c\geq 0$ such that

$$
\begin{aligned}
& c\left(\pi_{t}-\pi^{*}\right)^{2}+d \\
= & \min_{i_{t}} E\left[a\left(\pi_{t}-\pi^{*}\right)^{2}+bx_{t}^{2}+\delta\left(c\left(\pi_{t+1}-\pi^{*}\right)^{2}+d\right)\right] \\
= & \min_{i_{t}} E\left[a\left(\pi_{t}-\pi^{*}\right)^{2}+bx_{t}^{2}+\delta\left(c\left(\alpha x_{t}+\pi_{t}-\pi^{*}+e_{t+1}\right)^{2}+d\right)\right] \\
= & \min_{i_{t}} E\left[a\left(\pi_{t}-\pi^{*}\right)^{2}+\left(b+\delta c\alpha^{2}\right)x_{t}^{2}+2\alpha\delta c x_{t}\left(\pi_{t}-\pi^{*}\right)+\delta c\left(\pi_{t}-\pi^{*}\right)^{2}+\delta c\sigma_{e}^{2}+\delta d\right] \\
= & \min_{i_{t}}\left[\left(b+\delta c\alpha^{2}\right)\left(\beta^{2}\left(i_{t}-\pi_{t}\right)^{2}+\sigma_{u}^{2}\right)-2\alpha\delta c\beta\left(i_{t}-\pi_{t}\right)\left(\pi_{t}-\pi^{*}\right)\right] \\
& +a\left(\pi_{t}-\pi^{*}\right)^{2}+\delta c\left(\pi_{t}-\pi^{*}\right)^{2}+\delta c\sigma_{e}^{2}+\delta d.
\end{aligned}
$$

Differentiating with respect to i_t to get the minimum implies that the optimal policy satisfies

$$
i_{t}=\pi_{t}+m^{q}\left(\pi_{t}-\pi^{*}\right) \tag{A1}
$$

where

$$m^q = \frac{\alpha\delta c}{\beta\left(b + \delta c\alpha^2\right)}$$

and

$$\begin{aligned} & c\left(\pi_t - \pi^*\right)^2 + d \\ = \quad & -\frac{\alpha^2\delta^2 c^2}{\left(b + \delta c\alpha^2\right)}\left(\pi_t - \pi^*\right)^2 + a\left(\pi_t - \pi^*\right)^2 + \delta c\left(\pi_t - \pi^*\right)^2 \\ & + \left(b + \delta c\alpha^2\right)\sigma_u^2 + \delta c\sigma_e^2 + \delta d. \end{aligned}$$

Our conjecture holds and the Bellman equation is satisfied if left hand and right hand sides of this expression are the same function of π_t, that is provided the coefficients of $\left(\pi_t - \pi^*\right)^2$ are the same so

$$c = a - \frac{\alpha^2\delta^2 c^2}{\left(b + \delta c\alpha^2\right)} + \delta c$$

or equivalently

$$c = a + \frac{b\delta c}{\left(b + \delta c\alpha^2\right)} \tag{A2}$$

and

$$d = \left(b + \delta c\alpha^2\right)\sigma_u^2 + \delta c\sigma_e^2 + \delta d.$$

or equivalently

$$d = \frac{\left(b + \delta c\alpha^2\right)\sigma_u^2 + \delta c\sigma_e^2}{1 - \delta}.$$

Equation A2 can be written as

$$f^q\left(a, b, c\right) = \left(c - a\right)\left(b + \delta c\alpha^2\right) - b\delta c = 0.$$

Considering $f^q\left(a, b, c\right)$ as a function of c it is quadratic with a positive coefficient on c^2 whilst $f^q\left(a, b, 0\right) = -ab < 0$ if a and b are both positive. Then there is a unique positive root c^q and $\partial f^q/\partial c > 0$ when $c = c^q$. We have indeed got a solution to the Bellman equation. The optimal policy is given by A1 where

$$m^q = \frac{\alpha\delta c^q}{\beta\left(b + \delta c^q \alpha^2\right)} \tag{A3}$$

and c^q is the unique positive root of $(c - a)\left(b + \delta c \alpha^2\right) - b\delta c = 0$.

