How square is the policy frontier ?

SGB Henry, M Satchi and D Vines*

First Version December 2000

This version January 2002

Centre for International Macroeconomics

Department of Economics, University of Oxford

Abstract:

This paper assesses the implications of discounting on a result derived by Bean (1998): that in a model of monetary policy where policy acts with a lag, the outcomes of monetary policy are very similar for a wide range of weightings of the (non-discounting) monetary authority's objective function, with respect to inflation stability versus output stability. We show that when the authority discounts the future, outcomes become more sensitive to preferences, and that it is important to take the discount rate into account when examining the question of how the authority's remit should be specified.

Keywords: Monetary Policy, Policy Frontier, Discounting

JEL Nos: E53, E58, E61

* We are grateful to Stephen Hall, Adrian Pagan and John Muellbauer for helpful comments and suggestions. All errors remain the responsibility of the authors. We are grateful to the Leverhume Trust for financial support under grant no. F/108 519A, "EMU and European Macroeconomic Policy in a Global Context."

1. Introduction

Here we address an important finding derived by Bean (1998). This finding, when applied to the new UK monetary arrangements, implies that "it does not matter that the inflation remit does not specify the relative weight that the Bank of England should place on output volatility *vis-à-vis* inflation volatility." The result is obtained under the assumption that the authority does not discount the future. In practice, however, some discounting is likely to occur and planning horizons might be relatively short; for the U.K., members of the Monetary Policy Committee are appointed for a term of three years, and so the average remaining term of existing members is one and a half years. Given that a single time period in Bean (1998) is interpreted as corresponding to around one year, the assumption of no discounting by the central bank should not be taken for granted. We examine how the above finding alters when discounting is introduced.

Recent research in monetary policy has incorporated models with a more realistic treatment of the lag structure with which policy takes effect. Svennson (1997), Ball (1997) and Bean (1998) are recent examples. In these models, optimal policy typically takes the form of a Taylor rule.[1] Bean (1998) demonstrates in this setting, that the weighting of the loss function of the monetary authority with respect to "output volatility *vis-à-vis* inflation volatility" does not affect the optimal policy rule in such a way as to greatly alter the outcome of stabilisation policy. This can be seen in the shape of the Bean policy frontier (the set of efficient outcomes) which is relatively 'rectangular.' Figure (1), which will be described further below, shows the policy frontier and the optimal outcomes for two different weightings of the loss function. These are 'close,' both in terms of their absolute distance relative to the magnitude of the variance of supply and demand shocks, and also in terms of the ratio of output volatility to inflation volatility obtained in each case.

We find that this result is very robust to variations in the estimates of the parameters of the model. This holds to such an extent that it is almost independent of the empirical analysis undertaken by Bean himself in the original paper – we find that the property of rectangularity holds for any set of plausible estimates. His analysis is therefore likely to apply to cases other than the U.K. It can also be shown that except in the unrealistic case where policy affects inflation contemporaneously, the rectangularity result is also robust to simple variations in the lag structure used. Our main finding is that the rectangularity result, however, is dependent on the planning horizon of the policymaker. We argue that when discounting is introduced, the weighting of the loss function becomes increasingly important in affecting the trade-off.

[1] In a model without lags, the interest rate is typically a linear function of current shocks. When policy acts with a lag, the role of policy is to control the effects of past shocks which remain in the system via persistence. Policy

The essence of this argument goes as follows: here as is typical in models of monetary policy, the relative control of output or inflation deviations depends on the extent to which supply shocks are controlled – a tight control of supply shocks corresponds to tight control of inflation and lax control of output, and vice versa. However, the variances of both future output and inflation depend positively on the amount of 'inflationary pressure' already within the system. A myopic authority takes this as given and policy is sensitive to its preferences over output and inflation stability. However, a far-sighted authority, regardless of its preferences, has an incentive to control inflation now since this will lower the future variance of both output and inflation. This results in policy being less sensitive to preferences.

Dynamic monetary policy models generally rely on some form of persistence in either output or inflation. For instance Lockwood, Miller and Zhang (1998) discuss persistence in employment while Mash (2000) discusses persistence in inflation. Both of these models are characterised by inflation bias, which results from the authority targeting a level of output higher than the natural rate. The level of the discount rate is then shown to have implications for the consequent issues of time inconsistency and delegation; for instance Mash (2000) shows that an authority with long planning horizons offsets the problem of time inconsistency.

Persistence can also be explored in the case where the authority targets the natural rate of output and hence there is no inflation bias. Clark, Goodhart and Huang (1999) provide a discussion of the difference between policy made under commitment and discretion in the case where there is no inflation bias but there is persistence of inflation. Bean also adopts a framework with persistence in inflation and without inflation bias (i.e. in which the authority targets the natural rate of output).[2] In comparison with the framework put forward by Clark et al., Bean places a greater emphasis on the lag structure and a lesser emphasis on the treatment of expectations.[3] What we show here is that in the Bean framework, planning horizons can have a significant impact on the rectangularity argument i.e. that if planning horizons are short, the relative weighting of inflation and output variability in the loss function does have a much greater impact on the outcome of optimal policy.

---

then acts by controlling persistence, and so the optimal interest rate is then often a function of the persisting variables rather than the shocks. With simple lag structures, this can take the form of a Taylor rule.

[2] Bean argues that independence of the central bank is likely to achieve this situation, since it is then immune to the political pressures which cause output to be targeted at a level higher than the natural rate. He also suggests that lengthening the term of MPC members would further insulate the MPC from short-term political pressures. Our analysis broadly supports this measure since it could also be expected to reduce discounting.

[3] This is of course entirely natural given the different issues that the models analyse. For example, Clark et al. assume policy affects both output and inflation contemporaneously. However, taking into account that the policy instrument is the nominal rate, they allow a role for inflationary expectations in the evolution of output unlike the models of Ball (1997), Bean (1998) and Svensson (1997).

Section 2 provides an outline of the theoretical analysis presented by Bean (1998) and Ball (1997), and section 3 provides the corresponding analysis in the case of discounting. Svensson (1997) also solves this problem in the case of discounting; nevertheless, we present a solution here which stylistically is more in keeping with that of Ball (1997) and Bean (1998) since it allows a clearer discussion of rectangularity argument.[4] Section 4 discusses the implications of discounting on rectangularity and section 5 concludes.

