Revenue maximizing inflation[☆]

Kent P. Kimbrough*

Department of Economics, Duke University, Box 90097, Durham, NC 27708, USA

Received 6 May 2005; received in revised form 20 July 2005; accepted 21 July 2005

Available online 11 October 2006

Abstract

A classic monetary policy result is that revenue maximization entails setting the inflation tax rate equal to the inverse of the interest semi-elasticity of the demand for money. The standard approach underlying “Cagan’s rule” is partial equilibrium in nature, treating money demand as being given from outside the model and abstracting from the real effects of inflation. This paper reconsiders the question of the revenue maximizing inflation rate in a general equilibrium framework with a labor-leisure choice, where money is held because it reduces transactions costs. In this framework, the revenue maximizing inflation tax rate is lower than that implied by Cagan’s rule.

© 2006 Elsevier B.V. All rights reserved.

JEL classifications: E; E4; E5; E6

Keywords: Revenue maximizing inflation; Inflation tax; Seigniorage; Cagan money demand function; Cagan’s rule

1. Introduction

In his pioneering study of hyperinflation, Cagan (1956) introduces the semi-log or Cagan money demand function that has been a workhorse in monetary economics ever since. One especially useful result that Cagan establishes is that the revenue maximizing inflation rate is equal to $1/\alpha$, where α is the interest semi-elasticity of the demand for money. The “Cagan rule,” as this result has come to be known, is the keystone of the literature on revenue maximizing inflation. However, Cagan’s rule is partial equilibrium in nature in two important respects. First, it takes the demand for money as given from outside the model

[☆]I would like to thank the editors and the referee for their comments, which significantly improved the paper.

*Corresponding author. Tel.: +1 919 489 0826; fax: +1 919 684 8974.

E-mail address: kent@econ.duke.edu.

rather than deriving it from the underlying environment. Second, it ignores the potential impact of changes in the inflation tax rate on real variables such as output, consumption, and the real interest rate. The upshot is that one parameter, the interest semi-elasticity of the demand for money, pins down the revenue maximizing inflation tax rate. No other features of the economy and no other features of the money demand function matter.¹

This paper broadens the approach to the revenue maximizing inflation tax rate by (i) deriving the money demand function from the consumer's optimization problem and (ii) allowing for changes in the inflation tax rate to have real effects. The model is similar to that used by Kimbrough (1986a) to study the optimality of Friedman's (1969) rule when distorting taxes must be used to raise government revenues. Consumers face a labor-leisure choice and they hold money because it economizes on the amount of resources that must be devoted to transacting in the goods market. It is shown that, because changes in the inflation tax rate have real effects in such an economy, the revenue maximizing inflation tax rate is lower than that implied by Cagan's rule. The intuition for this result is quite simple. As with any other tax, maximizing revenues from the inflation tax requires setting the tax rate so that the marginal increase in tax revenues from increasing the tax rate is just equal to the marginal loss in revenues from the erosion of the tax base induced by the higher tax rate. In the standard approach underlying Cagan's rule, this marginal loss in revenues arises solely because money holdings per unit of consumption/output fall as higher inflation tax rates lead consumers to substitute out of money and into other assets. In the transactions costs approach adopted here there is another effect as well: As the inflation tax rate rises, consumption falls due to the combined effects of rising transactions costs and increased distortions in the labor market thus further reducing the demand for real balances. Put more succinctly; in the standard approach higher inflation reduces the tax base solely through a substitution effect whereas in the transactions cost approach there is both a substitution effect *and* a scale effect. This means that the marginal loss in revenues from increasing the inflation tax rate is higher in the transactions costs approach and, as a result, the revenue maximizing inflation tax rate is lower than that implied by Cagan's rule. Section 1 of the paper demonstrates this result for the general case. Section 2 looks at the special case where the underlying transactions cost function gives rise to the Cagan money demand function. Qualitative results are discussed and back-of-the-envelope type calculations are presented to try and assess the quantitative significance of the results. Concluding remarks are presented in Section 3.

