# A 'Timeless Perspective' on Optimality in Forward-Looking Rational Expectations Models

Andrew P. Blake*

*National Institute of Economic and Social Research*
*2 Dean Trench Street, Smith Square, London SW1P 3HE, UK*

October 11, 2001

**Abstract**

This paper discusses the 'timeless perspective' optimisation concept with reference to a much-studied forward-looking rational expectations model. We establish that this policy, as usually described, is not always superior to a time consistent alternative on the basis of the stochastic equilibrium. We derive an alternative 'timelessly optimal' rule which is globally optimal with respect to the unconditional variance and is therefore more supportable as a time consistent equilibrium.

*JEL classifications*: E52, E58, C61

*Keywords*: Monetary Policy, Optimality, Time Consistency, Timeless Perspective

*Email*: ablake@niesr.ac.uk, *Tel.*: +44 20 7654 1924

© NIESR 2001

*This research was made possible by funding from the ESRC's 'Understanding the Evolving Macroeconomy' program. I am grateful to James Sefton for helpful discussions which led to considerable clarification of the argument. Errors and omissions remain my responsibility.

# 1 Introduction

Recently the 'timeless-perspective' optimality concept proposed by Michael Woodford (see Woodford 1999a, Woodford 1999b) has received a great deal of attention in the inflation targeting literature. Whilst it involves no new solution concepts, the 'timelessly optimal' control rule emphasises the stochastic equilibrium to the monetary policy problem. This solution has been used by, *e.g.*, King and Wolman (1999), Clarida, Galí, and Gertler (1999), Steinsson (2000), McCallum and Nelson (2000) and Walsh (2001) as well as further papers by these authors and others.[1] The policy proposal is simple to outline. What makes an optimal policy time-inconsistent is the separate treatment of initial conditions and policy in the longer term. By ignoring those initial conditions and sticking to a constant policy, a policy can be implemented that has desirable long-run stabilisation properties—for example the absence of inflationary bias—which monetary policy makers may be able to sustain as a consistent policy. This is the 'rules versus discretion' debate revisited with the well known result that a rule to which a policymaker can commit may outperform the discretionary equilibrium, except that the particular proposed rule is the one associated with the optimal but time inconsistent policy at a 'mature' stage.

The many proponents of this as a solution to the optimal monetary policy problem argue that it corresponds to the policy which the monetary authorities would have wished to implement given the opportunity to make a binding choice in the distant past to current behaviour. They argue further that this conforms to the notion of time consistency implicit in much of the early literature. The popularity of this solution as a viable description of policy is enhanced by its analytical tractability, particularly for the model considered in this paper where troublesome 'pre-determined' Lagrange multipliers can be eliminated to yield a readily interpretable policy rule.

However there is a significant weakness to this proposition as a description of equilibrium monetary policy. In this paper we argue that the consistency argument is generally false. There exist superior policies in stochastic equilibrium that the policymaker will necessarily prefer in the absence of enforced discretion. If, in the long run, a policy which is at least as good as the usual 'timelessly optimal' one exists then policymakers will seek to adopt this as the rule they would have wished to commit themselves to in the distant past. In order to demonstrate the contrary one only needs to find an alternative policy which is better. This motivates the investigation undertaken by McCallum and Nelson (2000).

However, we can go further. If a policy could be found that is optimal in the metric usually used to evaluate the 'timelessly optimal' policy then

[1] In particular, the stochastic problem is treated in Svensson and Woodford (2000) and Svensson and Woodford (2001). We return to their analysis below.

this may fulfil the criteria for time consistency. We derive a policy which directly minimises the asymptotic loss function and discuss its properties. It can be derived in a similar fashion to the timeless-perspective optimal policy although, of course, the actual objective function is rather different. Because of this inherent superiority it is necessarily more supportable as a time consistent policy. However the method of derivation indicates very clearly that a *short-term* incentive to depart from the announced rule exists as it is not the solution to the originally posited objective function. The key to sustaining any 'timelessly optimal' policy as time consistent rests with convincing agents that the long-term is the true objective of policy.

The analysis of these issues and attempts to resolve the time inconsistency problem go back a long way. The time inconsistency of the optimal policy was recognised by Kydland and Prescott (1977). Barro and Gordon (1983) emphasised a static concept of time inconsistency, where the inability of the policymaker to commit to a particular monetary policy induces an inflationary bias. In a repeated game this bias is mitigated although not altogether eliminated. The dynamic rational expectations optimal policy problem was thoroughly investigated by a number of authors in the 1980s. See, *e.g.*, the collected articles in Buiter and Marston (1985) and Currie and Levine (1993).[2] This literature emphasised a second form of time inconsistency, nicely illustrated with the temptation to re-optimise captured by a co-state vector, which we discuss further below.[3] This dynamic aspect is critical to our analysis of the inconsistency problem. It turns out that introducing commitment technologies to the optimal but time inconsistent policy may in the long run be counter-productive. Any short-run gains from commitment may become longer term losses where it would be preferable to 'renege' and adopt some alternative policy rule. As the timeless optimality solution reflects the solution in equilibrium it is these longer-term costs which are relevant in assessing the sustainability of the policy regime. It may therefore be that the long-run associated with the discretionary equilibrium is to be preferred to the commitment solution.

The recent papers which analyse the timeless perspective policy mostly adopt a simple forward-looking Phillips curve. A graph of the trade-off between the unconditional variance of inflation and the output gap (in this paper $\sigma_\pi^2$ and $\sigma_y^2$) or a table of the weighted variances are often used to illustrate the impact of the concern given to anti-inflation priorities by the central bank for the policy regimes under consideration (see, *e.g.*, Clarida, Galí, and Gertler 1999). This would seem to be a relevant comparison, given that for stochastic initial conditions (which, we note below, is equivalent to

[2]More recently, Söderlind (1999) has presented some more general solutions to these problems and provided software solutions for other investigators, often used in the recent literature.

[3]An important link between these two sources of time inconsistency was Cohen and Michel (1988), a good example of a paper which analyses both.

the timeless-perspective optimality concept) the expected value of a usual discounted quadratic welfare loss function is proportional to a weighted average of these quantities. In this paper we argue strongly that potential policy rules and stochastic equilibria should be compared across the correct metric. Once this is adopted a fully optimal policy *for that problem* can be derived.

Initially two questions are addressed in this paper. Firstly, and as investigated by McCallum and Nelson (2000), is it necessarily the case that the 'timelessly optimal' policy is better than a consistent alternative on the basis of their stochastic equilibria? We find that it is not. Secondly, is there an optimal policy on the basis of this metric? We find that there is. In answering both of these questions we are able to offer a number of important insights into the nature of time inconsistency in a stochastic equilibrium. In particular, it becomes clear that the social discount factor plays an important role in determining the best policy.

The paper is organised as follows. In Section 2 we describe the model and characterise the policy problem. In Section 3 we derive the two candidate policies already discussed in the literature in the context of an asymptotic loss function. In Section 4 an alternative timelessly optimal policy is derived and all three equilibria are compared and a number of important aspects of the solutions discussed. Conclusions are drawn in Section 5. Three appendices give additional first-order conditions for the fully optimal solution for the model considered in the main text, a discussion of the form of the policy rules implemented and a brief treatment of the general LQG problem.