2. An Outline of the Bean Model

Bean (1998) and Ball (1997) use the following equations for aggregate demand and supply. Aggregate demand is given by

$$y = -\mu r_{-1} + \lambda y_{-1} + \eta \qquad \text{with } \eta \sim \mathrm{N}\,(0,\sigma_\eta^2) \tag{D1}$$

where r, the real interest rate, is the policy instrument. Aggregate supply is given by an 'accelerationist' Philips curve:[5]

$$\pi = \pi_{-1} + \alpha y_{-1} + \varepsilon \qquad \text{with } \varepsilon \sim \mathrm{N}\,(0,\sigma_\varepsilon^{\,2}) \tag{S1}$$

So real interest rates affect output with a lag of one period, output affects inflation with a further lag, and one period is interpreted as corresponding to a time span of around one year. The loss function of the non-discounting monetary authority is given by:

$$L = Var(\pi) + \beta Var(y) \tag{L1}$$

Bean provides a clear derivation of the optimal rule. The key observation is that the rule must be of the form:

$$Ey_{+1} = -\rho E\pi_{+1} = -\rho(\pi + \alpha y) \tag{1}$$

The constant ρ is referred to as the 'feedback parameter.' Equation (1) implies that output moves in such a way that as far as possible past demand shocks, which increase the volatility of both output and inflation, are removed from the system. The level of ρ determines to what extent supply shocks, which have opposing effects on output and inflation volatility, are controlled – a high ρ corresponds to a tight control of supply shocks and so inflation. In effect, by imposing (1) we are choosing an efficient 'sequencing' matrix described below in

[4] In particular of the feedback parameter, defined below. Svennson's solution, in many ways more direct, is ideal for his consideration of the inflation forecast as an intermediate target.

[5] This the reduced form of the equation $\pi = (1-\omega)\mathrm{E}\pi + \omega\pi_{-1} + \zeta y_{-1} + \varepsilon$ where $\alpha = \zeta/\omega$.

(3). To see the mathematical basis for equation (1) - see Svensson (1997) - we simply recognise that Ey_{t+1} can be considered as the control variable for this dynamic problem. Maximising the loss function (which due to the linear-quadratic structure of the model is a quadratic function of expected inflation deviations) subject to the constraint (S1) gives a first order condition that is a rule of the form (1). Equation (1) then determines the form of the optimal interest rate rule, which is a Taylor rule,

$$r = \{[\lambda + \alpha\rho]y + \rho\pi\}/\mu \tag{2}$$

The current values of inflation and output in terms of last period's values are then given by the following matrix, which we shall refer to as the 'sequencing' matrix:

$$\begin{bmatrix} y \\ \pi \end{bmatrix} = \begin{bmatrix} -\alpha\rho & -\rho \\ \alpha & 1 \end{bmatrix} \begin{bmatrix} y_{-1} \\ \pi_{-1} \end{bmatrix} + \begin{bmatrix} \eta \\ \varepsilon \end{bmatrix} \tag{3}$$

It is important to note that the sequencing matrix is of rank 1 – in this system, this is a consequence of following an efficient policy given by (1). As will be described below, this rank condition can be seen as the essential cause of rectangularity in the case when there is no discounting.

Equation (3) is used to calculate the unconditional variances:[6]

$$\begin{bmatrix} Var(y) \\ Var(\pi) \end{bmatrix} = \frac{1}{[2-\alpha\rho]} \begin{bmatrix} 2 & \rho/\alpha \\ \alpha\rho & [(1/\alpha\rho)+2-\alpha\rho] \end{bmatrix} \begin{bmatrix} \sigma_\eta^2 \\ \sigma_\varepsilon^2 \end{bmatrix} \tag{4}$$

These expressions for the variances are substituted into the objective function (L1) and this allows us to calculate the value of the feedback parameter ρ that minimises the loss. This is the positive root of the equation:

$$\beta\rho^2 + \alpha\rho - 1 = 0 \tag{5}$$

The issue of the rectangularity of the policy frontier is easily described using (L1). Once the optimal rule has been implemented, for a wide range of values of β, there is a small difference in the final values of $Var(y)$ and $Var(\pi)$ that are obtained. Figure (1) shows the shape of the policy frontier which, replicating that in Bean, is shown in the space of the

[6] Bean uses the result that if $X=BX_{-1}+E$ where B is a conformable matrix and E is a vector noise process, then $\mathrm{Vec}[\mathrm{Var}(X)] = [1 - (B \otimes B)]^{-1}\mathrm{Vec}[\mathrm{Var}(X)]$.

standard deviations of output and inflation. The outcomes for two different values of β, β=0.3333 and β=3, are shown and can be seen to be 'close' as described in the introduction.

2.1 The Effect of Parameters Estimates and Lag Structure on Rectangularity

In figure 1, we use the set of parameter estimates obtained by Bean (1998) in his original empirical investigation. These are $\sigma_\eta = 1.55$ and $\sigma_\varepsilon = 1.25$ for the demand and supply shocks respectively and $\alpha = 0.492$. We find that the rectangularity property is robust to realistic alterations in the parameter values and lag structure of the model (accordingly, we use Bean's parameter estimates in all subsequent figures[7]). To examine the effect of lag structure, we consider the following alternative equations for demand and supply respectively:

$$y = -\mu r + \lambda y_{-1} + \eta \qquad \text{(D2)}$$

$$\pi = \pi_{-1} + \alpha y + \varepsilon \qquad \text{(S2)}$$

In (D2) the lag in interest rates has been removed so policy affects output contemporaneously, while in (S2) the lag in output has been removed so output affects inflation contemporaneously. If we consider the two economies (S1) & (D2) and (S2) & (D1), we find that the property of rectangularity remains. In both of these economies, policy affects inflation with a one period lag: this is enough to ensure that the policy frontier becomes vertical as the preferences of the authority tend to pure inflation targeting, since policy cannot mitigate to any extent the contemporaneous supply shock. Only in the unrealistic model (S2) & (D2), where policy acts without lags and the fact that policy can control the contemporaneous supply shock implies that the policy frontier does not become vertical in the limit, is the shape of the policy frontier noticeably less rectangular.[8] Figure (2) shows the shapes of the policy frontier for all these various models together with the outcomes for β=0.3333 and β=3.

