

HYPERBOLIC DISCOUNTING AND POSITIVE OPTIMAL INFLATION

LIAM GRAHAM

University College London

DENNIS J. SNOWER

Kiel Institute for the World Economy

Christian-Albrechts University

CEPR

and

IZA

The Friedman rule states that steady-state welfare is maximized when there is deflation at the real rate of interest. Recent work by Khan, King, and Wolman [*Review of Economic Studies* 10 (4), 825–860] uses a richer model but still finds deflation optimal. In an otherwise standard New Keynesian model we show that, if households have hyperbolic discounting, small positive rates of inflation can be optimal. In our baseline calibration, the optimal rate of inflation is 2.1% and remains positive across a wide range of calibrations.

Keywords: Optimal Monetary Policy, Inflation Targeting, Phillips Curve, Nominal Inertia, Monetary Policy

1. INTRODUCTION

The Friedman rule states that optimal monetary policy in the long run requires deflation at the real rate of interest. More recent work by Khan et al. (2003) uses a richer model in which inflation or deflation has costs arising from price dispersion. They find the optimal rate of inflation to be higher than the Friedman rule, but still negative. Nevertheless policy makers have chosen positive inflation targets based on other criteria.

The central result of this paper is that, if households have hyperbolic discounting and there are multiperiod nominal wage contracts, low positive inflation rates can be optimal. Our baseline calibration implies an optimal rate of 2.1%, and we show that low positive rates between 1% and 4% are optimal over a wide range of calibrations.

To understand intuitively why hyperbolic discounting has this effect, it is important to realize that inflation redistributes the costs and benefits of employment through time. In particular, because a worker's nominal wage is constant over the wage contract period, whereas the price level rises continually, the real wage falls

We thank the Editor and two anonymous referees for their comments, which greatly improved the paper. Address correspondence to: Liam Graham, Department of Economics, University College London, Gower Street, London WC1E 6BT, UK; e-mail: Liam.Graham@ucl.ac.uk.

over the contract period. Thus the worker's disutility of labor is distributed from the present (when the real wage is relatively high and working hours correspondingly low) to the future (when the wage is low and working hours are high). Under hyperbolic discounting, the future receives much less weight than the present. Because the worker attaches relatively little importance to the future disutility of labor, she is inclined to set a wage that implies a higher average level of labor over the contract period.

The greater the inflation rate, the more the disutility of labor is shifted into the future. Thus the myopia inherent in hyperbolic discounting has a stronger effect, and so aggregate employment rises: there is a long-run positive relation between inflation and employment. This relation can provide a useful long-run role for monetary policy. Specifically, under imperfect competition and other market imperfections (such as income taxes), the equilibrium level of output and employment is inefficiently low. But by letting the money supply grow and thereby generating inflation, the monetary authority can stimulate long-run macroeconomic activity, thereby offsetting the imperfections. It turns out that, on this account, the optimal long-run inflation rate becomes positive.

For a positive rate of inflation to be optimal two conditions must be satisfied.

- (1) The zero-inflation steady state must be inefficient. In our model this arises from two sources: imperfect competition among firms and a labor income tax. It is this inefficiency that monetary policy can potentially address.
- (2) There exists, at least over some range of inflation above zero, a positive trade-off between inflation and real macroeconomic activity, such as production or employment. There are a number of possible theoretical justifications for the existence of such a trade-off, among them are downward nominal wage rigidities; "greasing the wheels" of the labor market; the optimality of seigniorage when only distortionary taxes are available; and the zero lower bound on nominal interest rates. Reviews can be found in Sinclair (2003) or Billi and Khan (2008).¹ In this paper we focus on the role of discounting.

Conventional New Keynesian models take account of the first condition (through their assumption that prices and/or wages are set under monopolistic competition), but not the second. The fact that discounting leads to a positive tradeoff is well known in the literature [Romer (1990); Ascari (1998); Graham and Snower (2004); Levin and Yun (2007)]. Yet with exponential discounting, the discount rate is close to zero, and this positive tradeoff is very small and only exists over a very small range of inflation. In short, macroeconomic activity is inefficiently low, but monetary policy is virtually powerless to do anything to raise it.

We show that the second condition becomes important under hyperbolic discounting, whereby households have a strong preference for a payoff today over a payoff tomorrow, but a much weaker preference for a payoff in a year's time over a payoff in a year and a day's time [Laibson (1997); Barro (1999)].²

The optimal rate of inflation is a function of the relative strengths of the different channels by which inflation affects real activity. In this paper, we present a simple New Keynesian model with households that choose between consuming

goods by paying cash, which is subject to the inflation tax, or with credit, which incurs a time cost. Households supply differentiated labor to firms that face a production function involving a Dixit–Stiglitz aggregate of labor types. Households set nominal wages by means of Taylor contracts.³

In this model there are five channels (described in detail in Section 4) by which inflation affects real variables. The first two of these channels are sources of the inflation tax: (i) “the wedge of monetary inefficiency” (the part of the inflation tax due to interest foregone on cash held) and (ii) “shopping time” (inflation raising the nominal interest rate and thereby inducing households to substitute credit for cash, thereby incurring a shopping time cost). These imply a negative relation between welfare and the rate of inflation. If only these two channels are present, the Friedman rule applies and the optimal monetary policy is to deflate at the real rate of interest.

In the presence of inflation, nominal wage contracts imply relative price instability in the form of wage dispersion. The third and fourth channels represent the effect of this wage dispersion on firms and households, respectively: (iii) “employment cycling” (inflation generates inefficient employment fluctuations through fluctuations in households’ real wages, in the presence of staggered nominal wage contracts) and (iv) “labor supply smoothing” (the employment fluctuations make households worse off, because of rising disutility of labor). These channels imply a negative relation between welfare and the absolute value of inflation. Considering only these two channels, the optimal inflation rate is zero. Thus in a model with the first four channels the optimal inflation rate will lie between zero and the Friedman rule. This is effectively the model of Khan et al. (2003).

The final channel is (v) “discounting” (a higher discount rate leads to households’ attaching less weight to disutility from future work and, given a rising profile of labor over the wage contract, leads households to set a lower wage, which means they supply more labor). As we noted earlier, discounting leads to a positive relation between inflation and output and employment. However, the welfare effects of this are theoretically ambiguous as increased output increases welfare but increased employment decreases welfare. In this paper we show that, under our baseline calibration, whereas optimal inflation without discounting is -1.2% (corresponding to Khan et al.’s 2003 result), if we introduce hyperbolic discounting the optimal rate becomes 2.1% .

Our result relies crucially on the interaction of nominal wage contracts and hyperbolic discounting. There is considerable empirical support for both of these features of the economy. For a review of the literature on hyperbolic discounting, see Laibson (1997). We review the evidence for the existence of nominal wage contracts in Section 6.2.

The paper is organized as follows. Section 2 presents the model. Section 3 explains the five channels, by which inflation affects real activity in the model. Section 4 gives our choice of parameters. Section 5 presents our benchmark result, explains the contribution of each of the five channels, and conducts sensitivity analysis. In Section 6 we discuss our assumptions on the nature of the wage

contract and draw some implications for monetary policy. Section 7 concludes. The Appendix contains the solution of the household’s problem. Further derivations are available in the working paper version, available from the corresponding author’s website.

2. MODEL

We present a simple dynamic general equilibrium model consisting of three types of agent: households, firms, and a government. There is no uncertainty, and in our *Results* section we focus on steady states of the model.

A large number S of households supply differentiated labor to firms and consume goods. Households can choose whether to purchase goods with cash or with nominal credit, which incurs a time cost. In the spirit of Taylor (1979), we group households into N wage-setting cohorts of equal size, each of which sets a nominal wage contract for N periods. Different cohorts set wages at different times, uniformly staggered.

Firms use all labor types to produce a homogeneous consumption good. The government imposes a tax on labor income, prints money, and rebates the proceeds to households as a lump sum.

2.1. Firms

The representative firm uses all types of labor in a production function with a constant elasticity of substitution [Dixit and Stiglitz (1977)] to produce a homogenous consumption good,

$$y_t = \left[\sum_{h=0}^S n_t(h)^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}, \tag{1}$$

where y_t is output, $n_t(h)$ is the amount of labor hired from household h , and θ is the elasticity of substitution between different labor types.

There is perfect competition in the product market. The firm’s cost minimization implies a standard demand function for each household’s labor,

$$n_t(h) = w_t(h)^{-\theta} y_t, \tag{2}$$

where $w(h)$ is the real value of the nominal wage set by household h . Note that because firms are perfect competitors, their markup will be zero and the real wage index constant and equal to unity.

