Optimal Monetary Policy When Lump-Sum Taxes Are Unavailable:

A Reconsideration of the Outcomes Under Commitment and Discretion*

Martin Ellison

Dept of Economics
University of Warwick
Coventry CV4 7AL
UK

m.ellison@warwick.ac.uk

Neil Rankin

Dept of Economics
University of Warwick
Coventry CV4 7AL
UK

n.rankin@warwick.ac.uk

August 2005

* We acknowledge helpful conversations with many people: Berthold Herrendorf, Henrik Jensen, Jonathan Thomas and Ted To, to name just a few. Thanks also to two anonymous referees, and to participants in numerous seminars. All errors and opinions are of course our own. Martin Ellison acknowledges support from an ESRC Research Fellowship, 'Improving Monetary Policy for Macroeconomic Stability in the 21st Century' (RES-000-27-0126).

Abstract

We re-examine optimal monetary policy when lump-sum taxes are unavailable. Under commitment, we show that, with alternative utility functions to that considered in Nicolini's related analysis, the direction of the incentive to cheat may depend on the initial level of government debt, with low debt creating an incentive towards surprise deflation, but high debt the reverse. Under discretion, we show that the economy will not necessarily tend to the Friedman Rule, as Obstfeld found. Instead it may tend to the critical debt level at which there is no cheating incentive under commitment, and inflation and could well be positive here.

Keywords

time consistency, optimal inflation-tax smoothing,
discretion, commitment, Friedman Rule

JEL Classification

E52, E61

1. Introduction

As is well known, optimal, welfare-based monetary policy, even in a flex-price, efficiently functioning economy, is subject to a time consistency problem if the government does not have lump-sum taxes or transfers at its disposal. This is because 'surprise' inflation may be a substitute source of non-distorting revenue, alleviating the problem of being unable to reach the first-best, 'Friedman Rule' outcome of a zero nominal interest rate. In this world, monetary policy must be implemented by open-market swaps of money and government debt, and the optimal second-best policy is to use debt to smooth intertemporally the distortions caused by inflation. Surprise inflation can help reduce these distortions.

Although this problem is familiar, significant puzzles about it remain. First is that the optimal monetary policy is likely to be degenerate, in that the best surprise rate of inflation to pick is infinity, because this maximises the lump sum of revenue appropriated by the government. Attention was drawn to this feature by Lucas and Stokey (1983). However, a well-defined optimum can be obtained if money enters the economy in such a way that there is a welfare cost of current inflation, since this cost must then be balanced against the benefits. This feature was introduced in an interesting contribution by Nicolini (1998). The presence of the cost, on the other hand, leads to a second puzzle: it may now be the case that the time consistency problem takes the form of an incentive to create surprise *deflation* (i.e. inflation lower than expected). This inverts all our usual ideas about time inconsistency in monetary policy. A third puzzle arises if we assume policy is conducted under discretion, rather than (as implicit in the discussion so far) under commitment. In this case lack of trust in the government's projected policy might be conjectured to lead to higher expected and actual inflation; but in fact it has been argued that discretion will lead in the long run to *lower* inflation – in particular to convergence to the Friedman Rule where inflation is negative. Such an argument is presented in two papers by Obstfeld (1991, 1997).[1]

[1] Here we refer to the case of Obstfeld's analysis where the authorities' objective function is as close as possible to private utility functions.

In this paper we reconsider this optimal monetary policy problem. We do so using Nicolini's model of a simple cash-in-advance economy, thereby avoiding the first of the above-mentioned puzzles. We extend his analysis of the case of commitment a little, but our main contribution is to the case of discretion, which he did not study. The model enables us to conduct a pure welfare-based analysis of optimal policy under discretion, avoiding Obstfeld's need to postulate an ad hoc objective function for the policymaker. Our analysis shows that the second and third of the above-mentioned puzzles are related. Specifically, we find that under discretion it is not necessarily true that in the long run the economy will converge on the Friedman Rule. Depending on consumer preferences, it may converge on a different steady state where inflation is above the Friedman-Rule level, and quite possibly positive. We call this the 'time-consistent steady state'. The force which makes such a steady state possible turns out to be the incentive which exists towards surprise deflation under commitment. This latter counteracts the more familiar incentive towards surprise inflation, and makes it possible that, under commitment, a critical level of inherited government debt exists such that the two incentives exactly cancel out, leading to no temptation to behave in a time-inconsistent manner. It transpires that this critical debt level is the same as the one associated with the 'time consistent steady state' to which the economy may converge under discretion.

The paper is organised as follows. Section 2 describes the structure of the economy. We examine optimal monetary policy under commitment in Section 3. In Section 4 we study optimal monetary policy under discretion, providing both analytical results and numerical computations. Section 5 concludes.

2. The Structure of the Economy

The economy consists of many identical households who consume the single type of output good, supply a single type of labour, and hold money - motivated by a cash-in-advance constraint - and bonds. Markets are perfectly competitive and all prices are flexible. The government issues money and bonds. Bonds are taken to be 'real', or indexed, as in the papers already cited.[2] A key constraint on the government is that it does not have access to lump-sum taxes and transfers: the money supply must be changed through open market operations, i.e. purchases or sales of bonds in exchange for money. Since our focus is purely on monetary policy, we shall ignore conventional distorting taxes, and treat government spending on goods and services as exogenous.

Technology takes the form $y_t = n_t$, where y_t is output and n_t is labour input. Hence competitive firms will set a price equal to the wage, and make zero profits. For the representative household, the optimisation problem is:

$$\text{maximise} \qquad \Sigma_{t=0}^{\infty} \beta^t \left[U(c_t) - \alpha n_t \right] \qquad (U' > 0,\ U'' < 0)$$

$$\text{s.t.} \qquad p_t c_t \leq M_t ,$$

$$M_t + b_t p_t (1 + R_t) + p_t n_t = p_t c_t + M_{t+1} + b_{t+1} p_{t+1} , \quad \text{for } t = 0, \ldots, \infty,$$

$$\text{with } M_0, b_0 \text{ given.} \tag{1}$$

M_t, b_t are, respectively, the quantities of money and real bonds held at the start of period t; R_t is the nominal interest rate between t-1 and t; c_t is consumption; n_t is labour supply; and p_t is the price level. The key feature of this problem is that M_0 and b_0 are given. In this way the model incorporates a 'welfare cost of current inflation': a rise in p_0 (and equal proportional rise in all p_t, $t > 0$) depresses the real value of the household's initial cash holdings, so that (provided the cash-in-advance constraint is binding) current consumption is squeezed. This liquidity squeeze occurs because the household is assumed to be unable to visit the asset market at the start of the period: in any period, goods markets open before asset markets. Such a timing assumption was introduced by Svensson (1985), and is the reverse of the more common timing assumption in cash-in-advance models associated with Lucas (e.g. Lucas

[2] Although it would be more 'realistic' to assume nominal bonds, this would obviously make it easier to generate a 'surprise inflation' result, so to assume real bonds cannot be said to facilitate our main conclusion. An interesting related analysis which does assume nominal bonds is by Diaz-Gimenez et al. (2002).

(1982)). Its value in the present context is that it captures the idea that current inflation has a cost because agents cannot instantaneously adjust their portfolios.[3]

From the first-order conditions for the household's problem we can readily derive:

$$U'(c_t) = \alpha[1 + R_t], \qquad t = 1,\ldots,\infty, \tag{2}$$

$$[1 + R_{t+1}]\frac{p_t}{p_{t+1}} = \frac{1}{\beta}, \qquad t = 0,\ldots,\infty, \tag{3}$$

(in which it has been assumed that $R_t > 0$, so that the cash-in-advance constraint binds). (3) indicates that the real interest rate is a constant equal to the inverse subjective discount rate. The nominal interest rate thus moves one-to-one with the expected inflation rate. (2) indicates that consumption is uniquely and negatively related (since $U'' < 0$) to the nominal interest rate, or equivalently to the expected inflation rate. This reflects the distorting effect of the inflation tax: the earnings from an increase in current labour supply cannot be spent immediately, owing to the cash-in-advance constraint, and their future value is hence eroded by expected inflation, which therefore provides a disincentive to labour supply and so to consumption. Writing $m_t \equiv M_t/p_t$, we note that $m_t = c_t$ provided that the cash-in-advance constraint binds, which reminds us that expected inflation equivalently reduces the demand for real balances.