2. Transactions costs and revenue maximizing inflation

The representative consumer chooses time profiles for consumption, c_t , labor effort, l_t , money holdings, M_t^d , and bond holdings, b_t , to maximize lifetime utility

$$U = \sum_{t=0}^{\infty} \beta^t U[c_t - g(l_t)], \quad (1)$$

¹In the general case, where the money demand function is not of the Cagan type, the partial equilibrium approach continues to imply that the revenue maximizing inflation tax rate is given by the inverse of the interest semi-elasticity of the demand for money. However, the interest semi-elasticity is no longer a constant but varies with the inflation tax rate.

where $U(\cdot)$ exhibits positive and diminishing marginal utility, $0 < \beta < 1$, and $g', g'' > 0$. Output is given by the production function $y_t = f(\ell_t)$ which is characterized by a positive and diminishing marginal product. Money is held because it helps reduce transactions costs as given by the transactions cost function²

$$T_t = T(m_t), T'(m_t) < 0, T''(m_t) > 0, \quad (2)$$

where $m_t \equiv M_t^d / P_t c_t$, P_t is the period t price level, and T_t is transactions costs per unit of consumption.

The consumer thus maximizes lifetime utility (1) subject to the available production technology, the transactions technology (2), and the sequence of budget constraints

$$[1 + T(m_t)]c_t + \frac{M_t^d}{P_t} + b_t = y_t + \frac{M_{t-1}^d}{P_t} + (1 + r_{t-1})b_{t-1} + \tau_t,$$

where r_t is the real interest rate and τ_t is a lump-sum transfer payment from the government.

It is straightforward to show that the first-order conditions from the consumer's problem imply that, at an optimum, real balances are determined by the condition

$$-T'(m_t) = i_t \quad (3)$$

for each period, where $i_t \equiv R_t / (1 + R_t)$ and R_t is the nominal interest rate which satisfies the Fisher equation. The first-order condition (3) states that consumers choose their money holdings so that the marginal benefit in terms of reduced transactions costs equals the opportunity cost of holding money. This yields the consumer's money "demand" function³

$$\frac{M_t^d}{P_t} = k(i_t)c_t, k'(i_t) = \frac{-1}{T''} < 0. \quad (4)$$

The consumer's first-order conditions can also be shown to imply that in each period optimal labor supply satisfies

$$g'(\ell_t) = \frac{f'(\ell_t)}{1 + T[k(i_t)] + i_t k(i_t)}, \quad (5)$$

where use has been made of (3) and (4). Condition (5) reflects the fact that inflation distorts the consumer's labor-leisure choice by driving a wedge between the opportunity cost of supplying additional labor, $g'(\ell_t)$, and the marginal product of labor, $f'(\ell_t)$. It follows immediately from (5) that increases in the inflation tax rate, i_t , reduce labor effort by increasing the magnitude of this distortion. Given the production technology, $y_t = f(\ell_t)$ and the implicit relation between optimal labor supply and the inflation tax rate, $\ell_t = \ell(i_t)$, implied by (5), it follows that output is given by

$$y_t = y(i_t), y'(i_t) = \frac{-f'[\ell(i_t)]g'[\ell(i_t)]k(i_t)}{\{1 + T[k(i_t)] + i_t k(i_t)\}g''[\ell(i_t)] - f''[\ell(i_t)]} < 0. \quad (6)$$

The government is the only other actor in the economy. They issue debt and use revenues from the inflation tax in order to finance lump-sum transfers to consumers

²The results that follow also hold for more general forms of the transactions cost function.

³Demand is in quotes because consumption must be solved for in order to get a proper demand function. The quotes are dropped from here on.

according to the sequence of budget constraints

$$\tau_t = \frac{M_t - M_{t-1}}{P_t} + b_t - (1 + r_{t-1})b_{t-1},$$

where M_t is the period t money supply. Discounting by $d_t \equiv \prod_{j=1}^t (1 + r_{j-1})^{-1}$ for $t = 1, 2, \dots$ and letting $d_0 = 1$, the government's intertemporal budget constraint is thus

$$\sum_{t=0}^{\infty} d_t \tau_t = \sum_{t=0}^{\infty} d_t i_t \frac{M_t}{P_t} - \frac{M_{-1}}{P_0}. \quad (7)$$

The government's goal is choose the time profile for the inflation tax rate, i_t , in order to maximize its revenues from the inflation tax. In the partial equilibrium problem, the government treats consumption, output and the real interest rate as given. In this case, assuming the money demand function still takes the form given by (4), it follows from the government's intertemporal budget constraint (7) that the first-order condition for revenue maximization is $k(i_t) + i_t k'(i_t) = 0$. The revenue maximizing inflation tax rate thus satisfies $i_t = -k(i_t)/k'(i_t) \equiv 1/\alpha(i_t)$, where $\alpha(i_t)$ is the interest semi-elasticity of the demand for money. That is, the revenue maximizing inflation rate is determined by Cagan's rule. In the special case where the money demand function is of the semi-log form, $\alpha(i_t) = \alpha$ and Cagan's rule implies that the revenue maximizing inflation rate is $1/\alpha$.⁴