As $f^q(a, b, c)$ is homogeneous of degree 2 in (a, b, c) if $f^q(a, b, c) = 0$ then $f^q(\lambda a, \lambda b, \lambda c) = 0$, so the root c^q is homogeneous of degree 1 in (a, b). Then equation A3 implies that the coefficient m^q of the policy response of the real interest rate to deviations of inflation from its target value is homogeneous of degree 0 in (a, b), only the ratio of the weights a/b matters to policy. When $a = 0$ so no weight is given to inflation $c = 0$ solves $f^q(a, b, c) = 0$, $m^q = 0$, the optimal policy is to set $i_t = \pi_t$ so the real interest rate is zero which minimises the output gap. When $b = 0$ no weight is given to output, the optimal policy response has $m^q = \frac{1}{\alpha\beta}$, so $i_t - \pi_t = \frac{1}{\alpha\beta}(\pi_t - \pi^*)$ which makes expected inflation at $t + 1$ equal to the target π^*. We now show that a/b increases from zero to infinity the coefficient m^q increases, interest rate policy becomes more aggressive. As only the ratio a/b matters this can be done by showing that m^q is an increasing function of a. As $\partial c^q/\partial a = -\left(\partial f^q/\partial a\right)/\left(\partial f^q/\partial c\right)$ at $c = c^q$, and we have already argued that $\partial f^q/\partial c > 0$ when $c = c^q$, it is enough to show that $\partial f^q/\partial a = -\left(b + \delta c \alpha^2\right) < 0$.

A.3 Proof of Proposition 3

We are seeking to solve the dynamic programming problem of minimising the expectation of

$$E_t \mathcal{L}_t^b = 1 - E_t \exp\left(-\sum_{\tau=0}^{\infty} \delta^\tau \left[a\left(\pi_{t+\tau} - \pi^*\right)^2 + b x_{t+\tau}^2\right]\right)$$

by choosing the interest rate $i_{t+\tau}$ at dates $t + \tau$ when

$$x_t = -\beta\left(i_t - \pi_t\right) + u_t$$

and

$$\pi_{t+1} = \alpha x_t + \pi_t + e_{t+1}.$$

We conjecture that the value function is of the form $v(\pi_t) = 1 - \exp\left(-c(\pi_t - \pi^*)^2 - d\right)$ in which case the Bellman equation is satisfied if

$$\begin{aligned}
& \exp\left(-c(\pi_t - \pi^*)^2 - d\right) \\
= \;& \max_{i_t} E\left[\exp\left(-a(\pi_t - \pi^*)^2 - bx_t^2 - \delta\left(c(\pi_{t+1} - \pi^*)^2 + d\right)\right)\right] \\
= \;& \max_{i_t} E\left[\exp\left(-a(\pi_t - \pi^*)^2 - bx_t^2 - \delta c(\alpha x_t + \pi_t - \pi^* + e_{t+1})^2 - \delta d\right)\right]
\end{aligned}$$

which from Proposition 1 is equal to

$$\max_{i_t} \frac{1}{\sqrt{1 + 2\delta c\sigma_e^2}} E\left[\exp(-Z)\right]$$

where

$$Z = a(\pi_t - \pi^*)^2 + bx_t^2 + \widehat{\delta} c(\alpha x_t + \pi_t - \pi^*)^2 + \delta d$$

and

$$\widehat{\delta} = \frac{\delta}{1 + 2\delta c\sigma_e^2}. \tag{A4}$$

Since $x_t = -\beta(i_t - \pi_t) + u_t$ the expression Z can be written as

$$\begin{aligned}
Z \;=\;& a(\pi_t - \pi^*)^2 + \left(b + \widehat{\delta} c\alpha^2\right) x_t^2 + 2\alpha\widehat{\delta} c x_t(\pi_t - \pi^*) + \widehat{\delta} c(\pi_t - \pi^*)^2 + \delta d \\
=\;& \left(b + \widehat{\delta} c\alpha^2\right)\left(-\beta(i_t - \pi_t) + u_t + \frac{\alpha\widehat{\delta} c(\pi_t - \pi^*)}{\left(b + \widehat{\delta} c\alpha^2\right)}\right)^2 - \frac{\alpha^2\widehat{\delta}^2 c^2(\pi_t - \pi^*)^2}{\left(b + \widehat{\delta} c\alpha^2\right)} \\
& + a(\pi_t - \pi^*)^2 + \widehat{\delta} c(\pi_t - \pi^*)^2 + \delta d.
\end{aligned}$$