Equilibrium in the private sector of the economy may now be determined. At this point, suppose an arbitrary monetary policy defined by a sequence of monetary growth rates, $\{\mu_t\}_0^\infty$ (where $\mu_t \equiv [M_{t+1}\text{-}M_t]/M_t$). Then equilibrium consumption must satisfy:

$$c_t \quad = \quad \frac{\beta}{\alpha}\frac{1}{1+\mu_t}U'(c_{t+1})c_{t+1}, \tag{4}$$

which has been obtained from (2), (3) and $p_{t+1}/p_t = (1+\mu_t)c_t/c_{t+1}$ (the latter being implied by $m_t = c_t$). Under perfect foresight (4) determines c_t in a pure forward-looking manner, as a

[3] As Nicolini (1998) shows, it can be interpreted as a more compact version of the 'limited participation' model of Grossman and Weiss (1983) and Rotemberg (1984). There, agents are divided into two groups who are only allowed to visit the bank in alternate periods.

function of current and future monetary growth rates $\{\mu_{t+s}\}_{s=0}^{\infty}$. With consumption tied down by (4), output is determined via the goods market clearing condition:

$$y_t \quad = \quad c_t + g \tag{5}$$

where g is the level of government spending.

A key role is played in what follows by the government's budget constraint. The single-period constraint may be written:

$$b_t p_t(1+R_t) + p_t g \quad = \quad M_{t+1} - M_t + b_{t+1} p_{t+1} \,. \tag{6}$$

Note the absence of lump-sum taxes or transfers. Under our timing assumptions, (b_t, M_t) are predetermined in period t, and the only policy action open to the government (since we treat g as exogenous) is an open-market sale or purchase of bonds, which raises or lowers b_{t+1}, lowering or raising M_{t+1}, respectively. We may integrate (6) forwards, incorporating (3) (and appealing to a 'No Ponzi Game' condition), to obtain:

$$b_0(1+R_0) \quad = \quad \Sigma_{t=0}^{\infty} \beta^t \mu_t m_t \; - \; \frac{1}{1-\beta} g \,. \tag{7}$$

This is the government's intertemporal budget constraint expressed in terms of 'cash-flow seigniorage', i.e. $\mu_t m_t$. We may also re-write it as:

$$b_0(1+R_0) \quad = \quad -m_0 \; + \; \Sigma_{t=1}^{\infty} \beta^t R_t m_t \; - \; \frac{1}{1-\beta} g \,. \tag{8}$$

(8) gives the same constraint in terms of 'opportunity-cost seigniorage', i.e. $R_t m_t$.[4] Notice that m_0 enters (8) differently from m_1, m_2, etc. Algebraically speaking, this is the source of time inconsistency, as we discuss further below.

[4] See Herrendorf (1997) for a useful discussion of these two definitions. To obtain (8) from (7), note that $[M_{t+1} - M_t]/p_t$ can be re-written as $m_{t+1}\beta[1+R_{t+1}] - m_t$.

3. Optimal Monetary Policy Under Commitment

In this section we assume that there is some commitment mechanism which obliges the government to adhere to the monetary policy which it chooses in period 0, $\{\mu_t\}_0^\infty$. Its optimisation problem is then:

$$\text{maximise w.r.t. } \{\mu_t\}_0^\infty \qquad \Sigma_{t=0}^\infty \beta^t [U(c_t) - \alpha c_t - \alpha g]$$

$$\text{s.t.} \qquad (1+R_0)b_0 + \frac{1}{1-\beta} g \;=\; -c_0 + \Sigma_{t=1}^\infty \beta^t \left[\frac{1}{\alpha} U'(c_t) - 1\right] c_t ,$$

$$c_t \;=\; \frac{\beta}{\alpha} \frac{1}{1+\mu_t} U'(c_{t+1}) c_{t+1} \quad \text{for } t = 0,\ldots,\infty,$$

$$\text{with } (1+R_0)b_0,\ g \text{ given.} \tag{9}$$

We have substituted out R_t using (2), m_t using $m_t = c_t$, and n_t using the production function and (5). It is clear that, rather than treat the μ_t's as the control variables, we can equivalently treat the c_t's as the control variables, leaving the μ_t's to be determined residually by the difference equation constraints. This reduces the problem to:

$$\text{maximise w.r.t. } \{c_t\}_0^\infty \qquad \Sigma_{t=0}^\infty \beta^t [U(c_t) - \alpha c_t - \alpha g]$$

$$\text{s.t.} \qquad -(1+R_0)b_0 - \frac{1}{1-\beta} g \;=\; c_0 - \Sigma_{t=1}^\infty \beta^t \left[\frac{1}{\alpha} U'(c_t) - 1\right] c_t . \tag{10}$$

We easily see that this problem has an *unconstrained* maximum where:

$$U'(c_t) = \alpha \qquad \text{for all } t. \tag{11}$$

This is the first-best, 'Friedman Rule' solution, where $R_t = 0$. The government's intertemporal budget constraint in general prevents the attainment of this outcome unless initial government debt is sufficiently negative. Thus, in a second-best situation, c_t will be less than its Friedman-Rule level, and an optimal policy must trade off the deviation of c_t from this with the deviation of any other c_s.

It is straightforward to derive the following set of first-order conditions for the problem (10):

$$U'(c_0) - \alpha \quad = \quad \frac{U'(c_t) - \alpha}{1 + [\sigma(c_t) - 1]U'(c_t)/\alpha} \qquad \text{for } t = 1,\dots,\infty. \tag{12}$$

Here, $\sigma(c_t)$ is the 'relative risk aversion' measure of curvature of $U(c_t)$, i.e. $-c_t U''(c_t)/U'(c_t)$. It is immediate from this that c_t is the same for all $t = 1,\dots,\infty$. It is further apparent that this common c_t ($\equiv c$, say) is in general different from c_0. Hence an optimal time path of c_t typically takes the form of either a 'step up' or a 'step down' (see Figure 1). Inspection of (12) further shows that if $\sigma < 1$ then we obtain $c_0 < c$ (a 'step up'), and if $\sigma > 1$ then we obtain $c_0 > c$ (a 'step down'). In the special case where $\sigma = 1$, the optimal time path is 'flat'.[5]

The above results are as in Nicolini (1998). Nicolini however restricts attention to utility functions of the 'constant relative risk aversion' (CRRA) class, so that σ is an exogenous parameter. Here we will explore the consequences of more general utility functions. A visual aid to doing so is provided by Figure 2. The diagram exploits the fact that, under an optimal policy, $c_t = c$ for $t = 1,\dots,\infty$, in order to represent the problem in reduced form. The indifference curves over (c_0,c) are given by the lifetime utility function written as $U(c_0) - \alpha c_0 - \alpha g + \beta[1-\beta]^{-1}[U(c) - \alpha c - \alpha g]$, and the government's intertemporal budget constraint contracts to:

$$c_0 \quad = \quad \frac{\beta}{1-\beta}\left[\frac{1}{\alpha}U'(c)c - c\right] \; - \; (1+R_0)b_0 \; - \; \frac{1}{1-\beta}g\,. \tag{13}$$

Figure 1 depicts this as a backward-bending curve. This reflects the 'seigniorage Laffer curve': as future inflation rates (and thus nominal interest rates) are raised, c falls, but nevertheless future seigniorage revenue at first rises, permitting less revenue to be raised through current seigniorage, and thus permitting higher c_0. However (depending on the shape of $U(c)$) the Laffer curve may have a peak, such that revenue starts to fall as future inflation continues to be increased and c continues to be reduced. Initial government debt, $(1+R_0)b_0$,

[5] Corresponding to these time paths of consumption, it is easy to show that the time paths of debt are a 'step down', a 'step up', and 'flat', respectively.

acts as a parameter which fixes the position of the budget constraint: higher initial debt shifts it to the left. By varying initial debt we thus trace out a locus of points of tangency with the indifference curves - the ‘expansion path’. The equation of the expansion path is just the first-order condition (12), with ‘c_t’ substituted by c.

Consider now the shape of the expansion path. First, it clearly passes through the point ($\bar{c}$,$\bar{c}$), where $\bar{c}$ is the Friedman Rule consumption level given by (11). Second, from the results presented earlier, the expansion path must lie above (below) the 45° line at levels of c such that $\sigma(c) < 1$ (>1). It follows that if $U(c)$ is such that there exists a value c_c at which $\sigma(c) = 1$, then the expansion path must cross the 45° line at this critical value of c. (For this to be of interest, we obviously need $c_c < \bar{c}$). Third, we would intuitively expect the expansion path to be upward-sloping (i.e. under an optimal policy lower initial government debt would be used to raise both c_0 and c)[6]. Hence, fourth, if c_c exists then the question arises as to whether the intersection with the 45° line occurs ‘from above’ or ‘from below’. Differentiating (12) and evaluating where $c_0 = c = c_c$, we obtain:

$$\left.\frac{dc}{dc_0}\right|_{C_c} = \frac{\alpha}{\alpha + c\sigma'[U' - \alpha]} . \tag{14}$$

This is clearly less than or greater than one as $\sigma'(c_c)$ is (respectively) positive or negative. That is, the expansion path cuts the 45° line from above if the relative risk aversion parameter is increasing in consumption at the intersection point, and from below if it is decreasing.[7]

Figure 2 illustrates the case where an intersection of the expansion path and the 45° line exists, and where it occurs from above. In this case there must also exist a critical value of $(1+R_0)b_0$ such that, if initial debt happened to take this value, the optimal policy would be to choose $c_0 = c = c_c$. The associated budget constraint is the one which passes through the point C. If initial debt were slightly higher than this level, the budget constraint would lie slightly farther to the left, and the optimum would be at a point like A, where $c_0 < c$. Conversely if

[6] The condition for this is that the RHS of (12) be decreasing in c. It turns out that the same condition is necessary for the optimisation problem’s second-order conditions to be satisfied, so we need this condition to hold. It holds provided that $\sigma(c)$ is increasing, constant, or not too strongly decreasing, in c.