However, in the context of the general equilibrium model being considered here, the government must take account of the real effects of inflation when determining the revenue maximizing inflation rate. Formally this amounts to taking account of the fact that the government's revenue maximization problem is subject to the constraints implied by the optimizing behavior of consumers and the equilibrium conditions in goods and money markets. These conditions require that

$$\frac{M_t}{P_t} = \frac{k(i_t)y(i_t)}{1 + T[k(i_t)]}, \quad (8)$$

where use has been made of the fact money demand is given by (4), output is given by (6), money market equilibrium requires $M_t = M_t^d$, and the goods market clears, so that $[1 + T(k(i_t))]c_t = y_t$. In addition, the real interest rate is determined by the intertemporal optimality conditions from the consumer's problem. In order to abstract from the complications this entails, it is assumed from here on that the function $U(\cdot)$ is linear so that the real interest rate is constant at $(1/\beta) - 1$.⁵ In this case the discount factors in (7) are given by $d_t = \beta^t$. Additionally, it is well know that the government has an incentive to engineer surprise inflations in an attempt to garner additional revenues through capital levies on money holdings. This time-consistency issue is abstracted from by assuming the government is able to follow a time-consistent policy that precludes price level jumps.

Under these conditions, substituting (8) into the government's budget constraint (7), it can be shown that the first-order conditions for maximizing revenues from the inflation tax

⁴Bailey (1956) also derives this result.

⁵Alternatively it could be assumed that the economy in question is a small open economy facing a constant world real interest rate.

require that the inflation tax rate be set in all periods so that

$$\frac{\{1 + T[k(i_t)]\}[k(i_t) + i_t k'(i_t)]}{\{1 + T[k(i_t)]\}^2} - \frac{i_t k(i_t) T'[k(i_t)] k'(i_t)}{\{1 + T[k(i_t)]\}^2} - \frac{\{1 + T[k(i_t)]\} k(i_t) \phi(i_t)}{\{1 + T[k(i_t)]\}^2} = 0, \quad (9)$$

where $\phi(i_t) \equiv -i_t y'(i_t)/y(i_t) > 0$ is the elasticity of output with respect to the inflation tax rate. Formally, it can be shown with some manipulation that (9) requires that the revenue maximizing inflation tax rate satisfies

$$i_t = \frac{1}{\alpha(i_t)} - \frac{\alpha(i_t)[1 - \sigma(i_t)] + \psi(i_t)\sigma(i_t)}{\alpha(i_t)[\alpha(i_t) + \psi(i_t)\sigma(i_t)]} < \frac{1}{\alpha(i_t)},$$

where

$$0 < \sigma(i_t) \equiv \frac{1 + T[k(i_t)]}{1 + T[k(i_t)] + i_t k(i_t)} \leq 1$$

and $\psi(i_t) \equiv -y'(i_t)/y(i_t) > 0$ is the interest semi-elasticity of output. This establishes that in the general equilibrium setting considered here, the revenue maximizing inflation tax rate is lower than that dictated by Cagan's rule.

To better understand the intuition behind this result, note that the left-hand side of the first-order condition (9) is, of course, the marginal effect of a change in the inflation tax rate on revenues. The first term is what the partial equilibrium approach underlying Cagan's rule accounts for: The increase in the inflation tax rate raises tax revenues but the increase in the opportunity cost of holding money reduces the tax base, and hence tax revenues, through a substitution effect as consumer's substitute out of money and into other assets. However, in general equilibrium there are two other effects that need to be considered. Both arise because increases in inflation lower equilibrium consumption. This scale effect, captured by the second and third terms on the left-hand side of (9), serves to reduce the revenue maximizing inflation tax rate below that implied by Cagan's rule. To see this more clearly, notice that if the government were to set the inflation tax rate according to Cagan's rule, the first term on the left-hand side of (9) would be zero. The government would be balancing the increase in tax revenues from a higher tax rate against the loss in tax revenues arising from the erosion of the tax base *given* the level of consumption. However, in the presence of transactions costs, the resulting inflation rate is too high since it fails to account for the loss of revenues arising from the additional erosion of the tax base due to the decline in equilibrium consumption. That is, in the presence of transactions costs, setting the inflation tax rate so that $k(i_t) + i_t k'(i_t) = 0$ assures that the government ends up on the far side of the Laffer curve. In order to avoid this, the government must set the inflation tax rate lower than the level indicated by Cagan's rule.