3. The Optimal Rule with Discounting

We will now use a loss function with discount factor δ:

$$L_t = E\left[\sum_{s=0}^{\infty} \delta^s (\pi_{t+s+1}{}^2 + \beta y_{t+s+1}{}^2)\right] \qquad \text{(L2)}$$

[7] A rectangular frontier is obtained from wide ranges of parameter values. Details available from the authors.
[8] In the model without lags, and for small values of β, decreasing persistence in inflation in the model further reduces rectangularity. This is very much not the case in the model with lags, see footnotes 1 and 9.

Note that when the monetary authority minimises (L1) the long run mean of both variables is zero, and so this is the same as minimising (L2) in the limit as $\delta \to 1$. Our analysis in this section proceeds in three stages. First we identify two mechanisms by which the introduction of discounting can affect optimal policy. Secondly we show that as an authority becomes more patient, both of these mechanisms cause the authority to be more concerned with controlling inflation regardless of its initial preferences. Thirdly we describe the effect of this on the sensitivity of policy to the weighting of the objective function. The sensitivity of policy to the weighting of the objective function is now primarily what we refer to when we discuss property of rectangularity, rather than the shape of the frontier; the shape of the policy frontier in terms of the long-run variances of output and inflation will not change when discounting is introduced. However, the shape of policy frontier in the space of short-run variances of output and inflation is much less rectangular, as seen below.

With discounting, the optimal rule must still take the same form (1) as above, but the feedback parameter ρ will now be a function of both β and δ. Allowing discounting alters optimal policy in two ways. In this system, interest rates affect output with a one period lag and inflation with a two period lag. Since policy now affects output and inflation with different lags, changing the discount factor acts to change the relative weight on output in the loss function: for a given β, as the authority becomes more myopic, it becomes more concerned with output stabilisation since it can only influence inflation at a longer horizon. We will call this mechanism the 'weighting mechanism.' Discounting alters policy by another mechanism described below which we shall refer to as the 'investment mechanism.' To separate the effects of the two mechanisms, we make the following reparametrisation $\beta = \tilde{\beta}\delta$. Varying δ while holding $\tilde{\beta}$ constant then allows us to isolate the effect of the investment mechanism, since the relative weighting in the loss function on output and inflation deviations at the horizon at which they are affected by current policy does not then change.

Since the optimal rule under discounting remains of the same form (1), the sequencing matrix also takes the same form (3) with respect to the feedback parameter. Because of the lag structure, at time t the interest rate decision affects $E[{y_{t+1}}^2]$ and $E[{\pi_{t+2}}^2]$. From (3) we have,

$$E[{y_{t+1}}^2] = \rho^2[\alpha y_t + \pi_t]^2 + C_1 \tag{6}$$

$$E[{\pi_{t+2}}^2] = (1-\rho\alpha)^2[\alpha y_t + \pi_t]^2 + C_2 \tag{7}$$

where the constants $C_1 = \sigma_\eta^2$ and $C_2 = \alpha^2\sigma_\eta^2 + 2\sigma_\varepsilon^2$. Noting that $\alpha y_t + \pi_t = E[\pi_{t+1}]$ we can rewrite (6) and (7) as

$$E[y_{t+1}{}^2] = \rho^2 E[\pi_{t+1}]^2 + C_1 \tag{8}$$

$$E[\pi_{t+2}{}^2] = (1-\rho\alpha)^2 E[\pi_{t+1}]^2 + C_2 \tag{9}$$

Equations (8) and (9) provide the key to understanding the other effect of discounting on rectangularity (the 'investment mechanism'). If we take $E[\pi_{t+1}]$ as given, the relative deviations of output and inflation will be determined by the feedback parameter: a high feedback parameter controls inflation but leaves output far from its target, whereas a low feedback parameter does the opposite. This is the trade-off that a myopic (δ=0) authority faces. However, the trade-off appears different for a long-sighted authority. In the future, output and inflation deviations will depend not only on the value of the feedback parameter, but *both* depend on the expected deviation of inflation in preceding periods. By getting inflation to target now, the authority reduces the expected deviations of *both* inflation and output in the future. Reducing inflation now can be seen as an 'investment' which creates long-term stability of both inflation and output. So as the authority looks to the future, regardless of whether it is primarily concerned with output or inflation stability, it becomes more concerned with controlling inflation in the present. Hence the policies of authorities with different preferences 'converge' as they become more far-sighted. [9]

Note that in general, if the 'sequencing' matrix is of rank 1, there will be a particular linear combination of (expected) output and inflation on which both inflation and output depend. As an authority becomes more far-sighted, its policy will become more concerned with controlling this. Here, a far-sighted authority will control inflation more tightly than a myopic one and have a higher feedback parameter. We now verify this. It is shown in the appendix that using the above reparametrisation the feedback parameter is given by the equation:

$$\tilde{\beta}\delta\rho^2 + \left[\tilde{\beta}\left(\tfrac{1}{\alpha}\right)(1-\delta) + \alpha\right]\rho - 1 = 0 \tag{10}$$

For the optimal policy rule we take the positive root of (10), and it can be seen that this lies in the range $[0, \tfrac{1}{\alpha}]$. Note that for the case δ=1 equation (10) gives the feedback parameter given by (5) above. Differentiating (10) with respect to δ while holding $\tilde{\beta}$ constant:

[9] This 'investment' argument depends on there being some persistence in inflation. One might then imagine that removing persistence from the system would reduce rectangularity. This is not so: while reducing persistence in inflation decreases the strength of the investment mechanism, in a model with lags it also reduces the trade-off between output and inflation. Without any persistence, the policy frontier is upward sloping and policy completely insensitive to preferences. To see this, make the following observation. If a policy maker cares only about output stability she just uses interest rates to control demand. If she cares about inflation as well, she uses output to reduce the persistence in inflation (see footnote 1) thereby increasing the volatility of output. This is the trade-off. If there is no persistence in inflation however, she has no need to do this; minimising output volatility then minimises inflation volatility. Persistence causes the trade-off, but also introduces a mechanism by which discounting alters the optimal policy outcome.

$$\rho'\left[2\tilde{\beta}\delta\rho + \tilde{\beta}\left(\tfrac{1}{\alpha}\right)(1-\delta) + \alpha\right] + \tilde{\beta}\rho(\rho - \alpha) = 0 \tag{11}$$

Since the term in the square brackets is strictly positive and $\rho \in [0, \tfrac{1}{\alpha}]$, it follows that ρ is an increasing function of δ. Unless the weighting objective function is such that it corresponds to a case of pure inflation targeting or pure output stabilisation, ρ is a strictly increasing function of δ. The investment mechanism augments the weighting mechanism: the more patient the authority the tighter is the policy in controlling inflation.

