DISCUSSION PAPER SERIES

No. 2664

TIME CONSISTENCY WHEN OPEN MARKET OPERATIONS ARE THE MONETARY POLICY INSTRUMENT: IS THERE REALLY A DEFLATION BIAS?

Neil Rankin

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

An online version of this Paper can be found at www.cepr.org/pubs/dps/DP2664.asp

ISSN 0265-8003

TIME CONSISTENCY WHEN OPEN MARKET OPERATIONS ARE THE MONETARY POLICY INSTRUMENT: IS THERE REALLY A DEFLATION BIAS?

Neil Rankin, University of Warwick and CEPR

Discussion Paper No. 2664
January 2001

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 2664

January 2001

ABSTRACT

Time Consistency When Open Market Operations are the Monetary Policy Instrument: Is There Really a Deflation Bias?*

We re-examine optimal monetary policy in a dynamic general equilibrium model where open market operations are the only policy instrument. The government optimizes purely over private agents' welfare. We use a money in the utility function approach with a welfare cost of 'current' inflation. Under commitment, for the most plausible specification time inconsistency takes the form of surprise inflation, if there is high initial government debt. Although 'orthodox', this result contradicts Nicolini's related analysis, in which surprise deflation is the main finding. Under discretion, we find that the long-run inflation rate is quite likely to be positive, not negative as in Obstfeld's related analysis.

JEL Classification: E52, E61
Keywords: inflation bias, monetary policy, open market operations, optimal seigniorage, time consistency

Neil Rankin
Department of Economics
University of Warwick
Coventry CV4 7AL
UK
Tel: (44 24) 76523 470
Fax: (44 24) 76523032
Email: n.rankin@warwick.ac.uk

* I acknowledge helpful conversations with many people: Berthold Herrendorf, Henrik Jensen, Jonathan Thomas and Ted To, to name just a few. Thanks also to the participants in numerous seminars. All errors and opinions are of course my own.

Submitted 20 November 2000

NON-TECHNICAL SUMMARY

The time consistency aspect of monetary policy, and the perceived inflation bias that it explains, have been the most researched aspect of monetary policy in the last two decades. Two basic sources of time inconsistency have been appealed to: first, the short-run Phillips curve trade-off, usually interpreted as embodying temporary nominal rigidities (a classic reference being Barro and Gordon (*Journal of Political Economy* 1983)); and, second, the use of surprise inflation as a means of generating non-distorting tax revenue when conventional lump sum taxes are impossible (a classic reference being Lucas and Stokey (*Journal of Monetary Economics* 1983)). Work on the first is much better known, but it can be criticised because it typically uses *ad hoc* models and *ad hoc* government objective functions. Work on the second, by contrast, belongs in the public finance tradition, where general equilibrium models are used, and private agents' utility, i.e. welfare, provides the government's objective function. Recently, the growing use of dynamic general equilibrium (DGE) models in macroeconomics has spawned some interesting attempts to study the Barro-Gordon problem using welfare as the government's objective function (for example Ireland (*Journal of Economic Dynamics & Control*, 1997) and Neiss (*Journal of Money, Credit & Banking*,1999)). These come close to being pure 'public finance' treatments. Nevertheless, a feature which prevents them from completely treating the problem as a public finance one is that they assume the possibility of lump sum taxes or transfers: the instrument for changing the money supply is taken to be lump sum cash handouts to (or levies on) private agents. Although this is standard in many DGE analyses of monetary policy, it is unrealistic, since in practice the instrument for changing the money supply is open-market operations, i.e. purchases or sales of government debt in exchange for money. Moreover, if lump sum taxation is possible, why not use it to finance subsidies to remove the underlying distortions to the real economy which give rise to the Barro-Gordon problem in the first place? These observations give additional motivation to the Lucas-Stokey strand of research. Before the two strands can be merged, however, a major puzzle concerning the latter remains.

Although the original finding of Lucas and Stokey was that a government would have an incentive to choose an inflation rate higher than it had previously announced, thereby acquiring real revenue without distorting agents' decisions at the margin, Nicolini (*Journal of Monetary Economics*, 1998) has recently suggested that the incentive may instead be to create *lower* inflation than previously announced: surprise *deflation* rather than surprise inflation. This is the sense in which his result is one of 'deflation bias'. The government's problem, as analysed in this setting, is to allocate optimally over time the distortions which inflation causes to agents' holdings of real money balances, subject to the present value of seigniorage revenue adding up to equal the outstanding initial government debt, plus the present value of any exogenous government spending obligations. Government debt is thus a medium for ‘inflation tax smoothing’. Nicolini modifies Lucas and Stokey’s set-up of the problem to avoid an extreme feature of the latter, namely that the optimal monetary policy is to drive the current inflation rate to infinity. His modification is to change the timing assumptions of the cash-in-advance model so as to introduce a welfare cost of current inflation, something which is lacking in Lucas-Stokey. The welfare cost is what restrains the government from setting current inflation to infinity, but it also has the unorthodox consequence that, for the most plausible parameter values, the government’s incentive is to renege on its previously promised inflation rate in the *downwards* direction. This inverts all the usual conclusions about time inconsistency in monetary policy, even though it has the great virtue of making the policy problem well defined without needing extra *ad hoc* assumptions. Hence it naturally raises the question of how robust such a finding is.

In this Paper we argue that the surprising ‘deflation bias’ finding is not very robust, and that an alternative specification of the problem, while still close in spirit to Nicolini’s, restores the possibility of a more orthodox ‘inflation bias’ finding. We also offer more intuitive insight into what is driving the possibility of ‘deflation bias’. Further, we show that the initial level of outstanding government debt is important for whether or not there is inflation bias in our setting, unlike in Nicolini’s. For the most reasonable specification, a low level of initial debt may generate deflation bias, but a high level generates inflation bias. It also emerges that there exists a critical level of initial debt such that there is *no* time inconsistency.

Nicolini concerns himself only with the case where monetary policy is fully credible. We begin with this case, but then proceed to study what happens in the opposite case where the government has no credibility, so that monetary policy is completely discretionary. Under discretionary policy, the public’s forecasts of inflation are based on the observed level of government debt, not on the government’s announced path of the future money supply. We find that, for the most reasonable specification, the economy converges over time to a steady state at the critical debt level referred to above. This part of our analysis is comparable to work by Obstfeld (*Journal of Economic Dynamics & Control*, 1991; *Macroeconomic Dynamics*, 1997). Obstfeld, however, is obliged to introduce an *ad hoc* modification of private agents’ utility function in order to obtain a government objective function which generates a well defined problem. This is not necessary under our approach, which is purely welfare-based. Obstfeld finds that in the long run the economy converges on the ‘Friedman Rule’, where the inflation rate is negative and government debt is also negative. Thus he, too, obtains a curious ‘deflation bias’ result (though in a somewhat different sense from earlier): a result which says that in the long run there is no need to worry about inflation, even though the government cannot commit its policy. Our finding is that this result is also not robust: in our analysis there is no reason why the long-run inflation rate has to be negative, or indeed the level of government debt. The orthodox view that inability to commit policy may have inflationary consequences, even in the long run, is restored in our analysis. It is moreover restored while avoiding *ad hoc* elements in the government's objective function.

In summary, then, this Paper helps to resolve two 'deflation bias' puzzles which exist in recent work on time consistency in monetary policy when open market operations are the policy instrument. More broadly, it is also a contribution to the objective of bringing the analysis of the time consistency problem into the world of DGE models, and of subjecting it to the discipline of the public finance approach to policy.

The modelling innovation, which underlies the Paper, is a variant on that used by Nicolini. Nicolini used a cash in advance framework and modified the timing assumptions so that agents could be temporarily unable to adjust their cash balances; we make a similar modification using the money in the utility function framework. Our assumption is the same as that used by Neiss: that it is money held at the *beginning*, not the end, of each period which enters, deflated by the price, the utility function for that period. This apparently minor alteration of the specification causes current inflation to have a welfare cost, since, if agents' beginning-of-period money holdings are predetermined, a rise in the current price level causes a shortage of the real balances needed for making current goods purchases, and thus a loss of liquidity services. The difference between our results and Nicolini's arises because the money in the utility function framework produces a more flexible money demand function.

1. Introduction

The time-consistency aspect of monetary policy, and the perceived inflation bias which it explains, has been the most researched aspect of monetary policy in the last two decades. Two basic sources of time inconsistency have been appealed to: first, the short-run Phillips curve trade-off, usually interpreted as embodying temporary nominal rigidities (a classic reference being Barro and Gordon (1983)); and, second, the use of surprise inflation as a means of generating non-distorting tax revenue when conventional lump-sum taxes are impossible (a classic reference being Lucas and Stokey (1983)). Work on the first is much the better known, but it can be criticised because it typically uses ad hoc models and ad hoc government objective functions. Work on the second, by contrast, belongs in the public finance tradition, where general equilibrium models are used, and private agents' utility, i.e. welfare, provides the government's objective function. Recently, the growing use of dynamic general equilibrium (DGE) models in macroeconomics has spawned some interesting attempts to study the Barro-Gordon problem using welfare as the government's objective function (for example Ireland (1997) and Neiss (1999)). These come close to being pure 'public finance' treatments. Nevertheless, a feature which prevents them from completely treating the problem as a public finance one is that they assume the possibility of lump-sum taxes or transfers: the instrument for changing the money supply is taken to be lump-sum cash handouts to (or levies on) private agents. Although this is standard in many DGE analyses of monetary policy, it is unrealistic, since in practice the instrument for changing the money supply is open-market operations, i.e. purchases or sales of government debt in exchange for money. Moreover, if lump-sum taxation is possible, why not use it to finance subsidies to remove the underlying distortions to the real economy which give rise to the Barro-Gordon problem in the first place? These observations give additional motivation to the Lucas-Stokey strand of research. However, before the two strands can be merged, a major puzzle concerning the latter remains.

Although the original finding of Lucas and Stokey was that a government would have an incentive to choose an inflation rate higher than it had previously announced, thereby acquiring real revenue without distorting agents' decisions at the margin, Nicolini (1998) has recently suggested that the incentive may instead be to create *lower* inflation than previously announced: surprise *deflation* rather than surprise inflation. This is the sense in which his result is one of 'deflation bias'. The government's problem, as analysed in this setting, is to allocate optimally over time the distortions which inflation causes to agents' holdings of real money balances, subject to the present value of seigniorage revenue adding up to equal the outstanding initial government debt, plus the present value of any exogenous government spending obligations. Government debt is thus a medium for 'inflation tax smoothing'. Nicolini modifies Lucas and Stokey's set-up of the problem to avoid an extreme feature of the latter, namely that the optimal monetary policy is to drive the current inflation rate to infinity. His modification is to change the timing assumptions of the cash-in-advance model so as to introduce a welfare cost of current inflation, something which is lacking in Lucas-Stokey. The welfare cost is what restrains the government from setting current inflation to infinity, but it also has the unorthodox consequence that, for the most plausible parameter values, the government's incentive is to renege on its previously promised inflation rate in the *downwards* direction. This inverts all the usual conclusions about time inconsistency in monetary policy, even though it has the great virtue of making the policy problem well-defined without needing extra ad hoc assumptions[1]. Hence it naturally raises the question of how robust such a finding is.