As already noted, the scale effect appearing in (9) is a composite of two effects. The second term on the left-hand side of (9) illustrates how the scale effect arises because increases in inflation lower money holdings relative to consumption thereby resulting in higher per unit transactions costs and lower consumption for any given level of output. However, increases in inflation also distort the labor-leisure choice further, resulting in reduced labor supply, output, and ultimately consumption. This second component of the scale effect is captured by the third term on the left-hand side of (9). Note that it depends

crucially on the sensitivity of output to changes in the inflation tax rate as measured by the elasticity $\phi(i_t)$. As can be seen from (6), this elasticity depends not only on the properties of the transactions technology (2), which underlies the money demand function (4), but also on tastes and the production technology through its dependence on the slopes of the labor supply and demand schedules as captured by g' and f' . Therefore, not only is the revenue maximizing inflation tax rate lower than that implied by Cagan's rule but in general equilibrium it also depends on features of the economy other than just the interest semi-elasticity of the demand for money.⁶

3. Transactions costs and the Cagan money demand function

This section of the paper provides an explicit solution for the model outlined in the previous section when the underlying transactions cost function (2) is such that the consumer's money demand function (4) is of the Cagan type and thus takes the form

$$\frac{M_t^d}{P_t} = B e^{-\alpha i_t} c_t. \quad (10)$$

This exercise is undertaken in order to provide a back-of-the envelope quantitative assessment of the importance for the revenue maximizing inflation tax rate of taking account of general equilibrium considerations. For simplicity, it is assumed that labor is supplied inelastically so that output is given exogenously in any period. The revenue maximizing inflation tax rate is still given by (9) but $\phi(i_t) = 0$ so that third term on the left-hand side of (9) can be ignored.

Brock (1989) shows that the transactions cost function underlying the Cagan money demand function is

$$T(m_t) = \theta + \frac{1}{\alpha} m_t [\ln(m_t) - \ln(B) - 1], \quad (11)$$

where θ is a constant of integration set to assure that transactions costs are always non-negative.^{7,8} In equilibrium consumers are on their money demand curve so that real balances are given by (10). Using this in (11), it can be seen that in equilibrium transactions costs per unit of consumption are

$$T[k(i_t)] = \theta - \frac{1}{\alpha} B e^{-\alpha i_t} (1 + \alpha i_t). \quad (12)$$

⁶In the general case where there is a wealth effect on labor supply, the third term on the left-hand side of (9) could be positive if there is a backward bending supply of labor at the aggregate level. However, even in this case the second and third terms in the numerator of (9) taken together must be negative since consumption must fall when the inflation tax rate is increased. See Kimbrough (1986b) for an illustration of this point in a slightly different environment.

⁷Brock (1989) uses a shopping time model to study the importance of reserve requirements for the revenue maximizing inflation tax rate. Most of his analysis focuses on the case where all shopping time comes out of leisure, labor supply remaining constant. Because higher inflation doesn't reduce output or consumption in his setup, Brock concludes that the revenue maximizing inflation tax rate satisfies Cagan's rule assuming that the government is free to simultaneously set reserve requirements to maximize revenues.

⁸To derive (11), take logarithms of (10) and use the first-order condition for optimal money holdings, (3), to eliminate i_t from the resulting expression to see that the transactions cost function must satisfy the condition $T'(m_t) = (1/\alpha)[\ln(m_t) - \ln(B)]$. Integrating by parts yields (11).

Notice that for transactions costs to be non-negative at a nominal interest rate of zero, i.e., when $i_t = 0$, $B/\alpha \leq \theta$ must hold. Here it is assumed that $B/\alpha = \theta$, which amounts to assuming that transactions costs are zero when Friedman's rule is followed. This reduces the number of parameters to be considered to just two: The interest semi-elasticity of the demand for money, α , and the scale parameter θ .