## 2 Model, policy description and solution

We consider the model analysed by (amongst others) Clarida, Galí, and Gertler (1999) and McCallum and Nelson (2000). The model is a very simple one. A forward-looking Phillips curve is given by:

$$\pi_t = \beta \pi^e_{t+1} + \alpha y_t + \varepsilon_t \tag{1}$$

where $\pi_t$ is inflation at time $t$, $y_t$ is the output gap, and $\varepsilon_t$ is an i.i.d. shock process with variance $\sigma^2$. $\alpha$ and $\beta$ are positive parameters, with the latter also the social discount factor. This particular model can be derived from optimising behaviour (we refer to Clarida, Galí, and Gertler (1999) for details) and is often described as New Keynesian.[4] In anticipation of the results derived below, we assume that monetary policy is set using:

$$\pi_t = \theta_0 y_t + \theta_1 y_{t-1} \tag{2}$$

[4]A second dynamic IS relationship is sometimes added, but makes no material difference to the results because the Lagrange multiplier is always zero.

so that inflation is related to the output gap and its lag.[5]

We assume that the model can be solved using an undetermined coefficient method.[6] Write:

$$\pi_t = s_{11}y_{t-1} + s_{12}\varepsilon_t \tag{3}$$

$$y_t = s_{21}y_{t-1} + s_{22}\varepsilon_t \tag{4}$$

Lead (3) and take expectations to give $\pi^e_{t+1} = s_{11}y_t$. Using this and (4), the undetermined coefficient equation for $y_t$, in (1) gives:

$$\pi_t = (\beta s_{11} + \alpha)s_{21}y_{t-1} + ((\beta s_{11} + \alpha)s_{22} + 1)\varepsilon_t \tag{5}$$

and substituting (4) in (2) gives:

$$\pi_t = (\theta_0 s_{21} + \theta_1)y_{t-1} + \theta_0 s_{22}\varepsilon_t. \tag{6}$$

To solve for the rational expectations equilibrium we can equate coefficients across equations. Doing so between (3) and (6) gives:

$$s_{12} = \theta_0 s_{22} \tag{7}$$

and:

$$s_{11} = \theta_0 s_{21} + \theta_1 \tag{8}$$

and between (3) and (5), using the solution for $s_{12}$ just derived, gives:

$$s_{22} = \frac{1}{\theta_0 - \alpha - \beta s_{11}} \tag{9}$$

and:

$$s_{21} = \frac{\gamma + \sqrt{\gamma^2 + 4\beta\theta_0\theta_1}}{2\beta\theta_0} \tag{10}$$

where $\gamma = \theta_0 - \alpha - \beta\theta_1$.

For what follows we need to derive the unconditional variances of $y$ and $\pi$ so that we can compare them across policy regimes. Note that from (4):

$$E(y_t^2) = s_{21}^2 E(y_{t-1}^2) + s_{22}^2 E(\varepsilon_t^2) \tag{11}$$

[5]It is perhaps more natural to treat the output gap as the policy instrument rather than the inflation rate, as is the case in Steinsson (2000). However, it is more convenient to preserve the inflation rate as the policy variable in common with the rest of the literature, in particular the static, Barro-Gordon approach. In Appendix B we discuss the form of implied 'simple rule' involved in each policy equilibrium and compare with the results derived by Levine and Currie (1987).

[6]In the numerical examples given below we use the Blanchard and Kahn (1980) method to calculate the expectational equilibria. This yields the same result as the undetermined coefficients solution.

as $E(y_{t-1}\varepsilon_t) = 0$. Let:

$$p = \frac{s_{22}^2}{1 - s_{21}^2} \tag{12}$$

which allows us to write $\sigma_y^2$ as:

$$\sigma_y^2 = p\sigma^2. \tag{13}$$

Using this and (3) we can write $\sigma_\pi^2$ as:

$$\sigma_\pi^2 = (ps_{11}^2 + s_{12}^2)\sigma^2. \tag{14}$$

If we consider a weight $\omega$ which indicates the degree of concern with output stabilisation, then a weighted average of the two asymptotic variances is:

$$c = \omega\sigma_y^2 + \sigma_\pi^2 = ((\omega + s_{11}^2)p + s_{12}^2)\sigma^2 = \left( \frac{(\omega + s_{11}^2)s_{22}^2}{1 - s_{21}^2} + s_{12}^2 \right) \sigma^2. \tag{15}$$

This could, of course, be interpreted as an appropriate asymptotic cost function. Indeed, it is often used as a comparison in the literature, either by the use of a 'trade-off' diagram, where $\sigma_\pi^2$ is plotted against $\sigma_y^2$, or as a direct report of the value of $c$ for a given policy. For example, candidate policy regimes can be compared where '...the asymptotic social loss function is evaluated under both precommitment and discretion. As Jensen (1999) and McCallum and Nelson (2000) have previously shown, precommitment achieves a lower value of the loss function than does discretion' (Walsh 2001, p. 4). It is the second part of this quote that we take as the starting point for the current paper and we turn to in the next section. We discuss whether this method of comparing equilibria is appropriate further once we derive the standard equilibria.

## 3 'Timeless perspective' and 'time consistent' optimal policies

In this section we derive the fully optimal and time consistent optimal policies for the model and discuss the timeless perspective modification. Such policies are usually derived by considering a discounted infinite horizon loss function that represents the social welfare function.

Compare (15) with the usual discounted quadratic loss function:

$$V_0 = \min_\pi E_0 \sum_{t=0}^{\infty} \beta^t \left( \pi_t^2 + \omega y_t^2 \right). \tag{16}$$

This is often taken as a representative social welfare loss function. See the discussion in, *e.g.*, Steinsson (2000). Optimisation proceeds usually by the method of Lagrange multipliers. Define:

$$H_k = \pi_k^2 + \omega y_k^2 + \lambda_k(\beta\pi_{k+1} + \alpha y_k + \varepsilon_k - \pi_k) \tag{17}$$

so that:

$$\hat{V}_0 = \min_{\pi} E_0 \sum_{t=0}^{\infty} \beta^t H_t \tag{18}$$

is the objective function (15) subject to the constraint (1).

First-order conditions for (18) are:

$$E_0(2\omega y_t + \alpha\lambda_t) = 0, \text{ for } t \geq 1, \tag{19}$$

$$E_0(2\pi_t - \lambda_t + \lambda_{t-1}) = 0, \text{ for } t \geq 1, \tag{20}$$

$$E_0(2\pi_0 - \lambda_0) = 0. \tag{21}$$

Given that (19) is a static relationship between $\lambda_t$ and $y_t$, so $\lambda$ can be eliminated from (20) and (21).[7] Thus we can write the 'timelessly optimal' (TP) policy as:

$$\pi_t = -\frac{\omega}{\alpha} y_t + \frac{\omega}{\alpha} y_{t-1} \tag{22}$$

for $t \geq 1$. In the initial period, notice from (21) that:

$$\pi_0 = -\frac{\omega}{\alpha} y_0 \tag{23}$$

so that a different policy rule is implemented when the policy is announced. The time inconsistency argument is nicely illustrated by the two-part policy rule. There is a clear incentive to renege in each period: It is better to announce a new regime is in place, and that in the current period (23) will be implemented and (22) will be followed for all subsequent periods. For this to be optimal it must be believed; agents recognise this is not the only time the policymaker will face such an incentive; the policy is time inconsistent.