In this paper we argue that the surprising 'deflation bias' finding is not very robust, and that an alternative specification of the problem, while still close in spirit to Nicolini's, restores the possibility of a more orthodox 'inflation bias' finding. We also offer more intuitive insight into what is driving the possibility of 'deflation bias'. Further, we show that the initial level of outstanding government debt is important for whether or not there is inflation bias in our setting, unlike in Nicolini's. For the most reasonable specification, a low level of initial debt may generate deflation bias, but a high level generates inflation bias. It

[1] A number of studies of the optimal seigniorage problem have used ad hoc assumptions to make it well-defined, such as postulating a government objective function different from private agents' utility. See, for example, Obstfeld (1991, 1997) (also discussed below), or Barro (1983). Nicolini's study is notable because it is the first to obtain a non-degenerate solution within a pure public finance approach.

also emerges that there exists a critical level of initial debt such that there is *no* time inconsistency.

Nicolini concerns himself only with the case where monetary policy is fully credible. We begin with this case, but then proceed to study what happens in the opposite case where the government has no credibility, so that monetary policy is completely discretionary. Since government debt is a state variable in this economy, discretionary policy is modelled using Markov-perfect equilibrium. We find that, for the most reasonable specification, the economy converges over time to a steady state at the critical debt level referred to above. This part of our analysis is comparable to work by Obstfeld (1991, 1997). Obstfeld, however, is obliged to introduce an ad hoc modification of private agents' utility function in order to obtain a government objective function which generates a well-defined problem. This is not necessary under our approach, which is purely welfare-based. Obstfeld finds[2] that in the long run the economy converges on the 'Friedman Rule', where the inflation rate is negative, and government debt is also negative. Thus he, too, obtains a curious 'deflation bias' result (though in a somewhat different sense from earlier): a result which says that in the long run there is no need to worry about inflation, even though the government cannot commit its policy. Our finding is that this result is also not robust: in our analysis there is no reason why the long-run inflation rate has to be negative, nor indeed the level of government debt. The orthodox view that inability to commit policy may have inflationary consequences, even in the long run, is restored in our analysis. It is moreover restored while avoiding ad hoc elements in the government's objective function.

In summary, then, this paper helps to resolve two 'deflation bias' puzzles which exist in recent work on time consistency in monetary policy when open market operations are the policy instrument. More broadly, it is also a contribution to the objective of bringing the analysis of the time consistency problem into the world of DGE models, and of subjecting it to the discipline of the public finance approach to policy.

---

[2] Here we refer to the case of his analysis which uses a version of the government's objective function which is as close as possible to private utility functions.

The modelling innovation which underlies the paper is a variant on that used by Nicolini (1998). Nicolini used a cash-in-advance framework and modified the timing assumptions so that agents could be temporarily unable to adjust their cash balances; we make a similar modification using the money-in-the-utility-function framework. Our assumption is the same as that used by Neiss (1999): that it is money held at the *beginning*, not the end, of each period which enters, deflated by the price, the utility function for that period. This apparently minor alteration of the specification causes current inflation to have a welfare cost, since, if agents' beginning-of-period money holdings are predetermined, a rise in the current price level causes a shortage of the real balances needed for making current goods purchases, and thus a loss of liquidity services. The difference between our and Nicolini's results arises because the money-in-the-utility-function framework produces a more flexible money demand function, as we explain in more detail in the body of the paper.

The body of the paper is organised as follows. Section 2 describes the structure of the economy; Section 3 looks at optimal monetary policy under precommitment; Section 4 looks at optimal monetary policy under discretion; and Section 5 concludes.

2. Structure of the Economy

The economy is populated by identical infinitely-lived households who consume the single type of output good, supply a single type of labour to firms, and hold money and bonds. Firms produce output using labour as the sole input. Markets are perfectly competitive and all prices are flexible. The government issues money and bonds. Bonds are taken to be 'real', or indexed, bonds, as in the papers already cited.[3] 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. To keep the focus purely on monetary policy, we shall also ignore conventional distorting taxes, and government spending on goods and services.

In more detail, the optimisation problem of the representative household is:

[3] Although it would be more 'realistic' to assume nominal bonds, this would obviously make it easier to generate a 'surprise inflation' result, so assuming real bonds cannot be said to favour our main conclusion.

$$\underset{\{c_t, l_t, M_t, d_t\}_0^\infty}{\text{maximise}} \quad \Sigma_{t=0}^{\infty}\beta^t u(c_t, l_t, m_t) \qquad (m_t \equiv M_{t-1}/p_t) \tag{1}$$

$$\text{s.t.} \quad d_{t-1} + M_{t-1}/p_t + w_t l_t + \pi_t = c_t + M_t/p_t + q_t d_t, \qquad t = 0,\ldots,\infty, \tag{2}$$

M_{-1}, d_{-1} given.

The variables here are: M_t, the stock of money held at the end of period t; d_t, the stock of government bonds held at the end of period t; w_t, the real wage; l_t, labour supply; π_t, real profits received from firms; c_t, consumption; p_t, the price of goods; and q_t, the price of a bond, where the latter is a promise of one unit of goods in one period's time. As in the tradition stretching back to Sidrauski (1967), real balances are included in the utility function to represent their role in helping consumers to buy goods. The less common feature of our specification is that the real balances entering the utility function are defined as money held at the *start* of period t, deflated by p_t, rather than as money held at the end of period t, deflated by p_t.[4] This is the feature which we share with Neiss (1999), who also uses the money-in-the-utility function approach. Moreover it is analogous to the definition of 'cash' in the cash-in-advance approach used by Nicolini (1998).[5] The definition captures the idea that there may be a temporary inability on the part of consumers to adjust their portfolios when confronted by an increase in the price level. Since M_{t-1} is predetermined in period t, a rise in p_t lowers m_t, imposing a welfare cost of current inflation on consumers because they find themselves short of cash with which to make goods purchases. A more developed version of the same idea also used by Nicolini is one in which consumers are divided into two groups, each group being allowed to visit the bank only in alternate periods. Here we shall work with the compact version, thereby avoiding the complications which would be introduced by heterogeneity amongst consumers.

The first-order conditions for this problem can readily be derived as:

[4] This second definition is nearly universal in the literature. Just to cite one example out of many, see the recent graduate text by Walsh (1998).

[5] Here we refer to the 'compact' version of the model suggested by Nicolini, rather than to the full version which allows for two groups of agents, as explained below.

$$u_c(t) = \frac{\beta}{q_t} u_c(t+1) \,, \tag{3}$$

$$w_t = -\frac{u_l(t)}{u_c(t)} \,, \tag{4}$$

$$u_m(t) = \left[\frac{1}{q_{t-1}} \frac{p_t}{p_{t-1}} - 1 \right] u_c(t) \,. \tag{5}$$

(3) is the standard consumption Euler equation (note that $1/q_t$ is the gross real interest rate), and (4) is the equally standard condition that the real wage equal the marginal rate of substitution between goods and leisure (where $u_c(t)$, etc., indicates the partial derivative of utility with respect to c in period t). In (5), inside the term in square brackets we have the product of the gross real interest rate and the gross inflation rate, which is therefore the gross nominal interest rate, i_{t-1}. Thus we can rewrite (5) as:

$$u_m(t) = [i_{t-1} - 1] u_c(t) \qquad \text{where} \qquad i_{t-1} \equiv \frac{1}{q_{t-1}} \frac{p_t}{p_{t-1}} \,. \tag{5'}$$

In this form it says that real balances should be held to the point where their marginal utility just equals the marginal utility of their opportunity cost in terms of lost consumption, the latter resulting from the fact that real balances, unlike bonds, do not pay interest.

Turning to the representative firm, its optimisation problem is to choose (l_t, y_t) to maximise profits $\pi_t = y_t - w_t l_t$ subject to the production function $y_t = f(l_t)$. This leads to the standard first-order condition $w_t = f'(l_t)$. In labour market equilibrium, using this together with the labour supply condition (4), and adding the assumption (which we shall maintain hereafter) that the utility function is additively separable, we have $f'(l_t) = -u_l(l_t)/u_c(c_t)$. Since, in goods market equilibrium, $c_t = y_t = f(l_t)$, we can write:

$$f'(l_t) = -\frac{u_l(l_t)}{u_c(f(l_t))} \,. \tag{6}$$

This equation determines l_t, and thus the equilibrium employment level in the economy. From it we can see the standard result that, under additive separability, the economy has a ‘natural rate’ of employment which is independent of monetary policy. Henceforth we denote this, and the corresponding output level, by $(\bar{l}, \bar{y})$. Note that this further implies $u_c(t) = u_c(\bar{y})$ for all t, whence (3) implies $q_t = \beta$ for all t, i.e. the equilibrium gross real interest rate is the constant, $1/\beta$.

It remains to determine equilibrium in the monetary sector of the economy. For now, suppose an arbitrary monetary policy, defined by a given sequence of money supplies $\{M_t\}_{t=0}^{\infty}$, or, equivalently, monetary growth rates $\{\mu_t\}_{t=0}^{\infty}$ (where $\mu_t \equiv M_t/M_{t-1}$). Using (5) to eliminate q_t from (4), and also the foregoing definitions and results, we arrive at the following difference equation in real balances, which gives the basic law of motion of the private-sector equilibrium:

$$m_t = \frac{\beta}{\mu_t} m_{t+1}\left[1 + \frac{u_m(m_{t+1})}{u_c(\bar{y})}\right]. \tag{7}$$

Since the endogeneity of m_t reflects the endogeneity of p_t, m_t is a non-predetermined variable. (7) must hence be solved in a forward-looking manner, in which (subject to an appropriate saddlepoint stability condition being satisfied), the current value of m_t is determined as the value which causes the resulting sequence $\{m_{t+s}\}_{s=0}^{\infty}$ to be non-divergent. In general, m_t thereby becomes a function of all current and future monetary growth rates, $\{\mu_t\}_{t=0}^{\infty}$.