With these assumptions, it follows from the goods market equilibrium condition, $[1 + T(k(i_t))]c_t = y_t$, and (12) that equilibrium consumption is given by

$$c_t = \frac{1}{1 + \theta - \theta e^{-\alpha i_t}(1 + \alpha i_t)} y_t. \quad (13)$$

Using this and (8), it can be seen that the inflation tax revenue for period t appearing in the government's intertemporal budget constraint (2) is

$$i_t \frac{M_t}{P_t} = \frac{i_t \alpha \theta e^{-\alpha i_t}}{1 + \theta - \theta e^{-\alpha i_t}(1 + \alpha i_t)} y_t. \quad (14)$$

After maximizing the government's inflation tax revenue, the resulting first-order condition can be shown, after some manipulation, to reduce to

$$1 - \alpha i_t = \frac{\theta}{1 + \theta} e^{-\alpha i_t}. \quad (15)$$

Condition (15) implicitly determines the revenue maximizing inflation tax rate as a function of the underlying parameters of the transactions technology α and θ .

The determination of the revenue maximizing inflation tax rate is shown in Fig. 1. The $v(i_t)$ schedule is defined by $v(i_t) \equiv 1 - \alpha i_t$. It intersects the y -axis at 1, has a slope of $-\alpha$, and intersects the x -axis when the inflation tax rate is $1/\alpha$. The $v(i_t)$ schedule reflects the marginal gain in revenues from increasing the inflation tax rate *less* the marginal loss in revenues from the erosion of the tax base as consumers reduce their money holding per unit of consumption. These are the factors resulting in Cagan's rule for the revenue

Fig. 1. Determination of the revenue maximizing inflation tax rate for the Cagan money demand function.

maximizing inflation tax rate. The $h(i_t)$ schedule is defined by $h(i_t) \equiv [\theta/(1 + \theta)]e^{-\alpha i_t}$. The $h(i_t)$ schedule reflects the marginal revenue loss that occurs when the inflation tax rate is increased because increasing transactions costs causes consumption to fall relative to output. It intersects the y -axis at $\theta/(1 + \theta) < 1$. It can also be seen that $h'(i_t) < 0$ and $h''(i_t) > 0$. Additionally, $h(i_t)$ is bounded below by zero. These facts guarantee that there is a unique revenue maximizing inflation rate shown by $i_{\max}$. This rate depends on the properties of the underlying transactions technology as captured by the parameters α and θ .⁹

Although a closed form solution for the revenue maximizing inflation tax rate cannot be obtained, it is apparent from (15) and the preceding discussion that the solution takes the form $\alpha i_{\max} = s(\theta)$, where $0 < s(\theta) < 1$ and $s'(\theta) < 0$.¹⁰ The *level* of the revenue maximizing tax rate is thus $i_{\max} = (1/\alpha)s(\theta)$. It depends on both α and θ . However, the revenue maximizing inflation tax rate in the transactions cost approach, $i_{\max}$, is a fraction, $s(\theta)$, of the rate $1/\alpha$ implied by Cagan's rule, and this fraction depends solely on the parameter θ . From this and (12)–(14), it is evident that the corresponding equilibrium values of inflation tax revenues, consumption, and transactions costs at the revenue maximizing inflation tax rate, all as shares of output, are also uniquely determined by the parameter θ . Note also from (12)–(14) that, when the inflation tax rate is set to $1/\alpha$, as implied by Cagan's rule, the equilibrium values of inflation tax revenues, consumption, and transactions costs depend only on θ . This means that the equilibrium for the transactions cost economy when the inflation tax rate is set at its revenue maximizing level, as given in (15), and the equilibrium for the same economy when the inflation tax rate is set to according to Cagan's rule can be directly compared for different values of θ without worrying about the particular value of the interest semi-elasticity of the demand for money, α .¹¹

This serves as the basis for the quantitative assessment of the importance of the general equilibrium considerations for the revenue maximizing inflation tax rate that follows. Intuitively, the idea underlying these calculations is this: Suppose that the economy were as described here but that, in an attempt to maximize revenues, policymakers mistakenly were to set the inflation tax rate at the rate $1/\alpha$ implied by Cagan's rule. How different would this equilibrium be from the one that would be achieved if policymakers actually knew the true structure of the economy and instead of setting the inflation tax rate to $1/\alpha$ set it at the true revenue maximizing rate $i_{\max}$? The preceding discussion implies that the difference between these two equilibria with respect to inflation tax revenues, consumption, and transactions costs (as shares of output) depends entirely on the parameter θ . The parameter α matters only when comparing the *level* of the inflation tax rate that the policymaker will set in the two cases.