By contrast a time consistent (TC) policy can be obtained by always setting the Lagrange multiplier associated with the forecast to zero, which gives:

$$\pi_t = -\frac{\omega}{\alpha} y_t \tag{24}$$

for all $t$ which is a policy which simply minimises the intra-period cost without regard to the dynamics of the problem. Note that this is simply (23) implemented in every period.

[7] The elimination of the Lagrange multiplier in this way is possible because of the lack of structural dynamics except through expectations. More generally the elimination requires a discounted feedback on each lag of the predetermined state. This was shown by Levine and Currie (1987), and more recently by Svensson and Woodford (2000).

We now turn to the stochastic implications of potential policy rules. Pearlman (1992) showed that the optimal but time inconsistent policy for a class of linear rational expectations models with quadratic objective function is *ex ante* certainty equivalent.[8] We need to be careful about the interpretation of this statement. At time $t$ the optimal time inconsistent policy expressed as a feedback rule is certainty equivalent. Thus in the face of shocks unknown at time $t$ there is no policy that can be implemented that is *expected* to be better given a feedback representation of the policy rule. However, as we have just shown, such a policy from the perspective of time $t+1$ *cannot* be optimal as it is time inconsistent.

By contrast, the timeless perspective (TP) policy as suggested by Woodford (1999a) ignores the constraint on the initial condition and implements (22) in every period. There is therefore a symmetry about the TP and TC policies. One ignores the first-order condition (21), the TP policy, and one ignores (20), the TC policy.

We need to motivate why 'timelessness' is an attractive equilibrium concept in a stochastic context. A possible description of the *intention* of a TP policy is that it is set to minimise:

$$V = \min_{\pi} E \sum_{t=0}^{\infty} \beta^t \left( \pi_t^2 + \omega y_t^2 \right) \tag{25}$$

where the *unconditional* expectation of the loss function is considered. Note that some authors have described it as the policy that minimises $V_{t+k}$ for some $k$ such that the state is in stochastic equilibrium. This treatment of initial conditions is described by King and Wolman (1999), who '... consider the behavior of an economy after the effects of an initial "start-up" period have worn away.' (King and Wolman 1999, p. 376).[9] This, of course, amounts to assuming not that the initial conditions are *ignored*, but rather that they are treated as stochastic.[10]

This has the implication that the initial state cannot be relevant to any cost comparison and so the stochastic equilibrium must be compared. Note that:

$$V_{0,\ k} = E_0 V_k = \sum_{t=k}^{\infty} \beta^t c = \frac{\beta^k}{1-\beta} c \tag{26}$$

for $k$ large enough that the initial conditions have indeed 'worn away'. Thus the relevant costs are proportional to $c$ and we can compare the value of

[8] Similar results are derived in Aoki (1998), Svensson and Woodford (2000) and Svensson and Woodford (2001).

[9] Of course, in a deterministic context any policy rule that stabilises the target variables at their target values is as good as any other, given the definition of timelessness.

[10] This distinction is much more obvious for a more general control problem where not just one lag appears in the policy rule. See Appendix C for details.

$c$ obtained for each policy regime directly. We note below the important implications of discounting on the optimal policy.[11]

An immediate consequence of this is that the policy derived by ignoring the first-order condition associated with the very first period is not necessarily the policy that minimises (15). It is, after all, time inconsistent, and is necessarily suboptimal in subsequent periods, particularly in a stochastic equilibrium where continuing shocks mean that the minimum of the objective function is never reached, despite being certainty equivalent at $t = 0$. It may be that the TC policy could dominate. This is the same argument as whether the time inconsistent policy is sustainable. After some time has passed, if the 'cost-to-go' for the time inconsistent policy is greater than for a candidate time consistent one, then the time inconsistent policy is obviously unsustainable from then onwards. Indeed, it is unclear why a policymaker would try to sustain it: Switching to the TC policy at that point would improve welfare. Such effects are particularly bad news for the TP policy, which is by definition the time inconsistent policy after 'sufficient' time has elapsed.

Table 1: Values of $c \times 10^5$ for $\beta = 0.975$

| | | $\omega$ | | | | |
|---|---|---|---|---|---|---|
| $\alpha$ | Policy | 0.01 | 0.10 | 0.25 | 0.50 | 1.00 |
| 0.005 | TP | 2.433 | 2.512 | 2.522 | 2.527 | 2.529 |
| | TC | 2.494 | 2.499 | 2.500 | 2.500 | 2.500 |
| 0.010 | TP | 2.315 | 2.477 | 2.503 | 2.515 | 2.522 |
| | TC | 2.475 | 2.498 | 2.499 | 2.500 | 2.500 |
| 0.020 | TP | 2.094 | 2.402 | 2.457 | 2.485 | 2.503 |
| | TC | 2.404 | 2.490 | 2.496 | 2.498 | 2.499 |
| 0.030 | TP | 1.894 | 2.327 | 2.409 | 2.451 | 2.481 |
| | TC | 2.294 | 2.478 | 2.491 | 2.496 | 2.498 |

It turns out that it is not possible to show that the TP policy is always better than the TC policy for the simple reason that it is not. We give some representative numbers in Table 1 for $\sigma = 0.005$ and $\beta = 0.975$. In Figure 1 we plot points where $c$ is equal for the TP and TC policies, as functions of the three parameters, $\alpha$, $\omega$ and $\beta$. Points above the surface are where the TP policy is superior and below the TC is superior. The range and scale are carefully chosen to include points where each policy dominates. If we consider an arbitrary value for $\alpha$, then a greater concern

[11]For example, Walsh (2001) reports $c/(1-\beta)$ as his measure of asymptotic cost. Interestingly, this is never a feasible value of the expected cost. We discuss this further in Appendix C.

Figure 1: Points of equal 'cost' for TP and TC policies

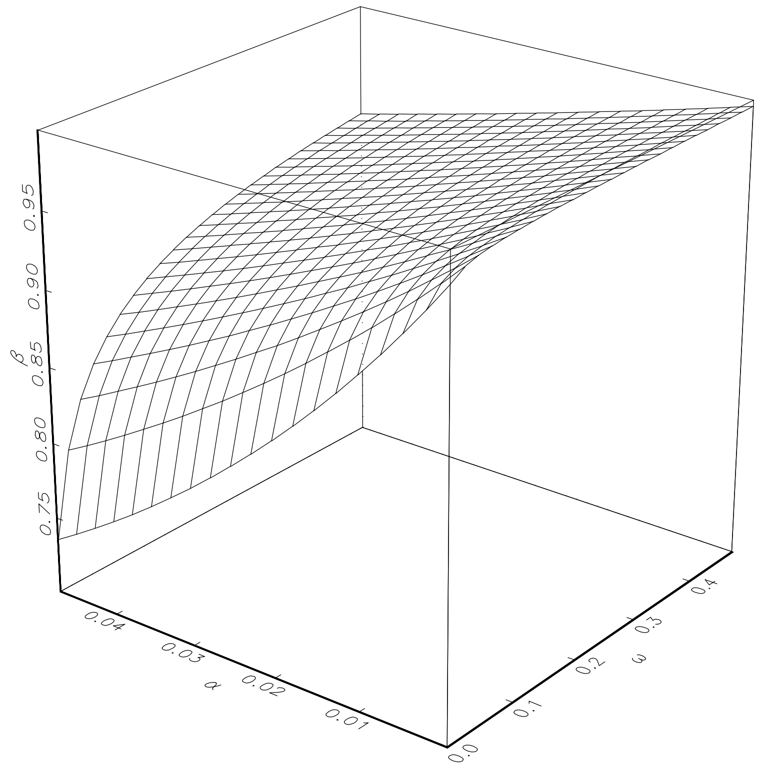


with output stabilisation (greater $\omega$) or a lower discount factor favours the TC policy. The points of equality on the graph are close to what might be regarded as empirically relevant, as they are only slightly outside experience, in particular requiring fairly low values of $\alpha$ for reasonable $\beta$.[12]

## 4 Timelessness and optimality

If neither the TP or TC policies are globally dominant, a logical question to ask is what constitutes the truly optimal policy? The optimal policy to minimise (15) turns out to be very simple and intuitive. However, to find the minimum to (15) directly—even for this simple model—turns out to be very unwieldy analytically. For example, we can substitute in to rewrite

[12] By contrast McCallum and Nelson (2000) do not find any points where the TC policy dominates. However, they do not choose a value for $\alpha$ that is small enough or, crucially, vary the discount factor.