2.2. Households

A typical household h maximizes its lifetime utility, defined over consumption c and leisure l :

$$U_t(h) = \sum_{i=0}^{\infty} \beta_{t,t+i} u[c_{t+i}(h), l_{t+i}(h)]. \tag{3}$$

We write the discount factor between periods t and $t + i$ in the general form $\beta_{t:t+i}$ to allow different specifications.

In what follows we restrict ourselves to utility functions of the following form:

$$u(c, l) = \log c + \zeta \frac{l^{1-\eta}}{1-\eta}. \quad (4)$$

Households maximize their utility subject to a budget constraint, written in real terms as

$$m_t(h) + \frac{1}{1+R} b_{t+1}(h) = \frac{b_t(h) - d_t(h)}{1+\pi} + (1-\tau) w_t(h) n_t(h) + \chi_t, \quad (5)$$

where m is real money balances, d is credit used to buy goods, χ is lump-sum transfers from government, and n is time spent working. R and π , the nominal rate of interest and the inflation rate, respectively, we assume to be constant.⁴ Holdings of one-period nominal bonds, which pay a unit on redemption, are given by b .

Households face a choice between buying goods with cash, and buying them with nominal credit. If the proportion of goods bought with credit is ξ , the households' real money balances are given by

$$m_t(h) = [1 - \xi_t(h)] c_t(h) \quad (6)$$

and holdings of nominal consumption debt, repayable in the next period, are

$$d_{t+1}(h) = \xi_t(h) c_t(h). \quad (7)$$

To model the split between cash and credit goods, we follow Khan et al. (2003) by assuming that buying goods with credit involves a cost in terms of "shopping time," n^s . Then if we normalize the total time endowment to unity, the household's time constraint is

$$l_t(h) + n_t(h) + n_t^s(h) = 1. \quad (8)$$

We follow Khan et al. (2003) in interpreting the single consumption good as a continuum of goods, and modeling the cash/credit split by assuming that if credit is used, the purchase involves a random fixed time cost ϖ that is only known after the consumer has chosen how much of the good to buy, but before the decision to buy it with cash or credit. If F is the cumulative distribution function of time costs, total time costs borne by household h are then

$$n_t^s(h) = \int_0^{F^{-1}[\xi_t(h)]} \varpi dF(\varpi). \quad (9)$$

The household also faces the downward-sloping demand curves for the household's labor type (2) and the constraint that it can only change its wage every N periods.

2.3. Government

Government revenue arises from a proportional tax τ on labor income and from seigniorage. There is no government spending, and proceeds are rebated to households equally by lump-sum transfers χ , so that the government budget constraint is

$$\tau \sum_{h=0}^S w_t(h) n_t(h) + \left(m_t - \frac{m_{t-1}}{1 + \pi} \right) = N\chi_t. \tag{10}$$

The government follows an exogenous rule for the money supply with a constant growth rate g and this gives steady state inflation, $\pi = \Delta m = g$.

2.4. Aggregates

In the absence of government spending, the aggregate budget constraint is

$$y_t = c_t. \tag{11}$$

We define aggregate labor as

$$n_t = \sum_{h=0}^S n_t(h) \tag{12}$$

and the aggregate wage index in terms of efficiency labor is

$$w_t = \left[\sum_{h=0}^S w_t(h)^{1-\theta} \right]^{\frac{1}{1-\theta}}. \tag{13}$$

The real interest rate is given by

$$r = R - \pi, \tag{14}$$

which is constant in a given steady state because R and π are assumed constant.

2.5. Equilibrium

A competitive equilibrium for the above economy is a sequence of plans for

- allocations of households⁵ $\{m_t(h), d_{t+1}(h), b_{t+1}(h), n_t^s(h), n_t(h)\}_{t=1:\infty}^{h=1:\infty}$
- prices $\{w_t(h), r\}_{t=1:\infty}^{h=1:\infty}$
- aggregate labor input $\{n_t\}_{t=1:\infty}$

such that

- (1) Given prices, the allocations solve the utility maximization problem. If preferences are time-inconsistent, the allocations are also Nash equilibria of the game-theoretic problem of each household.

- (2) $\{w_t\}_{t=1:\infty}$ is the marginal product of aggregate labor (12).
- (3) Product, money, and bond markets clear.

2.6. Discounting

With exponential discounting, the household’s discount factor is constant through time, $\beta_{t:t+i} = \beta^i$. To model hyperbolic discounting, we follow Laibson (1996, 1997) and much of the subsequent literature in approximating the hyperbolic discount function by a “quasi-hyperbolic” discount function in which the discount factors from the perspective of the current period are $1, \gamma\delta, \gamma\delta^2, \gamma\delta^3, \dots$. The discount factor between the first two periods is $\gamma\delta$; but that between all subsequent periods is δ . Thus we can think of γ as a short-run discount factor, and δ as a long-run discount factor.

In this case the household’s behavior is time-inconsistent: in future periods the household will have an incentive to change plans made in the current period. Because the household cannot commit itself to a plan beyond the current period and because a household with rational expectations will take this into account when making its decisions, it is standard practice [e.g., Laibson (1996)] to formulate the household’s problem as a game played by the household at time t (self t) against the household in future periods (future selves).

2.7. Solving the Household’s Problem

We show in the Appendix that the equilibrium strategy⁶ of an infinitely lived household with preferences (3) facing constraints (2), (5), and (8) and able to reset its nominal wage every N periods is

$$\begin{aligned}
 s_t(h) &= \{c_t^*(h), \xi_t^*(h)\} : t \neq iN, \\
 s_t(h) &= \{c_t^*(h), \xi_t^*(h), w_t^*(h)\} : t = iN,
 \end{aligned}
 \tag{15}$$

where the real value of the nominal wage the household chooses when it can change its wage is

$$w_t^*(h) = \omega \frac{\sum_{i=0}^{N-1} \beta_{t:t+i} l_{t+i}^{*-\eta}(h) n_{t+i}^*(h)}{\frac{1}{Nc_t^*(h)(1+R(1-\xi_t^*(h)))} \sum_{i=0}^{N-1} \frac{n_{t+i}^*(h)}{(1+\pi)^i}},
 \tag{16}$$

where

$$\omega = \frac{\zeta\theta(1 - (1+r)^{-1})}{(\theta - 1)(1 - \tau)(1 - (1+r)^{-N})}.$$

The proportion of credit goods is defined implicitly by

$$\xi_t^*(h) = F \left(\frac{R}{\zeta l_t^*(h)^{-\eta} \{1 + R[1 - \xi_t^*(h)]\}} \right), \tag{17}$$

and consumption is a constant fraction of a_t , the household’s lifetime wealth,

$$c_t^*(h) = \frac{r}{1 + R[1 - \xi_t^*(h)]} a_t(h). \tag{18}$$

This differs from the standard permanent income result that consumption is the annuity value of wealth, $ra(h)$, because of the distortionary effects of inflation, equivalent to a proportional tax on consumption, discussed in Section 3.5. Comparing (17) and (18) with the equivalent equations (34) and (35) in Khan et al. (2003) shows that the equilibrium strategy is of the same form as the optimal choice in a time-consistent world.

2.8. The Steady State

For the remainder of the paper we consider the properties of a symmetric steady state in which all households consume the same amount in all periods⁷ and, when they can reset their wage, choose a nominal wage that obtains the same real wage.

Households set nominal wages every N periods. Because contracts are set in staggered fashion, under nonzero inflation the N different cohorts charge N different real wages (the highest charged by the cohort that sets its wage this period; the lowest charged by the cohort that set its wage $N - 1$ periods ago). In short, the economy is characterized by real wage dispersion.

An important feature of this steady state is that the *intertemporal* distribution of real wages across the contract period is identical to the *intra*temporal distribution of wages between cohorts. This can be seen from a simple example involving $N = 4$ with four cohorts setting wage contracts lasting four periods. Whenever a cohort resets its wage, it chooses a nominal wage to give a real wage w^* . In the next period, the real value of this is $\frac{w^*}{1+\pi}$, and so on. Table 1 shows the resulting nominal and real wage distributions.