Since lump-sum taxes are ruled out, the government does not have a completely free choice of the μ_t sequence: the μ_t's must satisfy its intertemporal budget constraint. First look at the government's single-period budget constraint:

$$d_{t-1} = \frac{M_t - M_{t-1}}{p_t} + q_t d_t \quad \left(\equiv [\mu_t - 1]m_{t-1} + q_t d_t\right). \tag{8}$$

This says that outstanding debt at the start of period t must be financed either by issuing money or by issuing new debt. By repeatedly using (8) advanced one period to substitute out d_t, d_{t+1},…, etc. (also noting $q_t = \beta$), and applying the 'No Ponzi Game' condition that in the limit debt cannot grow faster than the real interest rate $1/\beta$, we obtain:

$$d_{t-1} = \Sigma_{s=0}^{\infty}\beta^s[\mu_{t+s} - 1]m_{t+s}. \tag{9}$$

This form of the intertemporal constraint states that the present value of 'cash-flow' seigniorage must equal initial outstanding debt, where 'cash-flow seigniorage' is the real value of the increase in the money supply during a period, i.e. $[M_t\text{-}M_{t\text{-}1}]/p_t$ or $[\mu_t\text{-}1]m_t$. An alternative version of (9) is in terms of 'opportunity cost' seigniorage, where the latter is defined as real balances times the net nominal interest rate, $[i_{t\text{-}1}\text{-}1]m_t$: [6]

$$d_{t-1} = -m_t + \Sigma_{s=1}^{\infty}\beta^s[i_{t+s-1} - 1]m_{t+s}\,. \tag{9'}$$

(9') helps to reveal the source of the time consistency problem, because it shows how 'current' real balances enter differently from 'future' real balances. Since in equilibrium each period's real balance level and nominal interest rate is associated with that period's inflation rate, this translates into a different impact of current and and anticipated inflation on the present value of seigniorage. Thus, as time passes and 'anticipated inflation' becomes 'current inflation', an incentive may arise for the government to alter its earlier choice for inflation. The study of this change in incentive will be a central concern of the paper.

3. Optimal Monetary Policy Under Commitment

The structure of the government's optimisation problem

In this section, a given plan for monetary policy, $\{\mu_t\}_0^{\infty}$, is assumed to be fully credible by the public, so that the public's expectations of future μ_t's simply equal the μ_t's in the plan. Under this assumption, the government's optimisation problem from the perspective of period 0 may be stated as:

$$\underset{\{\mu_t\}_0^{\infty}}{\text{maximise}} \quad \Sigma_{t=0}^{\infty}\beta^t u(\bar{y},\bar{l},m_t)$$

$$\text{s.t.} \quad d_{-1} = -m_0 + \Sigma_{t=1}^{\infty}\beta^t \frac{u_m(m_t)}{u_c(\bar{y})}m_t\,, \ \ d_{-1} \text{ given,}$$

[6] See Herrendorf (1997) for a useful discussion of these two definitions. To obtain (9') from (9), note that $[M_t\text{-}M_{t\text{-}1}]/p_t$ can be rewritten as $m_{t+1}\beta i_t\text{-}m_t$; regrouping the terms in the sum then yields (9').

$$m_t = \frac{\beta}{\mu_t} m_{t+1}\left[1 + \frac{u_m(m_{t+1})}{u_c(\bar{y})}\right],\ t = 0,\ldots,\infty. \tag{10}$$

Note that we have substituted out i_{t-1} from the government's budget constraint using (5'). It is clear from the structure of the problem (10) that, rather than treat $\{\mu_t\}_0^\infty$ as the choice variables, we can equivalently treat $\{m_t\}_0^\infty$ as the choice variables, leaving the difference-equation constraints to determine $\{\mu_t\}_0^\infty$ residually. Thus we can rewrite the problem more compactly as:

$$\underset{\{m_t\}_0^\infty}{\text{maximise}} \quad \Sigma_{t=0}^\infty \beta^t u(\bar{y}, \bar{l}, m_t)$$

$$\text{s.t.} \quad d_{-1} = -m_0 + \Sigma_{t=1}^\infty \beta^t \frac{u_m(m_t)}{u_c(\bar{y})} m_t,\ d_{-1} \text{ given.} \tag{10'}$$

The nature of the problem can first be described in intuitive terms. Notice straight away that, if lump-sum taxes (τ_t, say) were permitted, the RHS of the constraint in (10') would be augmented by $\Sigma_{t=0}^\infty \beta^t \tau_t$. With these extra instruments available, there would be no obstacle to achieving the first-best solution to the problem, which is to set $m_t = \overline{m}$ for all t, where $\overline{m}$ is the 'satiation level' of real balances such that $u_m(\overline{m}) \equiv 0$ (assuming the utility function to possess such a level). This is Friedman's (1969) 'Optimum Quantity of Money' Rule, in which the nominal interest rate is zero, and there is deflation of prices ($[p_t - p_{t-1}]/p_{t-1} = \beta - 1$) and negative monetary growth ($\mu_t - 1 = \beta - 1$). In the absence of lump-sum taxation, however, this outcome is not feasible, since the revenue needed to finance continual withdrawals of money from the economy is not available[7]. The government must trade off the objective of getting m_t as close as possible to $\overline{m}$ with the objective of getting m_s as close as possible to $\overline{m}$, for any $s \neq t$. The cost of a high value of m_t is a lower value of m_s: the mechanism through which this occurs is that, to raise m_t, the government must lower the inflation rate, p_t/p_{t-1}, which lowers seigniorage revenue in period t; to compensate it must then raise seigniorage in period s by

[7] If we were to allow marginal taxation (other than the 'inflation tax'), which for simplicity we have excluded, other revenue would be available, but it would have a distortionary cost which would still make it generally suboptimal to follow the Friedman Rule, as was originally pointed out by Phelps (1973). Some authors have challenged Phelps's conclusion (for example, Chari et al. (1996)) but, to be valid, the challenge requires particular assumptions about preferences over goods and leisure. In this paper we focus on pure monetary policy and leave this debate to one side.

increasing the inflation rate p_s/p_{s-1}, which lowers m_s. For t=0, the negative relationship of inflation to real balances derives simply from the fact that M_{-1} in $m_0 \equiv M_{-1}/p_0$ is predetermined. For t≥1, it derives from private agents' money-holding first-order condition (5) (recalling that $q_t = \beta$ and $u_c(t) = u_c(\bar{y})$).

To look at this another way, observe that (5) or (5') can be inverted to obtain the 'demand for money' function:

$$m_t = u_m^{-1}\left(u_c(\bar{y})[i_{t-1} - 1]\right). \tag{11}$$

This shows that m_t is a unique function of i_{t-1}-1 (or equivalently of p_t/p_{t-1}); and it is moreover decreasing, since $u_{mm} < 0$ is required for convexity of preferences. A condition which is obviously necessary for the tradeoff just described to arise is that seigniorage revenue $[i_{t-1}-1]m_t$ be increasing in i_{t-1}-1, or, equivalently, decreasing in m_t. Intuitively, this is the condition that the economy be on the upward slope of the 'seigniorage Laffer curve'. For it to hold, we need that the net-interest elasticity of the money demand function (11) be less than one.

More formally, now, we may derive the following first-order conditions for the problem (10'):

$$\rho(m_t) = 1 + \frac{u_c(\bar{y})}{u_m(m_0)}, \ t = 1,\ldots,\infty, \qquad \text{where} \qquad \rho(m_t) \equiv -\frac{m u_{mm}(m_t)}{u_m(m_t)}. \tag{12}$$

ρ is a measure of the curvature of the subutility function over real balances (the 'relative risk aversion' parameter, although of course in the present context there is no risk). It is also the inverse of the net-interest elasticity of demand for money, as is easily shown by differentiating (11). A number of points can be seen from (12). First, an optimal monetary policy under commitment requires that all *future* levels of real balances, $(m_1,m_2,\ldots)$, be set at the same level. This is the standard 'tax smoothing' result, first noted for conventional taxes by Barro (1979). Second, *current* real balances will generally be set at a different level from future real balances. This is clear from (12) because m_0 enters in a different way from $(m_1,m_2,\ldots)$. The optimal path of real balances as a function of time thus has a 'step' shape, as we show in Figure 2 below: either a 'step up' or a 'step down'. Mathematically speaking, the reason for this is obviously that m_0 and $(m_1,m_2,\ldots)$ enter the government's intertemporal budget constraint in different ways. This is a key point, and we will discuss why they enter in different ways below. Third, (12) can only be satisfied if $\rho(m_t) > 1$: this is the condition that the optimum must occur where money demand is net-interest-inelastic, as already noted.

Since an optimal policy involves only two levels of real balances, m_0 and m_1 (where 'm_1' stands for the common m_t, $t\geq 1$), we can summarise it by the (m_0,m_1) pair which satisfies the following two equations:

$$d_{-1} = -m_0 + \frac{\beta}{1-\beta}\frac{u_m(m_1)}{u_c(\bar{y})}m_1, \tag{13}$$

$$\rho(m_1) = 1 + \frac{u_c(\bar{y})}{u_m(m_0)}. \tag{14}$$

We can also depict the government's problem in reduced form on an indifference curve diagram, as in Figure 1:

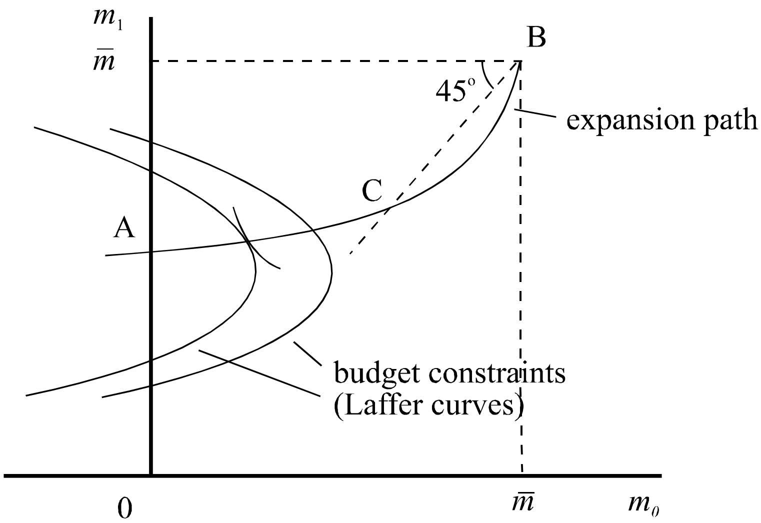

Figure 1

Here the indifference curves are implicitly defined by $U = u(\bar{y},\bar{l},m_0) + \beta[1-\beta]^{-1}u(\bar{y},\bar{l},m_1)$, the budget constraint by (13), and the locus AB by (14). Note that the budget constraint takes its shape from the seigniorage Laffer curve: as m_1 falls (corresponding to a rise in future inflation), future seigniorage revenue at first rises, enabling less current seigniorage to be raised, and hence higher m_0 to be attained. However, it may happen that beyond a certain point the Laffer curve peaks, reversing the sign of the relationship. For a given d_{-1}, the optimal (m_0,m_1) are given by the intersection of the budget constraint for that d_{-1} and AB. AB can also be interpreted as the 'expansion path', i.e. the locus of optimal (m_0,m_1) pairs traced out as initial debt, d_{-1}, varies. Observe that as d_{-1} rises, the budget constraint shifts leftwards in parallel fashion, and AB picks out the the points of tangency with the indifference curves.