There is plenty of evidence suggesting plausible values for the interest semi-elasticity α . See, for example, Lucas (2000). A plausible range for the crucial scale parameter θ was pinned down by noting that if there were no transactions costs, inflation tax revenues for period t would be given by $i_t(M_t/P_t) = i_t\alpha\theta e^{-\alpha i_t}y_t$, where use has been made of the fact that money demand is given by (10) with $B = \alpha\theta$ and goods market equilibrium implies $c_t = y_t$. As a share of output, inflation tax revenues would thus be $i_t\alpha\theta e^{-\alpha i_t}$. Since Cagan's rule would maximize revenues under these conditions, maximum revenues from the

⁹In the general case it would also depend on labor supply considerations and the production technology.

¹⁰This follows from the fact that the semi-elasticity, α , and the inflation tax rate, i_t , always appear as αi_t in (15).

¹¹The equilibrium for other values of the inflation tax rate depends on both α and θ .

Table 1
Revenues and welfare under alternative inflation tax rates

θ	$i_{\max}/(1/\alpha)$	$\text{Rev}(i_{\max})$	$c(i_{\max})$	$T(i_{\max}) = w(i_{\max})$	$\text{Rev}(1/\alpha)$	$c(1/\alpha)$	$T(1/\alpha) = w(1/\alpha)$
0.136	0.954	4.8	96.7	3.4	4.8	96.5	3.6
0.204	0.933	7.1	95.3	4.9	7.1	94.9	5.4
0.272	0.914	9.4	94.1	6.3	9.3	93.3	7.2
0.340	0.897	11.5	92.9	7.7	11.5	91.8	8.9
0.408	0.880	13.7	91.7	9.0	13.5	90.3	10.7

Note: All numbers are as a percent of output except those in columns 1 and 2.

The values for θ are associated with maximum revenues from the inflation tax of 5%, 7.5%, 10%, 12.5%, and 15% of GDP in a partial equilibrium setting where there are no real effects from inflation and Cagan's rule is valid. Throughout the table $x(i)$ denotes that the variable x is being evaluated at the inflation tax rate $i = i_{\max}, 1/\alpha$.

inflation tax would thus be θe^{-1} . The parameter θ can thus be thought of as a scale parameter determining the maximum revenues that might be obtained from inflationary finance when Cagan's rule applies. Previous studies concerning the maximum level of inflation tax revenues as a share of output can thus be used to pin down reasonable values for θ .

Cagan (1956) finds that for the seven hyperinflations he studies, governments typically collected inflation tax revenues of less than 10% of output. The one exception is the Hungarian hyperinflation of 1923–1924 where revenues amount to somewhere between 13% and 19% of output. However, in every case the average monthly inflation rate exceeded the revenue maximizing inflation rate implied by Cagan's rule and by his estimate of the interest semi-elasticity of the demand for money. Barro (1997) and Romer (1996) both suggest that the available evidence implies that the maximum amount of revenue a government might hope to collect from the inflation tax is no more than 10% of GDP. Bali and Thurston (2000) have recently estimated Laffer curves for 30 countries.¹² Their estimates are based on the Cagan money demand function and thus are directly applicable here. Except for Peru and Zaire, which are extreme outliers on the high end, they find that the maximum revenues from the inflation tax range from a low of one quarter of one percent of GDP for Canada to a high of slightly over 17% of GDP for Argentina. In light of this evidence, θ is set here so that, if Cagan's rule is valid, maximum inflation tax revenues range between 5% and 15% of GDP. The results for values of θ consistent with maximum revenues less than 5% of GDP will be evident from the discussion that follows.

Results are presented in Table 1 for five values of the parameter θ ranging between 0.136 and 0.408. These values imply maximum revenues from the inflation tax of 5%, 7.5%, 10%, 12.5%, and 15% of GDP in a partial equilibrium setting where there are no real effects from inflation and Cagan's rule is valid. In the transactions cost approach, which explicitly accounts for the real effects of inflation, inflation tax revenues at both the revenue maximizing inflation tax rate, $i_{\max}$ and at the rate $1/\alpha$ are necessarily below these levels. All numbers in the table are expressed as a percent of output except those in column 1 for θ and those in column 2 for the revenue maximizing inflation tax rate relative to the

¹²Actually they estimate Laffer surfaces since they consider the combination of the inflation tax rate and reserve requirements that maximize revenues.

rate $1/\alpha$, $i_{\max}/(1/\alpha) = s(\theta)$. (Throughout the table $x(i)$ denotes that the variable x is being evaluated at the inflation tax rate $i = i_{\max}, 1/\alpha$.)