3.1 *Discounting in an Alternative Framework*

As discussed above, the weighting mechanism relies on the fact that policy affects output and inflation at different horizons. Suppose we wanted to consider the effect of discounting on a system where policy affected both output and inflation at the same horizon. Consider for example the system (S2) & (D1) which as seen in figure (2) has the property of rectangularity when there is no discounting. We refer to this as the 'one-lag' case. The optimal rule now has the form $Ey_{+1} = -\rho\pi$ which gives us the analogues of equations (8) and (9):

$$E[{y_{t+1}}^2] = \rho^2{\pi_t}^2 + C_1 \tag{12}$$

$$E[{\pi_{t+1}}^2] = (1-\rho\alpha)^2{\pi_t}^2 + C_2 \tag{13}$$

where $C_1 = \sigma_\eta^2$ and $C_2 = \alpha^2\sigma_\eta^2 + \sigma_\varepsilon^2$

Here there is no weighting mechanism but the investment mechanism remains the same: it can be shown that the feedback parameter is given by an equation identical to (10) save for the reparametrisation made above:[10]

$$\beta\delta\rho^2 + \left[\beta\left(\tfrac{1}{\alpha}\right)(1-\delta) + \alpha\right]\rho - 1 = 0 \tag{14}$$

3.2 *Policy Equivalence Curves*

We can summarize as follows: the system of equations for output and inflation determines how policy affects the sequencing matrix. This dictates the form of the optimal rule in terms of a feedback parameter that is constant over time. Once the form of the optimal rule has been determined, it only remains to solve for the feedback parameter to specify policy precisely. The feedback parameter is a function of the parameters β and δ, respectively preferences over inflation and output stability, and the level of discounting.

[10] Details available from the authors.

Ignoring the weighting mechanism and considering the one-lag case for the moment, the feedback parameter is given by equation (14). Then treating ρ as fixed in (14) we can then plot curves in (β, δ) space which are the locus of points which yield equal values of ρ and so equivalent policy rules. We refer to these as 'policy equivalence' curves. Figure (3A) shows a series of such curves for equation (14). Bearing in mind that these curves are contour lines representing equal values of ρ and that contours with lower values of ρ correspond to higher values of β, figure (3A) demonstrates graphically the result derived above, that for a given value of β the feedback parameter is increasing in δ i.e. that policy becomes tighter as the authority becomes more far-sighted.

We can also see that as the authority becomes more far-sighted and δ increases, changes in β result in smaller changes in the feedback parameter. As argued above, when the authority is far sighted, the weighting of the objective function matters less in the formation of policy. Figure (3B) shows the results in the original Bean framework using equation (10) where the weighting mechanism is taken into account; we can see that both these features remain, only the effect is more dramatic.

4. The Implications for Rectangularity

How do we assess the implications of these results for rectangularity? As an example, we can calculate the completely myopic feedback parameter in the one-lag case when δ=0 i.e. when the authority just looks one period ahead:

$$\rho = \frac{\alpha}{\alpha^2 + \beta} \qquad (15)$$

The short run policy frontier given by (12) and (13) in { $E[y_{t+1}{}^2]$, $E[\pi_{t+1}{}^2]$ } space depends on the current level of inflation, particularly its magnitude relative to the standard deviation of the supply and demand shocks. It is easy to see that this will not be rectangular. Figure 4A shows how the short-run policy frontier shifts outwards as the current level of inflation increases, with the policy outcomes marked for β=0.3333 and β=3.

Suppose a social planner who did not discount the future wanted to consider the impact of not specifying the weighting of the objective function of her myopic monetary authority. Would similar outcomes result from a wide range of weightings? To address this question, we use the myopic feedback rule for different values of β to calculate the long-run variances of inflation and output. Note that in the space of long-run variances, since the form of the optimal rule is the same except for the value of the feedback coefficient, the shape of the policy frontier will remain unchanged. The effect of myopia is shown on the distance between outcomes for two different values of β; these can be seen for β=3 and β=0.3333 in figure 4B for the one-lag case.[11] The filled circles show the outcomes for a perfectly far-sighted authority, whereas the hollow circles show the outcome for myopia, and as can be seen these are appreciably further apart. When the authority is myopic outcomes are more sensitive to differences in the weighting.

In his original analysis Bean provides an analysis of the 'excess loss' that results from the authority choosing the 'wrong' β in calculating the feedback coefficient. Supposing for instance that a social planner has a particular value of β, say β=1, the excess loss diagram shows the percentage excess loss, from the point of view of the planner, when policy is implemented by an authority with a weighting β^*. In the case with discounting, we can augment this analysis by supposing that the planner and the authority each have their own discount rate. Supposing that the planner does not discount the future (i.e. δ=1), figure 5A shows the excess loss from the planner's point of view both for an authority with δ=1 and for an authority with δ=0.65 in the one-lag case[12]; figure 5B shows the corresponding diagram for the two-lag case. In both diagrams, the dashed line shows the loss function when the authority does not discount, and the continuous line when it does. We can see from the policy equivalence curves in figure 4, that when the authority and the social planner have differing discount rates, the excess loss of the social planner will be zero for a particular weighting of the authority's objective function β^* with $\beta^*\neq 1$. This is verified in figure 5. As we can see, consideration of the relative discount rates can be important in the excess loss rates; when the authority does not discount the future excess loss rates are relatively low for a wide range of values of β^* whereas when the authority does discount, the excess loss can be much higher.

5. Conclusions

Whilst it is true in general that a 'rectangular' policy frontier implies that similar policy outcomes are obtained from a wide range of weightings of the objective functions, it should be emphasised that *an important implication of rectangularity of the policy frontier is that it makes sense for the central bank to concern itself with output volatility to at least some strictly positive extent*. As can be seen in figure 1, a central bank that purely targets inflation will not exploit the rectangularity of the policy frontier; it will obtain an outcome at the point *A* whereas it might be considered socially desirable to be in the general vicinity of *B*. From

[11] It is not helpful to consider a completely myopic authority in the two-lag case, since the authority will only care about output, and so we obtain ρ=0 and an infinite variance of inflation regardless of the value of β. In general, however, because of the additional impact of the weighting mechanism, the sensitivity of outcomes to β is further heightened in the two-lag case when δ is small. In the one-lag case discounting has relatively little impact on policy when β is small but non-zero; this is not true in the two-lag case because of the weighting mechanism.

point A, large gains in output stability can be obtained by compromising to a small degree on inflation volatility, since the slope of the policy frontier becomes vertical at point A.