A condition for when a 'step up' or 'step down' in the time path of real balances is optimal may now be given. Notice that 'step up' behaviour occurs at points on the expansion path to the left of the 45^0 line, and 'step down' behaviour at points to the right. Suppose AB intersects the 45^0 line (as in Figure 1), and moreover suppose for the moment that d_{-1} is such that the optimum occurs exactly at the point of intersection (C, in Figure 1). In this case $m_0 = m_1$. From (14), m_1 then satisfies:

$$\rho(m_1) = 1 + \frac{u_c(\bar{y})}{u_m(m_1)}. \tag{16}$$

Now, the private first-order condition for money holding, (5'), implies that the RHS of (16) equals $i_0/[i_0-1]$, whence, at point C, $[1/\rho(m_1)][i_0/(i_0-1)] = 1$. But $[1/\rho(m_1)][i_0/(i_0-1)]$ is just the gross-interest elasticity of demand for real balances m_1 (i.e. $[-\partial m_1/\partial i_0][i_0/m_1]$), so this tells us that if the optimal path for real balances is 'flat', this occurs where the gross-interest elasticity of future real balances is unity. From this it is not hard to deduce that if, instead, d_{-1} is such that the optimum is at a point on AB to the left of the 45^0 line (indicating 'step up' behaviour), the gross-interest elasticity of future real balances here must be greater than unity; and conversely if to the right. Next note that the elasticity of future real balances with respect to the gross interest rate is the same as with respect to the gross inflation rate (since $i_0 = (1/\beta)p_1/p_0$), and that the elasticity of current real balances with respect to the gross inflation rate p_0/p_{-1} is one, since it derives solely from the fact that M_{-1} is predetermined ($m_0 \equiv M_{-1}/p_0$). Thus:

Lemma 1 *Under optimal monetary policy with commitment, future real balances are less (greater) than current real balances according as the gross-inflation elasticity of future real* *balances is less (greater) than the gross-inflation elasticity of current real balances (the latter being unity).*

Roughly speaking, this corresponds to the familiar idea in optimal tax theory that the good which should be taxed most heavily, and therefore demand for which is likely to be most depressed, is the one in the most inelastic supply.[8]

We end with three further remarks. First, (14) shows that the condition for the expansion path to be upward-sloping is that $\rho'(m_t) > 0$. This is necessary if (m_0,m_1) are to be 'normal goods' for the government, i.e. if a reduction in initial debt d_{-1} is to raise both real balance levels. In fact $\rho'(m_t) > 0$ is also necessary to ensure that, in the choice between any m_t and m_s for $t,s \geq 1$, the local second-order conditions are satisfied, as we show in the Appendix. Hence we shall maintain this assumption in what follows. Second, if the utility function possesses a satiation level, $\overline{m}$, then, under mild conditions, the expansion path will tend to the point $(\overline{m},\overline{m})$ at its upper end (as drawn in Figure 1).[9] Intuitively, this 'first-best' outcome will occur if initial debt d_{-1} is sufficiently negative, i.e. if the debt is in fact a sufficiently large stock of assets, since then the government can use the (real) interest earned on these assets to finance continual monetary withdrawals, and so implement the Friedman Rule. Third, we remark that the solution defined by (11) exhibits 'path dependence', in the sense that the long-run value of m_t under an optimal policy (i.e., 'm_1') depends on the given initial debt, d_{-1}. This is different from what is found under optimal discretionary policy, as will be seen below.

The incentive to renege

We have just discussed the optimal policy from the perspective of period 0. Let the values of (m_0,m_1) chosen in the optimal plan made in period 0 be denoted $(m_{0|0},m_{1|0})$. Once period 1 has arrived, the government may have an incentive to choose an m_1 different from

---

[8] The correspondence is only 'rough', because the standard rule is about tax rates rather than quantities, and because in our case nothing is in fact implied about the relative size of current and future 'inflation tax rates': the optimal current inflation rate depends on p_{-1}, whereas the optimal level of m_0 does not. The optimal monetary policy problem (10') cannot be made completely equivalent to a standard optimal commodity tax problem, because 'current seigniorage revenue' is $-m_0$, not $[i_{-1}-1]m_0$.

[9] As can be seen from (14), $(\overline{m},\overline{m})$ will lie on the expansion path if $\rho(m_1)$ tends to infinity as m_1 tends to $\overline{m}$. In the Appendix we give a sufficient condition for this to be the case.

$m_{1|0}$.[10] Its optimal choice from the perspective of period 1 will be denoted $m_{1|1}$. We now investigate whether, and in which direction, $m_{1|1}$ differs from $m_{1|0}$. Since $m_1 \equiv M_0/p_1$ and M_0 is predetermined in period 1, clearly 'surprise inflation' corresponds to the case where $m_{1|1}$ - $m_{1|0} < 0$ and 'surprise deflation' to the case where $m_{1|1}$ - $m_{1|0} > 0$.

Consider first how debt evolves over time under plan 0. The 'step' property of the optimal plan implies that the government budget constraint in (10') applied to $\{m_{t|0}\}_0^\infty$ can be written:

$$d_{-1} = -m_{0|0} + \frac{\beta}{1-\beta}\frac{u_m(m_{1|0})}{u_c(\bar{y})}m_{1|0}. \tag{17}$$

The same budget constraint applied to the 'continuation' of this plan, i.e. to $\{m_{t|0}\}_1^\infty$, gives:

$$d_0 = -m_{1|0} + \frac{\beta}{1-\beta}\frac{u_m(m_{1|0})}{u_c(\bar{y})}m_{1|0}. \tag{18}$$

Subtracting,

$$d_0 - d_{-1} = m_{0|0} - m_{1|0}. \tag{19}$$

So debt rises (falls) during the first period of the plan if it prescribes a 'step down' ('step up') in the path of real balances. Next, (10') applied to plan 1 gives:

$$d_0 = -m_{1|1} + \frac{\beta}{1-\beta}\frac{u_m(m_{2|1})}{u_c(\bar{y})}m_{2|1}.$$

Equating this expression for d_0 with that in (14), we obtain:

$$m_{1|1} - m_{1|0} \quad = \quad \frac{\beta}{1-\beta}\frac{u_m(m_{2|1})m_{2|1} - u_m(m_{1|0})m_{1|0}}{u_c(\bar{y})}. \tag{20}$$

[10] Standard remarks apply to the interpretation of this type of investigation. If policy is truly committed, as a result of some mechanism which ties the government in, then reneging is not actually feasible, and the discussion here is hypothetical: it concerns what the government *would* do, *if* it could alter its plan. If, on the other hand, reneging is feasible, then the discussion here is not hypothetical, but the expectations held by the public in period 0 were not rational – the public were foolish to believe the government's plan. Which is the case is not our particular concern in this section; in the next section we will look explicitly at what happens when the government cannot commit, and the public know this.

As noted, surprise inflation or deflation occurs accordingly as this expression is negative or positive. Now, as seen, if plan 0 involves a ‘step down’, $d_0 > d_{-1}$. Given that real balances are ‘normal goods’ for the government, the rise in d implies that $m_{2|1} < m_{1|0}$. Recalling that $\rho > 1$, whence $u_m(m)m$ is decreasing in m, it follows that the RHS of (16) is positive. Thus:

Lemma 2 *Under optimal monetary policy with commitment, in period 1 the government will be tempted to choose a higher (lower) inflation rate than that which was planned in period 0 (i.e. to create surprise inflation [deflation]), if the plan specified a step up (step down) in the time path for real balances (i.e. specified higher [lower] real balances in period 1 than period 0).*

A visual summary of the two possible relationships between plan 0 and plan 1 is presented in Figure 2.

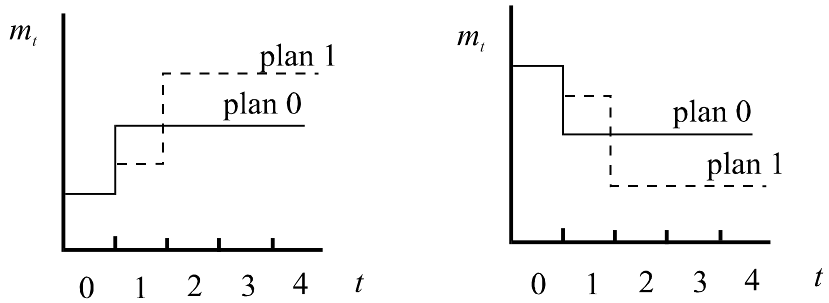

Figure 2

Linking this to our earlier analysis of the ‘expansion path’, it can be seen that if, in period 0, the economy is at a point on the expansion path to the left of the 45^0 line, then, in period 1, the government will have an incentive to create surprise inflation; while if it is to the right of the 45^0 line, the government will have an incentive to create surprise deflation. The position and shape of the expansion path are therefore key to whether surprise inflation or surprise deflation is the more likely. A possibility of particular interest is that the expansion path may cut the 45^0 line (as in Figure 1), in which case whether there is surprise inflation or surprise deflation depends on where on the path the economy lies, and thus on the level of initial government debt, d_{-1}. In the case in Figure 1, the expansion path cuts the 45^0 line 'from above'. Thus, for a sufficiently low level of initial debt, d_{-1}, the economy will be at a point on the segment CB, where in period 1 the government will have an incentive to create surprise deflation, as Nicolini (1998) found. However, for a high level of debt, the economy will be at a point on the segment AB, where in period 1 the government will be tempted to create surprise inflation. This latter is the orthodox finding in the literature on time consistency in monetary policy. In our model, then, if the expansion path appears as in Figure 1, a sufficiently high level of initial debt yields the orthodox, surprise inflation, result. By contrast, in Nicolini's model, the size of the debt is irrelevant to the sign of the inflation surprise. For the most plausible parameter values, his model always predicts surprise deflation.[11] It seems fairly intuitive that, as we find here, surprise inflation should be associated with high, rather than low, initial debt: high debt means that the need for revenue is strong, and surprise inflation has traditionally been regarded as a measure to which a government is likely to resort when hard-pressed to raise revenue.