In the face of this nonsmooth path of income, households exchange bonds to smooth their consumption,⁸ and all households consume the same amount in every period,

$$c^* = \frac{1}{N} \sum_{i=0}^{N-1} \frac{w^* n^{i*}}{(1 + \pi)^i}, \tag{19}$$

where n^i is the labor supplied by a typical household in the i th period of its contract. The equation for the aggregate wage index, (13), implies that when households can reset their nominal wage they choose it so that its real value is

$$w^* = \left[\frac{1}{N} \frac{1 - \pi^{N(\theta-1)}}{1 - \pi^{\theta-1}} \right]^{\frac{1}{\theta-1}}. \tag{20}$$

TABLE 1. Nominal and real wage distribution in the steady state

		Nominal wage of cohort . . .			
Period	Price level	1	2	3	4
1	1	w^*	$\frac{w^*}{(1 + \pi)^3}$	$\frac{w^*}{(1 + \pi)^2}$	$\frac{w^*}{1 + \pi}$
2	$1 + \pi$	w^*	$w^*(1 + \pi)$	$\frac{w^*}{(1 + \pi)^2}$	$\frac{w^*}{1 + \pi}$
3	$(1 + \pi)^2$	w^*	$w^*(1 + \pi)$	$w^*(1 + \pi)^2$	$\frac{w^*}{1 + \pi}$
4	$(1 + \pi)^3$	w^*	$w^*(1 + \pi)$	$w^*(1 + \pi)^2$	$w^*(1 + \pi)^3$

		Real wage of cohort . . .			
Period		1	2	3	4
1		w^*	$\frac{w^*}{(1 + \pi)^3}$	$\frac{w^*}{(1 + \pi)^2}$	$\frac{w^*}{1 + \pi}$
2		$\frac{w^*}{1 + \pi}$	w^*	$\frac{w^*}{(1 + \pi)^3}$	$\frac{w^*}{(1 + \pi)^2}$
3		$\frac{w^*}{(1 + \pi)^2}$	$\frac{w^*}{1 + \pi}$	w^*	$\frac{w^*}{(1 + \pi)^3}$
4		$\frac{w^*}{(1 + \pi)^3}$	$\frac{w^*}{(1 + \pi)^2}$	$\frac{w^*}{1 + \pi}$	w^*

Combining (16) and (19) and using the government's budget constraint (10) to eliminate transfers gives a labor supply relation⁹

$$\sum_{i=0}^{N-1} \beta_{0:i} (1 - n^{*i})^{-\eta} n^{*i} = \frac{1}{\omega[1 + R(1 - \xi^*)]}. \quad (21)$$

This sets the present value of the marginal disutility of labor with respect to the wage (on the left-hand side) equal to the present value of the marginal utility of the extra consumption resulting from a higher wage (on the right-hand side). Because of consumption smoothing, the right-hand side is independent of the wage because an increase in the wage leads to an increase in consumption from (19) and thus a fall in its marginal utility, which, with log utility, exactly offsets the increase in income resulting from the higher wage.

Given an inflation rate π , we can use (20) to solve for the real equivalent w^* of the nominal reset wage. Given a level of aggregate labor, (2) then gives the path of labor across the contract period. Exploiting the symmetry between the intratemporal and intertemporal properties of the model, we can then calculate consumption c from (18) and obtain the proportion of credit goods ξ from (17) and the allocation between labor n , shopping time n^s , and leisure l from (8) and (9). We then iterate on the quantity of aggregate labor until (21) is satisfied.

3. FIVE CHANNELS BY WHICH INFLATION AFFECTS REAL ACTIVITY

In our model there are five channels by which the rate of inflation (equal to the rate of money growth in the steady state) affects real activity. In this section we describe each of these channels in turn.

3.1. Employment Cycling

The real wage dispersion arising from staggered contracts means that, given the form of the production function (1), firms substitute toward the labor of relatively cheap households and away from relatively expensive households. This is inefficient in the sense that for a given level of aggregate labor, the higher the degree of real wage dispersion, the lower will be the level of output. Thus real wage dispersion leads to a negative relation between the *absolute* value of inflation and real variables. This effect is familiar in the New Keynesian literature on optimal monetary policy. For Woodford [2003, Section 2.1, 393], cycling between different types of labor is the key distortion that leads to a stabilizing role for monetary policy.

3.2. Labor Supply Smoothing

With strictly concave preferences over leisure ($\eta > 0$), the fluctuations in hours arising from employment cycling make households worse off. In response, they supply less labor at a given wage over the contract period. This leads to a negative relation between the *absolute* value of inflation and real variables.

3.3. Discounting

When choosing their wage, households compare the present value (over the contract period) of the marginal disutility of labor (MDL) with the present value of the marginal utility of consumption (MUC). Because the latter is constant as a result of consumption smoothing and, under positive inflation, the former is rising, an increase in the discount factor means that the MDL falls by more than the MUC, so, other things being equal, the household will set a lower wage, which means it supplies more labor.

Working through the steps in detail:

- (1) Households set their wage so that the present value (over the contract period) of the marginal disutility of labor with respect to the wage is equal to the present value of the marginal utility of consumption from the labor income generated by a wage change. This is the labor supply relation (21).
- (2) When inflation is positive, labor supply increases over the contract period because of employment cycling, and thus the MDL increases over the contract period. In contrast, the MUC remains constant through time, because consumption is constant, given by (19).

- (3) Because of discounting, future marginal utilities and disutilities receive less weight than current ones. Because the MDL increases over the contract period, whereas the MUC remains constant, an increase in discounting reduces the present value of the MDL relative to the present value of the MUC. In response, the household sets a lower wage that raises its labor supply, thereby raising the present value of the MDL to bring the two present values back into equality. So, given inflation, as the discount rate increases, the labor supply increases.
- (4) An increase in inflation means that employment cycling gets stronger, so the MDL increases more over the contract period, and the discounting effect decreases the present value of the MDL more relative to the present value of the MUC. Consequently the household sets a wage that increases its labor supply to bring the two present values back into equality. So, given discounting, as the inflation rate increases, the labor supply increases.

3.4. Shopping Time

As inflation rises, and hence the nominal interest rate rises as well, households optimally substitute away from cash and toward credit goods, as given by (17). This involves direct costs in terms of shopping time, given by (9). The time constraint (8) then means that less time is available for market activities, so this generates a negative relation between the level of inflation and real activity.

3.5. Wedge of Monetary Inefficiency

The inflation tax raises the cost of consumption of cash goods, and we define the wedge of monetary inefficiency, following Khan et al. (2003), as the difference between the actual cost of consumption and the cost when the nominal interest rate is zero. From inspection of (18), this wedge is equal to $R(1 - \xi_t^*)$, the product of the nominal interest rate and the proportion of cash goods in consumption ($1 - \xi$). In other words, steady state inflation acts as a proportional tax on consumption, with the tax rate being $R(1 - \xi_t^*)$. This generates a negative relation between inflation and real activity.

3.6. Comparison

Table 2 compares the effects of five channels on real activity.¹⁰

4. CHOICE OF PARAMETERS

We choose a benchmark set of parameters, shown in Table 3, to allow us to explore the determinants of the optimal inflation rate in our model. We perform a sensitivity analysis for all the important parameters in Section 5.3.

TABLE 2. Five channels linking inflation and real activity

	Effect on real activity of an increase in π	
	$\pi > 0$	$\pi < 0$
Employment cycling	↓	↑
Labor supply smoothing	↓	↑
Discounting	↑	↑
Shopping time	↓	↓
Wedge of monetary inefficiency	↓	↓

4.1. Wage Setting

We take the length of a contract to be one year, following Taylor (1999), and assume it remains constant as inflation varies. We discuss the empirical rationale for this in Section 6.2.

4.2. Period of Analysis

In our model, the period of analysis is important for two reasons. First it determines the time over which the high short-run discount factor applies and thus the strength of the discounting effect—this is discussed in the next section. Second, it affects average money balances and so the importance of the inflation tax. We choose a period of analysis of one month, and this, along with our assumption of a contract length of one year, implies that there are twelve wage-setting cohorts, $N = 12$.

4.3. Preferences

We have already restricted households' utility functions (4) to be separable in consumption and leisure, implying logarithmic preferences over consumption, and we further choose logarithmic preferences over leisure, $\eta = 1$. The weight on

TABLE 3. Parameters

Parameter	Value
η	1
ζ	2.3
r	0.04
γ	0.7
N	12
θ	5
τ	0.35
n	0.2

leisure in the utility function, ζ , is chosen so that households spend 20% of their time working (standard in the RBC literature).

In Appendix A.8 we show that the three time preference parameters δ , γ , and r are related by

$$\delta = \frac{1}{1 + \gamma r}, \quad (22)$$

so given our calibration of the real interest rate, we only need to pick one of the discounting parameters.