From Proposition 1

$$E\exp(-Z) = \frac{1}{\sqrt{1 + 2\left(b + \widehat{\delta} c\alpha^2\right)\sigma_u^2}} \exp(-Y)$$

where

$$\begin{aligned}
Y \;=\;& \frac{\left(b + \widehat{\delta} c\alpha^2\right)}{1 + 2\left(b + \widehat{\delta} c\alpha^2\right)\sigma_u^2}\left(-\beta(i_t - \pi_t) + \frac{\alpha\widehat{\delta} c(\pi_t - \pi^*)}{\left(b + \widehat{\delta} c\alpha^2\right)}\right)^2 \\
& - \frac{\alpha^2\widehat{\delta}^2 c^2(\pi_t - \pi^*)^2}{\left(b + \widehat{\delta} c\alpha^2\right)} + a(\pi_t - \pi^*)^2 + \widehat{\delta} c(\pi_t - \pi^*)^2 + \delta d.
\end{aligned}$$

The optimal interest rate policy minimises Y by setting

$$i_t = \pi_t + m^b (\pi_t - \pi^*) \tag{A5}$$

where

$$m^b = \frac{\alpha \widehat{\delta} c}{\beta \left(b + \widehat{\delta} c \alpha^2 \right)}$$

implying that

$$Y = -\frac{\alpha^2 \widehat{\delta}^2 c^2 (\pi_t - \pi^*)^2}{\left(b + \widehat{\delta} c \alpha^2 \right)} + a (\pi_t - \pi^*)^2 + \widehat{\delta} c (\pi_t - \pi^*)^2 + \delta d$$

so the Bellman equation is satisfied if

$$\begin{aligned} & \exp\left(-c (\pi_t - \pi^*)^2 - d \right) \\ = \; & \max_{i_t} \frac{1}{\sqrt{1 + 2\delta c \sigma_e^2}} E\left[\exp(-Z) \right] \\ = \; & \frac{1}{\sqrt{\left(1 + 2 \left(b + \widehat{\delta} c \alpha^2 \right) \sigma_u^2 \right) (1 + 2\delta c \sigma_e^2)}} \exp(-Y) \, . \end{aligned}$$

Equating coefficients of $(\pi_t - \pi^*)^2$ implies

$$c = a - \frac{\alpha^2 \widehat{\delta}^2 c^2}{\left(b + \widehat{\delta} c \alpha^2 \right)} + \widehat{\delta} c \tag{A6}$$

or equivalently

$$c = a + \frac{b \widehat{\delta} c}{\left(b + \widehat{\delta} c \alpha^2 \right)}$$

whilst equating the other terms

$$\exp(-d) = \frac{1}{\sqrt{\left(1 + 2 \left(b + \widehat{\delta} c \alpha^2 \right) \sigma_u^2 \right) (1 + 2\delta c \sigma_e^2)}} \exp(-\delta d) \, . \tag{A7}$$

Using A4 equation A2 can be written as

$$c = a + \frac{b\widehat{\delta c}}{\left(b + \widehat{c\delta}\alpha^2\right)} = a + \frac{b\delta c}{b\left(1 + 2\delta c\sigma_e^2\right) + c\delta\alpha^2}$$

or

$$f\left(a, b, c\right) = \left(c - a\right)\left(b\left(1 + 2\delta c\sigma_e^2\right) + \delta c\alpha^2\right) - b\delta c = 0.$$

Considered as a quadratic function of c has a positive coefficient on c^2 whilst if $c = 0$ $f = -ab < 0$ if a and b are both positive, in which case there is a unique positive root c^b. Equations A4 and A7 imply

$$d = \frac{1}{2\left(1 - \delta\right)}\left[\ln\left(1 + 2\left(b + \frac{\delta c\alpha^2}{1 + 2\delta c\sigma_e^2}\right)\sigma_e^2\right) + \ln\left(1 + 2\delta c\sigma_e^2\right)\right]$$

Our conjecture is satisfied, we have a solution to the Bellman equation. From A1 the optimal policy at t can be written as

$$i_t = \pi_t + m^b\left(\pi_t - \pi^*\right)$$

where

$$m^b = \frac{\alpha\delta c^b}{\beta\left(b\left(1 + 2\delta c^b\sigma_e^2\right) + \delta c^b\alpha^2\right)} \tag{A8}$$

and c^b is the unique positive root of $\left(c - a\right)\left(b\left(1 + 2\delta c\sigma_e^2\right) + \delta c\alpha^2\right) - b\delta c = 0$. The optimal policy minimises the expectation of

$$\exp\left[-\left(a\left(\pi_t - \pi^*\right)^2 + bx_t^2 + \delta c^b\left(\pi_{t+1} - \pi^*\right)^2\right)\right].$$

When $a = 0$ so no weight is given to inflation $c = 0$ solves $f^b\left(a, b, c\right) = 0$, $m^b = 0$, the optimal policy is to set $i_t = \pi_t$ so the real interest rate is zero which minimises the output gap. When $b = 0$ no weight is given to output, the optimal policy response has $m^b = \frac{1}{\alpha\beta}$, so $i_t - \pi_t = \frac{1}{\alpha\beta}\left(\pi_t - \pi^*\right)$ which makes expected inflation at $t+1$ equal to the target π^*. In both these polar cases the optimal policy is the same with the bell and the quadratic loss functions. For all other cases the policy differs.