[7] (14) is not the correct expression for the slope at the Friedman-Rule point. The slope there is instead given simply by σ, and so is greater or less than one as σ is greater or less than one.

initial debt were slightly lower, the optimum would be at a point like B, where $c_0 > c$. In the case in Figure 2, then, the shape of the optimal time path of consumption depends on the initial level of debt: 'high' debt gives rise to a 'step up' shape, and low debt to a 'step down' shape. With the CRRA utility function used by Nicolini (1998), on the other hand, the expansion path lies either always above, or always below, the 45° line, and the magnitude of initial government debt is therefore irrelevant, qualitatively speaking, to the shape of the optimum time path of consumption.

In order to assess whether there is a potential time consistency problem with the optimal policy, it must be compared with the optimal policy if the government were permitted to re-optimise in period 1. This re-optimisation is only hypothetical, since we have assumed that there is a commitment mechanism to prevent the government from actually deviating from its period-0 plan. Let us denote the optimal choice of c_t from the perspective of period s ($s \leq t$) as $c_{t|s}$. We are particularly interested in whether $c_{1|1}$ is smaller or greater than $c_{1|0}$. Since $c_1 = M_1/p_1$ and M_1 is predetermined, $c_{1|1} < c_{1|0}$ indicates that the government in period 1 would wish to set p_1 higher than was planned in period 0, i.e. that there is an incentive to 'surprise inflation'. Conversely, if $c_{1|1} > c_{1|0}$, the incentive is to 'surprise deflation'. Following Nicolini's method we may readily show that, if the period-0 optimal time path takes the form of a 'step up', then in period 1 there is an incentive to surprise inflation; while if it takes the form of a 'step down', then in period 1 there is an incentive to surprise deflation. A summary sketch of the two possible relationships between the period-0 and period-1 optimal time paths for consumption is thus as in Figure 1.

Together with the earlier results, this then implies that if there exists a critical consumption level (strictly less than the Friedman Rule level) at which the coefficient of 'relative risk aversion' in consumption equals one, then associated with this is a critical level of government debt such that, if initial debt happens to take this value, there is no time inconsistency. Moreover, if relative risk aversion is increasing (decreasing) in consumption at the critical level, then initial debt above the critical debt level will be associated with a temptation to create surprise inflation (deflation); and initial debt below, with a temptation to create surprise deflation (inflation). This is an extension of Nicolini's (1998) main finding. It says that if relative risk aversion is not a constant, then the direction of time inconsistency may depend on the level of initial government debt. An intuitively plausible outcome, we would argue, arises in the case of increasing relative risk aversion. Here, while the rather unorthodox incentive to create surprise deflation dominates for low levels of debt, the more conventional incentive to create surprise inflation dominates for high levels. This outcome occurs if the expansion path has the shape shown in Figure 2.[8]

A natural follow-up question concerns how likely it is to find a utility function possessing the properties just highlighted. Although the CRRA function is a common utility function with some convenient features, there is of course no shortage of alternatives. Consider, for example, quadratic utility, $U(c) = c[\hat{c} - c/2]$. (With this, $\hat{c} \geq \alpha$ is needed to ensure a non-negative Friedman-Rule consumption level.) $\sigma(c)$ = $c/[\hat{c} - c]$ for this function, so a value of c such that $\sigma = 1$ clearly exists, namely at $c = \hat{c}/2$. Moreover, σ is clearly everywhere increasing in c. The associated expansion path then looks exactly like that in Figure 2, provided that $\hat{c} > 2\alpha$.[9] Another common functional form, the 'constant absolute risk aversion' (CARA) function, can similarly yield an expansion path like that in Figure 2. We hence conclude that the preferences required in order for the above result to apply are not especially unusual. It is true that CRRA preferences are widely used in macroeconomics, partly because they have the convenient property of being consistent with balanced growth. However, quadratic preferences are also widely used in consequence of some other convenient properties: for example, in the presence of uncertainty they generate the exact consumption-as-a-random-walk result, and they underpin mean-variance portfolio analysis. In the present paper our aim is to explore the implications of some familiar utility functions as a theoretical exercise, rather than to claim any one function fits the facts well. Clearly

[8] A referee points out that, in an extension to his main analysis, Nicolini (1998) also notes that there may be a critical level of initial debt such that there is no temptation to time inconsistency. This is where $\sigma > 1$ and debt is *nominal*. Nominal debt generates an incentive towards surprise inflation which could exactly offset the incentive towards surprise deflation. However Nicolini does not develop the analysis of this case.

[9] This condition ensures that $\hat{c}/2$ is less than the Friedman-Rule consumption level, $\hat{c}-\alpha$. If it is violated, the expansion path lies everywhere above the 45^{o} line. Even though quadratic utility by itself does not therefore guarantee an expansion path like that in Figure 2, we can still say that, with quadratic utility, if an incentive to surprise deflation is ever to exist (i.e. if part of the expansion path is ever to lie below the 45^{o} line), then a critical consumption level as defined in the main text must also exist.

much more elaboration of our bare-bones model would be needed before it could be confronted with the data.

A further natural question concerns why the direction of time inconsistency should depend on the size of σ. As noted, in a CIA model with Lucas's timing assumption, the temptation to 'cheat' in monetary policy is always towards surprise inflation. In the present framework, the incentive to raise all current and future prices equiproportionally in order to appropriate revenue in a 'non-distorting' way is counterbalanced by the incentive not to generate a large welfare cost of current inflation through the liquidity-squeeze mechanism explained earlier. The direction of the temptation to cheat depends on how these incentives change, relative to the incentives for the setting of subsequent periods' inflation, as time advances. First, observe that a temptation to reduce current inflation relative to what was planned would arise if an increase in current inflation (increase in p_0/p_{-1}, from the perspective of period 0) were to provide a smaller gain in the PV of seigniorage and a larger loss of lifetime utility[10] than a projected increase in future inflation (increase in p_1/p_0, from the perspective of period 0). This is because, when the future 'arrives', the seigniorage benefit would be smaller than it was from the vantage point of the period before and the utility cost would be larger, so the government would perceive an incentive to deviate downwards from what it had planned. Such changes in incentive occur when $\sigma > 1$ because future consumption demand is then inelastic with respect to future inflation (or, equivalently, $1+R_t$), as can be seen from (2); while current consumption demand is always unit-elastic with respect to current inflation, as follows from the cash-in-advance constraint. Hence a 1-unit increase in p_0/p_{-1} causes a larger loss of lifetime utility than a $(1/\beta)$-unit increase in p_1/p_0. It also causes a smaller gain in the PV of seigniorage, because the inelastic future consumption demand means a weaker dampening effect on revenue of the shrinkage in the tax base as the inflation rate is raised, so that seigniorage increases more strongly with future than with current inflation.

[10] Through its 'direct' effect on utility, ignoring the indirect effect via its budgetary impact

4. Optimal Monetary Policy Under Discretion

If the government is unable to commit to a given policy plan made in period t, $\{\mu_{t+s|t}\}_{s=0}^{\infty}$, then it is not rational for households' forecasts of future μ_{t+s}'s, as of time t (denote these by $\mu^e_{t+s|t}$), to equal those in the plan. Instead, forecasts should be based only on variables observable at time t. The government's inherited stock of debt is one obvious 'observable' on which to condition forecasts, since it is clear from the previous section that initial debt is a key determinant of the government's optimal policy choices. Other observables could also be used, such as current and past values of μ_t. However, this would introduce an element of 'reputation-building' behaviour into policy. Since we wish to focus on purely 'discretionary' behaviour, we use only debt as the basis for private forecasts. Formally, the concept of equilibrium employed will be that of Markov-perfect equilibrium.[11]

First, it is useful, as Obstfeld does, to define the concept of government 'commitments'

$$k_t \equiv b_t(1+R_t) + \Sigma_{s=0}^{\infty}\beta^s g_{t+s}. \tag{15}$$

'Commitments' is just the spending, in present-value terms, over which the government does not have discretion. Formally, we will use k_t rather than debt as the state variable in what follows, although, given our assumption that g is time-invariant, k_t is simply debt plus the constant, $g/(1\text{-}\beta)$. Note that the government's budget constraint in terms of k_t is (cf. (6)):

$$k_t = \mu_t m_t + \beta k_{t+1}. \tag{16}$$

We now suppose that, in period t, households forecast (μ_{t+1}, k_{t+2}) using the rules:

$$\mu^e_{t+1|t} = \hat{\phi}(k_{t+1}), \tag{17}$$

$$k^e_{t+2|t} = \hat{\psi}(k_{t+1}). \tag{18}$$

---

[11] This is the same basic idea as in Obstfeld (1991, 1997). His 1991 paper studies a small open economy version of the optimal inflation tax smoothing problem, while his 1997 paper does the same for a closed economy. The key difference from the present analysis, as noted in the Introduction, is that Obstfeld's model does not include a 'welfare cost of current inflation'. This obliges him to use a government objective function which differs from the private utility function, being obtained by adding on an ad hoc 'cost of current inflation' term to the latter.