(15) as:

$$c = \left[ \frac{\omega + \theta_0^2 + \theta_1^2 + \frac{\theta_1}{\beta}}{\left(\gamma - \frac{\phi}{2}\right)^2 \left(1 - \frac{1}{\theta_0^2}\left(\frac{\phi}{2\beta}\right)^2\right)} \right] \sigma^2 \tag{27}$$

where:

$$\phi = \gamma + \sqrt{\gamma^2 + 4\beta\theta_0\theta_1}$$

with $\gamma$ as before. First-order conditions with respect to $\theta_0$ and $\theta_1$ can be derived. However, even a computer algebra tool such as *Maple* (see, *e.g.*, Heck 1996) is of more use for checking a solution rather than deriving it. This we do in Appendix A. In what follows we find the solution to a related problem and argue that it solves the original one and establish it in the appendix. We, of course, use numerical examples to illustrate the solution.

Remember that the timeless perspective policy is analytically straightforward and certainty equivalent from the perspective of $t = 0$ (although, of course, not optimal). If the approach is to find an alternative policy that solves the original problem more tractably then it is sensible to start with a timelessly optimal one. Consider the *undiscounted* constrained control problem:

$$\tilde{V}_0 = \min_{\pi} E_0 \sum_{t=0}^{\infty} H_t \tag{28}$$

where $H_t$ is defined as before. What if we consider this from a 'timeless' perspective? Clearly, the value of any loss function for a stochastic control problem is usually infinity without discounting. However, if we consider a finite horizon $T$-period problem with arbitrary initial and terminal periods and seek to minimise the expected value in each period we will obtain the required optimal policy that minimises (15). Note that minimising the quadratic loss function is then equivalent to minimising $T \times c$ over that interval for stochastic initial conditions and so must minimise $c$. The TP-based welfare loss is (by definition) proportional to the weighted steady-state variances but it is not *necessarily* the optimal policy that minimises the weighted steady state variance. Neither is the TC policy. However, under the comparison made above (really the only meaningful comparison that we can make) the undiscounted policy must dominate both. Notice that a policy calculated *with* discounting necessarily concentrates the action of that policy to the immediate future at the expense of the stochastic equilibrium. Remember, we are essentially trying to solve (26) for arbitrary $k$ and therefore eliminating the short-run policy effects.

First-order conditions for (28) are (19) and (21) but instead of (20) we obtain:

$$E_0(2\pi_t - \lambda_t + \beta\lambda_{t-1}) = 0 \tag{29}$$

for $t \geq 1$. That $\beta$ appears in the undiscounted problem and not in the discounted one reflects its role as a parameter in the model (1). From a 'timeless' perspective, we would again ignore (21) and assume that policy is set the same way in every period. This yields the optimal (OP) policy rule:

$$\pi_t = -\frac{\omega}{\alpha} y_t + \beta \frac{\omega}{\alpha} y_{t-1}. \tag{30}$$

In Appendix A we show how *Maple* can be used to check that this is indeed the minimum. Having derived three possible policies to be implemented we compare them on the basis of their asymptotic welfare losses in Table 2. Notice that the OP, TP, and TC feedback rules in the form of (2) can be written:

$$\theta_0 = -\frac{\omega}{\alpha} \tag{31}$$

and:

$$\theta_1 = \mu \frac{\omega}{\alpha} \tag{32}$$

where $\mu = 1$, for the TP policy, $\mu = 0$ for the TC policy and $\mu = \beta$ for the OP policy. Given this parameterisation it is apparent that a lower discount factor should favour the TC policy over the TP as the superior OP policy has less lagged feedback. Note that for $\beta = 1$ the OP policy and the TP policy coincide. Equally, for $\beta = 0$ the problem is static and all three solutions coincide. Note that this reflects the role of the discount factor as a parameter of the model.

We illustrate the three policies with numerical examples. In Table 2 we give the values for the OP, TP and TC policies for several values of $\alpha$ and $\omega$ with $\beta = 0.98$ and $\sigma = 0.005$, differing from Table 1 by the larger discount factor which gives a greater relative advantage to the timelessly optimal policy. For this model and the reported parameterisation the gains for OP over TP are modest and are certainly smaller than the proportionate gains available for the TP policy over the TC policy, *e.g.* for $\alpha = 0.1$, $\omega = 0.1$. Of course, the OP policy reflects this advantage, having at least as much gain.

The existence of this policy raises a number of important issues. Woodford (1999a, p. 294) argues that any policymaker using a 'timeless' perspective '...chooses to act as one *would have* wished to commit oneself to act at a date far in the past, not as one *actually did* commit oneself to act at any such distant past date'. Taken together with the constraint of an expectational equilibrium, this is a perfect description of time consistency. Woodford suggests that the TP policy fulfils these criteria. However, we have found a policy which dominates it in stochastic equilibrium. Thus the TP policy rule *cannot* be the policy 'one would have wished to commit oneself to': The OP policy is. Only this policy is the one to which a central bank would still choose to commit after recomputing an optimal strategy from a timeless perspective (see Woodford 1999a, pp. 293–294).

Table 2: Values of $c \times 10^5$ for $\beta = 0.98$

| $\alpha$ | Policy | $\omega$ 0.01 | 0.10 | 0.25 | 0.50 | 1.00 |
|---|---|---|---|---|---|---|
| | OP | 2.416 | 2.486 | 2.494 | 2.497 | 2.498 |
| 0.005 | TP | 2.423 | 2.503 | 2.515 | 2.519 | 2.522 |
| | TC | 2.494 | 2.499 | 2.500 | 2.500 | 2.500 |
| | OP | 2.301 | 2.456 | 2.479 | 2.489 | 2.494 |
| 0.010 | TP | 2.305 | 2.467 | 2.494 | 2.507 | 2.515 |
| | TC | 2.475 | 2.498 | 2.499 | 2.500 | 2.500 |
| | OP | 1.547 | 2.172 | 2.301 | 2.369 | 2.416 |
| 0.050 | TP | 1.547 | 2.174 | 2.305 | 2.373 | 2.423 |
| | TC | 2.000 | 2.439 | 2.475 | 2.488 | 2.494 |
| | OP | 0.965 | 1.855 | 2.083 | 2.209 | 2.301 |
| 0.100 | TP | 0.966 | 1.856 | 2.084 | 2.211 | 2.305 |
| | TC | 1.250 | 2.273 | 2.404 | 2.451 | 2.475 |

Of course, the policy (30) is very specific to this problem and alternative 'timeless perspective' optimal policies need to be computed for a wider range of models. It may be that the gap between the equivalent TP and OP policies might be much greater, or indeed the frequency that it is bettered by a time-consistent alternative may be much greater. The main point to emphasise is this: The TP policy is not *necessarily* even a good policy, let alone the best.