The implications for rectangularity of discounting may be summarised as follows. When the authority is myopic, it focuses on the short-run policy frontier which is not rectangular. The authority cares less about inflation the more it discounts the future, and its actions become increasingly sensitive to the weighting of its objective function. From the point of view of a social planner who does not discount the future, large excess losses can be generated by not obviously implausible combinations of β and δ for the authority.

What does this say about the way in which an authority's remit should be specified? Suppose society wants to achieve on outcome at A on the long-run policy frontier. This can be achieved easily – the authority should engage in pure inflation targeting and whether it discounts the future or not is irrelevant since discounting has no effect on the actions of an authority that purely targets inflation. However, if society wants to exploit the rectangularity of the long-run frontier and achieve an outcome in the vicinity of B, then it does become important to consider the discount rate of the authority. Given mixed targeting and no discounting, rectangularity implies that we are fairly likely to end up at B. Of course discounting *per se* does not imply an outcome away from B; were the authority targeting inflation excessively at the expense of output stabilisation, increased discounting could help achieve an outcome at B. With discounting however, policy outcomes become more sensitive to the value of β, and again many not obviously implausible combinations of β and δ can result in an outcome far from B. This is particularly true of the two-lag case where the 'weighting mechanism' plays an additional role.

Since the relative weighting of a monetary authority's objective function is often not specified in its remit, as in the U.K. case, agents in the economy in general may not have a precise idea of the authority's β. When it discounts the future less, the actions of the authority become less sensitive to the value of β and ought to be more predictable, so stability is likely to be improved when the authority is patient. So, provided the authority does not target inflation excessively, arrangements that decrease myopia[13] ought to result in greater long run stability.

Bibliography

Ball L. 1997 "Efficient Rules for Monetary Policy" NBER Working Paper W5952

---

[12] The reason we consider that such a low discount factor (or high discount rate) might be plausible is for the reasons mentioned in the introduction: that members of the MPC serve a term of three years, so the average remaining term is one and a half years; and that a period in the model corresponds to around one year.
[13] For instance for the U.K., by increasing the length of service on the MPC.

Bean C. 1998 "The new U.K. monetary arrangements" Economic Journal 108: 1795-1809

Clark P.B., Goodhart C.A.E. and Huang H. 1999 "Optimal monetary policy rules in a rational expectations mode of the Phillips curve" Journal of Monetary Economics 43: 497-520

Lockwood B., Miller M. and Zhang L. 1997 "Designing monetary policy when unemployment persists" Economica 65: 327-45

Mash R. 2000 "The time inconsistency of monetary policy with inflation persistence" mimeo University of Oxford

Svensson L.E.O. 1997 "Inflation forecast targeting: Implementing and monitoring inflation targets" European Economic Review 1997 41: 1111-1146

Appendix

Define $V_t = E\left[\sum_{s=0}^{\infty} \delta^s (\pi_{t+s+2}{}^2 + \tilde{\beta} y_{t+s+1}{}^2)\right]$. Then the loss function (L2) can be written:

$$L_t = \delta V_t + E[\pi_{t+1}{}^2] = \delta V_t + \left((\alpha/\omega) y_t + \pi_t\right)^2 + \sigma_\varepsilon{}^2 \tag{A1}$$

Since ρ only enters (A1) via V_t, by the principle of optimality,

$$V_t = \min_\rho E[\pi_{t+2}^2 + \tilde{\beta} y_{t+1}^2 + \delta V_{t+1}] \tag{A2}$$

From equations (8) and (9) in the main text, we can see that V_t will have the form:

$$V_t = A_0 + A_2 E[\pi_{t+1}]^2 \tag{A3}$$

Substituting equations (8), (9) and (A2) into the right hand side of (A2) we get:

$$V_t = \min_\rho \left[\left((1 - \rho\alpha/\omega)^2 (1 + A_2\delta) + \tilde{\beta}\rho^2\right) E[\pi_{t+1}]^2 + \delta A_0 + \tilde{\beta} C_1 + (1 + \delta A_2) C_2\right] \tag{A4}$$

Using the envelope theorem, we can then obtain the first order condition for ρ :

$$\tilde{\beta}\rho = \alpha/\omega \left(1 - \rho\alpha/\omega\right)\left(1 + A_2\delta\right) \tag{A5}$$

Comparing the coefficient of $E[\pi_{t+1}]^2$ in (A3) and (A4) we then obtain

$$A_2 = \left(1 - \rho\alpha/\omega\right)^2 (1 + A_2\delta) + \tilde{\beta}\rho^2$$

Substituting for $\tilde{\beta}\rho$ from (A5) gives

$$A_2 = \left(1 - \rho\alpha/\omega\right)\left(1 + A_2\delta\right) \tag{A6}$$

Comparing (A5) and (A6) we can see that $\rho = \left(\alpha/\omega\tilde{\beta}\right) A_2$. Substituting this into (A5) gives

$$\tilde{\beta}\delta\rho^2 + \left[\tilde{\beta}\left(\omega/\alpha\right)(1 - \delta) + \left(\alpha/\omega\right)\right]\rho - 1 = 0 \tag{A7}$$