(Note that k_{t+1} is observable by households in period t - since they can observe $\mu_t m_t$ and thus use (16) - so that it is appropriate to base their forecasts upon it.) $\hat{\phi}(.)$ and $\hat{\psi}(.)$ are for the moment treated as arbitrary functions, but they will later be determined by imposing a rationality requirement as part of the conditions of Markov-perfect equilibrium. To generate an s-period ahead forecast, (18) may be used repeatedly in (17):

$$\mu^e_{t+s|t} \quad = \quad \hat{\phi}(\hat{\psi}^{s-1}(k_{t+1})), \tag{19}$$

where $\hat{\psi}^n(.) \equiv \hat{\psi}(\hat{\psi}(...\hat{\psi}(.)...))$ denotes the nth iterate of the function $\hat{\psi}(.)$.

These forecasting rules can next be used to determine equilibrium consumption in period t. Recall that the equilibrium value of c_t is given by the private-sector law of motion, (4), solved in a forward-looking manner. The relevant μ_t's to use in this equation are now the expected values as given by (19), rather than the values from the government's policy plan, since what counts for determining the actual c_t are households' expectations. It is helpful to consider the determination of c_t in two stages. First, given k_{t+1} and thus a sequence of expected values $\{\mu^e_{t+s|t}\}^{\infty}_{s=1}$ generated by (19), we use (4) for periods t+1 onwards to solve for c_{t+1}. This yields households' forecast, as of period t, of the equilibrium value of c_{t+1}. Since it is contingent on the observed k_{t+1}, we may write it as $c^e_{t+1|t} = \hat{\theta}(k_{t+1})$. Second, using $\hat{\theta}(k_{t+1})$ in (4) for period t, we obtain the equilibrium value of current c_t:

$$\begin{aligned} c_t \quad &= \quad \frac{\beta}{\alpha}\frac{1}{1+\mu_t}U'(\hat{\theta}(k_{t+1}))\hat{\theta}(k_{t+1}) \\ &\equiv \quad \frac{e(k_{t+1})}{1+\mu_t}. \end{aligned} \tag{20}$$

Current c_t thus depends on two variables: the currently observed μ_t, and the currently observed k_{t+1}, whose effect operates via influencing expectations about future μ_{t+s}'s. The function $e(.)$ captures this expectations effect. The form of $e(.)$ derives from the form of $\hat{\theta}(.)$, which is in turn derived from the forms of $\hat{\phi}(.)$ and $\hat{\psi}(.)$.

The government now treats (20) as part of the economy's structure, and thus as a given. Its optimisation problem under discretion is therefore to choose $\{\mu_t\}^{\infty}_{0}$ to maximise lifetime utility of the typical agent subject to (20) and (16) for t = 0,...,∞, and to given k_0. By substituting (20) into (16) and into the maximand, we can express this as:

$$\text{maximise w.r.t. } \{\mu_t\}_0^\infty \qquad \Sigma_{t=0}^{\infty}\beta^t\left[U\left(\frac{e(k_{t+1})}{1+\mu_t}\right)-\alpha\frac{e(k_{t+1})}{1+\mu_t}-\alpha g\right]$$

$$\text{s.t.} \qquad k_t = \mu_t\frac{e(k_{t+1})}{1+\mu_t}+\beta k_{t+1}, \quad \text{for } t = 0,...,\infty,$$

$$\text{with } k_0 \text{ given.} \tag{21}$$

(21) reveals that the government's dynamic optimisation problem under discretion - unlike that under commitment - has a standard recursive form. That is, every period, the new value of the state (k_{t+1}) depends only on the current value of the state (k_t) and on the current value of the control (μ_t), while the flow maximand also depends only on these same two variables (although it is k_{t+1} which appears in the flow maximand, k_{t+1} is an implicit function of (k_t,μ_t) via the constraint). This structure means that the problem can in principle be solved by dynamic programming. In turn, dynamic programming ensures that the solution is time consistent.[12]

The dynamic programming perspective also makes clear that the solution to our problem can be expressed as a pair of feedback rules on the state; for example

$$\mu_t = \phi(k_t), \tag{22}$$

$$k_{t+1} = \psi(k_t). \tag{23}$$

(22)-(23) define the government's optimal monetary policy having taken as our starting point the public's arbitrary forecasting rules, (17)-(18). We notice that (22)-(23) relate the same variables as (17)-(18); hence, for (17)-(18) to provide 'rational' forecasts by the public, we need the functions $\hat{\phi}(.)$ and $\phi(.)$, and also $\hat{\psi}(.)$ and $\psi(.)$, to be the same. In this case households will forecast correctly no matter what the value of k_t. The discretionary, or

[12] Given that the function e(.) has as yet unknown properties, there is a question as to whether the optimisation problem (21) is well defined. Here we proceed as if this is the case, but we return to the question below.

Markov-perfect, equilibrium is thus a pair of forecasting rules which have the property that they reproduce themselves in the guise of optimal government policy rules.

It is convenient, as earlier, to rewrite the optimisation problem in order to treat the c_t's rather than μ_t's as the controls. To do this we use (20) to substitute out the μ_t's from (21), so that the problem is transformed to:

$$\text{maximise w.r.t. } \{c_t\}_0^\infty \qquad \Sigma_{t=0}^{\infty}\beta^t[U'(c_t)-\alpha c_t-\alpha g]$$

$$\text{s.t.} \qquad k_t \;=\; e(k_{t+1})-c_t+\beta k_{t+1}, \quad \text{for } t=0,...,\infty,$$

$$\text{with } k_0 \text{ given.} \tag{24}$$

The first-order conditions for this are then easily derived:

$$U'(c_{t+1})-\alpha \;=\; [U'(c_t)-\alpha][1+\frac{1}{\beta}e'(k_{t+1})] \quad \text{for } t=0,...,\infty. \tag{25}$$

Repeating here the constraint from the problem (24),

$$k_t \;=\; e(k_{t+1})-c_t+\beta k_{t+1}, \tag{26}$$

we see that (25)-(26) constitute a first-order dynamical system in (k_t,c_t), which determines the evolution of the economy under the optimal discretionary policy. Given that k_t is a predetermined variable whereas c_t is not, c_0 will generally be tied down in relation to k_0 by the transversality condition. This hence determines a particular solution of the system (25)-(26) which constitutes the solution of the optimisation problem. We denote this solution as:

$$c_t=\theta(k_t)\,. \tag{27}$$

Equivalently, (27) is the optimal feedback rule of the control upon the state variable. If the optimum is such that the economy converges on a steady state, then (27) is also the 'saddlepath' solution of (25)-(26).

Although $\theta(.)$, like $e(.)$, is still at present an unknown function, once it has been determined we can substitute it back into (26) to get:

$$e(k_{t+1}) + \beta k_{t+1} \quad = \quad k_t + \theta(k_t)\,. \tag{28}$$

This implicitly determines the optimal feedback rule for k_{t+1}, (23). The optimal feedback rule for μ_t, (22), is similarly recoverable by substituting (27) and (23) into (20). The functions $e(.)$ and $\theta(.)$ thus play a central role in the Markov-perfect equilibrium, since, if they can be determined, the other unknown functions $\phi(.)$ and $\psi(.)$ follow. It is also worth noting that in equilibrium the function $\theta(.)$ must turn out to be the same as the function $\hat{\theta}(.)$. This is because if households' forecasting rules are always to yield correct predictions, their forecast of c_{t+1} contingent on k_{t+1} must coincide with the government's optimal, and thus actual, choice of c_{t+1} contingent on k_{t+1}.