The choice of γ , the parameter measuring the size of short-run subjective discounting, is more difficult. Values of between 60% and 70% are standard in the hyperbolic discounting literature, for example, Laibson et al. (2007), but applying to annual data. There are some attempts to estimate quasi-hyperbolic discount functions over shorter time horizons, all using the pattern of consumption between paydays. Shapiro (2005) finds a daily discount factor of around 90%; Huffman and Barenstein (2005) find a range from 87% to 97%; Mastrobuoni and Weinberg (2009) obtain broadly similar results. Note that these studies are incompatible with Laibson's work because they imply an annual discount factor much higher than 70%. There are also a number of serious methodological issues in comparing different studies.¹¹

The only survey we are aware of is Fredrick et al. (2002). Their findings on the short-run discount factor, in their Table 1a, suggest that it is somewhere between 0% and 80%. Given the lack of a consensus, we choose $\gamma = 70\%$ for our baseline calibration. This choice implies via (22) a value for the long-run discount parameter δ of 0.998. So our calibration is broadly compatible with Laibson's because our annual discount factor will be $\gamma\delta^{11} \approx 68\%$.

A further choice is whether the period of analysis should be the same as the period over which the high short-run discount factor applies. Making the period of analysis longer than the period of high discounting would have no great effect. Making it shorter would involve an extra parameter and would require assumptions about how to "compound down" the short-run discount factor. In the absence of any empirical guidance on these issues, we choose to keep the period of analysis the same as the period over which the short-run discount factor applies; i.e., the "kink" in the quasi-hyperbolic discount function is after one period.

4.4. Money Demand

Demand for real balances in our model is given by $m_t = (1 - \xi_t)c_t$ and ξ , the proportion of credit goods, depends on the distribution of costs via (17). Because we take our model of credit goods from Khan et al. (2003), we can also use their estimates for the form of the cost distribution function F . They posit that F has a generalized beta distribution,

$$F(x) = \xi_L + \xi_H B\left(\frac{x}{\kappa}; b_1, b_2\right), \quad (23)$$

where B is the beta distribution, and then use U.S. data to estimate the five parameters, $\xi_L, \xi_H, \kappa, b_1, b_2$, based on two data sets, a “short sample” starting in 1959 and a “long sample” starting in 1948.¹² For our benchmark calibration, we use the short sample estimates and report results for the long sample estimates in our sensitivities section.

4.5. Other Parameters

For the elasticity of labor substitution, θ , we choose a value of 5.¹³ We set the level of labor taxation τ to be 35%, corresponding roughly to the average for the OECD; see Mendoza et al. (1994).

4.6. The Zero-Inflation Steady State

At zero inflation, there are no relative price distortions because the real wage is constant over the contract period. Given the above choice of parameters, households buy 65% of goods with credit, and hence the ratio of money balances to consumption is 35%. Households spend 0.003% of their time endowment shopping. This is similar to the value in Khan et al. (2003) and is consonant with the observation by Lucas (2000) that households spend “several hours per year” avoiding the inflation tax.

5. THE OPTIMAL RATE OF INFLATION

Optimal monetary policy in the long run involves choosing a steady state inflation rate (money growth rate) to maximize social welfare. However, defining social welfare is not straightforward, because with time-inconsistent preferences, different selves will not necessarily agree on rankings of utility flows [see İmrohoroğlu et al. (2003)]. One approach is to define social welfare as a sum of the utility of all selves at all points in time, choosing a backward discount factor to measure how a self at some particular time values utility flows at earlier times. The choice of a backward discount factor is complex [see İmrohoroğlu et al. (2003) and Caplin and Leahy (2004)], but it can be assumed to be between 0 and 1 (i.e., utility in the past is valued less than utility today). It can be shown that the social welfare function, defined as the sum of the utility of all N selves in all N contracts, is proportional to

$$SW = \sum_{i=0}^{N-1} u(c^i, l^i), \quad (24)$$

where the constant of proportionality is a function of the backward discount factor and the steady state rate of interest; i.e., it is independent of all the parameters we will vary.¹⁴

Output and hence consumption in the zero-inflation steady state are below the first best because of the presence of distorting labor taxes and imperfect

TABLE 4. Selectively removing channels

	Optimal inflation rate (%)
Benchmark result	2.1
Eliminate	
1. Discounting	-1.3
2. Discounting, shopping time	-0.88
3. Discounting, wedge	-0.55
4. Discounting, shopping time, wedge	0.0
5. Discounting, employment cycling, labor supply smoothing	-3.9

competition. This means that if there is a positive trade-off between real variables and inflation, at least in some range from zero inflation upward, the monetary policy maker can potentially improve social welfare by choosing higher inflation.

Given our choice of parameters, the annual inflation rate that maximizes social welfare is 2.1%. In the next two sections we explain the role of the five channels in this result, and show how it depends on our choice of parameters.

5.1. Selectively Removing Channels

How much does each of the five channels by which inflation affects real variables (described in Section 3) contribute to our benchmark result? Table 4 reports the effect of selectively eliminating combinations of the channels.¹⁵

Discounting is the only one of the five channels which leads to a positive relation between inflation and real activity, and eliminating it means that deflation becomes optimal. This mirrors the result of Khan et al. (2003) that, in the absence of discounting, the other channels imply that mild deflation is optimal (Khan et al.'s equivalent optimal rate is -0.76%).

Shopping time and the wedge of monetary inefficiency represent the effects of the inflation tax in terms of shopping time and the cost of consumption, respectively. Eliminating either of these therefore will raise the optimal inflation rate, and lines (2) and (3) of the table show this. If both of these effects and discounting are eliminated, the only channels that remain are employment cycling and labor supply smoothing. Both of these imply a negative relation between the *absolute* value of inflation and real variables: hence taken on their own they imply that zero inflation is optimal, line (4) of the table.

If we eliminated the effects of relative price instability by "switching off" employment cycling, the only effects of inflation would be shopping time and the wedge of monetary inefficiency. Setting the nominal interest rate to zero eliminates both of these (as can be seen from inspection of (17) and (21)), so the Friedman rule is optimal. With $R = 0$, the Fisher relation implies that $1 + \pi = (1 + r)^{-1}$, and given our choice of $r = 4\%$, this means $\pi = -3.8\%$.

5.2. Monetary Policy and Inefficiency

The level of optimal inflation will depend on the degree to which the steady state is inefficient. In our model the zero-inflation steady state is inefficient because of imperfect competition in the labor market and the tax on labor income, and we can demonstrate the following striking result: the greater the inefficiency in our economy—so that the lower is steady-state output relative to optimal output—the greater is the optimal inflation rate. In short, the optimal inflation rate is not independent of the degree of economic inefficiency. On the contrary, the optimal inflation rate is positive because it can compensate for an inefficient steady state level of real activity.

The degree of imperfect competition is given by the parameter θ , which represents the level of market power of a particular labor type. From (1), the higher is θ , the more substitutable are labor types, so the lower is market power. If we weaken the degree of imperfect competition by raising θ from the baseline of 5 to a value of 7, the optimal rate of inflation falls to 1.8%. Similarly, if we increase the degree of imperfect competition by choosing $\theta = 3$, the optimal rate of inflation increases to 2.3%.

The proportional labor income tax is a rate τ , chosen to be 35% in our baseline calibration. Reducing the tax rate to 25%, thus making the steady state less inefficient, reduces the optimal rate of inflation to 1.7%. Increasing the tax rate to 45% increases the optimal inflation rate to 2.6%.

This result has wide-ranging implications, and we discuss some of these in Section 6.3.

5.3. Sensitivity Analysis

In this section we describe how variations in other model parameters affect the optimal rate of inflation. Although there is a high degree of uncertainty about many of the parameters of our model, this section shows that our main result, that a positive rate of inflation is optimal, is robust to variations in our benchmark calibration.

Discounting and Period of Analysis. Table 5 shows how the optimal rate of inflation varies with the short-run discount factor and the period of analysis.¹⁶ We choose short-run discount factors ranging from 60% to 95%, and a period of analysis from semiannual to weekly (semiannual contracts are the longest that are compatible with overlapping wage contracts a year in length). Looking along the rows, we see that the higher the short-run discount factor, the greater the optimal rate of inflation. For optimal inflation to be positive when equilibrium output is inefficiently low, there must be a positive relation between inflation and real economic activity. The only channel in our model that leads to such a positive relation is discounting. The more heavily households discount the short-term future, the greater is the positive effect of inflation on real economic activity and

TABLE 5. Sensitivity of optimal inflation rate to discount factor and period of analysis

<i>N</i>	0.6	0.7	0.8	0.9	0.95
52	1.4%	1.0%	0.74%	0.56%	0.47%
12	3.7%	2.1%	0.92%	-0.02%	-0.41%
4	8.8%	5.6%	2.8%	-1.9%	-2.8%
2	10%	6.8%	2.2%	-3.9%	-3.9%

thus the greater is the optimal inflation rate. Note that the macroeconomic literature on hyperbolic discounting typically uses values in the first three columns of the table.