Figure 1 The Bean Policy Frontier

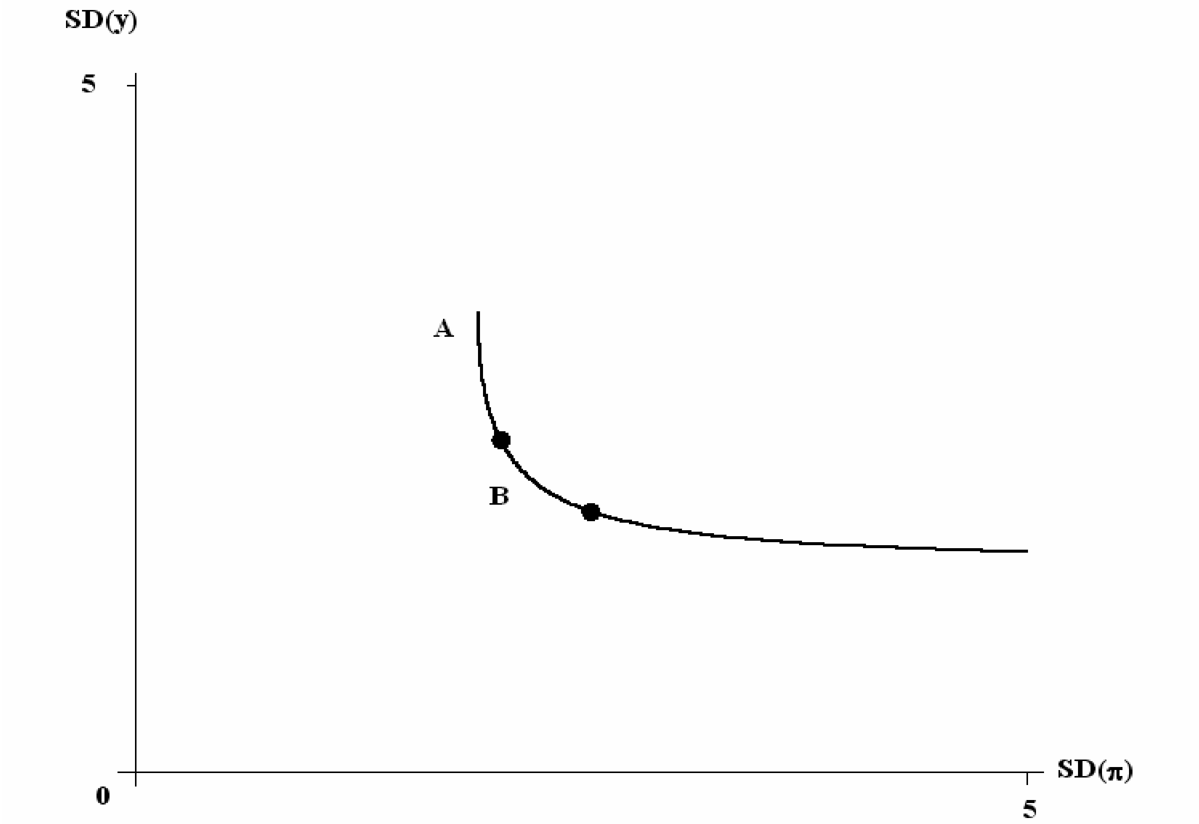

Figure 2 Lag Structure

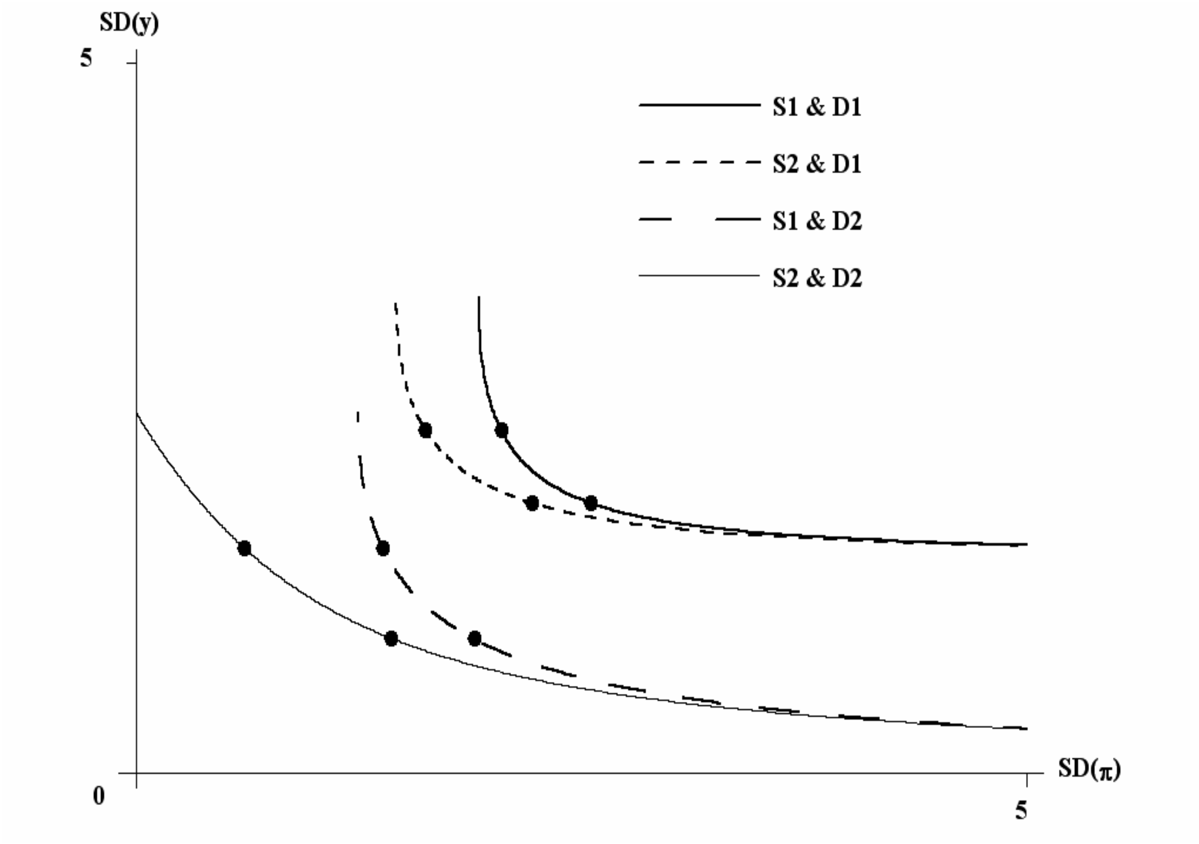

Figure 3 A. Policy Equivalence Curves

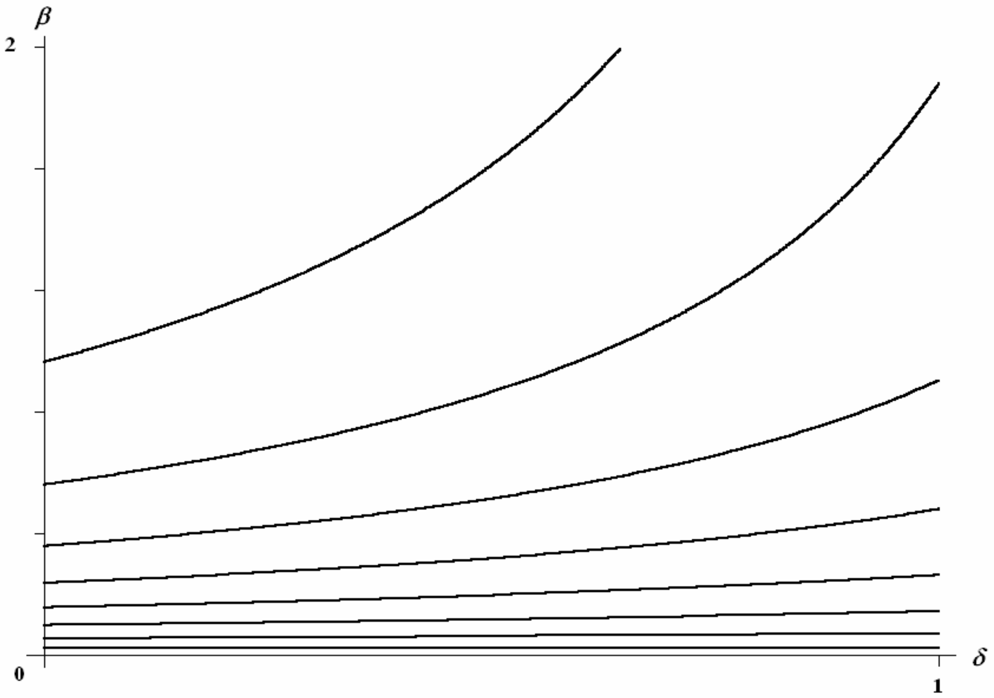