It also must be noted that this policy does not altogether convincingly solve the time inconsistency problem. At any given time an incentive exists to implement the optimal time inconsistent policy to minimise (16) from the perspective of that time. We require agents to support the OP policy as much as the TP policy, in the same way that Barro and Gordon (1983) argue that an *almost* optimal policy can be supported. They argue that both agents and the policymaker recognise that agents will punish time inconsistent actions. The mechanism assumed is that agents impose the requirement that an inferior time consistent policy is followed for a period once a policymaker is observed to renege on a prior promised policy. Of course, greater support exists for the OP policy than for the TP policy, as there is more to lose because the OP policy derived here is the global optimum to the policy problem. If the social welfare function *is* represented by the stochastic equilibrium, then the incentive to renege is absent. This is the support that agents need to give to the OP policy to sustain it as an equilibrium: They need to believe that the monetary policymaker will eschew all short-run gains on the basis of (16) and instead will always act to minimise (15).

## 5 Conclusions

In this paper we have evaluated a policy proposal for monetary policy that has gained considerable popularity recently in the context of a simple dynamic rational expectations model.

We have established two results for our simple model, both of which generalise. Firstly, it is clear that neither the TP or TC policy necessarily dominates on the basis of the asymptotic cost function, with the points at which they are equal shown in Figure 1. Secondly, there is an alternative 'timeless perspective' optimal policy, not one of the standard equilibria, that is better. As it must be that the time inconsistent policy dominates both from the perspective of the origin of that policy (*i.e.* $t = 0$) it may also be that there is a policy that dominates them from the same perspective used to derive the TP rule. This turns out to be the solution to the undiscounted 'timeless' control problem.

Two sets of constraints *ensure* that a policy is time consistent. Firstly, that the policy is optimal subject to the rational expectation implied by that policy[13] and, secondly, that it is the best policy taking into account the reactions of agents. This second constraint often corresponds to taking bygones as bygones, so that a continuing policy is required to be independent of what went before. Thus the dynamic programming principle that minimises the 'cost-to-go' stage-wise is adhered to.[14] This would seem to rule out the timeless perspective equilibrium as a potentially time consistent policy: No such constraint is imposed.

However, by concentrating on the asymptotic loss function the policy problem is translated to something akin to a static one. The same (*i.e.* single set of) first order conditions are required to be satisfied in all periods by the 'timeless perspective' constraint. As demonstrated above, there *is* a best policy that satisfies this constraint that is not necessarily derived from the optimal time inconsistent policy. Nevertheless, it *is* time inconsistent in exactly the same way as the static Barro and Gordon (1983) problem. If the monetary policymaker treats the expected inflation rate next period as parametric, then in the current period reoptimisation can yield a better outcome with a lower expected welfare loss achieved. For the model in the text, because the problem reduces to a static one for exogenous expectations, the best time consistent *and* timeless perspective optimal policy is simply the TC policy.

[13]This, incidentally, is what rules out the policy often describes as 'perfect cheating' where the policymaker plans to renege each period and implement the optimal time inconsistent policy afresh each period and is believed by agents. This violates the rational expectations assumption and is exactly what makes the policy time inconsistent in the first place.

[14]We should be careful about our definition of time consistency. As Fershtman (1989) showed, dynamic programming is not required to ensure time consistency, but the stronger condition of subgame perfection. For present purposes we take these to be equivalent.

Ultimately the importance of time inconsistency is an empirical question. That we can find a superior policy to the TP policy may only of marginal importance relative to the overall improvement in comparison with a time consistent policy. The ability to commit to a rule—almost any rule—might be the most important element. However, the fact that the OP policy dominates both standard equilibria may indicate that richer models could admit wider classes of potentially good rules that improve on discretion whilst retaining simplicity and transparency.

# A First-order conditions

In this appendix we outline first-order conditions for the direct optimisation of (15) using *Maple*. These are included for completeness. Given the OP solution derived above we can parameterise either $\theta_0$ or $\theta_1$ and differentiate with respect to the other and evaluate the resulting expressions at the assumed optimum. A *Maple* worksheet that does this is available on request.

As we have two policy parameters, and they are functions of each other, we can set $\theta_1 = -\mu\theta_0$, and then check the optimality conditions with respect to $\mu$ and $\theta_0$ to verify the solution. If we initially assume that $\theta_1 = -\beta\theta_0$, we can differentiate $c$ with respect to $\theta_0$ and verify that at $\theta_0 = -\frac{\omega}{\alpha}$ the first order condition is zero. The following output from a *Maple* session illustrates this.

Firstly, we need to set up the model. We first set up equations 8, 9, 7 and 10 for an arbitrary policy rule:

```
>   s[11] := theta[0]*s[21] + theta[1];
```

$$s_{11} := \theta_0 s_{21} + \theta_1$$

```
>   s[22] := 1/(theta[0]-alpha-beta*s[11]);
```

$$s_{22} := \frac{1}{\theta_0 - \alpha - \beta(\theta_0 s_{21} + \theta_1)}$$

```
>   s[12] := theta[0]*s[22];
```

$$s_{12} := \frac{\theta_0}{\theta_0 - \alpha - \beta(\theta_0 s_{21} + \theta_1)}$$

```
>   g := theta[0]-alpha-beta*theta[1];
```

$$g := \theta_0 - \alpha - \beta\theta_1$$

```
>   s[21] := (g+(g^2+4*beta*theta[0]*theta[1])^(1/2))/(2*beta*theta[0]);
```

$$s_{21} := \frac{1}{2} \frac{\theta_0 - \alpha - \beta\theta_1 + \sqrt{(\theta_0 - \alpha - \beta\theta_1)^2 + 4\beta\theta_0\theta_1}}{\beta\theta_0}$$

The value of $p$ in (12) is then:

```
>   p := s[22]^2/(1-s[21]^2);
```

$$p := 1 \Big/ \left( \left( \theta_0 - \alpha - \beta \left( \frac{1}{2} \frac{\theta_0 - \alpha - \beta\theta_1 + \sqrt{(\theta_0 - \alpha - \beta\theta_1)^2 + 4\beta\theta_0\theta_1}}{\beta} + \theta_1 \right) \right)^2 \left( 1 - \frac{1}{4} \frac{\left(\theta_0 - \alpha - \beta\theta_1 + \sqrt{(\theta_0 - \alpha - \beta\theta_1)^2 + 4\beta\theta_0\theta_1}\right)^2}{\beta^2\theta_0^2} \right) \right)$$

and the cost function (15) is:

```
>   c := (omega + s[11]^2)*p + s[12]^2;
```

$$c := \frac{\omega + \left(\frac{1}{2}\frac{\%1}{\beta} + \theta_1\right)^2}{\left(\theta_0 - \alpha - \beta\left(\frac{1}{2}\frac{\%1}{\beta} + \theta_1\right)\right)^2 \left(1 - \frac{1}{4}\frac{\%1^2}{\beta^2\theta_0^2}\right)} + \frac{\theta_0^2}{\left(\theta_0 - \alpha - \beta\left(\frac{1}{2}\frac{\%1}{\beta} + \theta_1\right)\right)^2}$$

$$\%1 := \theta_0 - \alpha - \beta\theta_1 + \sqrt{(\theta_0 - \alpha - \beta\theta_1)^2 + 4\beta\theta_0\theta_1}$$