For surprise inflation to be associated with high initial debt, we need the expansion path to cut the 45^0 line 'from above'. Earlier it was shown that points on the expansion path to the left of this line exhibit a gross-interest elasticity of demand for future real balances greater than one, and points to the right, an elasticity less than one. Given this, we can summarise the main finding of this section in the following way:

Proposition 1 *Under monetary policy with commitment, if a critical level of future real balances exists whose gross-interest elasticity of demand equals one, then there will be an associated critical initial debt level such that, if initial debt takes this value, there is no time inconsistency. If, moreover, the elasticity is increasing (decreasing) in the interest rate at this point, then initial debt above the critical level will be associated with a temptation to create surprise inflation (deflation), and initial debt below the critical level will be associated with a temptation to create surprise deflation (inflation).*

[11] The key parameter value for Nicolini is the size of the intertemporal elasticity of substitution of consumption. If this is less than one – as most econometric studies suggest – surprise deflation rather than surprise inflation occurs in his model.

The likely shape of the expansion path

To assess how likely the condition of increasing gross-interest elasticity is to hold, one approach is to consider particular functional forms for subutility of real balances. For example, suppose utility takes the quadratic form:

$$m_t[\overline{m} - m_t / 2]. \tag{21}$$

With this, it can readily be verified that behaviour is qualitatively identical to that depicted in Figure 1. Notice that $u_m(m_t) = \overline{m} - m_t$, so that the money demand function is linear: $m_t = \overline{m} - u_c(\bar{y})[i_{t-1}-1]$. As with any linear demand function, the price-elasticity of demand is increasing in the price, and this is true here whether we use either the net or the gross interest rate as the 'price'.[12] Another example is that of the 'constant relative risk aversion' (CRRA) functional form:

$$[m_t^{1-\rho} - 1]/[1-\rho]. \tag{22}$$

This has often been used in DGE models - especially in the limiting case where ρ tends to one, when it becomes ln m_t. However it is clear from the foregoing that CRRA utility causes problem here, since it violates the condition that ρ be increasing in m_t. It generates an expansion path which is a vertical line (as can easily be seen from (14)), making m_0 a 'non-normal good' for the government. This further implies that the relationship of the incentive to renege to initial debt d_{-1} is reversed: low debt is now associated with surprise inflation, and high debt with surprise deflation – a somewhat counter-intuitive result. This could, then, be taken as evidence that increasing gross-interest elasticity is not particularly likely. However, it is clear that the CRRA functional form has a significant limitation in the current context, because it fails to incorporate a satiation level of real balances. It may be appropriate for monetary policy experiments involving small or moderate changes in real balances, but

[12] To ensure that the gross interest elasticity passes through one for real balances in the feasible range, the restriction $u_C(\bar{y}) < \overline{m}$ is necessary. Otherwise the expansion path lies everywhere to the left of the 45^0 line, so that the incentive is always to create surprise inflation.

inappropriate for experiments involving optimisation over the global range of what is feasible.

A second approach to assessing how likely the increasing-gross-interest-elasticity condition is to hold, is to consider particular functional forms for the money demand function, especially those which have found favour in econometric studies. Simple early studies often just regressed the level of real balances on the level of the nominal interest rate. This linear form is derivable from the quadratic utility function, as noted above, and hence it supports the increasing-gross-interest-elasticity condition. Another commonly used form is the semi-logarithmic, in which the log of real balances is regressed on the level of the nominal interest rate (see, for example, Hendry and Ericsson (1991)). It is straightforward to check that this also implies an increasing gross-interest elasticity.

Clearly these observations are not completely conclusive, but they do provide some evidence supporting the type of situation depicted in Figure 1.

The source of the incentive to create surprise deflation

Notwithstanding our argument that surprise deflation is relatively unlikely, it can still arise in our model, and since it is an unorthodox form of time inconsistency, it is worth seeking an intuitive understanding of where it comes from. In general the time consistency problem arises because, although a small change in real balances has the same effect on utility (after allowing for discounting) irrespective of whether it is current or anticipated, it has a *different* effect on the present value of seigniorage depending on whether it is current or anticipated. This can be seen by considering once more the government's optimisation problem (10'), noting again that m_0 and m_t, $t \geq 1$, enter the budget constraint in different ways. Therefore as time passes and 'future' real balances become 'current', although the perceived utility benefit of a small increase in the level of real balances does not change, the perceived loss-of-revenue cost *does* change. Intuitively, if this cost goes down, there is an incentive to raise current real balances above their previously planned level: this is the case of surprise deflation. Alternatively it may go up, producing instead surprise inflation. Hence the key is to understand why current and anticipated seigniorage are different functions of real balances.

To probe into this, consider the changes in monetary growth rates necessary to achieve, on the one hand, a 1-unit rise in m_0, and, on the other, a $(1/\beta)$-unit rise in m_1, with all other real balance levels unchanged. Using the private-sector law of motion, (7), it is straightforward to see that a 1-unit rise in m_0 requires a reduction in μ_0 and no change in any other μ_t. Its effect on the present value (PV) of seigniorage is obviously to reduce it by 1 unit. Note that we can also view this 1-unit reduction using the 'cash-flow seigniorage' version of the government budget constraint, (9), which shows that it has to be equal to the change in $[\mu_0\text{-}1]m_0$. Turning to m_1, and again using (7), we can see that a 1-unit rise in m_1 requires the same reduction in μ_1 and again no change in any μ_t, $t>1$; however, it now also requires some alteration in μ_0, since if μ_0 is not changed, the anticipation of the rise in m_1 will alter m_0 (as can be seen by writing (7) for $t=0$). It is the need for an offsetting change in μ_0 which explains why a $(1/\beta)$-unit rise in m_1 has a different effect on the PV of seigniorage from a 1-unit rise in m_0. To see this, refer again to (9), where, now, there is the same 1-unit fall in the PV resulting from the fall in $[\mu_1\text{-}1]m_1$, but in addition there is a change resulting from an alteration in $[\mu_0\text{-}1]m_0$ – positive or negative depending on whether μ_0 has to be raised or lowered. Clearly, it is the 'forward-looking' nature of the private-sector equilibrium which is responsible for this difference.

Since surprise deflation occurs when a higher loss of seigniorage is associated with an increase in m_1 than with an increase in m_0, it follows from the foregoing that surprise deflation occurs in situations where to achieve a 1-unit increase in m_1 without changing m_0 requires a reduction (rather than an increase) in μ_0. In turn this implies that, if μ_0 were not reduced, the reduction in μ_1 and thence in p_1 would cause p_0 to rise. Put another way, surprise deflation is associated with situations where there is 'perverse transmission' from expected future monetary policy: an anticipated future money supply decrease[13] causes the current price level to *increase*. Intuitively, perverse transmission is possible because there are two conflicting effects of a fall in p_1 on p_0. First, at the old (p_0, i_0), the real interest rate has risen, inducing, in the goods market, an attempt to substitute from current to future consumption.

[13] Note that a decrease in μ_1 alone is an equiproportionate decrease in $M_1, M_2, \ldots$, with M_0 unchanged.

This puts downward pressure on current price level, and is the normal mechanism whereby anticipated future deflation induces current deflation. Second, however, the fall in p_1 raises real balances M_0/p_1. Since M_0 enters the utility function deflated by p_1, this means there is an excess supply of real balances at the old i_0, i.e. money market disequilibrium, and to remove this requires a fall in i_0. This tends to offset the rise in the real interest rate, moderating the need for p_0 to fall. If the fall in i_0 is sufficiently strong, p_0 may need to rise, generating 'perverse transmission'. This effect of p_1 in disturbing period 0's money market equilibrium is a consequence of the use of 'beginning-of-period' real balances as the liquidity variable. With the more conventional 'end-of-period' real balances, i.e. M_0/p_0, in the utility function, it would not arise, and perverse transmission could not occur. Note also that it is more likely the lower is the interest-elasticity of money demand, since then a bigger fall in i_0 is needed to clear the money market: this accords with our earlier finding about the role of interest-elasticity.[14]

4. Optimal Monetary Policy Under Discretion

Definition of the discretionary equilibrium

If the government is known to be unable to commit to a given policy plan made in period t, $\{\mu_{t+s|t}\}_{s=1}^{\infty}$, then it is not in general rational for households' forecasts of future μ_{t+s}'s, as of time t (denote these subjectively expected values by $\mu^{e}_{t+s|t}$), simply to equal the values in the plan, as has been assumed up to now. Instead, forecasts should be based only on 'observables' at time t. In principle these observables could include all current and past variables. One relevant variable is clearly the stock of government debt, whose value at the end of period t is d_t.[15] As we have seen, the government's inherited level of debt is a key

[14] The explanation for surprise deflation given here can also be adapted to explain Nicolini's (1998) surprise deflation result. Nicolini uses a cash-in-advance model, where the interest-elasticity of money demand, for *given* consumption, is necessarily zero (velocity of circulation must be unity). However, since consumption is endogenous, the intertemporal substitution of consumption in response to a change in the real interest rate implies that real balances can respond to the interest rate. When the elasticity of intertemporal substitution of consumption is less than unity, this implies a low interest-elasticity of money demand; and it is the former condition which Nicolini shows to give rise to surprise deflation.

[15] d_t is observable to the public in period t because the current monetary growth rate μ_t, which determines it via (8), is instantaneously observable.

determinant of its optimal choices under precommitment, so that the same is likely to be true under discretion, and it is hence rational for households to base their forecasts on d_t. We could also assume that they base their forecasts on current and past values of μ_t, but this would introduce an element of reputation-building behaviour into optimal government policy. Since we wish to focus on purely 'discretionary' behaviour, we exclude everything but d_t from households' forecasting rules.