We now aim to study the properties of the discretionary equilibrium. We start with the steady states. Inspection of (25) suggests two ways in which a stationary solution of the dynamical system (25)-(26) may occur. First, (25) is clearly satisfied at the Friedman Rule, where $U' - \alpha = 0$ for all t. We shall refer to this as the 'Friedman Rule steady state' (FRSS), It is intuitively clear that if initial government debt is sufficiently negative that the underlying 'second-best' problem is absent, then time inconsistency is removed, and a government which started in this fortunate position would have no incentive to move away from it. Obstfeld (1991, 1997) similarly identifies the Friedman Rule allocation as a steady state of the discretionary equilibrium in his analysis. However, whereas for Obstfeld the Friedman Rule allocation is the only steady state, this is not necessarily true here. A second way in which a stationary solution of (25) could occur is if there exists a k_{t+1} at which $e'(k_{t+1}) = 0$. More specifically, we might conjecture that this would be true at the value of k_{t+1} corresponding to the 'critical' debt level as we defined it for monetary policy under commitment. The motivation for such a conjecture is that we know from Section 3 that there is no time inconsistency if initial debt happens to equal the critical value, and that, under commitment, if the government started with this amount of debt it would choose to stay there. Thus it might be hypothesised that, with this critical amount of initial debt, the optimal policy under discretion would be the same as under commitment. This second type of steady state we shall refer to as the 'time consistent steady state' (TCSS).

To prove that a TCSS exists under discretion if, in the problem under commitment, there exist critical consumption and debt levels as defined in Section 3, consider again the definition of the function *e*(.) (given in (20)). Differentiating this function with respect to *k*, we obtain:

$$e'(k) \quad = \quad \frac{\beta}{\alpha}[1-\sigma(c)]U'(c)\theta'(k) \tag{29}$$

(in which we have equated $\hat{\theta}(.)$ with $\theta(.)$, for the reason explained). Although $\theta'(k)$ is unknown, we do know from Section 3 that $\sigma(c) = 1$ at the critical consumption level c_c. This is therefore sufficient to prove that $e' = 0$ at the corresponding critical commitments level (call this k_c). Hence a stationary solution of (25) does indeed arise at k_c.

The level of inflation at the TCSS is higher than the negative inflation rate (β-1) which prevails at the Friedman Rule. Its value depends on the utility function: from (2) and (3), the general expression for the inflation rate is $(\beta/\alpha)U'(c) - 1$, and at c_c this could be positive or negative. Hence, if we can show that there are conditions under which the discretionary equilibrium converges on the TCSS (for the general case in which the initial $k_0 \neq k_c$) then it follows that the long-run destination of the economy under discretion is not necessarily a situation of deflation. We may also note that the level of debt at the TCSS is higher than the negative level required to sustain the Friedman Rule, and the same is true of the level of commitments. Whether they are negative or positive in the absolute again depends on the utility function, and also on *g*.[13]

The evolution of the economy under discretion is governed by (23), which we saw to take its form implicitly from (28). Differentiating (28), and evaluating at a generic steady state *k*, we have:

$$\frac{dk_{t+1}}{dk_t} = \psi'(k) \quad = \quad \frac{1+\theta'(k)}{\beta + e'(k)}. \tag{30}$$

Local convergence to either the FRSS or the TCSS clearly requires that $-1 < \psi' < 1$. Hence we shall proceed by attempting to solve for $\theta'(k)$ and $e'(k)$ at each type of steady state, in

[13] Specifically, setting $c_0 = c = c_c$ in (13), we have $k_c = (1-\beta)[(\beta/\alpha)U'(c_c)-1]c_c$, and $b_c(1+R_c) = k_c - g/(1-\beta)$.

order to determine whether this condition can be satisfied. An expression for e' has already been obtained in (29). Differentiating (29) for a second time yields the following expression for e'', which will be useful below:

$$e'' \;=\; (\beta/\alpha)[-\sigma' U'(\theta')^2 + (1-\sigma)U''(\theta')^2 + (1-\sigma)U'\theta''\theta']. \tag{31}$$

In order to find an expression for θ', recall that $\theta(k_t)$ is the saddlepath solution of the dynamical system (25)-(26). By taking a linear approximation to this system about a steady state, we may find an expression for θ'. Such a calculation yields:

$$1+\theta'$$

$$= \frac{1}{2\beta}\left\{\left\{\frac{U'-\alpha}{U''}e'' + (\beta+e')^2 + \beta \pm \sqrt{\left[\frac{U'-\alpha}{U''}e'' + (\beta+e')^2 + \beta\right]^2 - 4\beta(\beta+e')^2}\right\}\right. . \tag{32}$$

We now observe that (29), (31) and (32) constitute a system of three simultaneous equations in the four unknowns $(e', e'', \theta', \theta'')$. Although we cannot solve it as it stands, if we proceed to evaluate it at either the FRSS or TCSS, it turns out that we obtain additional restrictions which are sufficient to make solution possible.

Consider first the local dynamics of the FRSS. Setting $U' = \alpha$, notice that e'' drops out of (32). (29) and (32) then constitute a system of two equations in just two unknowns, e' and θ', from which we may hope to solve for e' and θ'. Appendix A presents the relevant calculations. We show that the system can be reduced to a quadratic equation in θ' as the single unknown. Once θ' has been determined, e' follows from (29) and ψ' from (30). Since the quadratic implies two possible solutions for θ', there are also two possible solutions for e' and ψ'. The point of particular interest is whether either of the solutions for ψ' have absolute value less than one. The key result demonstrated in Appendix A is that if $\sigma < 1$, there is exactly one solution for ψ' with absolute value less than one; but if $\sigma > 1$ there are no solutions with absolute value less than one. Combining this with our earlier finding, we thus conclude that:

Proposition 1 *The Friedman-Rule consumption level and associated negative debt level constitute a steady state of the discretionary equilibrium. Moreover, if the coefficient of 'relative risk aversion' in consumption is less (greater) than one at this steady state, then, within a neighbourhood of it, a discretionary equilibrium which converges on it exists (does not exist).*

In the case $\sigma < 1$ we also find, more specifically, that ψ' lies in (0,1) (ensuring that convergence is monotonic), and that θ' and e' are negative. Thus, as the Friedman Rule is approached, debt steadily falls (or equivalently - once debt becomes negative - government assets steadily rise) and consumption steadily rises. Inflation also steadily falls (becomes more negative). This is the outcome obtained by Obstfeld (1991, 1997). Such an outcome differs from what happens under commitment, where, if $\sigma < 1$, consumption also rises along an optimal time path, but only between periods 0 and l. Debt correspondingly falls between periods 0 and 1, but remains permanently above its Friedman Rule level. The clue as to why a government acting under discretion goes farther in reducing its debt over time than a government acting under commitment, lies in the negative e', as Obstfeld pointed out. $e' < 0$ means that lower debt (lower k_{t+1}) induces higher current consumption through affecting private agents' expectations (recall (20)). From the government's perspective, $e' < 0$ thus increases the 'return' to a marginal reduction in the debt: not only does lower debt next period mean the future inflation-tax revenue needed is lower, but also, by lowering private expectations of future inflation[14], it raises current consumption. By contrast, under commitment, private agents do not use the level of government debt as the basis for their forecasts - instead, they trust the government to carry out today's optimal plan. Hence a change in the debt level per se does not have this added 'Obstfeld effect'.[15]

A numerical illustration of a set of discretionary equilibria (one for each initial value of k_t) converging on the FRSS is given in Figure 3. This is calculated for the CRRA utility

---

[14] To see this, note that since θ' is negative, and recalling that $\theta' = \hat{\theta}'$, a reduction in k_{t+1} raises private agents' forecasts of c_{t+1}. c_{t+1} is negatively related to expected inflation p_{t+1}/p_t by (2)-(3).

[15] Although, under this outcome, utility is higher in the long run under discretion than under commitment, lifetime utility from the perspective of period 0 is lower, because consumption in the short run can be shown to be lower. This is as it should be, because inability to commit acts as a constraint on the optimal policy.

function with $\sigma = 0.5$, $\alpha = 5$ and $\beta = 0.95$. (We have also set $g_t = 0$ for all t.) A description of the algorithm used for the computations is provided in Appendix C. The Friedman-Rule consumption level in this example is $\overline{c} = 0.04$, and the corresponding level of 'commitments' needed to sustain this is $k = -0.04$. Panel (a) shows the $e(k_t)$ function for this example. This confirms the negative effect of the stock of debt (to which 'commitments' are here equivalent), operating via private-sector expectations, on c_{t-1}. Panel (b) shows the $\theta(k_t)$ function (the line labelled 'c_t'), indicating the negative total effect of the stock of debt on the government's optimal choice of c_t. Note that the range of c_t values considered extends from the maximum, Friedman-Rule, value down to approximately one quarter of that value, so that the picture is not just concerned with a small neighbourhood of the FRSS. The function is clearly quite close to being linear. Panel (c) plots the $\psi(k_t)$ function and also, for reference, the 45^{o} line. As Proposition 1 predicts, the slope is less than one - in fact, it is about 0.9 - which confirms that the economy does indeed converge monotonically on the FRSS for any initial level of debt above -0.04. For comparison, panel (b) also plots the two consumption levels which result from the problem under commitment. The line 'c_0' gives consumption in the first period of an optimal plan as a function of the initial debt k_0; while the line 'c' gives consumption in the second and all later periods, again as a function of k_0. We thus see that for the same inherited level of debt, in this example consumption in the short run is chosen to be lower under discretion than under commitment and, correspondingly, inflation is chosen to be higher.