Notice that $c$ should really be multiplied by $\sigma^2$, but we can drop this as it is exogenously given. Set $\theta_1$ to be related to $\theta_0$ by our assumed solution:

```
>   theta[1] := -beta*theta[0];
```

$$\theta_1 := -\beta\theta_0$$

and evaluate the first order condition with respect to $\theta_0$:

```
>   dct0 := diff(c, theta[0]);
```

$$dct0 := 2\frac{\%5\left(\frac{1}{2}\frac{\%3}{\beta} - \beta\right)}{\%6^2\left(1 - \frac{1}{4}\frac{\%4^2}{\beta^2\theta_0^2}\right)} - 2\frac{(\omega + \%5^2)\left(1 - \beta\left(\frac{1}{2}\frac{\%3}{\beta} - \beta\right)\right)}{\%6^3\left(1 - \frac{1}{4}\frac{\%4^2}{\beta^2\theta_0^2}\right)} - \frac{(\omega + \%5^2)\left(-\frac{1}{2}\frac{\%4\%3}{\beta^2\theta_0^2} + \frac{1}{2}\frac{\%4^2}{\beta^2\theta_0^3}\right)}{\%6^2\left(1 - \frac{1}{4}\frac{\%4^2}{\beta^2\theta_0^2}\right)^2} + 2\frac{\theta_0}{\%6^2} - 2\frac{\theta_0^2\left(1 - \beta\left(\frac{1}{2}\frac{\%3}{\beta} - \beta\right)\right)}{\%6^3}$$

$$\%1 := \theta_0 - \alpha + \beta^2\theta_0$$

$$\%2 := \%1^2 - 4\beta^2\theta_0^2$$

$$\%3 := 1 + \beta^2 + \frac{1}{2}\frac{2\%1(1 + \beta^2) - 8\beta^2\theta_0}{\sqrt{\%2}}$$

$$\%4 := \theta_0 - \alpha + \beta^2\theta_0 + \sqrt{\%2}$$

$$\%5 := \frac{1}{2}\frac{\%4}{\beta} - \beta\theta_0$$

$$\%6 := \theta_0 - \alpha - \beta\%5$$

and substitute in the hypothesised solution:

```
>  theta[0] := -omega/alpha;
```

$$\theta_0 := -\frac{\omega}{\alpha}$$

which simplifies to:

```
>  simplify(dct0);
```

$$0$$

confirming the optimum if $\theta_0 = -\frac{\omega}{\alpha}$.

Similarly, if we assume that $\theta_1 = -\mu\theta_0$, we can differentiate $c$ with respect to $\mu$ and verify that at $\mu = \beta$ the first order condition is zero. The following *Maple* session does this:

```
>  s[11] := theta[0]*s[21] + theta[1];
>  s[22] := 1/(theta[0]-alpha-beta*s[11]);
>  s[12] := theta[0]*s[22];
>  g := theta[0]-alpha-beta*theta[1];
>  s[21] := (g+(g^2+4*beta*theta[0]*theta[1])^(1/2))/(2*beta*theta[0]);
>  p := s[22]^2/(1-s[21]^2);
>  c := (omega + s[11]^2)*p + s[12]^2;
```

This gives the same output as above. Now instead we set:

```
>  theta[0] := -omega/alpha;
```

$$\theta_0 := -\frac{\omega}{\alpha}$$

and set the parameter on the lagged output value as:

```
>  theta[1] := -mu*theta[0];
```

$$\theta_1 := \frac{\mu\omega}{\alpha}$$

where $\mu$ is a dummy parameter which we expect to be equal to $\beta$ at the optimum. Taking the derivative with respect to the 'dummy' parameter gives:

```
> dcmu := diff(c, mu);
```

$$dcmu := 2\frac{\%5\left(\frac{1}{2}\frac{\%3}{\beta}+\frac{\omega}{\alpha}\right)}{\%6^2\left(1-\frac{1}{4}\frac{\%4^2\alpha^2}{\beta^2\omega^2}\right)}+2\frac{(\omega+\%5^2)\beta\left(\frac{1}{2}\frac{\%3}{\beta}+\frac{\omega}{\alpha}\right)}{\%6^3\left(1-\frac{1}{4}\frac{\%4^2\alpha^2}{\beta^2\omega^2}\right)}$$
$$+\frac{1}{2}\frac{(\omega+\%5^2)\%4\alpha^2\%3}{\%6^2\left(1-\frac{1}{4}\frac{\%4^2\alpha^2}{\beta^2\omega^2}\right)^2\beta^2\omega^2}+2\frac{\omega^2\beta\left(\frac{1}{2}\frac{\%3}{\beta}+\frac{\omega}{\alpha}\right)}{\alpha^2\%6^3}$$
$$\%1 := -\frac{\omega}{\alpha}-\alpha-\frac{\beta\mu\omega}{\alpha}$$
$$\%2 := \%1^2-4\frac{\beta\omega^2\mu}{\alpha^2}$$
$$\%3 := -\frac{\beta\omega}{\alpha}+\frac{-\frac{\%1\beta\omega}{\alpha}-2\frac{\beta\omega^2}{\alpha^2}}{\sqrt{\%2}}$$
$$\%4 := -\frac{\omega}{\alpha}-\alpha-\frac{\beta\mu\omega}{\alpha}+\sqrt{\%2}$$
$$\%5 := \frac{1}{2}\frac{\%4}{\beta}+\frac{\mu\omega}{\alpha}$$
$$\%6 := -\frac{\omega}{\alpha}-\alpha-\beta\%5$$

Now set $\mu$ equal to the hypothesised optimum:

```
> mu := beta;
```

$$\mu := \beta$$

with the final simplification again returned as:

```
> simplify(dcmu);
```

$$0$$

confirming the optimum.

# B An optimal simple rule interpretation

As a brief digression, which resolves a number of conflicts in the literature, we discuss various *simple rule* interpretations of the optimal policies. Simple rules have over the years been proposed by a many authors for many reasons.

However, we specifically have in mind here the literature on Taylor rules (see, *e.g.*, Taylor 1999, Svensson 2001). This literature often considers only simple proportional rules. We consider more general proportional-plus-integral-plus-derivative (PID) control rules of the form:

$$u_t = \kappa_P e_t + \kappa_I \sum_{i=0}^{\infty} e_{t-i} + \kappa_D \Delta e_t \tag{33}$$

or in differenced form as:

$$\Delta u_t = \kappa_P \Delta e_t + \kappa_I e_t + \kappa_D \Delta^2 e_t \tag{34}$$

where $u$ is the control instrument, $e$ the deviation of the target variable from target value, $\kappa_P$ is the so-called proportional control coefficient, $\kappa_I$ is the integral control coefficient and $\kappa_D$ is the derivative control coefficient. These are needed to encompass the rules derived in this paper. Integral control is typically included to ensure that the target is met and derivative control to improve the stability properties of a rule. See Franklin, Powell, and Emami-Naeini (1994) for a typical treatment of feedback control.