The first thing to note is that as the parameter θ increases, and as maximum inflation tax revenues become larger, the discrepancy between the revenue maximizing inflation tax rate and the rate implied by Cagan's rule also increases. When the maximum revenue from the inflation tax is low, the revenue maximizing inflation tax rate is 0.954 times the rate implied by Cagan's rule. When the maximum revenue from the inflation tax is high, the revenue maximizing inflation tax rate falls to only 0.88 times the rate implied by the Cagan's rule. To see what this means for *levels* of the revenue maximizing inflation tax rate and the associated nominal interest rate, $R_{\max}$ (recall $i_t = R_t/(1 + R_t)$), suppose that θ is 0.272. This corresponds to inflation tax revenues of 10% of GDP when Cagan's rule is valid. Given this value of θ , consider the implications of different values of the interest semi-elasticity of demand for money for the inflation tax rate and the nominal interest rate when they are set according to Cagan's rule as opposed to being set at their revenue maximizing values. When the demand for money is highly sensitive to changes in the nominal interest rate, $\alpha = 7$ for example, the inflation tax rate and nominal interest rate implied by Cagan's rule are fairly low at 14.3% [= 100(1/7)] and 16.7%, respectively. The revenue maximizing inflation tax rate and nominal interest rate are one to two percentage points lower at 13.1% and 15%, respectively. When $\alpha = 3$ the inflation tax rate and the nominal interest rate are 33% and 50% under Cagan's rule whereas the revenue maximizing values are 30.4% and 43.8%. Finally, when the demand for money is much less sensitive to changes in the nominal interest rate, $\alpha = 4/3$, the inflation tax rate and nominal interest rate implied by the Cagan's rule are considerably higher at 75% and 300%, respectively. The revenue maximizing inflation tax rate and nominal interest rate are much lower in absolute terms now, coming in at 68.6% and 218%. The differences for monetary policy and money growth rates implied by following the revenue maximizing policy versus the Cagan's rule can thus be quite substantial in absolute terms.

Another striking feature of Table 1 is its implications for inflation tax revenues. One key point is that the maximum revenues in the general equilibrium environment (column 3) can be significantly below those in partial equilibrium where the real effects of inflation are ignored. Within the context of the general equilibrium, transactions cost approach, a second key point is how small the difference is between inflation tax revenues at the revenue-maximizing rate, $i_{\max}$, and at the rate $1/\alpha$. Indeed, the difference is never more than two tenths of one percent of output and for three of the five cases considered it is so small that given the rounding in the table the differences don't even show up. This suggests that in the neighborhood of the revenue maximizing tax rate the revenue function is fairly flat.

Evidently, for a government whose *sole* concern is with revenues from the inflation tax, it matters little whether they set the inflation tax rate at the revenue maximizing level or at the rate $1/\alpha$. However, if the government is also concerned with consumption and welfare, the choice becomes much more meaningful. To see this, following Lucas (2000), define $w(i_t)$ as the proportionate increase in period t consumption that is necessary to just induce a consumer to accept an inflation tax rate of i_t rather than zero. This welfare cost measure thus satisfies the condition $U\{[1 + w(i_t)]y_t/(1 + T[k(i_t)])\} = U(y_t)$, given that transactions costs are zero when the nominal interest rate is zero and real balances are at their satiation level. Obviously $w(i_t) = T[k(i_t)]$. That is, the welfare cost of an inflation tax rate of i_t is simply equal to the transactions costs borne by consumers because real balances are lower

and transactions costs are higher than they would be if Friedman's rule were adopted. Columns 5 and 8 of Table 2 thus provide a measure of the welfare costs of inflation at the revenue maximizing inflation tax rate and at the rate $1/\alpha$. As can be seen, the difference in welfare costs rises significantly as the potential maximum revenues from the inflation tax rise. When maximum revenues are at 5% of GDP or less, following Cagan's rule leads to welfare costs that are at most two tenths of a percentage point higher than those incurred when the inflation tax rate is set at the revenue-maximizing level. This difference climbs to a more substantial half a percentage point when maximum revenues are near 7.5%, is almost a full percentage point at revenues in the 10% range, and is close to two full percentage points when maximum revenues are high at 15% of GDP. These differences are large as welfare cost measures go. (The differences in welfare costs are mirrored by the differences in consumption shown in columns 4 and 7 of Table 1.)