When $\sigma > 1$, however, Proposition 1 indicates that it is not possible for the outcome obtained by Obstfeld to occur in our model. The force driving the economy away from the Friedman-Rule outcome in such a case is discussed below in the context of Proposition 2. This suggests that in this case the destination may instead be the time-consistent steady state, where the latter exists.

We now turn to the local dynamics of the TCSS. At the TCSS itself, $\sigma = 1$ and $e' = 0$, as shown previously. Hence θ'' drops out of (31), and the simultaneous system of (29), (31) and (32) can be reduced to the system of just (31) and (32), in the unknowns θ' and e''. From this we may hope to solve for θ' and e''. Appendix B presents the relevant calculations. We show that the system can be reduced to a quadratic equation in θ' as the single unknown, analogous to, but different from, the quadratic applying at the FRSS. Once θ' has been determined, ψ' follows from (30). Since the quadratic implies two possible solutions for θ', there are also two possible solutions for ψ'. The point of particular interest is whether either of the solutions for ψ' have absolute value less than one. The key result demonstrated in Appendix D is that if $\sigma' > 0$, there is exactly one solution for ψ' with absolute value less than one; but if $\sigma' < 0$ there are no solutions with absolute value less than one. Combining this with our earlier findings, we thus conclude that:

Proposition 2 *If the critical consumption and debt levels as defined for the problem under commitment exist, then a steady state of the discretionary equilibrium also exists at these consumption and debt levels. Inflation could be positive or negative at this steady state. Moreover if, under commitment, initial debt above the critical level was associated with a temptation to create surprise inflation (deflation), and below, with a temptation to create surprise deflation (inflation), then, under discretion, within a neighbourhood of the steady state, an equilibrium which converges on it exists (does not exist).*

As Proposition 2 emphasises, in the neighbourhood of the critical consumption level there is a close relationship between the optimal policy under commitment and that under discretion. Moreover, in the case where under commitment a temptation to surprise inflation is associated with debt above the critical level and vice versa (i.e. in the case where $\sigma' > 0$), we can show that $0 < \psi' < 1$, thus ensuring that convergence under discretion is monotonic; and we can also show that θ' and e'' are negative. If $k_0 > k_c$, consumption then steadily rises over time until it reaches the critical level. Along this path, inflation and debt are falling. If, on the other hand, $k_0 < k_c$, then consumption steadily falls over time, with accompanying rises in debt and inflation. By comparison, under commitment, paths with the same k_0 would involve a rise or a fall in consumption (respectively) between periods 0 and 1, and a corresponding fall or rise in debt during period 0, but they would not continue all the way to (c_c,k_c). The reason why the evolution is carried further under discretion is that $e' < 0$ at $k_0 > k_c$ and $e' > 0$ at $k_0 < k_c$, as follows from the fact that $e'' < 0$ at the TCSS. $e' < 0$ means that there is an ‘Obstfeld effect’, as described above, so that the government perceives additional benefits of debt decumulation when unable to commit. $e' > 0$, by contrast, implies that there is a ‘reverse Obstfeld effect’: now a marginal increase in end-of-period debt *raises* current consumption through the effect on agents’ expectations, giving the government an incentive to accumulate debt relative to the case in which it can commit. Intuitively, $e' > 0$ here because this is the region of k_0 in which, under commitment, the temptation is to create surprise *deflation*. In this region, for a government following the optimal committed policy, after one period the higher debt which it faces would create an incentive for it to choose consumption to be greater than in the plan: under discretion, this incentive is reflected in the way debt influences consumption through expectations.[16]

Figure 4 provides a numerical example of a set of discretionary equilibria (one for each initial value of k_t) converging on a TCSS. This is calculated for the quadratic utility function with $\hat{c} = 0.08$, $\alpha = 0.03619$ and $\beta = 0.95$. (Again, $g_t = 0$ for all t.) The algorithm used is again that described in Appendix C. The implied ‘critical’ value c_c, i.e. where $\sigma(c) = 1$, is 0.04, and correspondingly $k_c = 0.04$. The $e(k_t)$ function for this example is depicted in panel (a). As predicted by the theory, it shows that the gradient is zero at k_c and the shape is concave: i.e. small increases in debt here affect consumption via the public’s expectations in opposite ways, depending on whether debt is high or low. Panel (b) plots the total effect of debt on the government’s optimal choice of c_t. The dynamics of debt are illustrated in panel (c). To make the picture clearer, we plot k_{t+1}-k_t on the vertical axis, since k_{t+1} as a function of k_t turns out to be very close to the 45^o line. It can be seen that the TCSS at k_c is indeed stable, even if convergence is very slow. Panel (c) also shows the FRSS (at $k = -0.0438$) and so reveals how it is unstable. For comparison, in panel (b) we also plot the two consumption levels which result from the problem under commitment (cf. Figure 3). We see that while, for debt above

[16] Having established that $e' = 0$, $e'' < 0$ in the neighbourhood of the TCSS, it is straightforward to check and confirm that the second-order conditions for the government’s optimisation problem are satisfied when $e(.)$ has these properties. In the neighbourhood of the FRSS we cannot determine e'' analytically. However our numerical experiment confirms that a maximum exists for the chosen parameter values, and also suggests that under CRRA utility $e'' = 0$ more generally may not be a bad approximation. It is again straightforward to check and confirm that the second-order conditions are satisfied when $e(.)$ is linear.

k_c, consumption in the short run is lower under discretion than commitment, for debt below k_c the opposite is true.

Our results demonstrate the main claim made in the Introduction, which is that it is not inevitable that under discretion the optimal policy will converge on the Friedman Rule. As we have just argued, the force which may prevent the attainment of the Friedman Rule, or even drive the economy away from it, is the fact that under commitment there can be an incentive to surprise deflation rather than surprise inflation.

It may be remarked that, although the conditions for local convergence to the FRSS and TCSS are not the same, it turns out they appear similar when viewed in terms of the 'expansion path' diagram. For local convergence to be possible, at either steady state the expansion path must cut, or meet, the 45^o line from above. This is because the stability conditions in Propositions 1 and 2 are the same as those governing the slopes of the expansion path at the relevant points, as is clear from the results in Section 3. Thus, if the expansion path has the shape illustrated in Figure 2, then, under discretionary policy, local convergence to the FRSS is not possible, but local convergence to the TCSS is. If instead the expansion path lies everywhere above the 45^o line - except at the Friedman Rule point where it must meet it - then the FRSS is the only steady state under discretion, and local convergence to it is possible.[17]

5. Conclusions

Our main findings having been summarised in the Introduction, we here comment on their generality. First, it is not essential to the results that we have used a cash-in-advance framework. A common alternative is the money-in-the-utility function approach. If this is adopted, then, in order to capture a welfare cost of current inflation, an analogous modelling device to the use of 'Svensson's timing' in cash-in-advance is to use beginning-of-period nominal balances deflated by the current price as the real balances variable in the utility function (as is done, in a different application, by Neiss (1999)). In early work we employed

[17] These are the only two possible shapes for the expansion path under quadratic or CARA utility.

such an approach and obtained results which exactly parallel those presented here. In fact, it can be shown that the two approaches are mathematically equivalent. However the interpretation required of the functions is different: what matters in the money-in-the-utility function approach is the shape of the utility-of-real-balances function, rather than the shape of the utility-of-consumption function.

Second, the precise conditions which we have obtained in this paper for a time-consistent steady state of the discretionary equilibrium to exist, and for convergence to it to be possible, are ones which we would not expect to be robust to other ways of modelling the welfare cost of current inflation. The way the latter is represented here is simple and convenient but there are more sophisticated ways in which it could be captured, such as in the 'limited participation' models referred to earlier. Nevertheless we would still expect that a 'time-consistent steady state', as we have defined it, could occur in such models, albeit with modifications to the exact conditions for its existence. Hence we anticipate that our conclusion that optimal monetary policy under discretion does not necessarily converge on the Friedman Rule would still apply in more general models.