The last step in the proof is establishing that less weight is given to the future with bell loss function, that is $c^b < c^q$ and $m^b < m^q$. Given the definition of c^q as the positive root of

$$f_q(c) = (c - a)\left(b + \delta c\alpha^2\right) - b\delta c = 0$$

$c^q - a > 0$ and

$$(c^q - a)\left(b + \delta c^q\alpha^2\right) - b\delta c^q = 0$$

so subtracting this expression from $f_b(c^q) = (c^q - a)(b(1 + 2\delta c^q\sigma_e^2) + \delta c^q\alpha^2) - b\delta c^q$ implies that

$$f_b(c^q) = 2\delta c^q\sigma_e^2(c^q - a) > 0.$$

As $f_b(c)$ is a quadratic, with $f_b(0) < 0$ and a positive coefficient on c^2 this implies that $c^q > c^b$. From A8

$$m^b = \frac{\alpha\delta c^b}{\beta(b(1 + 2\delta c^b\sigma_e^2) + \delta c^b\alpha^2)} < \frac{\alpha\delta c^b}{\beta(b + \delta c^b\alpha^2)} < \frac{\alpha\delta c^q}{\beta(b + \delta c^q\alpha^2)} = m^q$$

since $\frac{\alpha\delta c}{\beta(b+\delta c\alpha^2)}$ is an increasing function of c. Hence $m^b < m^q$ interest rate policy with a bell loss function is less aggressive than with the corresponding quadratic loss function.

A.4 Proof of Proposition 4

If $\alpha\beta_0$ is positive and x is negative replacing x by $-x$ does not change x^2 and increases $(\beta_0 x - \alpha)^2$ so increases the expected loss, thus the optimal x is non-negative A similar argument implies that if $\alpha\beta_0 < 0$ the optimal x is non-positive. Assume temporarily that $\alpha\beta_0$ is positive. From equation 8

$$E\mathcal{L}^b = 1 - \frac{1}{\sqrt{1 + 2k\left(\sigma_\beta^2 x^2 + \sigma_u^2\right)}}\exp\left(-\frac{k(\beta_0 x - \alpha)^2}{1 + 2k\left(\sigma_\beta^2 x^2 + \sigma_u^2\right)}\right).$$

$$\frac{\partial E\mathcal{L}^b}{\partial x} = \frac{2k}{\left(1 + 2k\left(x^2\sigma_\beta^2 + \sigma_u^2\right)\right)^{5/2}}\exp\left(-\frac{k(\beta_0 x - \alpha)^2}{1 + 2k\left(x^2\sigma_\beta^2 + \sigma_u^2\right)}\right)f(x) \quad \text{(A9)}$$

where

$$f(x) = \left(1 + 2k\sigma_u^2 + 2k\sigma_\beta^2 x^2\right)\left(x\left(\sigma_\beta^2 + \beta_0^2\right) - \alpha\beta_0\right) - 2xk\left(\beta_0 x - \alpha\right)^2 \sigma_\beta^2. \tag{A10}$$

The first two terms on the right hand side of A9 are strictly positive, so the stationary points are the roots of the cubic function $f(x)$. When $x = \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}$ $f(x)$ is negative, and when $x = \frac{\alpha}{\beta_0}$ $f(x)$ is positive, so by continuity there is at least one positive root in the interval $\left(\frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}, \frac{\alpha}{\beta_0}\right)$. Let $\widehat{x}$ be the largest positive root. Expanding $f(x)$ gives

$$\begin{aligned} f(x) &= -\alpha\beta_0 - 2k\sigma_u^2\alpha\beta_0 + x\left(\sigma_\beta^2 + \beta_0^2 + 2k\sigma_u^2\sigma_\beta^2 + 2k\sigma_u^2\beta_0^2 - 2k\sigma_\beta^2\alpha^2\right) \\ &\quad +2k\sigma_\beta^2 x^2\alpha\beta_0 + 2k\sigma_\beta^4 x^3. \end{aligned} \tag{A11}$$