Suppose, then, that households forecast next period's monetary growth rate according to the rule:

$$\mu^e_{t+1|t} = \hat{\sigma}(d_t)\,, \tag{23}$$

where $\hat{\sigma}(.)$ is for the moment treated as an arbitrary function. Combining this with the government's budget constraint yields an associated rule for forecasting d_{t+1}:

$$d^e_{t+1|t} = \hat{g}(d_t)\,. \tag{24}$$

To generate an s-period-ahead forecast we can use (24) repeatedly in (23) to get:

$$\mu^e_{t+s|t} = \hat{\sigma}(\hat{g}^{s-1}(d_t))\,, \tag{25}$$

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

These forecasting rules can now be used to determine equilibrium real balances in period t. Recall that the equilibrium value of m_t is given by the saddlepath solution of the private-sector law of motion (7). However, the relevant μ_t's to use in this equation are here the forecast values $\{\mu^e_{t+s|t}\}_{s=1}^{\infty}$ as given by (25), rather than the values $\{\mu_{t+s|t}\}_{s=1}^{\infty}$ from the government's policy plan, since what counts for determining the actual m_t are households' expectations. It is helpful to think of the determination of m_t in two parts. First, given d_t and thus the households' sequence of forecasts as generated by (25), we use (7) for periods t+1,...,∞ to find the saddlepath solution for m_{t+1}. This gives households' forecast of the equilibrium value of m_{t+1}, and since it is contingent on the given d_t, we may write it as $m^e_{t+1|t}$ = $\hat{\theta}(d_t)$, where the function $\hat{\theta}(.)$ implicitly derives from the forecasting functions and the private-sector law of motion. Second, using this in (7) for period t, we obtain the equilibrium value of current m_t:

$$m_t \quad = \quad \frac{\beta}{\mu_t}\hat{\theta}(d_t)\left[1+\frac{u_m(\hat{\theta}(d_t))}{u_c(\bar{y})}\right] \quad \equiv \quad \frac{e(d_t)}{\mu_t}. \tag{26}$$

This shows that m_t depends on μ_t with an elasticity of –1, and on d_t through a function which we can summarise as $e(d_t)$. The form of the function $e(.)$ derives from that of $\hat{\theta}(.)$, and thus ultimately from that of $\hat{\sigma}(.)$, as well as from the form of the utility function.

Using the equilibrium-real-balances function (26), we can specify the government's optimisation problem:

$$\underset{\{\mu_t\}_0^\infty}{\text{maximise}} \quad \Sigma_{t=0}^{\infty}\beta^t u(\bar{y},\bar{l},m_t)$$

$$\text{s.t.} \quad d_{t-1} = [\mu_t - 1]m_t + \beta d_t,$$

$$m_t = e(d_t)/\mu_t, \qquad t = 0,...,\infty, \quad d_{-1} \text{ given,} \tag{27}$$

or, substituting out m_t:

$$\underset{\{\mu_t\}_0^\infty}{\text{maximise}} \quad \Sigma_{t=0}^{\infty}\beta^t u(\bar{y},\bar{l},e(d_t)/\mu_t)$$

$$\text{s.t.} \; d_{t-1} = [1-1/\mu_t]e(d_t) + \beta d_t, \quad t = 0,...,\infty, \quad d_{-1} \text{ given.} \tag{27'}$$

Note that the problem (27') has a standard 'recursive' form, in which the new value of the state (d_t) just depends on the current value of the state (d_{t-1}) and on the current value of the control (μ_t); while the flow maximand also just depends on the current value of the state (if we imagine rearranging the constraint to substitute d_t out of the utility function) and on the current value of the control. Therefore it can be solved by dynamic programming, which in turn ensures that the solution will be time consistent. We also know that the dynamic programming solution can be written as a pair of feedback rules on the state, which we can denote as:

$$\mu_t \;=\; \sigma(d_{t\text{-}1}), \tag{28}$$

$$d_t \;=\; g(d_{t\text{-}1}). \tag{29}$$

(28) and (29) define the government's optimal monetary policy taking as given the public's arbitrary forecasting function, (23). However, for (23) to imply 'rational' forecasts, it needs to coincide with the optimal feedback rule, (28), which is based on it. In this case households will forecast correctly, no matter what the value of d_t. The discretionary equilibrium (Markov-perfect equilibrium) of the model is thus a function $\hat{\sigma}(d_t)$ with the property that it generates a $\sigma(d_t)$ such that $\hat{\sigma}(d_t) = \sigma(d_t)$ for all d_t.

Characterising the discretionary equilibrium: steady states

It is not in general possible to solve explicitly for the equilibrium forecasting rule-cum-policy function $\sigma(d_t)$, and its associated functions $g(d_t)$, $\hat{\theta}(d_t)$ and $e(d_t)$. This would be possible if the government's optimisation problem were linear-quadratic, but the problem here is not linear-quadratic even if we assume a quadratic utility function, since the constraint involves a 'Laffer curve', which is inherently non-linear.[16] Hence our procedure will be to try to characterise the discretionary equilibrium by conjecturing its properties, and then seeking to verify or falsify these conjectures. It will turn out, however, that we are able to solve for an approximation to the equilibrium path in the neighbourhood of some steady states. We begin by examining possible steady states themselves.

First, observe that the problem (27) may be re-expressed so as to treat m_t rather than μ_t as the control variable:

$$\underset{\{m_t\}_0^\infty}{\text{maximise}} \quad \Sigma_{t=0}^{\infty}\beta^t u(\bar{y},\bar{l},m_t) \tag{30}$$

$$\text{s.t.} \quad d_{t-1} = e(d_t) - m_t + \beta d_t, \quad t = 0,\ldots,\infty, \quad d_{-1} \text{ given.} \tag{31}$$

The first-order condition for this problem is readily derived as:[17]

---

[16] Obstfeld (1997) does consider a linear-quadratic version of the problem, but he is only able to do this by ad hoc approximations of some of the underlying functions.

[17] Given that the function $e(d_t)$ has unknown properties, there is a question as to whether the optimisation problem is well defined. For the moment we proceed as if this is the case, but we return to the question below.

$$u_m(m_{t+1}) = \left[1 + \frac{1}{\beta} e'(d_t)\right] u_m(m_t) . \tag{32}$$

(31)-(32) constitute a pair of first-order difference equations in (m_t,d_t) which in principle describe the evolution of the economy along the optimal policy path. The steady state (m,d) of this system, if it exists, is thus defined by:

$$[1-\beta]d = e(d) - m , \tag{33}$$

$$u_m(m) = \left[1 + \frac{1}{\beta} e'(d)\right] u_m(m) . \tag{34}$$

Examination of (33)-(34) suggests two ways in which they might be satisfied. First, if $m = \overline{m}$ so that $u_m(\overline{m}) = 0$, then (34) is automatically satisfied, and, given m, (33) determines d. We will refer to this as the 'Friedman Rule' steady state, for reasons which are clear from the discussion in Section 3. Second, suppose there exists a d such that $e'(d) = 0$. Then (34) is clearly satisfied, and, given d, (33) determines m. We will refer to this as the 'time consistent' steady state. We now comment further on each of these possibilities.

The 'Friedman Rule' steady state is exactly the one which Obstfeld (1991, 1997) identifies as being the long-run outcome in his closely related analysis.[18] In this steady state government debt is negative (in fact equal to $-\overline{m}$, as is seen by setting $\mu = \beta$ in the budget constraint in (27)), and inflation is negative (equal to β-1). We have already noted that, under commitment, if initial debt were sufficiently negative the government would be able to achieve the first-best allocation by using the real interest on its assets to finance continual monetary withdrawals at the rate β-1, and this is what is happening in this steady state. In this extreme but happy situation, the optimal tax problem has been removed, as has the incentive for time-inconsistent behaviour. The question then arises as to whether there exists a discretionary equilibrium path which converges on this steady state, if initial debt d_{-1} starts above $-\overline{m}$. We investigate this below. If there does, it follows that, for the given d_{-1}, whereas long-run inflation under commitment is quite likely to be positive, long-run inflation under

[18] Here we refer to the case of his analysis which uses a version of the government's objective function which is as close as possible to private utility functions.

discretion is *negative*. This is Obstfeld's surprising result: that under discretion there is no long-run problem of inflation, and indeed there is a tendency towards deflation ('deflation bias').[19] Intuitively, the reason why the government may drive the economy to this extreme outcome when it cannot commit is that reducing debt is the only method it has of lowering inflation expectations. The government cannot lower expected inflation just by announcing low future monetary growth, since this will not be believed. Expected inflation is based only on the observed level of debt, so that to reduce expected inflation it must reduce debt. This fact means that the 'rate of return' to a unit reduction in debt, from the government's point of view, is higher under discretion than under commitment, with the effect that the optimal policy is to give greater emphasis to the reduction of debt.

Turning to the 'time consistent' steady state, the conjecture that such exists is motivated by our finding that, under commitment, a critical initial debt level may exist such that there is no incentive to create surprise inflation or deflation. At such a debt level, we might therefore expect that the outcome under discretion would be the same as under commitment. Hence, more specifically, we now conjecture that the debt level d at which $e'(d)$ is zero coincides with the value of d_{-1} which, under commitment, induces $m_0=m_1$. From (13)-(14) this latter, and the corresponding m value – call them (d_c,m_c) - are defined by:

$$d_c = -m_c + \frac{\beta}{1-\beta}\frac{u_m(m_c)}{u_c(\bar{y})}m_c\,, \tag{35}$$

$$\rho(m_c) = 1 + \frac{u_c(\bar{y})}{u_m(m_c)}\,. \tag{36}$$

If our conjecture is correct, and if the discretionary equilibrium can converge to (d_c,m_c) for values of d_{-1} in some neighbourhood of d_c (see below), then Obstfeld's result that under discretion the government will always drive the economy to the 'Friedman Rule' steady state is not robust. This moreover means that price deflation is not the inevitable long-run consequence of discretionary policy: in general the net inflation rate associated with m_c

[19] However, this does not imply that discretion is better than commitment from a welfare point of view. Short-run inflation will in general be positive and higher under discretion than under commitment. Taking the time path as a whole, welfare is still higher under commitment, as Obstfeld shows.

(namely $\beta - 1 + \beta u_m(m_c)/u_c(\bar{y})$) may as well be positive as negative. This outcome therefore removes the seemingly paradoxical and rather extreme conclusion that discretion is associated with long-run 'deflation bias'.

In fact it is straightforward to confirm that $e'(d)$ is zero where $d=d_c$. Given the definition of $e(d_t)$ in (26), we have:

$$e'(d_t) \quad = \quad \frac{d}{dm_{t+1}}\left(\beta m_{t+1}\left[1+\frac{u_m(m_{t+1})}{u_c(\bar{y})}\right]\right)\times\hat{\theta}'(d_t)\,. \tag{37}$$

On the RHS of this, although the derivative $\hat{\theta}'(d_t)$ depends on the unknown function $\hat{\theta}(d_t)$ and is hence unknown, the derivative with respect to m_{t+1} is easily calculated. Evaluating it at m_c as defined by (36), we readily show that it equals zero (see Appendix). Hence $e'(d_c)=0$. Thus (provided that d_c exists, i.e. that the expansion path in the problem under commitment cuts the 45^0 line), a 'time consistent' steady state of the discretionary equilibrium does exist. What remains to be established is whether, if d_{-1} starts away from d_c, convergence to this steady state can occur.[20]

Characterising the discretionary equilibrium: dynamics

We now attempt to solve for a linear approximation to the non-stationary discretionary equilibrium in the neighbourhood of each of the types of steady state just discussed, the aim being to see whether a time path exists which converges on the steady state. If the properties of the function $e(d_t)$ were fully known, the task would be straightforward. The complicating factor is that the properties of $e(d_t)$ are unknown (except for the fact that $e'(d_c)=0$, as just established), and that indeed we wish to use this exercise to tie down its properties further. In particular, at the time-consistent steady state, it would be of interest to determine the sign and magnitude of e''; and, at the Friedman Rule steady state, to determine the sign and magnitude of e' and e''. Our strategy is to make use of two sources of information. First, treating $e(d_t)$ as if it were known, we derive the approximated saddlepath solution of the system (31)-(32).