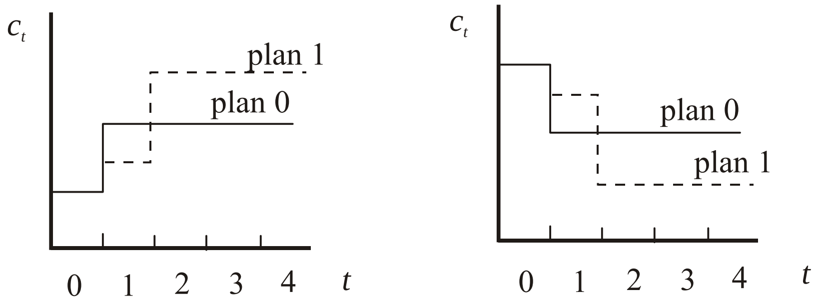

Figure 1 Optimum time paths of consumption under commitment

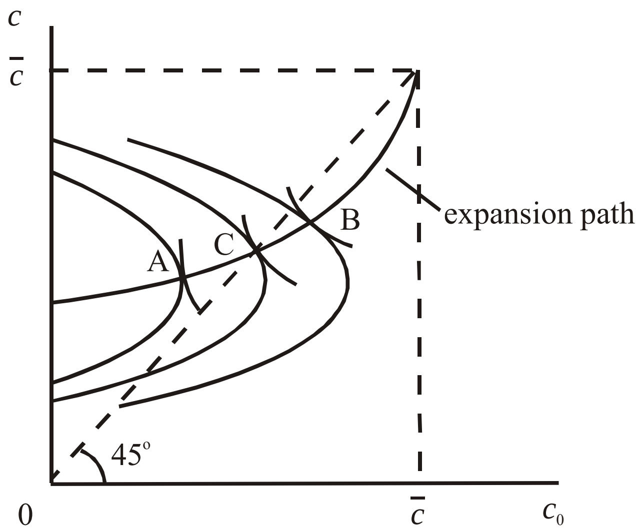

Figure 2 The government's optimisation problem under commitment in reduced form

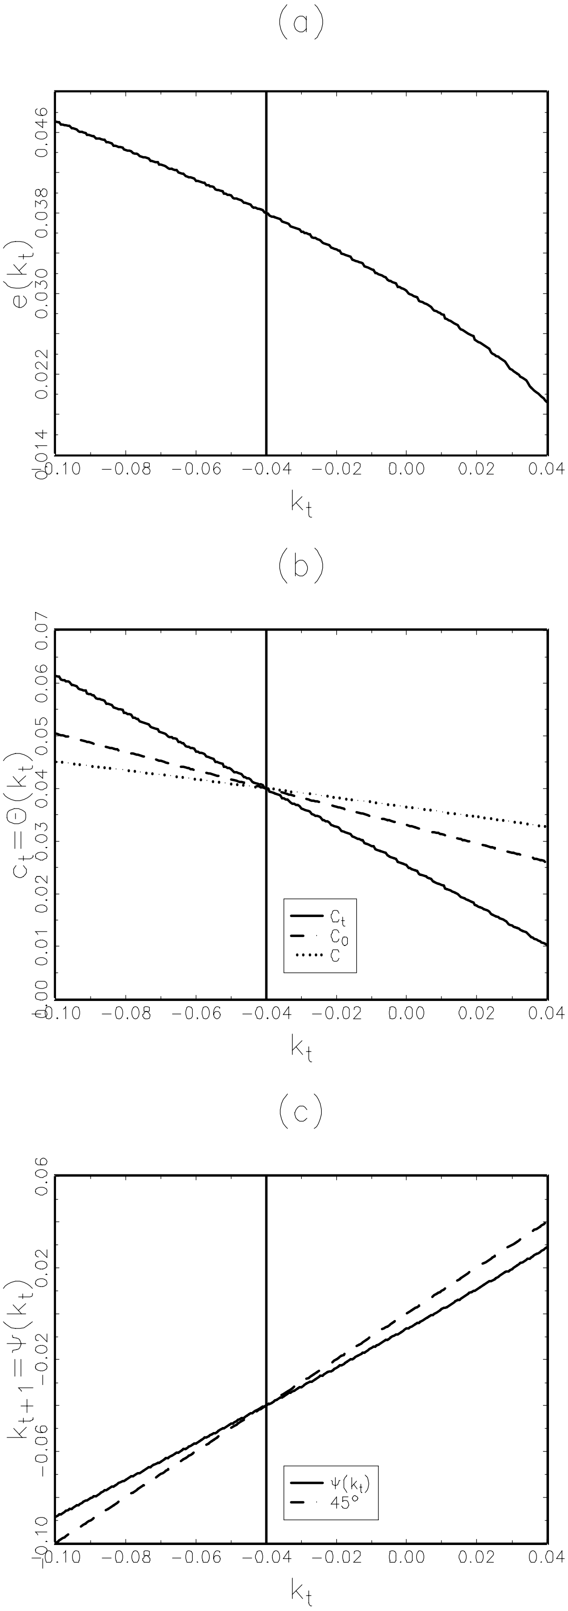

FIGURE 3 Discretionary equilibrium with CRRA utility
$\sigma = 0.5$, $\alpha = 5$, $\beta = 0.95$; FRSS value of $k_t = -0.04$

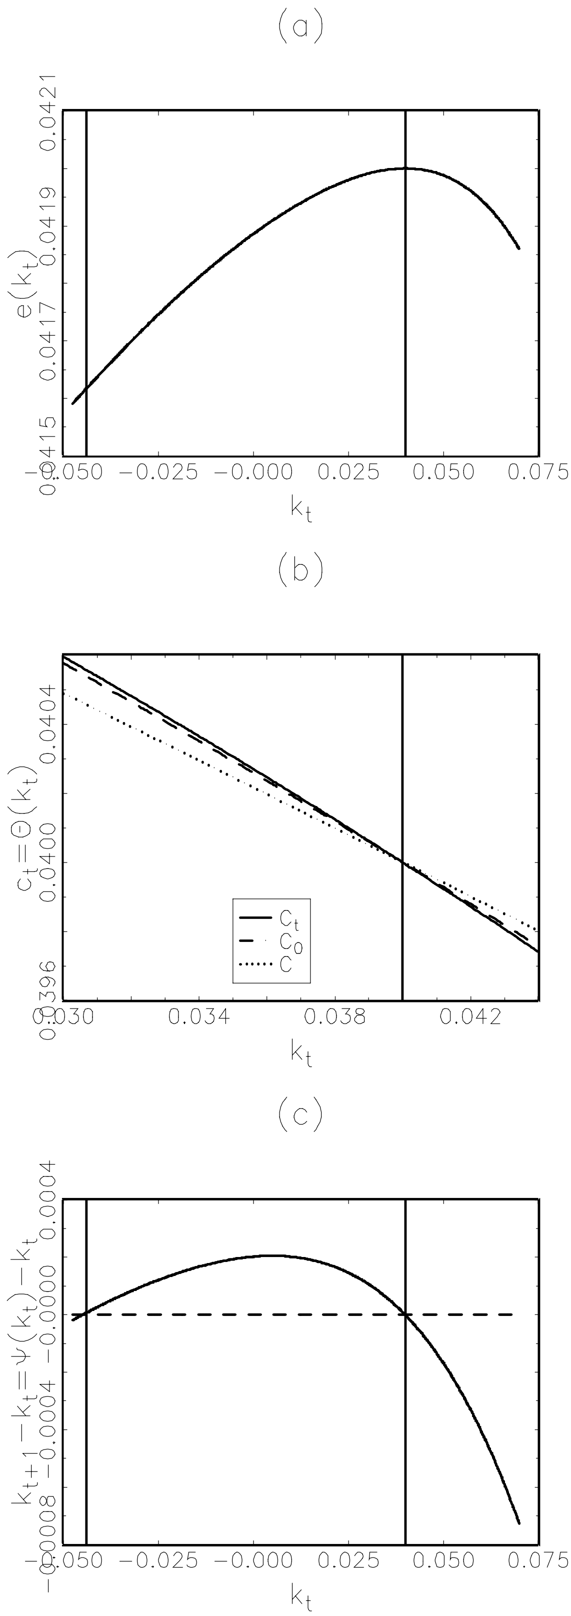

FIGURE 4 Discretionary equilibrium with quadratic utility

$\hat{c}$ =0.08, α = 0.03619, β = 0.95; TCSS value of k_t = 0.04; FRSS value of k_t = -0.0438

References

Diaz-Gimenez, J., Giovanetti, G., Marimon, R. and Teles, P. (2002) ‘Nominal Debt as a Burden on Monetary Policy’, unpublished paper, Dept of Economics, Universidad Carlos III de Madrid

Grossman, S. and Weiss, L. (1983) ‘A Transactions-Based Model of the Monetary Transmission Mechanism’, *American Economic Review* 73, 871-880

Herrendorf, B. (1997) ‘Time Consistent Collection of Optimal Seigniorage: A Unifying Framework’, *Journal of Economic Surveys* 11, 1-46

Lucas, R.E. (1982) ‘Interest Rates and Currency Prices in a Two-Country World’, *Journal of Monetary Economics* 10, 335-359

Lucas, R.E. and Stokey, N. (1983) ‘Optimal Fiscal and Monetary Policy in an Economy without Capital’, *Journal of Monetary Economics* 12, 55-93

Neiss, K.S. (1999) ‘Discretionary Inflation in a General Equilibrium Model’, *Journal of Money, Credit and Banking* 31, 357-374