B Equivalence Curves with the Weighting Effect

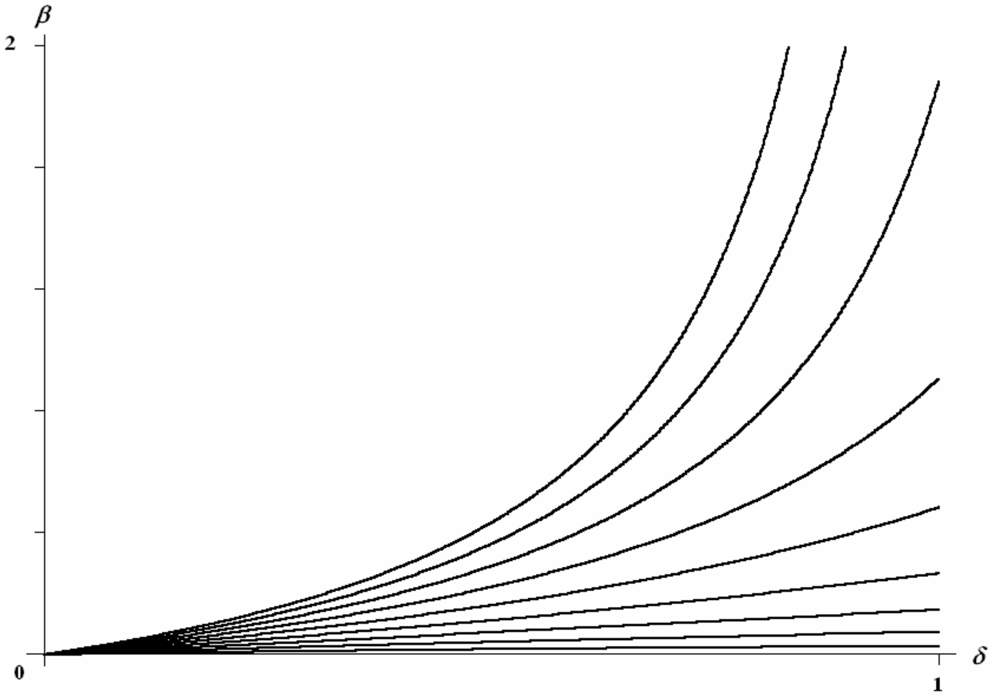

Figure 4 A. Short-Run Policy Frontiers

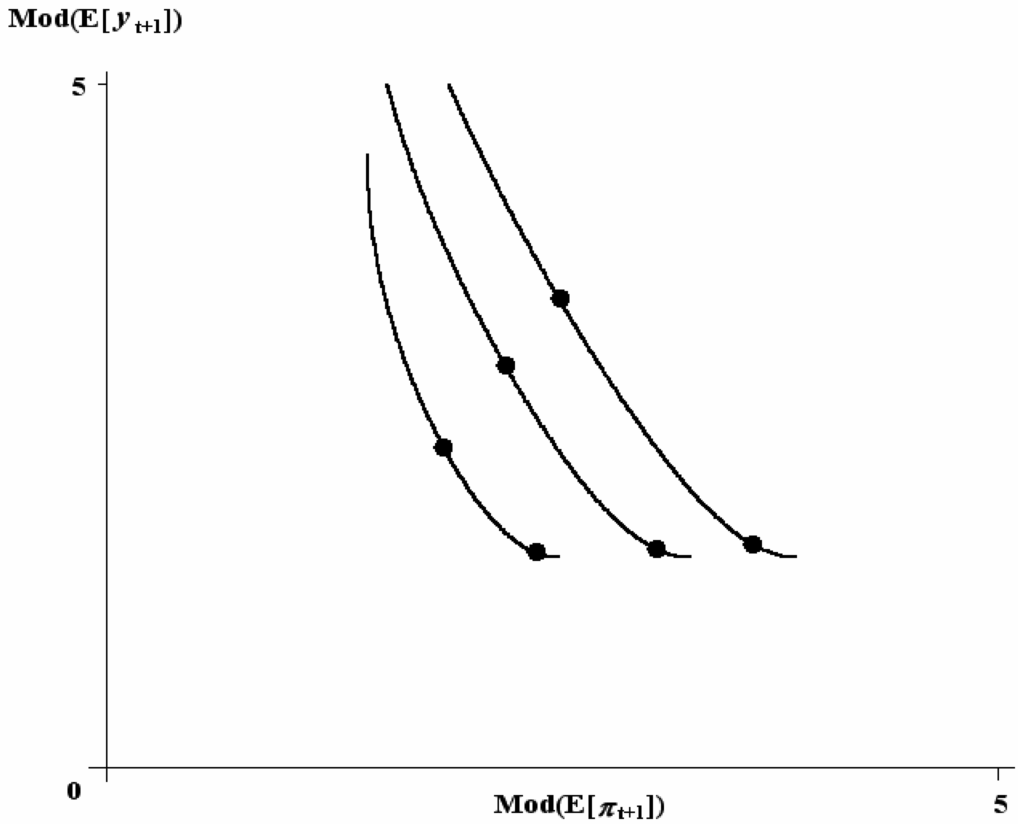

B. Long-Run Policy Frontiers – Myopic and Far-Sighted Outcomes