[20] At this point it is appropriate to return to the question of whether the government's optimisation problem (30)-(31) is well defined. In the Appendix we look at a local approximation to the full problem, exploiting the finding that $e'(d_c)=0$, and confirm that this version of the problem is well defined.

Equivalently, this is an approximation of the optimal feedback rule which solves the problem (30)-(31): let us denote this optimal feedback rule as:[21]

$$m_t = \theta(d_{t-1}). \tag{38}$$

Such an exercise yields θ' as a function of *e*' and *e*" (evaluated at the appropriate steady state). Secondly we use the definition of the function $e(d_t)$ as given by (26), plus the condition of discretionary equilibrium that the public's forecasting rule (23) must deliver correct forecasts for all d_t, which equivalently means that the public's rule for forecasting m_{t+1}, $\hat{\theta}(d_t)$, must coincide with the government's optimal feedback rule (38), i.e. $\theta(d_t)$. Differentiating this then yields *e*' and *e*" as functions of θ' and θ". Exploiting the restrictions which follow from evaluating these derivatives at either type of steady state, it turns out to be possible to solve for θ' at these steady states, simultaneously with the associated derivatives of *e*(.). This then allows us to find the conditions for locally convergent discretionary equilibrium paths to exist.

The details of the calculations just described are presented in the Appendix. Here we will summarise the main results. A preliminary remark is that, if we substitute (38) into (31), we obtain a first-order difference equation in d_t which describes the evolution of debt in the discretionary equilibrium:

$$d_{t-1} = e(d_t) - \theta(d_{t-1}) + \beta d_t.$$

The local stability of this at an arbitrary steady state, *d*, is governed by the eigenvalue:

$$\frac{dd_t}{dd_{t-1}} = \frac{1+\theta'(d)}{\beta + e'(d)}, \tag{39}$$

so that, for convergence, (39) must have an absolute value less than unity.

[21] Since we have already observed that the problem (27') (which is equivalent to the problem (30)-(31)), is soluble by dynamic programming, it follows that the optimal value of the control can be expressed as a feedback on the state.

Consider first behaviour in the neighbourhood of the time-consistent steady state, (m_c, d_c). In the Appendix we show that the discretionary equilibrium value of 1+θ' solves the following quadratic equation:

$$\left(1-\frac{m\rho'}{\rho(\rho-1)}\right)(1+\theta')^2 \quad + \quad \frac{m\rho'}{\rho(\rho-1)}(1+\theta') \quad - \quad \beta \quad = \quad 0\,, \tag{40}$$

(m, ρ, ρ' being evaluated at the steady state). The term $m\rho'/\rho(\rho-1)$ in this equation can easily be shown, by differentiating (14), to equal the inverse of the slope of the expansion path for the problem under precommitment, at the intersection with the 45^0 line. Now, since e'=0 at the time-consistent steady state, (39) implies that for stability we need 1+θ' to lie in the range (-β,β). To determine whether either of the two solutions of (40) lie in this range, note that when $m\rho'/\rho(\rho-1) < 1$, the LHS of (40) when plotted against 1+θ' is a parabola. This parabola cuts the vertical axis at -β, so that for at least one solution to lie in (-β,β), we need the LHS when evaluated at either –β, or β, or both, to be positive. Evaluating at -β, we obtain an unambiguously negative number, as is easily verified. Evaluating at β, we obtain

$$(\text{LHS of }(40))_{1+\theta'=\beta} \quad = \quad \beta(\beta-1)\left(1-\frac{m\rho'}{\rho(\rho-1)}\right), \tag{41}$$

which is also negative when $m\rho'/\rho(\rho-1) < 1$. Thus, when the expansion path cuts the 45^0 line 'from below' in the commitment problem, no path of the discretionary equilibrium which converges to the time-consistent steady state exists. When $m\rho'/\rho(\rho-1) > 1$, the LHS of (40) is an inverted parabola. It again cuts the vertical axis at -β, and is negative when 1+θ' = -β; at 1+θ' = β its value is given as before by (41), which is now positive. From this it follows that one solution lies in (0,β), and one in (β,∞). Hence, when the expansion path cuts the 45^0 line 'from above', a path of the discretionary equilibrium which converges to the time-consistent steady state exists (and is unique).[22]

This result shows that there is a correspondence between whether high initial debt is associated with surprise inflation or deflation under commitment, and whether the

[22] The relationship of e'' to θ' is $(\beta\rho/m)[1-m\rho'/\rho(\rho-1)](\theta')^2$, so that in this case e'' is negative. Intuitively, this means that above d_c, higher debt is associated with a forecast by the public of lower equilibrium real balances next period; while below d_c the reverse is the case.

equilibrium converges to, or diverges from (respectively), the steady state under discretion. In Section 3 we gave reasons for thinking that money-holding preferences were more likely to be such as to cause high debt to be associated with surprise inflation, which hence suggests that, under discretion, convergence to, rather than divergence from, the time-consistent steady state is more likely. Thus Obstfeld's (1991, 1997) conclusion that under discretion the economy is bound to end up at the Friedman Rule steady state, with consequent negative inflation, does not hold in our framework. Rather it is more probable that the economy will end up at an alternative and less extreme steady state, in which there is no reason why inflation should not be positive. We summarise this finding by:

Proposition 2 *Under optimal monetary policy with discretion, a steady-state equilibrium (the 'time consistent steady state') exists at the critical debt level defined in Proposition 1 (where such a debt level exists), and inflation could be either positive or negative there. If, moreover, the gross-interest elasticity of money demand is increasing (decreasing) in the interest rate at this steady state, then, within a neighbourhood of the steady state, a discretionary equilibrium which converges on it exists (does not exist).*

Consider, lastly, behaviour in the neighbourhood of the Friedman Rule steady state. In the Appendix we show that the discretionary equilibrium value of θ' is related to that of of e' by $e' = \beta\kappa\theta'$, where $\kappa \equiv 1+\overline{m}u_{mm}(\overline{m})/u_c(\overline{y})$ (note that κ could be either positive or negative). We also show that e' solves the following quadratic equation:

$$(e')^2 \quad + \quad (2\beta - 1/\kappa)e' \quad + \quad \beta(\beta - 1) \quad = \quad 0\,. \tag{42}$$

Moreover, we demonstrate that in this case the eigenvalue (39) can be expressed as:

$$\frac{dd_t}{dd_{t-1}} \quad = \quad 1 + e'/\beta\,. \tag{43}$$

Thus, for convergence, we need that the solution for e' should lie in the range (-2β,0). To determine whether either of the two solutions of (42) lie in this range, note that the LHS of (42) when plotted against e' is a parabola. It has a negative intercept with the vertical axis, whence one solution for e' is always positive, and so outside the range (-2β,0). Evaluating at *e'* = -2β, the LHS is $\beta(2/\kappa+\beta-1)$, which is positive or negative as κ is positive or negative, respectively (note $\kappa \leq 1$). Thus, a positive κ implies exactly one solution for *e'* in (-2β,0), while a negative κ implies no solutions in this range. It follows that the condition for a discretionary equilibrium to exist which converges locally to the Friedman Rule steady state is $\kappa > 0$. Hence we have:

Proposition 3 *Under optimal monetary policy with discretion, if the utility function has a satiation level of real balances ($\overline{m}$), then a steady-state equilibrium (the 'Friedman Rule steady state') exists at which real balances equal their satiation level, and inflation and debt are negative. If, moreover, the condition $-\overline{m}u_{mm}(\overline{m})/u_c(\overline{y})$ < 1 (>1) holds, then within a neighbourhood of the steady state, a discretionary equilibrium which converges on it exists (does not exist).*

The significance of this result is that it shows that, even where a satiation level of real balances exists, and thus a Friedman Rule steady state of the discretionary equilibrium exists, it is not guaranteed that under discretion the economy will converge to it despite starting close by. More broadly speaking, this supports our basic point above: Obstfeld's conclusion that under discretion negative inflation will be the outcome in the long run does not have generality.

A further conjecture might be that, where both a Friedman Rule steady state and a unique time-consistent steady state exist, and where the latter is locally stable, the Friedman Rule steady state will always be locally unstable. If this were true it would suggest that under discretion the economy would converge on the time-consistent steady state for all values of d_{-1} ($>\overline{m}$), not just those within some neighbourhood of the steady state. We can show that this is indeed the case if subutility of real balances is quadratic (given by (21)): in this case the condition $\kappa < 0$ is also the condition for the expansion path to intersect the 45^0 line in Figure 1 (moreover the intersection will always be 'from above'). However, the conjecture does not appear to hold for a general unspecified utility function. In particular, there is no general correspondence between the stability of the Friedman Rule steady state and the slope of the expansion path at ($\overline{m}$,$\overline{m}$) in Figure 1.

5. Conclusions

We have studied the time consistency problem in the realistic situation in which open market operations are the instrument of monetary policy, lump-sum taxes and transfers being prohibited. Our modelling of money demand captures a welfare cost of current inflation, as Nicolini (1998) did, and this enables a non-degenerate optimal monetary policy to exist even when the government's objective is purely to maximise the welfare of private agents. However we do not find a general presumption that time inconsistency will take the form of 'surprise deflation', unlike Nicolini: replacing the cash-in-advance by a money-in-the-utility function motive for holding money, we conclude that the conventional temptation to 'surprise inflation' is more likely. We have also studied optimal monetary policy under discretion. Here we conclude that it is not inevitable that in the long run inflation will end up negative, contrary to what Obstfeld (1991, 1997) found in a similar analysis but in which the government's objective was not pure private welfare maximisation. In our framework under discretion long-run inflation could very well be positive, with real balances far distant from their 'Friedman Rule', satiation levels.