As the model has one or two states, depending on the policy rule, we can interpret the optimal rules in the PID framework. The time consistent rule (24) is straightforwardly a purely proportional rule control rule, so that, in terms of (33):

$$\pi_t = \kappa_P y_t = -\frac{\omega}{\alpha} y_t \tag{35}$$

where there is no lagged state dependence. This reflects the usual requirement of time consistency that at any time $t$ optimality is calculated conditional on past behaviour being independent of policies implemented in the future. By contrast, the timeless perspective rule (22) is a purely *derivative* control rule:

$$\pi_t = \kappa_D \Delta y_t = -\frac{\omega}{\alpha} \Delta y_t \tag{36}$$

where the value of output lagged appears explicitly in the optimal policy. This has rightly been interpreted as an output growth rule, and inspired the analysis in Walsh (2001), for example. Our optimal rule (30) is a proportional-plus-derivative control rule:

$$\pi_t = -(1-\beta)\frac{\omega}{\alpha} y_t - \beta \frac{\omega}{\alpha} \Delta y_t. \tag{37}$$

Immediately there is an important departure here from previous literature, as this contrasts with the result in Levine and Currie (1987). They note that time inconsistent policies have a form of state-dependence which can be described as a form of integral control rule. Our time-inconsistent policies are by contrast described best as including *derivative* control elements. This still has a state-dependent aspect, but only through the path of the target rather than the control variable.

The resolution to this is the interpretation of the control instrument as $\pi$ whereas it should properly be $y$. We can simply invert (35) and (36) to yield the equivalent pure proportional and pure integral rules. (37) inverted does not yield an exact PID equivalent, but rather like (34):

$$(1-\beta L)y_t = -\frac{\alpha}{\omega}\pi_t \tag{38}$$

where $L$ is the lag operator. This is a modified integral control rule, very similar to the one found in Levine and Currie (1987).

# C The general LQG problem with stochastic initial conditions

In this appendix we discuss a more general stochastic linear rational expectations control problem. First we show how the asymptotic welfare loss can be described as an infinite sum. Write a general quadratic loss function as:

$$\begin{aligned} V_0 &= E_0 \sum_{t=0}^{\infty} \beta^t s_t' W s_t && (39) \\ &= E_0 \sum_{t=0}^{\infty} \beta^t \mathrm{tr}(s_t s_t' W) && (40) \\ &= \sum_{t=0}^{\infty} \beta^t \mathrm{tr}(P_t W) && (41) \end{aligned}$$

where $s_t$ is an $n$-vector of state variables, $W$ is a positive semi-definite weighting matrix, $\beta$ is the discount factor, $P_t = E_0[s_t s_t']$ is the forecast state variance, and we have used the result that $\mathrm{tr}(ABC) = \mathrm{tr}(BCA) = \mathrm{tr}(CAB)$ for conformable matrices. Now assume that we take expectations at the same date but start the cost calculation later, say $k > 0$. Write this:

$$\begin{aligned} V_{0,\ k} &= E_0 \sum_{t=k}^{\infty} \beta^t s_t' W s_t && (42) \\ &= \sum_{t=k}^{\infty} \beta^t \mathrm{tr}(P_t W). && (43) \end{aligned}$$

Couple this with a model:

$$s_{t+1} = A s_t + C\varepsilon_{t+1} \tag{44}$$

where $\varepsilon_t \sim N(0,\ \Omega)$. We return to parameterising the model later. Note that when $k$ is large enough, say $\bar{k}$, the state variance of $s_t$ for all $t \geq \bar{k}$ is such that:

$$\bar{P} = A\bar{P}A' + C\Omega C' \tag{45}$$

its asymptotic value, so we can write:

$$V_{0,\ \bar{k}} = \sum_{t=\bar{k}}^{\infty} \beta^t \mathrm{tr}(\bar{P}W) \tag{46}$$

$$= \frac{\beta^{\bar{k}}}{1-\beta} \mathrm{tr}(\bar{P}W). \tag{47}$$

This is necessarily where initial conditions have 'worn away'. The 'cost-to-go' from $\bar{k}$ onwards is therefore proportional to $\mathrm{tr}(\bar{P}W)$ with the proportion depending on the discount factor, $\beta$. Compare this with:

$$V = \frac{1}{1-\beta} \mathrm{tr}(\bar{P}W)$$

which is equivalent to the value reported by Walsh (2001). For this to be a feasible value of the loss function to report it would require that the inherited state was the minimum cost in future periods. If there is discounting this cannot be the case. The future stochastic equilibrium is less costly than current state-disequilibrium so this is reduced at the cost of greater potential variation in the state in the long term.

We turn to the derivation of the optimal policies described in the text. The problem is to find the optimal policy to minimise $V_{0,\ \bar{k}}$ subject to:

$$\begin{bmatrix} z_{t+1} \\ x^e_{t+1} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} z_t \\ x_t \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} u_t + \begin{bmatrix} \varepsilon_{t+1} \\ 0 \end{bmatrix} \tag{48}$$

where:

$$s_t = \begin{bmatrix} C_1 & C_2 \end{bmatrix} \begin{bmatrix} z_t \\ x_t \end{bmatrix} + Du_t. \tag{49}$$

In Pearlman (1992) and Svensson and Woodford (2000) it is demonstrated how the (time inconsistent) commitment solution is certainty equivalent. We can therefore solve the deterministic problem to find the optimal stochastic problem.

Write:

$$\begin{aligned} C_t &= y_t'C'WCy_t + 2y_t'C'WDu_t + u_t'D'WDu_t \\ &= y_t'Qy_t + 2y_t'Uu_t + u_t'Ru_t \end{aligned}$$

and define the Hamiltonian:

$$H_t = \beta^t C_t + 2\mu'_{t+1}(Ay_t + Bu_t - y_{t+1}) \tag{50}$$

so that:

$$\hat{V}_0 = \sum_{t=0}^{\infty} H_t. \tag{51}$$

First-order conditions are:

$$\frac{1}{2}\left(\frac{\partial H_t}{\partial y_t}+\frac{\partial H_{t-1}}{\partial y_t}\right) = \beta^t(Qy_t+Uu_t)+A'\mu_{t+1}-\mu_t=0 \quad (52)$$

$$\frac{1}{2}\frac{\partial H_t}{\partial u_t} = \beta^t(U'y_t+Ru_t)+B'\mu_{t+1}=0 \quad (53)$$

and the model (48).

The Hamiltonian for the *undiscounted* control problem is, of course:

$$\tilde{H}_t = C_t + 2\tilde{\mu}'_{t+1}(Ay_t+Bu_t-y_{t+1}) \quad (54)$$

with necessary first-order conditions are:

$$\frac{1}{2}\left(\frac{\partial \tilde{H}_t}{\partial y_t}+\frac{\partial \tilde{H}_{t-1}}{\partial y_t}\right) = Qy_t+Uu_t+A'\tilde{\mu}_{t+1}-\tilde{\mu}_t=0 \quad (55)$$

$$\frac{1}{2}\frac{\partial H_t}{\partial u_t} = U'y_t+Ru_t+B'\tilde{\mu}_{t+1}=0 \quad (56)$$

and the model (48).