Our baseline calibration takes the period of analysis to be a month. In standard models, the period of analysis has small effects, if any. In contrast, in our model the period has two offsetting effects. First, the longer the period of analysis, the greater are the real money balances that households hold, other things being equal, and thus the greater is the effect of the inflation tax and the lower is the optimal rate of inflation. Second, the period of analysis is also the length of time for which the high short-run discount rate applies, and tends to strengthen the discounting effect and thus raise the optimal rate of inflation. The table shows that the first effect dominates when discounting is weak; and the second when discounting is strong.

Shopping Time. The properties of shopping time and the monetary wedge are largely determined by the estimates of the shopping time function, which we take from Khan et al. (2003). They present two sets of estimates, one based on a short sample period and the other on a longer one. In our base calibration we use the former. If instead we use the latter, the optimal inflation rate becomes 2.6%. Although the optimal rate does not change much, the relative magnitudes of the effects of the monetary wedge and shopping time change considerably.¹⁷

Labor Supply Elasticity. As the elasticity of labor supply falls (η increases), households become more averse to nonsmooth paths of labor, so the positive trade-off between output and inflation weakens and the optimal rate of inflation falls. This effect is shown in Table 6.¹⁸

Summary. As discussed in Section 4, there is a high degree of uncertainty over all the important parameters of the model, in terms both of their measurement and of mapping what is a very stylized model onto the world. However, this section has shown that our result of a positive optimal rate of inflation is robust to plausible variations in the parameters, with the exception of the discount factor. This reemphasizes the importance of discounting as the only channel that leads to a positive relation between inflation and real variables.

TABLE 6. Sensitivity to labor supply elasticity

η	Optimal inflation rate (%)
0	6.7
1	2.1
2	1.4
5	0.72
10	0.45

6. DISCUSSION

6.1. Calvo Contracts

Calvo contracts are common in standard New Keynesian models, which are linearized around a zero inflation rate. However, they give anomalous results in models with nonzero steady state inflation. The problem is as follows. Calvo contracts imply that, if there are a large number of households, some of them never reset their nominal wage. With positive steady state inflation, this means the real wage of such households falls to zero. Given the simple CES form of firms' production function (1), firms substitute completely to these wage setters, meaning production falls toward zero. Bakhshi et al. (2003) address this issue more formally and show that Calvo contracts imply a maximum rate of steady state inflation that can be analyzed—given their calibration, this maximum value is 5.5%—and that they distort the steady state for any value of inflation substantially different from zero.

6.2. The Absence of Indexation

A key assumption of our analysis is that wages are set annually in nominal terms, and that within this contract period of one year there is no indexation. In this assumption we follow Khan et al. (2003), who specify a more complex process of price adjustment than us but still assume that it is exogenous. In this section we review the evidence for this.

In thinking about the nature of wage contracts, it is important to distinguish between unionized and nonunionized workers. Taylor (1999), reviewing the direct empirical evidence for the unionized sector in the United States, cites a large number of studies that suggests that annual contracts are the most common length of wage-setting interval. The wide variation across countries is captured by Groth and Johansson (2004) who update the study of Bruno and Sachs (1985) with data from 1985–1995. They assign countries an index that takes a value of 0 if the average length of union contract is a year or less, 1 if it is from 1 to 3 years, and 2 if it is greater than 3 years. The mean of this index across OECD countries was 0.94 for 1985–1995, suggesting a mean contract length greater than a year.

There is much less evidence on the nonunionized sector, partly because the most common source of data panel data sets (the PSID for the United States, or the BHPS for the United Kingdom), which typically collect their data annually and so necessarily miss any subannual wage setting. However, there does seem to be a consensus in the literature that most wages are set annually. For example, for the United States, Taylor (1999) writes, “setting nominal wages at a fixed level for more than several quarters and usually for as long as a year appears to be just as prevalent for workers who are not in unions, or do not have formal employment contracts, as for union workers with employment contracts.” Smith (2000), referring to the United Kingdom, notes that “pay negotiation in the United Kingdom typically occur at annual intervals, and pay awards are often made every 12 months even for workers who are not covered by collective bargaining,” and Brown et al. (2004) state that the annual wage settlement is “the principal source of pay change for most employees in Britain.” This is consonant with estimates of the degree of wage stickiness [reviewed in Taylor (1999)] using aggregate data, which typically find the average length of time between wage changes to be greater than one year.

Because we assume that the length of nominal contracts is fixed at one year, whereas we vary inflation, our model is subject to a version of the Lucas critique. Would agents not change their wage-setting behavior as steady state inflation varied? So an important empirical question is the degree to which the contract period changes with steady state inflation. Taylor (1999) writes that the average contract period fell to one year (our baseline value) during the great inflation of the 1970s, when average inflation was well above the range we consider in this paper. Looking across the OECD, Groth and Johansson (2004) find that their index variable was 0.78 in the period 1975–1985, still consistent with an average contract length above a year, and rose to 0.94 in the next 10 years, when average inflation was lower. So this suggests that modeling nominal contracts as fixed at 1 year over a range of inflation from 0 to 10% as we do in this paper is, if anything, a conservative assumption, and we could expect contracts to be longer now that average inflation is low.

Our assumption of Taylor contracts, without intracontract indexation, instead of Fisher contracts, is similarly open to the Lucas critique. First note that our model of annual nominal contracts nests annual indexation, because wages are optimally reset each year. Only indexation within the annual contract period would affect our results. The empirical evidence on indexation is striking. Card (1986), writing about a period when average inflation was well over the top of the range we consider, notes that “perhaps no more than 10% of all US workers are covered by cost-of-living provisions.” This surprising result arises because indexation is only observed in union contracts and, as Card (1983) reports, only in around 60% of these, and “escalation provisions are rare in the non-union sector.” Further, although one can find evidence of quarterly cost of living (COLA) adjustments, the vast majority seem to be annual [Vroman, (1985); Kaufman and Woglom (1986)]. More recently, Christofides and Leung (2003), looking at Canadian data

from 1976 to 1999, write that “very few contracts contain COLA clauses.” So although wage indexation is a feature of very high inflation rates [Marinakis (1997)], it seems largely absent over the range of inflation we consider.

6.3. Monetary Policy Making

Although our model is very stylized, an implication of our analysis is that the structure of the economy—insofar as it affects the degree of economic inefficiency—should be a factor that influences policy makers’ choice of an inflation target. At present, inflation targets are chosen to be around 2%, largely on the grounds that such a rate is consistent with price stability (in the sense that inflation of around 2% has little adverse impact on agents’ decisions) and sufficiently positive to avoid the nominal rate hitting the zero lower bound too often. However, our analysis suggests that the structure of the economy should also be taken into account on choosing an inflation target. Specifically, the more inefficient are equilibrium output and employment, the higher the inflation target should be. This implies that inflation targets should be reevaluated over time as the structure of the real economy changes. It also has obvious implications for the analysis of optimum currency areas. Different countries in a currency union will in general have different degrees of steady state inefficiency, so the mechanism in this paper gives a new source of differential welfare effects across countries in addition to those discussed in, for example, Benigno (2004).

We most certainly do not wish to argue that monetary policy should always compensate for inefficiencies in labor and product markets. Usually policy makers can deal with these inefficiencies more effectively through the use of structural and fiscal policies, e.g., policies that promote competition in product and labor markets, structural policies that reduce institutional rigidities, and taxes and subsidies that overcome market failures in these markets. The reason is that the latter policies may reduce the inefficiencies at lower cost than monetary policy. But for any given level of inefficiency that remains once fiscal and structural policies are at their optimal levels, monetary policy should not ignore the inefficiencies that remain.

7. CONCLUSION

We have shown that the interaction between nominal contracts and hyperbolic discounting can lead to a tradeoff between inflation and macroeconomic activity for which a small positive inflation rate is optimal. In our baseline calibration, this rate is 2.1% and remains positive for a wide range of calibrations. Our analysis is meant to narrow the gulf between monetary theory and practice. Thus far, monetary theory has suggested that the optimal inflation rate is negative—beginning with the Friedman rule and ranging to the more moderate, but still negative optimal inflation of Khan et al. (2003)—but central bankers have never taken this seriously.