The broader objective of this research has been to contribute towards bringing the analysis of optimal monetary policy under the roof of public finance theory. The model used here has been deliberately highly 'classical' and stripped-down in its features. It is to be hoped that in future work it would be possible to introduce features which would enable an integration with more traditional 'macroeconomic' analyses of monetary policy, such as temporary nominal rigidities.

References

Barro, R.J. (1979) ‘On the Determination of the Public Debt’, *Journal of Political Economy* 87, 940-971

Barro, R.J. (1983) ‘Inflationary Finance Under Discretion and Rules’, *Canadian Journal of Economics* 16, 1-16

Barro, R.J. and Gordon, D.B. (1983) ‘A Positive Theory of Monetary Policy in a Natural-Rate Model’, *Journal of Political Economy* 91, 589-610

Chari, V.V., Christiano, L.J. and Kehoe, P.J. (1996) ‘Optimality of the Friedman Rule in Economies with Distorting Taxes’, *Journal of Monetary Economics* 37, 203-223

Friedman, M. (1969) ‘The Optimum Quantity of Money’, in Friedman, M., *The Optimum Quantity of Money and Other Essays*, Chicago: Aldine

Hendry, D.F. and Ericsson, N.R. (1991) ‘Modeling the Demand for Narrow Money in the United Kingdom and the United States’, *European Economic Review* 35, 833-886

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

Ireland, P. (1997) ‘Sustainable Monetary Policies’, *Journal of Economic Dynamics and Control* 22, 87-108

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

Phelps, E. (1973) ‘Inflation in the Theory of Public Finance’, *Swedish Journal of Economics* 75, 67-82

Sidrauski, M. (1967) ‘Rational Choice and Patterns of Growth in a Monetary Economy’, *American Economic Review* 57 (P&P), 534-544

Walsh, C.E. (1998) *Monetary Theory and Policy*, Cambridge MA: MIT Press

Appendix

1. 2nd-order conditions for the commitment problem

Consider the government's optimal choice of a particular (m_t,m_{t+1}) pair ($t \geq 1$), treating other real balances as given. The government's (like households') indifference curve between m_t and m_{t+1} has slope $-u_m(m_t)/\beta u_m(m_{t+1})$. The government's budget constraint between m_t and m_{t+1} has a slope which can be calculated from (10') as:

$$\frac{dm_{t+1}}{dm_t} = -\frac{1}{\beta}\frac{u_m(m_t)}{u_m(m_{t+1})}\frac{\rho(m_t)-1}{\rho(m_{t+1})-1}. \tag{A1}$$

This constraint is clearly not necessarily convex. In the special case where ρ is a constant we see that ρ cancels from (A1), and hence the slope of the constraint everywhere coincides with that of the indifference curve. In this special case, therefore, the government is indifferent about the allocation of seigniorage across periods 1,2,...,∞. More generally, in order for (12) to define a strict maximum, we need that, as we slide up an indifference curve starting from a point of tangency, the budget constraint should become flatter than the indifference curve, and thus that the ratio of the (ρ−1)'s in (13) should decrease (see Figure A1). This clearly requires that ρ be increasing in *m*.

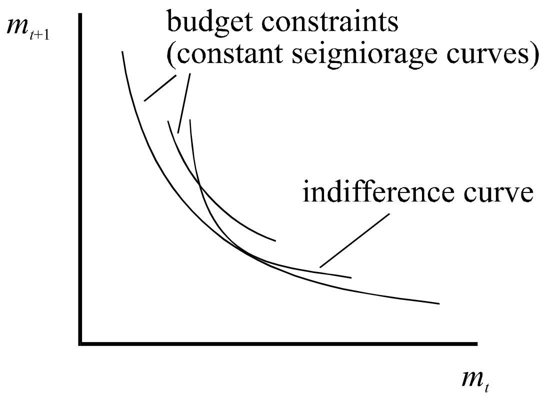

Figure A1

2. The value of ρ when real balances reach satiation

From the definition of ρ as $-mu_{mm}/u_m$ it is clear, except in the case where $u_{mm}(\overline{m})=0$, that ρ tends to infinity as $m \to \overline{m}$. In fact ρ still tends to infinity even if $u_{mm}(\overline{m})=0$, provided that utility remains continuous and continuously differentiable at $m=\overline{m}$, and has at least one non-zero derivative. To show this, note that when $u_{mm}(\overline{m})=0$ we can write the following Taylor series approximations:

$$u_{mm}(m) \;\approx\; u_{mmm}(\overline{m})[m-\overline{m}] \tag{A2}$$

$$u_{m}(m) \;\approx\; \frac{1}{2}u_{mmm}(\overline{m})[m-\overline{m}]^2 . \tag{A3}$$

Here we assume that $u_{mmm}(\overline{m})$ is non-zero, and neglect higher-order terms. Thus:

$$\frac{u_{mm}}{u_m} \;\approx\; \frac{1}{2[m-\overline{m}]} . \tag{A4}$$

This tends to minus infinity as m tends to $\overline{m}$, and hence ρ tends to infinity. If $u_{mmm}(\overline{m})=0$, the 2nd- and 3rd-order terms may be used instead in (A2) and (A3), and so on.

3. An approximation of the problem under discretion

At the time consistent steady state we know that $e'(d_c) = 0$. Therefore an approximation to the full problem (30)-(31) in the neighbourhood of this steady state can be made by treating $e(d_t)$ as simply a constant: $\overline{e}$, say, where $\overline{e} = e(d_c)$. This gives the problem:

$$\underset{\{m_t\}_0^\infty}{\text{maximise}} \quad \Sigma_{t=0}^{\infty}\beta^t u(\overline{y},\overline{l},m_t) \quad \text{s.t.} \quad d_{-1} = \Sigma_{t=0}^{\infty}\beta^t[\overline{e}-m_t] . \tag{A5}$$

Note that the assumed independence of $e(d_t)$ from d_t means that the sequence of single-period constraints can be integrated into a single intertemporal constraint. This version of the problem has similarities to the general problem under commitment, (10'). The differences are that current and future real balances here enter identically; and seigniorage revenue is linear in m_t. The latter means that the government's budget constraint between m_t and m_{t+1} is also linear, so that the choice over any (m_t,m_{t+1}) pair can be pictured in the same way as in Figure A1, but where the budget constraint is a straight line. It is clear, then, that (A5) is a well-defined optimisation problem.

4. Approximation of the discretionary equilibrium around the steady state

Our strategy for obtaining this solution was described in the main text. First, we derive a linear approximation to the saddlepath of the equation system (31)-(32). A linear approximation to (31)-(32) itself, about an arbitrary steady state (m,d), is:

$$\begin{bmatrix} m_t - m \\ d_t - d \end{bmatrix} = \begin{bmatrix} 1+e'/\beta & -me''/\rho\beta \\ 1/\beta & [1-me''/\rho\beta]/[\beta+e'] \end{bmatrix} \begin{bmatrix} m_{t-1} - m \\ d_{t-1} - d \end{bmatrix}. \tag{A6}$$

From this, we may solve for the eigenvalues of the coefficient matrix. In the main text we saw that, if θ' denotes the derivative of the optimal feedback rule, then the eigenvalue of the optimal path for debt is $(1+\theta')/(\beta+e')$ (see (39)). Since this must coincide with the stable eigenvalue (if such exists) of the matrix in (A6), then by equating the two we obtain an equation determining θ':

$$\frac{1+\theta'}{\beta+e'} = \frac{1}{2}\left\{1+\frac{e'}{\beta}+\frac{1}{\beta+e'}\left[1-\frac{1}{\beta}\frac{m}{\rho}e''\right] \pm \sqrt{\left(1+\frac{e'}{\beta}+\frac{1}{\beta+e'}\left[1-\frac{1}{\beta}\frac{m}{\rho}e''\right]\right)^2 - \frac{4}{\beta}}\right\}. \tag{A7}$$

Second, using $\theta(d_t)$ rather than $\hat{\theta}(d_t)$ in the definition of $e(d_t)$ (see (26)), we have:

$$e(d_t) = \beta\theta(d_t)\left[1+u_m(\theta(d_t))/u_c(\bar{y})\right].$$

The first and second derivatives of this, at an arbitrary value of d_t, are:

$$e' = \beta\{1+[1-\rho]u_m/u_c\}\theta', \tag{A8}$$

$$e'' = \beta\{1+[1-\rho]u_m/u_c\}\theta'' + \beta\{[1-\rho]u_{mm}/u_c - \rho' u_m/u_c\}(\theta')^2, \tag{A9}$$

(where u_c denotes the constant $u_c(\bar{y})$). (A7)-(A9) provide 3 equations in 4 unknowns, (θ',θ'',e',e''), and so cannot yet be solved. However, we now show that by making use of the restrictions known to apply at either the time consistent or the Friedman Rule steady state, solutions can be obtained.

At the time consistent steady state, $1+[1-\rho]u_m/u_c = 0$ (see (36)), and hence (A8) implies $e'=0$, as asserted in the main text. Thus e' drops out of (A7), which becomes a relationship between θ' and e" alone. Further, θ" drops out of (A9), which hence provides a second relationship between θ' and e". Using (A9) to substitute e" out of (A7) we obtain an equation in θ' alone. This equation is cubic in θ', but it has θ'=0 as one root. We may discard this root as being inconsistent with convergence, since it clearly implies that the eigenvalue given by (39) in the main text is unstable. The remaining equation is then quadratic, and may be arranged as equation (40) in the main text.

At the Friedman Rule steady state, by definition u_m=0 and m equals its satiation level, $\overline{m}$. In section 2 of the Appendix above, we showed that, under mild conditions, $\rho(\overline{m}) = \infty$. In this case e" drops out of (A7), which becomes a relationship between θ' and e' alone. Moreover, u_m=0 implies that (A8) can be written as $e' = \beta\kappa\theta'$, where $\kappa \equiv 1+\overline{m}\, u_{mm}(\overline{m})/u_c(\overline{y})$, as asserted in the main text. Thus we have two relationships between θ' and e', from which we may attempt to solve for both. In (A7), the term under the square root sign in (A7) is in fact an exact square, with the result that the RHS of (A7) simplifies to either $1/(\beta+e')$ or $1+e'/\beta$. Using $1/(\beta+e')$, (A7) reduces to θ'=0, and (A8) then implies e'=0. However convergence cannot occur under this solution, since the eigenvalue given by (39) in the main text is then clearly unstable. Therefore we turn to the case where the RHS of (A7) is $1+e'/\beta$. Combining this with $e' = \beta\kappa\theta'$, we arrive at equation (42) in the main text. Since the LHS of (A7) is just the eigenvalue (39), $1+e'/\beta$ is an equivalent expression for it, which explains (43) in the main text.