Another way to see the magnitude of these differences is to compare the welfare cost per unit of revenue raised, i.e., the average welfare cost of raising revenue, at the two tax rates. It is easy to show from Table 1 that, as θ increases, the average welfare cost per unit of revenue ranges from 4 cents per dollar to 13 cents per dollar higher when the inflation tax rate is set at $1/\alpha$ rather than at the revenue maximizing level $i_{\max}$. These results suggest the following conclusion: For economies where the maximum revenues from the inflation tax are lower than 5% of GDP, there is little difference between an inflation tax rate of $1/\alpha$ and the revenue-maximizing inflation tax rate. This is true either in terms of revenues or in terms of consumption and welfare. However, in economies where the maximum revenues from the inflation tax are 7.5% of GDP or higher the difference between following the Cagan's rule rather than setting the inflation tax rate at its revenue-maximizing level is economically substantial. The revenue differences are never big for the parameter values considered here, but the differences in consumption and welfare are sizable.

4. Conclusion

This paper has examined the revenue maximizing inflation rate in general equilibrium in an economy with a labor-leisure choice where money is held because it reduces transactions costs. It has been shown that in this environment the revenue maximizing inflation tax rate is below that implied by Cagan's rule. A stripped down version of the model was then used to provide back-of-the-envelope type calculations suggesting that the differences between the revenue maximizing inflation tax rate and that implied by Cagan's rule are economically significant.

A number of important extensions to this analysis readily suggest themselves. First, the theoretical model underlying the results abstracts from the government's ability to influence intertemporal relative prices and from related time consistency issues. It would be interesting to extend both the theory and the quantitative analysis to see how this would change the results. Second, the quantitative assessment of the results relies on a stripped down version of the model without a labor-leisure choice. The theory presented in Section 1 indicates that accounting for variable labor supply would only magnify the quantitative differences between the revenue maximizing inflation tax rate and that implied by Cagan's rule. It would be interesting to know by how much.

Although a transactions cost approach to money demand has been adopted here, the results appear to point to a general property of modern monetary models: They contain two channels by which inflation alters money holdings and thus affects the revenues from

money creation. One is the standard movement along a “money demand function” identified by Cagan. The other is a general equilibrium effect on scale, and perhaps other variables, that enter into the money demand function. The transactions cost framework is only one of many “more structural” models of money demand (relative to Cagan’s). However, it seems that any such model, appropriately constrained by the facts on money demand at the micro and/or macro levels, should have the two effects identified here. The example presented here thus illustrates key features that are likely to be recurrent in future work on the revenues from money creation that are grounded in modern monetary models.

References

- Bailey, M.J., 1956. The welfare cost of inflationary finance. *Journal of Political Economy* 64, 93–110.
- Bali, T.G., Thurston, T., 2000. Empirical estimates of inflation tax laffer surfaces: a 30-country study. *Journal of Development Economics* 63, 529–546.
- Barro, R.J., 1997. *Macroeconomics*. MIT Press, Cambridge.
- Brock, P.L., 1989. Reserve requirements and the inflation tax. *Journal of Money, Credit and Banking* 21, 106–121.
- Cagan, P., 1956. The monetary dynamics of hyperinflation. In: Friedman, M. (Ed.), *Studies in the Quantity Theory of Money*. University of Chicago Press, Chicago, pp. 25–117.
- Friedman, M., 1969. The Optimum Quantity of Money. In his *The Optimum Quantity of Money and Other Essays*. Aldine, Chicago, pp. 1–50.
- Kimbrough, K.P., 1986a. The optimum quantity of money rule in the theory of public finance. *Journal of Monetary Economics* 18, 277–284.
- Kimbrough, K.P., 1986b. Inflation, employment, and welfare in the presence of transactions costs. *Journal of Money, Credit and Banking* 18, 127–140.
- Lucas Jr., R.E., 2000. Inflation and welfare. *Econometrica* 68, 247–274.
- Romer, D., 1996. *Advanced Macroeconomics*. McGraw-Hill, New York.