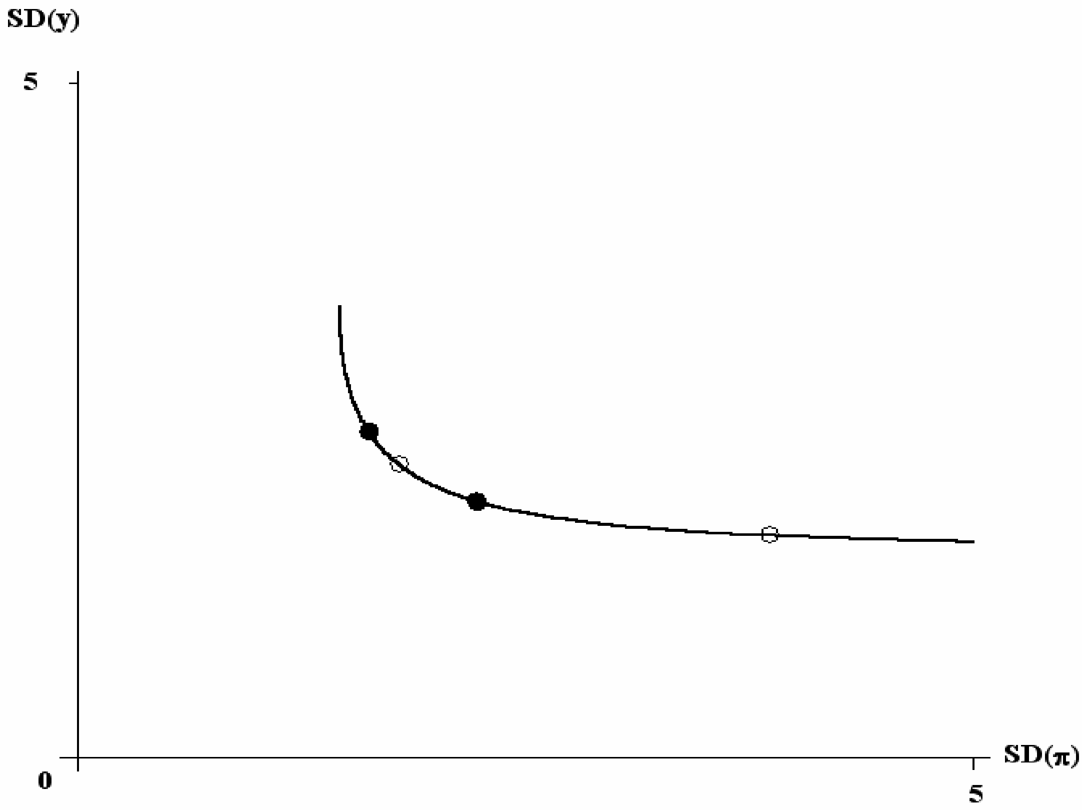

Figure 5 A. Excess Losses in the One-Lag Case with and without Discounting

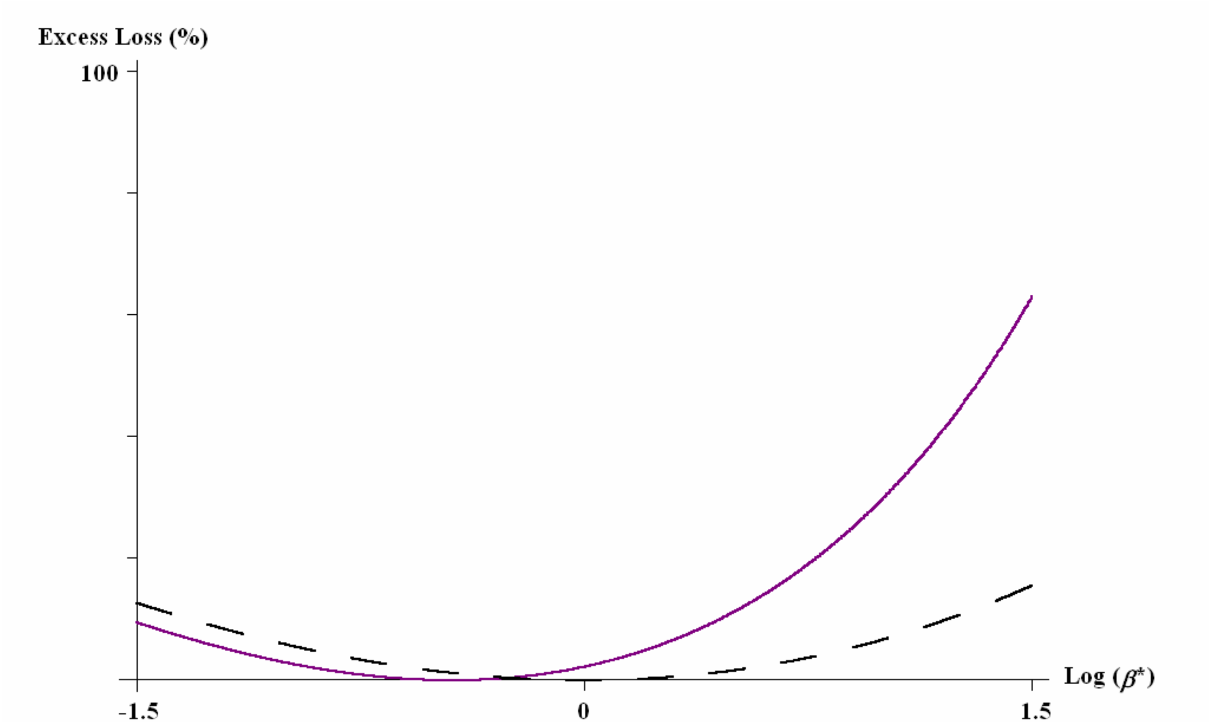

B. Excess Losses in the Two-Lag Case with and without Discounting