Our analysis is a first step in reconciling the optimal inflation literature with the practice of targeting small positive inflation rates.

The optimal inflation rate in our model is a result of the relative magnitudes of five channels by which inflation affects real activity, and we have analyzed the effects of these channels in our model: shopping time, the wedge of monetary inefficiency, employment cycling, labor supply smoothing, and discounting.

A general implication of our analysis is that the optimal rate of inflation depends on the structure of the economy in two ways. First, via the degree of structural inefficiency. We show that as the equilibrium becomes more inefficient, because of higher labor taxes or more imperfect competition, the optimal inflation rate rises. Second, via the mechanisms by which inflation affects real variables. An important topic for further research would be to examine a variety of different common inefficiencies and assess whether they have similar implications for monetary policy.

NOTES

1. There is also a literature finds a zero rate of inflation optimal by analysing the effects of steady state inflation on growth, for example, Gomme (1993) and Amano et al. (2008).

2. Barro (1999) shows that such high short-run discount factors are compatible with observed market returns and that there exists an “observational equivalence” between an economy with hyperbolic discounting and one with exponential discounting. Graham and Snower (2008) show that the presence of nominal wage contracts breaks this observational equivalence in the sense that an econometrician estimating both the Euler equation and the long-run Phillips curve could distinguish an economy with hyperbolic discounting from one with exponential discounting.

3. Another possibility would be Calvo contracts. We discuss our choice in Section 6.1.

4. Because our results will relate only to steady states there is no loss of generality here.

5. Equations (6) and (7) allow this to be written either in terms of money m and consumption debt holdings d or consumption c and the credit goods split ξ .

6. Equations (2), (8), and (9) show that choosing a strategy in terms of the wage and the proportion of goods to buy with credit $\{w, \xi\}$; is equivalent to choosing one in terms of the time to allocate to employment and shopping activities $\{n, n^s\}$. On this account, $l_t^*(h)$ is not listed as an argument of the household’s equilibrium strategy.

7. In the absence of redistribution, different cohorts would have (very slightly, under the range of inflation rates we consider in the paper) different present values of lifetime wealth because they start time in different periods of the wage contract. In Appendix B of the working paper version, available on the corresponding author’s website, we show the start-of-time redistribution necessary to remove this asymmetry.

8. We show that the bond market clears in Appendix B of the working paper version, available from the corresponding author’s website.

9. This expression is derived in Appendix B of the working paper version, available from the corresponding author’s website.

10. There is a close mapping between these five channels and the model of Khan et al. (2003). The key difference is that we assume sticky nominal wages, whereas Khan et al. (2003) assume sticky nominal prices. The real wage dispersion in our model corresponds to “relative price distortion” in theirs: we prefer the term “employment cycling” because it focuses on the effect of inflation on quantities directly relevant for welfare, rather than prices. “Labor supply smoothing” corresponds to their “markup distortion” —in our model inflation, via household preferences, effects the markup of

the real wage over the marginal disutility of labor. The discounting effect is present in their model, but is small because of their assumption of exponential discounting [for a comparison of exponential and hyperbolic discounting see Graham and Snower (2008)].

11. One issue is that discounting would be expected to vary across groups in society. Because the different studies use different data sets, they may be measuring the discounting of different groups. In the current paper we describe the behavior of wage setters who will be present in different data sets to different extents. Another issue is that discounting parameters are estimated jointly with the intertemporal elasticity of consumption and also depend on other assumptions such as the extent of credit constraints.

12. Details are in Khan et al. (2003), footnote 22.

13. Erceg et al. (2000) use a value of 4; Huang and Liu (2002) use 6.

14. See Appendix B of the working paper version, available from the corresponding author's website.

15. We eliminate the discount effect by setting $\beta_{t,t+i} = 1$ for all t, i ; labor supply smoothing by setting $\eta = 0$; the wedge by setting $R(1 - \xi)$ to zero in the labor supply equation; shopping time by forcing $n_t^s = 0$; employment cycling by forcing n to be constant over the contract period at its calibrated value.

16. When we change the short-run discount factor we also need to change the long-run discount factor according to (22) to keep the real interest rate at its calibrated value.

17. In our baseline calibration, shopping time increases by 1.9% from its value at zero inflation to its value at the optimal rate, and the proportion of credit goods by 0.02%. Using the "long sample" estimates of money demand, these values become 0.15% and 0.13% respectively.

18. The table includes the limiting case $\eta = 0$. This corresponds to the case of "indivisible labor" [Hansen (1985)] in the RBC literature but the justification for this is based on a lottery among heterogeneous agents. Such a justification is not valid in our model because we would expect the outcome of such a lottery to be different under hyperbolic discounting.

19. This is necessary to ensure that in the steady state households in all cohorts have the same lifetime wealth. For full details see Appendix B of the working paper version, available from the corresponding author's website.

REFERENCES

- Amano, Robert, Kevin Moran, Stephen Murchison, and Andrew Rennison (2007) Trend Inflation, Wage and Price Rigidities, and welfare. Working Papers 07-42, Bank of Canada.
- Ascari, Guido (1998) Superneutrality of money in staggered wage-setting models. *Macroeconomic Dynamics* 2, 383–400.
- Bakhshi, Hasan, Pablo Burriel-Llombart, Hashmat Khan, and Barbara Rudolf (2003) Endogenous Price Stickiness, Trend Inflation, and the New Keynesian Phillips Curve. Working paper 191, Bank of England.
- Barro, Robert J. (1999) Ramsey meets Laibson in the neoclassical growth model. *Quarterly Journal of Economics* 114(4), 1125–1152.
- Benigno, Pierpaolo (2004) Optimal monetary policy in a currency area. *Journal of International Economics*, 63(2), 293–320.
- Billi, Roberto M. and George A. Khan (2008) What is the optimal inflation rate? *Economic Review, Federal Reserve Bank of Kansas City Q II*, 5–28.
- Brown, Donna, Peter Ingram, and Jonathan Wadsworth (2004) Everyone's a Winner? Union Effects on Persistence in Private Sector Wage Settlements: Longitudinal Evidence from Britain. Department of Economics discussion papers 1104, Department of Economics, University of Surrey.
- Bruno, M. and Jeffery D. Sachs (1985) *The Economics of Worldwide Stagflation*. Oxford, UK: Blackwell.
- Caplin, Andrew and John Leahy (2004) The social discount rate. *Journal of Political Economy*, 112(6), 1257–1268.

- Card, David (1983) Cost of living escalators in major union contracts. *Industrial and Labor Relations Review* 37, 34–48.
- Card, David (1986) An empirical model of wage indexation provisions in union contracts. *Journal of Political Economy* 94(3), S144–S175.
- Christofides, Louis N. and Man Tuen Leung (2003) Nominal wage rigidity in contract data: A parametric approach. *Economica* 70(280), 619–638.
- Dixit, Avinash K. and Joseph E. Stiglitz (1977) Monopolistic competition and optimum product diversity. *American Economic Review* 67(3), 297–308.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000) Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics* 46, 281–313.
- Fredrick, Shane, George Loewenstein, and Ted O’Donoghue (2002) Time discounting and time preference: A critical review. *Journal of Economic Literature* 40, 351–401.
- Gomme, Paul (1993) Money and growth revisited: Measuring the costs of inflation in an endogenous growth model. *Journal of Monetary Economics* 32, 51–77.
- Graham, Liam and Dennis J. Snower (2004) The Real Effects of Money Growth in Dynamic General Equilibrium. Working paper 412, European Central Bank.
- Graham, Liam and Dennis J. Snower (2008) Hyperbolic discounting and the Phillips curve. *Journal of Money, Credit and Banking* 40(2–3), 427–448.
- Groth, Charlotta and Asa Johansson (2004) Bargaining structure and nominal wage flexibility. *European Economic Review* 48(6), 1349–1365.
- Hansen, G.D. (1985) Indivisible labor and the business cycle. *Journal of Monetary Economics* 16(3), 309–327.
- Huang, Kevin X.D. and Zheng Liu (2002) Staggered price setting, staggered wage setting and business cycle persistence. *Journal of Monetary Economics* 49, 405–433.
- Huffman, David and Matias Barenstein (2005) A Monthly Struggle for Self-Control? Hyperbolic Discounting, Mental Accounting, and the Fall in Consumption Between Paydays. IZA Discussion paper 1430.
- İmrohoroğlu, Ayse, Selahattin İmrohoroğlu, and Douglas H. Joines (2003) Time-inconsistent preferences and social security. *Quarterly Journal of Economics* 118(2), 745–784.
- Kaufman, Roger T. and Geoffrey Woglom (1986) The degree of indexation in major U.S. union contracts. *Industrial and Labor Relations Review*, 39(3), 439–448.
- Khan, Aubhik, Robert G. King, and Alexander L. Wolman (2003) Optimal monetary policy. *Review of Economic Studies* 10(4), 825–860.
- Laibson, David (1996) Hyperbolic Discount Functions, Undersaving, and Savings Policy. NBER working paper 5635.
- Laibson, David (1997) Golden eggs and hyperbolic discounting. *Quarterly Journal of Economic* 112(2), 443–477.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman (2007) Estimating Discount Functions with Consumption Choices over the Lifecycle. NBER working paper 13314.
- Levin, Andrew and Tack Yun (2007) Reconsidering the natural rate hypothesis in a New Keynesian framework. *Journal of Monetary Economics* 54(5), 1344–1365.
- Lucas, R.E., Jr. (2000) Inflation and welfare. *Econometrica* 68(2), 247–274.
- Marinakakis, Andrés E. (1997) Wage policy in high inflation countries: The role of indexation in Latin America during the 1980s. *Journal of Economic Studies* 24(6), 356–378.
- Mastrobuoni, Giovanni and Matthew Weinberg (2009) Heterogeneity in intra-monthly consumption patterns, self-control, and savings at retirement. *American Economic Journal: Economic Policy* 2, 163–189.
- Mendoza, Enrique G., Assaf Razin, and Linda L. Tesar (1994) Effective tax rates in macroeconomics: Cross-country estimates of tax rates on factor incomes and consumption. *Journal of Monetary Economics* 34(3), 297–323.
- Romer, David (1990) Staggered price setting with endogenous frequency of adjustment. *Economics Letters* 32(3), 205–210.