Nicolini, J.P. (1998) ‘More on the Time Consistency of Monetary Policy’, *Journal of Monetary Economics* 41, 333-350

Obstfeld, M. (1991) ‘A Model of Currency Depreciation and the Debt-Inflation Spiral’, *Journal of Economic Dynamics and Control* 15, 151-177

Obstfeld, M. (1997) ‘Dynamic Seigniorage Theory: An Exploration’, *Macroeconomic Dynamics* 1, 588-614

Rotemberg, J. (1984) ‘A Monetary Equilibrium Model with Transactions Costs’, *Journal of Political Economy* 92, 40-58

Svensson, L.E.O. (1985) ‘Money and Asset Prices in a Cash-in-Advance Economy’, *Journal of Political Economy* 93, 919-944

Appendices

A. Dynamics in the neighbourhood of the Friedman-Rule steady state

As noted in the main text, setting $U' = \alpha$ causes e'' to drop out of (32). The square root on the right-hand side of (32) can then be evaluated exactly, so that (32) reduces to:

$$1+\theta' \quad = \quad \frac{1}{\beta}(\beta+e')^2 \quad \text{or} \quad 1\,. \tag{A1}$$

The solution $1+\theta' = 1$ we discard, since it is inconsistent with local convergence. To see this, note that in combination with (29) it implies $e' = 0$, and hence (30) implies $\psi' = 1/\beta$, which has absolute value greater than one. We focus, then, on the other solution in (A1). Substituting out e' using (29) and re-arranging, we arrive at the following quadratic equation in θ' as the single unknown, as referred to in the main text:

$$[1-\sigma]^2(\theta')^2 \ + \ [2(1-\sigma)-1/\beta]\theta' \ + \ [1-1/\beta] \quad = \quad 0\,. \tag{A2}$$

Rather than study the implications of this for θ', it is of greater interest to study its implications for ψ'. Using the solution from (A1) in (30), we have:

$$\begin{aligned} \psi' \quad &= \quad (\beta+e')/\beta \\ &= \quad 1+(1-\sigma)\theta'\,, \end{aligned} \tag{A3}$$

where the second line again employs (29). We may now use (A3) to re-express (A2) as a quadratic equation in ψ':

$$\beta(\psi')^2 \ - \ \frac{1}{1-\sigma}\psi' \ + \ \frac{\sigma}{1-\sigma} \quad = \quad 0\,. \tag{A4}$$

Denoting this equation schematically as $a(\psi')^2 + b\psi' + c = 0$, we proceed to some standard tests to determine its implications for ψ'. First, the roots are real if $b^2-4ac > 0$, which we may easily verify always to be satisfied. Second, their sum is $-b/a = 1/\beta(1-\sigma)$ and their product is $c/a = \sigma/\beta(1-\sigma)$, whence $\sigma < 1$ implies two positive roots and $\sigma > 1$ implies one positive and one negative. Third, consider the two test parameters $(a+b+c)/(a-b+c)$ and $(c-a)/(a-b+c)$. If the first test parameter is negative, there is exactly one root with absolute value less than one; if it is positive, then there are either 0 or 2 roots with absolute value less than one as, respectively, the second test parameter is positive or negative. We have:

$$\frac{a+b+c}{a-b+c} \;=\; \frac{\beta-1}{\beta \;-\; (\sigma+1)/(\sigma-1)},$$

$$\frac{c-a}{a-b+c} \;=\; \frac{\sigma/(1-\sigma) \;-\; \beta}{\beta \;-\; (\sigma+1)/(\sigma-1)}.$$

If $\sigma < 1$, the first test parameter is negative, so there is exactly one root with absolute value less than one. (From the foregoing, we can moreover say that this root must be positive.) If $\sigma > 1$, the first test parameter is positive, and the second is also positive, so there are no roots with absolute value less than one. These are the results asserted in the main text.

Having seen that when $\sigma < 1$ there exists a $\psi' \in (0,1)$, it follows from (A3) that the associated θ' and e' are negative, as also asserted in the main text.

B. Dynamics in the neighbourhood of the time-consistent steady state

As noted in the main text, setting $\sigma = 1$ and $e' = 0$ causes θ'' to drop out of (31), so that (31) and (32) then constitute a system of two equations in the two unknowns e'' and θ'. Using (31) to substitute e'' out of (32) and re-arranging, we obtain a cubic equation in θ'. However, $\theta' = 0$ turns out to be one root of this equation. We discard this root, since it is inconsistent with local convergence. (To see this, note that $\theta' = 0$ together with $e' = 0$ imply, from (30), that $\psi' = 1/\beta$, which has absolute value greater than one.) The remaining equation in θ' is then the quadratic:

$$[U'/\alpha-1]c\sigma'(1+\theta')^2 \;-\; [(U'/\alpha-1)c\sigma'+1](1+\theta') \;+\; \beta \;=\; 0, \qquad \text{(B1)}$$

as referred to in the main text.

Rather than study the implications of this for θ', it is of greater interest to study its implications for ψ'. Using $e' = 0$ in (30) we have $\psi' = (1+\theta')/\beta$. This enables us to re-express (B1) as a quadratic equation in ψ':

$$[U'/\alpha - 1]c\sigma'\beta(\psi')^2 \; - \; [(U'/\alpha - 1)c\sigma' + 1]\psi' \; + \; 1 \;\; = \;\; 0\,. \qquad \text{(B2)}$$

Next, notice that the term $(U'/\alpha - 1)c\sigma' + 1$ in this equation is equal to the inverse slope of the expansion path at the critical consumption level c_c, as can be seen from (14). Hence, denoting the slope in (14) as s, we can further re-write (B2) as:

$$\beta(\psi')^2 \; - \; \frac{1}{1-s}\psi' \; + \; \frac{s}{1-s} \;\; = \;\; 0\,. \qquad \text{(B3)}$$

Our concern is now with the implications of (B3) for ψ', and especially with whether any of the roots of (B3) have absolute value less than one. At this point it may be observed that (B3) is identical in form to the equation (A4) analysed above, but with ‘s’ in (B3) replacing ‘σ’ in (A4). Therefore the conclusions reached above also apply here. That is, if $s < 1$, there is exactly one root with absolute value less than one (and we can moreover say that this root must be positive). If $s > 1$, there are no roots with absolute value less than one. Given that s is < 1 or > 1 as $\sigma' > 0$ or < 0, respectively, then, translated into statements about σ', the conclusions are that if $\sigma' > 0$, there is exactly one root with absolute value less than one; and if $\sigma' < 0$, there are no roots with absolute value less than one. These are the results asserted in the main text.

Having seen that when $\sigma' > 0$ there exists a $\psi' \in (0,1)$, it follows from $\psi' = (1+\theta')/\beta$ that the associated θ' is negative, and from (31) that e'' is negative, as also asserted in the main text.

C. The algorithm used for numerical computations

The value function for the policy problem (24) - call it $V(k_t)$ - must satisfy the Bellman equation:

$$V(k_t) \; = \; \max_{c_t} \left[U(c_t) - \alpha c_t - \alpha g + \beta V(k_{t+1}) \;\text{ s.t. }\; k_t = e(k_{t+1}) + \beta k_{t+1} - c_t \right].$$

$\theta(k_t)$ - (27) in the main text - is the associated policy function. If $e(k_{t+1})$ were a known, exogenous function we could compute $V(k_t)$ by the standard method of iterating on the value function. According to this method we begin by making a guess at the value function: $V_0(k_t)$, say. This function is approximated numerically on a grid, and the maximisation operation described on the RHS of the Bellman equation is carried out by a grid search routine. In this way a new value function, $V_1(k_t)$, is generated. The process is then repeated N times until the function has converged, as judged by an appropriate criterion of convergence.

In our case, however, $e(k_{t+1})$ is an unknown, endogenous function. It must satisfy (20) in the main text. This, then, suggests a 'double iteration' procedure. We first make a guess at $e(k_{t+1})$: call it $e^0(k_{t+1})$. Using this in the policy problem (24), we then employ the above value function iteration method to compute the corresponding value function - $V^0(k_t)$, say. Associated with this is a policy function, which we may denote $\theta^0(k_t)$. In the next stage, we generate a new version of $e(k_{t+1})$ by using (20):

$$e^1(k_{t+1}) = [\beta / \alpha] U'(\theta^0(k_{t+1}))\theta^0(k_{t+1})$$

$e^1(k_{t+1})$ is then used to generate a new version of the policy problem, and the process loops round again. This higher-level iteration continues until successive versions of the function $e(.)$ are judged to have converged. Helpful guesses for the initial function $e^0(k_{t+1})$ can be obtained from Propositions 1 and 2. In particular we can directly calculate a linear approximation to $e(k_{t+1})$ which is satisfied at either the FRSS or the TCSS, and which has the correct slope.

The above algorithm, implemented in Gauss, was used to generate the numbers underlying Figures 3 and 4.