implying that

$$\frac{\partial^2 f(x)}{\partial x^2} = 4k\sigma_\beta^2\alpha\beta_0 + 12k\sigma_\beta^4 x$$

which is positive for positive x. Hence $f(x)$ is a convex function for positive x. From 9 $f(0) = -\alpha\beta_0$ which is by our temporary assumption negative. For any x in the interval $(0, \widehat{x})$ convexity of f implies that

$$\begin{aligned} f(x) &= f\left(\frac{x}{\widehat{x}}\widehat{x} + \left(1 - \frac{x}{\widehat{x}}\right)0\right) \leq \frac{x}{\widehat{x}} f(\widehat{x}) + \left(1 - \frac{x}{\widehat{x}}\right) f(0) \\ &= \frac{x}{\widehat{x}}0 + \left(1 - \frac{x}{\widehat{x}}\right) f(0) < 0 \end{aligned}$$

so $f(x)$ cannot have any roots in the interval $(0, \widehat{x})$. As we assumed that $\widehat{x}$ is the largest positive roots, and there cannot be any smaller positive roots, $\widehat{x}$ must be the unique positive root of $f(x) = 0$. We have argued that given our temporary assumption that $\alpha\beta_0$ is positive, the optimal policy is positive; we have also argued that there is a root in the interval $\left(\frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}, \frac{\alpha}{\beta_0}\right)$, so this root must be the optimal policy. Use notation x^* for the optimal policy. As $f\left(\frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}\right) < 0$, and $f\left(\frac{\alpha}{\beta_0}\right) > 0$, $f_x(x^*) > 0$. Thus as

$$\frac{\partial x^*}{\partial k} = -\frac{f_k(x^*)}{f_x(x^*)}$$

$\frac{\partial x^*}{\partial k}$ has the opposite sign to $f_k(x^*)$. From A10

$$f(x) = \left(x\left(\sigma_\beta^2 + \beta_0^2\right) - \alpha\beta_0\right) + 2kg(x) \tag{A12}$$

where

$$g(x) = \left(\sigma_u^2 + \sigma_\beta^2 x^2\right)\left(x\left(\sigma_\beta^2 + \beta_0^2\right) - \alpha\beta_0\right) - x\left(\beta_0 x - \alpha\right)^2 \sigma_\beta^2. \tag{A13}$$

As $x^* > \frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}$ the term $x^*\left(\sigma_\beta^2 + \beta_0^2\right) - \alpha\beta > 0$, thus from A12, as $f(x^*) = 0$, $g(x^*) < 0$. But A12 and A13 also imply that $f_k(x^*) = 2g(x^*)$. Thus $f_k(x^*) < 0$, and so $\frac{\partial x^*}{\partial k} > 0$, the optimal policy is an increasing function of k. A similar argument implies that $\frac{\partial x^*}{\partial \sigma_\beta^2}$ is opposite in sign to $f_{\sigma^2}(x^*)$. From A12 and A13 $f_{\sigma^2}(x^*) = \left(x^*\left(\sigma_\beta^2 + \beta_0^2\right) - \alpha\beta\right) > 0$, so the optimal policy is a decreasing function of σ_β^2.

To get the limits as k tends to 0, note that the first order condition with $k = 0$ is $x\left(\sigma_\beta^2 + \beta_0^2\right) - \alpha\beta = 0$. As we have shown that the optimal policy x^* is increasing in k and bounded above by α/β_0, x^* must tend to a finite limit as k tends to infinity. At this limit the first order condition becomes $g(x) = 0$. Note from A13 that $g(\alpha/\beta_0) > 0$, so the upper limit is strictly less than α/β_0. Finally note that as σ_β^2 tends to zero, so the Brainard uncertainty disappears, the interval $\left(\frac{\alpha\beta_0}{\sigma_\beta^2+\beta_0^2}, \frac{\alpha}{\beta_0}\right)$ and thus the optimal policy collapses to the point α/β_0.

Now suppose that contrary to our temporary assumption $\alpha\beta_0$. As $\left(\beta_0 x - \alpha\right)^2 = \left(\beta_0(-x) - (-\alpha)\right)^2$ and $(-x)^2 = x^2$, we can consider $-x$ as the policy variable, and note that $-\alpha\beta_0$ is positive. Then everything we have proved about x and α applies to $-x$ and $-\alpha$, in particular the limiting arguments hold, and increasing functions become decreasing functions.