For either control problem, the solution can be expressed as:

$$u_t = -F_z z_t - F_\mu \mu_t^2 \quad (57)$$

$$x_t = -N_z z_t - N_\mu \mu_t^2 \quad (58)$$

$$\mu_{t+1}^2 = \Gamma_z z_t + \Gamma_\mu \mu_t^2 \quad (59)$$

see Currie and Levine (1993). The predetermined states and co-states under control can then be written:

$$z_{t+1} = \Phi_z z_t + \Phi_\mu \mu_t^2 \quad (60)$$

$$\mu_{t+1}^2 = \Gamma_z z_t + \Gamma_\mu \mu_t^2 \quad (61)$$

where $\Phi_z = A_{11} - A_{12}N_z - B_1F_z$ and $\Phi_\mu = -B_1N_\mu - B_1F_\mu$. Equations (60) and (61) are equivalent to (44).

In Levine and Currie (1987) it was pointed out that $\mu$ could be solved out as a function of the predetermined state given that $\mu_0^2 = 0$ to give:

$$\mu_t^2 = \sum_{j=0}^{\infty}\Gamma_\mu^j\Gamma_z z_{t-j-1}$$

so that:

$$u_t = -F_z z_t - F_\mu \sum_{j=0}^{\infty}\Gamma_\mu^j\Gamma_z z_{t-j-1}. \quad (62)$$

This is similar to (33) and for the scalar case can be written as a generalisation of the representative policy (38).

Note that the cost of the optimal control can be written as:

$$C_0 = s_0' S s_0 = \mathrm{tr}(\Sigma_0 S)$$

where $S$ is the value function and $\Sigma_0 = s_0 s_0'$. The value function works by evaluating the 'cost-to-go' associated with a given set of initial conditions. If the stochastic problem is certainty equivalent then the same control can be used. If the shocks appear additively as disturbances to the future state then their discounted expected value in each period is the expected disturbance to the state, and so the cost is written:

$$C_0 = \mathrm{tr}((\Sigma_0 + \frac{\beta}{1-\beta}\hat{\Omega})S)$$

where $\hat{\Omega} = C\Omega C'$. Note that similar to $V_{0,\ \bar{k}}$ above we can write:

$$C_{0,\ \bar{k}} = \frac{\beta^{\bar{k}}}{1-\beta}\mathrm{tr}(\hat{\Omega}S) \tag{63}$$

where, of course, initial conditions have been ignored.

We can compare (47) and (63) for a given control problem. $W$ is a sparse matrix of weights which has nothing to do with the specification of the model (other than the location of the weights) and $\bar{P}$ reflects both the model under control and the covariance structure of the shocks. $\hat{\Omega}$ mainly reflects the covariance structure of the shocks and static parts of the model under control and $S$ reflects all the non-stochastic parts—the model, control and loss function. The latter reflects the certainty equivalence of the control problem.

It should be clear now why (63) based on the discounted control problem is inferior to it based on the undiscounted one when applied to the asymptotic loss function. This is entirely due to the nature of the timeless optimality problem. The inherited state turns out to be important. It reflects a policy (that will be continued to be pursued in the future) where the asymptotic variance is is steady state. Because of this, a discounted control problem necessarily makes the focus of policy a rejection of disturbances at a faster rate than the non-discounted one—discounting means that disequilibria in the near future is more costly than further away. This is therefore likely to implement a control rule which minimises the short-term impacts of shocks at the cost of greater future variability. (63) for an *ex ante* optimal policy and the policy which minimises (47) need not coincide as this is *not* what certainty equivalence implies.

Practically, of course, the introduction of any timeless-perspective optimal rule, however calculated, is hugely problematic. Whilst the optimal policy rule for the simple model (1) can be seen as a straightforward growth targeting rule, that implied by (62) is a discounted lag function of the current state, the past values of which have evolved under a variety of different policy regimes. It is impossible to determine ahead of time what impact this would have—such a policy could be disastrous, depending on what policies had been followed in the past. However, implementing the time inconsistent policy initially, perhaps on the grounds that it would have predictable effects, is impossible. It is, after all, time inconsistent and we are back where we started.

## References


Aoki, K. (1998): "On the Optimal Monetary Policy Response to Noisy Indicators," Working Paper, Princeton University.

Barro, R., and D. Gordon (1983): "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12, 101–121.

Blanchard, O., and C. Kahn (1980): "The Solution of Linear Difference Models Under Rational Expectations," *Econometrica*, 48(5), 1305–1311.

Buiter, W. H., and R. C. Marston (eds.) (1985): *International Economic Policy Coordination*. Cambridge University Press.

Clarida, R. H., J. Galí, and M. Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

Cohen, D., and P. Michel (1988): "How Should Control Theory be Used to Calculate a Time-Consistent Government Policy?," *Review of Economic Studies*, 55, 263–274.

Currie, D., and P. Levine (1993): *Rules, Reputation and Macroeconomic Policy Coordination*. Cambridge University Press.

Fershtman, C. (1989): "Fixed Rules and Decision Rules: Time Consistency and Subgame Perfection," *Economics Letters*, 30, 191–194.

Franklin, G. F., J. D. Powell, and A. Emami-Naeini (1994): *Feedback Control of Dynamic Systems*. Addison Wesley, third edn.

Heck, A. (1996): *Introduction to Maple*. Springer-Verlag, second edn.

Jensen, H. (1999): "Targeting Nominal Income Growth or Inflation?," CEPR Discussion Paper No. 2167.

King, R. G., and A. L. Wolman (1999): "What Should the Monetary Authorities Do When Prices are Sticky?," in *Monetary Policy Rules*, ed. by J. B. Taylor. University of Chicago Press.

Kydland, F., and E. Prescott (1977): "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 473–492.

Levine, P., and D. Currie (1987): "The Design of Feedback Rules in Linear Stochastic Rational Expectations Models," *Journal of Economic Dynamics and Control*, 11, 1–28.

McCallum, B. T., and E. Nelson (2000): "Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models," NBER Working Paper No. 7915, http://www.nber.org/papers/w7915.

Pearlman, J. G. (1992): "Reputational and Nonreputational Policies Under Partial Information," *Journal of Economic Dynamics and Control*, 16(2), 339–357.

Söderlind, P. (1999): "Solution and Estimation of RE Macromodels with Optimal Policy," *European Economic Review*, 43, 813–823.

Steinsson, J. (2000): "Optimal Monetary Policy in an Economy with Inflation Persistence," Central Bank of Iceland, Working Papers, No. 11.

Svensson, L. E. (2001): "What Is Wrong With Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules," unpublished paper, available at: http://www.iies.su.se/leosven/papers/JEL.pdf.

Svensson, L. E., and M. Woodford (2000): "Indicator Variables for Optimal Policy," NBER Working Paper No. W7953, available at: http://www.princeton.edu/~woodford/swind011.pdf.

——— (2001): "Indicator Variables for Optimal Policy under Asymmetric Information," NBER Working Paper No. W8255, available at: http://www.princeton.edu/~woodford/swind102.pdf.

Taylor, J. B. (ed.) (1999): *Monetary Policy Rules*. University of Chicago Press.

Walsh, C. E. (2001): "The Output Gap and Optimal Monetary Policy," paper presented at NBER Summer Institute, July 19.

Woodford, M. (1999a): "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?," in *New Challenges for Monetary Policy: A Symposium Sponsored by the Federal Reserve Bank of Kansas City*, pp. 277–316. Federal Reserve Bank of Kansas City.

——— (1999b): "Optimal Monetary Policy Inertia," *The Manchester School*, 67(Supplement), 1–35.