Shapiro, Jesse (2005) Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of Public Economics* 89(2–3), 303–325.

Sinclair, Peter J.N. (2003) The Optimal Rate of Inflation: An Academic Perspective. Bank of England Quarterly Bulletin 2000 (Autumn), 343–351.

Smith, Jennifer C. (2000) Nominal wage rigidity in the United Kingdom. *Economic Journal* 110(462), C176–C195.

Taylor, John (1979) Staggered wage setting in a macro model. *American Economic Review* 69(2), 108–113.

Taylor, John (1999) Staggered price and wage setting in macroeconomics. In John. B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, chap. 15, pp. 1009–1050. Elsevier.

Vroman, Wayne (1985) Cost-of-living escalators and price–wage linkages in the U.S. economy, 1968–1980. *Industrial and Labor Relations Review* 38(2), 225–235.

Woodford, Michael (2003) *Interest and Prices*. Princeton, NJ: Princeton University Press.

APPENDIX: SOLVING THE HOUSEHOLD’S PROBLEM

Under hyperbolic discounting, the household’s preferences are time-inconsistent, so standard dynamic programming techniques cannot be used. Instead, we model the household at time t as playing a game against its future selves. Throughout this Appendix we are considering the problem of a single household, so we drop the (h) notation for brevity.

A.1. PRELIMINARIES

Household utility (4) is defined in terms of consumption c_t and leisure l_t . However, leisure is determined by (i) the proportion of credit goods ξ_t and (ii) the wage w_t , which determines the firm’s labor demand. We use the time constraint (8) and the expressions for labor demand (2) and shopping time (9) to write leisure in terms of the choice variables,

$$l_t = 1 - w_t^{-\theta} y_t - \int_0^{F^{-1}(\xi_t)} \varpi dF(\varpi), \tag{A.1}$$

and we use this to rewrite the utility function in terms of the choice variables (recall that y_t is exogenous to the household’s problem),

$$v(c_t, w_t, \xi_t) = \log c_t + \frac{\zeta}{1 - \eta} \left[1 - w_t^{-\theta} y_t - \int_0^{F^{-1}(\xi_t)} \varpi dF(\varpi) \right]^{1 - \eta}. \tag{A.2}$$

It will prove helpful to define the consumption function in terms of net assets. This section reconciles this approach with the budget constraint (5), reproduced here:

$$m_t + \frac{1}{1 + R} b_{t+1} = \frac{b_t - d_t}{1 + \pi} + (1 - \tau) w_t n_t + \chi_t. \tag{A.3}$$

- (1) We substitute for money and credit from (6) and then subtract d_{t+1} from both sides, using (7) to obtain

$$b_{t+1} - d_{t+1} = (1 + r)(b_t - d_t) + (1 + R) [(1 - \tau) w_t n_t + \chi_t] - (1 + R)(1 - \xi_t) c_t - \xi_t c_t.$$

- (2) We define net assets as net bonds less credit, $a_t = b_t - d_t$
 (3) We assume that at the start of time the household sells its lifetime income (which consists of labor income, transfers, and the initial redistribution of wealth, which is equivalent to spreading the household's income equally over the contract period¹⁹) in exchange for bonds,

$$a_0 = \sum_{j=0}^{m-1} \frac{1}{(1 + r)^{mN}} z_j, \tag{A.4}$$

where z_j is the discounted value of *average* contract income in the contract period starting in $t = jN$,

$$z_j = \frac{\varphi}{N} \sum_{t=jN}^{(j+1)N-1} [(1 - \tau) w_t n_t + \chi_t], \tag{A.5}$$

and

$$\varphi = \frac{1 - (1 + r)^{-N}}{1 - (1 + r)^{-1}}.$$

- (4) Then assets evolve according to

$$a_{t+1} = (1 + r) a_t - (1 + R)(1 - \xi_t) c_t - \xi_t c_t. \tag{A.6}$$

A.2. THE HOUSEHOLD'S PROBLEM

Consider a game of length mN where m is an integer. Time runs from $T = 0$ to $T = mN - 1$ and wages are set in periods $0, N, 2N, \dots, (m - 1)N$. The household's value function is

$$V(a_t, w_t) = \max [v(c_t, w_t, \xi_t) + \gamma \delta V(a_{t+1}, w_{t+1})]; 0 \leq t \leq T, \tag{A.7}$$

subject to asset evolution (A.6). v is as defined in (A.2).

Continuation utility is

$$V(a_{t+1}, w_{t+1}) = \sum_{i=1}^{T-t} \delta^{i-1} v(c_{t+i}, \xi_{t+i}, w_{t+i}). \tag{A.8}$$

Note the following:

- (1) Time inconsistency means that actual utility from the perspective of the household at time $t + 1$ will be different from continuation utility, because a different discount factor will be applied to the utility flows. This is why standard dynamic programming techniques cannot be used.
 (2) Wages are set every N periods, so w_t is in general the wage *prevailing* at time t and not necessarily a choice variable in that period.

(3) We do not allow the use of credit in the final period of the game.

The optimum strategy is found by backward induction. The problem is solved first in the final period T , and then in a general period $0 \leq t < T$.

A.3. THE FINAL PERIOD, $t = T$

In the final period we assume that credit goods cannot be used (because the credit would never be repaid), so $\xi_{mN-1} = 0$ and the household simply consumes the value of the remaining assets, which from (A.6) is

$$c_T = \frac{1+r}{1+R} a_T = \lambda_T a_T, \tag{A.9}$$

where $\lambda_T = \frac{1}{1+\pi}$.

A.4. OTHER PERIODS $0 \leq t < T$

The first-order conditions of the maximization problem given by (A.7) with respect to c_t , ξ_t , and w_t are

$$v_c = -\gamma \delta V_a \frac{da_{t+1}}{dc_t}; 0 \leq t < mN - 1, \tag{A.10}$$

$$v_\xi = -\gamma \delta V_a \frac{da_{t+1}}{d\xi_t}; 0 \leq t < mN - 1, \tag{A.11}$$

$$v_w = -\gamma \delta \left(V_a \frac{da_{t+1}}{dw_t} + V_w \right); t = nN, n \in \mathbf{Z}, 0 \leq n < m. \tag{A.12}$$

The first equation sets the marginal utility of consumption equal to its cost in terms of reduced future utility. The second sets the marginal disutility of an increase in the proportion of credit goods (in terms of longer shopping time and hence reduced leisure) equal to the benefit in terms of higher future consumption. The third sets the marginal benefit of a higher wage (in terms of increased leisure in the current period) against the future marginal costs in terms of reduced future consumption ($V_a \frac{da_{t+1}}{dw_t}$) and reduced future leisure (V_w). Recall that because wages are only set every N periods, the first-order condition for the wage is only valid for $t = 0, N, 2N, \dots$

A.5. THE STRATEGY FOR CONSUMPTION

Using the utility function (A.2),

$$v_c = \frac{1}{c_t}. \tag{A.13}$$

Conjecture that the optimal strategy for consumption is of the form

$$c_t = \lambda_t a_t. \tag{A.14}$$

Differentiate the asset evolution equation (A.6) to give

$$\frac{da_{t+1}}{dc_t} = -[1 + R(1 - \xi_t)]. \tag{A.15}$$

Combine these with the first-order condition for consumption (A.10) to give

$$\frac{1}{c_t} = \gamma\delta [1 + R(1 - \xi_t)] V_a. \tag{A.16}$$

Differentiating continuation utility (A.8) and using (A.14),

$$V_a = \sum_{i=1}^{T-t} \delta^{i-1} \frac{\lambda_{t+i}}{c_{t+i}} \frac{da_{t+i}}{da_{t+1}}. \tag{A.17}$$

With the help of (A.14), the asset evolution equation (A.6) can be written

$$a_{t+1} = \mu_t a_t, \tag{A.18}$$

where

$$\mu_t = (1 + r) - [1 + R(1 - \xi_t)] \lambda_t, \tag{A.19}$$

so

$$a_{t+i} = \prod_{j=1}^{i-1} \mu_{t+j} a_{t+1}; i > 1, \tag{A.20}$$

and then

$$\frac{da_{t+i}}{da_{t+1}} = \prod_{j=0}^{i-1} \mu_{t+j} = \frac{a_{t+i}}{a_{t+1}}. \tag{A.21}$$

Substituting this into (A.17),

$$V_a = \sum_{i=1}^{T-t} \delta^{i-1} \frac{\lambda_{t+i}}{c_{t+i}} \frac{a_{t+i}}{a_{t+1}} \tag{A.22}$$

$$= \sum_{i=1}^{T-t} \delta^{i-1} \frac{1}{a_{t+i}} \frac{a_{t+i}}{a_{t+1}} \tag{A.23}$$

$$= \frac{1}{a_{t+1}} \sum_{i=1}^{T-t} \delta^{i-1} \tag{A.24}$$

$$= \frac{1}{\mu_t a_t} \frac{1 - \delta^{T-t-1}}{1 - \delta}, \tag{A.25}$$

where the first step uses (A.14) and the third uses (A.18).

So (A.16) becomes

$$\frac{1}{c_t} = \frac{1 + R(1 - \xi_t)}{\Theta \mu_t a_t}, \tag{A.26}$$

where

$$\Theta = \frac{1 - \delta}{\gamma\delta (1 - \delta^{T-t-1})}.$$

Then using (A.14),

$$\lambda_t = \frac{\Theta \mu_t}{1 + R(1 - \xi_t)}, \tag{A.27}$$

and using (A.19),

$$\lambda_t = \frac{\Theta}{1 + \Theta} \frac{1 + r}{1 + R(1 - \xi_t)}. \tag{A.28}$$

A.6. THE STRATEGY FOR THE PROPORTION OF CREDIT GOODS

Using the utility function (A.2),

$$v_\xi = -F^{-1}(\xi_t) \zeta l_t^{-\eta}. \tag{A.29}$$

Differentiating the asset evolution equation (A.6),

$$\frac{da_{t+1}}{d\xi_t} = R c_t. \tag{A.30}$$

Combining these with (A.16) and the first-order condition (A.11) gives

$$F^{-1}(\xi_t) \zeta l_t^{-\eta} = \frac{R}{1 + R(1 - \xi_t)} \tag{A.31}$$

or

$$\xi_t = F \left\{ \frac{R}{\zeta l_t^{-\eta} [1 + R(1 - \xi_t)]} \right\}, \tag{A.32}$$

which defines implicitly the choice of ξ_t .

A.7. THE STRATEGY FOR THE WAGE

Recall that the wage can only be chosen in periods $t = 0, N, 2N, \dots$. If we are in one of these periods the following analysis applies. Otherwise the wage is given.

Using the utility function (A.2)

$$v_w = \theta \zeta \frac{n_t l_t^{-\eta}}{w_t}. \tag{A.33}$$

Differentiating continuation utility (A.8) gives

$$V_w = \sum_{i=1}^{N-1} \delta^{i-1} \left(\theta \frac{n_{t+i}}{w_t} \zeta l_{t+i}^{-\eta} \right). \tag{A.34}$$

Assets in $t + 1$ depend on the wage chosen in this period via the initial sale of labor income in (A.4), so

$$\frac{da_{t+1}}{dw_t} = (1 + r)^t \frac{da_0}{dw_t} \tag{A.35}$$

$$= (1 - \theta)(1 - \tau) \frac{\varphi}{N} \sum_{i=0}^{N-1} \frac{n_{t+i}}{(1 + \pi)^i}. \tag{A.36}$$

Substituting this, (A.34), and (A.16) into the first-order condition for the wage (A.12) gives

$$\frac{\theta \zeta}{w_t} \left(n_t l_t^{-\eta} + \gamma \delta \sum_{i=1}^{N-1} \delta^{i-1} n_{t+i} l_{t+i}^{-\eta} \right) = -(1-\theta)(1-\tau) \frac{1}{c_t [1+R(1-\xi_t)]} \frac{\varphi}{N} \sum_{i=0}^{N-1} \frac{n_{t+i}}{(1+\pi)^i}, \tag{A.37}$$

or

$$w_t = \frac{\theta \zeta}{\varphi (\theta - 1) (1 - \tau)} \frac{n_t l_t^{-\eta} + \gamma \delta \sum_{i=1}^{N-1} \delta^{i-1} n_{t+i} l_{t+i}^{-\eta}}{\frac{1}{c_t [1+R(1-\xi_t)]} \frac{1}{N} \sum_{i=0}^{N-1} \frac{n_{t+i}}{(1+\pi)^i}}. \tag{A.38}$$

Using the general discount factor ($\beta_{t:t} = 1$, $\beta_{t:t+i} = \gamma \delta^i$ for $i > 0$), the optimal choice of the wage can be written

$$w_t = \omega \frac{\sum_{i=0}^{N-1} \beta_{t:t+i} n_{t+i} l_{t+i}^{-\eta}}{\frac{1}{N c_t [1+R(1-\xi_t)]} \sum_{i=0}^{N-1} \frac{n_{t+i}}{(1+\pi)^i}}, \tag{A.39}$$

where

$$\omega = \frac{\theta \zeta}{\varphi (\theta - 1) (1 - \tau)}. \tag{A.40}$$

A.8. AN INFINITE HORIZON GAME

As $m \rightarrow \infty$ (recall $T = mN$) the expression for λ (A.28) becomes

$$\lambda_t = \frac{\frac{1-\delta}{\gamma \delta}}{1 + \frac{1-\delta}{\gamma \delta}} \frac{1+r}{1+R(1-\xi_t)}. \tag{A.41}$$

In the infinite horizon game, if assets are to neither collapse or explode, and ruling out oscillatory solutions, from (A.18)

$$\mu = 1, \tag{A.42}$$

which implies from (A.19) that

$$1 = (1+r) - [1+R(1-\xi_t)] \lambda_t \tag{A.43}$$

or

$$\lambda_t = \frac{r}{1+R(1-\xi_t)}, \tag{A.44}$$

and equating this with (A.41) gives a relation between the three time preference parameters:

$$\frac{1-\delta}{\gamma \delta} = r. \tag{A.45}$$

A.9. UNIQUENESS

If ξ_t and w_t were exogenous, (A.44) would give a unique value for λ_t and then (A.39) a unique value for w_t . However, in general, (A.32) shows that ξ_t depends on itself both directly and via leisure. Geometrically, the left-hand side of (A.32) is the 45° line, and because F is a CDF so defined at every value on the interval $[0, 1]$, Brouwer's fixed point theorem tells us there will always be at least one solution. However, because there is no closed form expression for consumption or output, nothing general can be said about whether there might be more than one solution. Given this, the only option is to check uniqueness numerically for each steady state (each value of inflation for each calibration) considered in the paper. This process is straightforward:

- (1) Take ξ exogenous over a fine grid and at each point of the grid evaluate the right-hand side of (A.32).
- (2) Check that it is weakly monotonically decreasing. If so, there will be a unique solution for ξ_t and hence a unique equilibrium strategy in the chosen steady state.

In practice, the right-hand side is found to be monotonically decreasing in ξ_t , implying a unique solution, for a far wider range of calibrations than that considered in the paper.