Implications of Habit Formation for Optimal Monetary Policy

Jeffery D. Amato
Bank for International Settlements
Thomas Laubach
Board of Governors of the Federal Reserve System*

First version: November 2000
Revised: May 2001

Abstract

We study the implications for optimal monetary policy of introducing habit formation in consumption into a general equilibrium model with sticky prices. Habit formation affects the model's endogenous dynamics through its effects on both aggregate demand and households' supply of output. We show that the objective of monetary policy consistent with welfare maximization includes output stabilization, as well as inflation and output gap stabilization. We find that the variance of output increases under optimal policy, even though it acquires a higher implicit weight in the welfare function. We also find that a simple interest rate rule nearly achieves the welfare-optimal allocation, regardless of the degree of habit formation. In this rule, the optimal responses to inflation and the lagged interest rate are both declining in the size of the habit, although super-inertial policies remain optimal.
JEL classification: E31, E32, E52.
Keywords: Habit formation, optimal monetary policy, interest rate rules.

---

*Amato: 4002 Basel, Switzerland, jeffery.amato@bis.org; Laubach: Washington, D.C. 20551, tlaubach@frb.gov.
We gratefully acknowledge helpful discussions with Marc Giannoni, Ed Nelson, John Williams, and Michael Woodford. All remaining errors are ours. The views expressed herein are those of the authors and do not necessarily reflect those of the Bank for International Settlements, the Board of Governors of the Federal Reserve System, or the staff of either institution.

1 Introduction

This paper investigates the implications of habit formation for optimal monetary policy. It is motivated by recent studies (Edge (2000), Fuhrer (2000)) that show including habit formation improves the degree to which small-scale business cycle models are able to fit certain aspects of U.S. time series.[1] For example, Fuhrer argues that it helps to explain the gradual response of output to shocks observed in VAR studies. Despite the large number of papers that examine the desirable empirical properties of models with habit formation, only McCallum and Nelson (1999) and Fuhrer (2000) have addressed the subject of monetary policy, and neither of these authors characterizes optimal policy in a model in which agents make optimal choices about both consumption and labor supply. This paper addresses this gap in the growing literature that evaluates optimal monetary policy and simple interest rate rules within the context of small, structural models.

There is considerable empirical evidence that the autocorrelations of detrended output and inflation are large and positive, and die out slowly, in most industrialized countries (Fuhrer and Moore (1995), Coenen and Wieland (2000)). Including habit formation improves the empirical performance of small-scale business cycle models because it introduces endogenous persistence into the structural equations. From the Euler equation of an optimizing household, habit formation implies that the marginal utility of current consumption depends upon both past and expected future consumption. Therefore, the IS equation derived from this Euler equation depends upon both expected future and lagged output. By contrast, the IS equation based on time-separable utility in consumption depends only upon current and forward-looking variables. Similarly, as we will show, habit formation alters the form of the Phillips curve. Current inflation depends upon both past and expected future output gaps, in addition to the current output gap and expected inflation as in the standard new-Keynesian Phillips curve. This occurs because the marginal utility of income of consumers affects the optimal pricing decisions of suppliers. The reduced-form processes for output and inflation thus change for two reasons: first, because inflation depends upon the output gap, whose dynamic properties have been affected by the habit through an altered IS equation; and, second, because the valuation of revenues by suppliers also depends upon the habit in consumption,

[1] Habit formation has also been used to explain various anomalies in the finance literature. For example, see the discussion in chapter 8 of Campbell, Lo and MacKinlay (1997). While habit formation may be useful in explaining various aspects of aggregate data, direct evidence for habit formation based on consumption data is hard to find, see Dynan (2000).

which leads to additional output gap terms appearing in the Phillips equation. Overall, a relatively large habit in consumption can lead to substantial persistence in both output and inflation.

It can be difficult to distinguish, however, whether the observed persistence in the data is derived from endogenous dynamics or exogenous shocks. Yet, optimal monetary policy may differ in important respects depending upon the source of persistence. This is suggested by the existing literature, which reaches different conclusions about the nature of desirable policies in part due to differences in the specification of endogenous dynamics. For example, in studies based on a standard sticky price model with optimizing agents, highly inertial policy is desirable and output terms get a low weight in optimal interest rate rules (e.g. Rotemberg and Woodford (1999)). Neither of these implications obtain from studies that use non-utility based models that include lagged endogenous terms (e.g. Rudebusch and Svensson (1999)). It is therefore of interest to investigate the implications for optimal policy of introducing habit formation – and, by implication, endogenous persistence – into a model for monetary policy analysis, while at the same time maintaining the advantages of working with an optimization-based model.

One of the advantages of working with an optimization-based model is that the representative agent's discounted utility function provides a natural measure to evaluate alternative policies. As in other recent analyses of optimal policy (Rotemberg and Woodford (1997, 1999) and Erceg, Henderson and Levin (2000)), we evaluate alternative policies according to a second-order approximation to agents' utility. In their model with time-separable utility, Rotemberg and Woodford show that welfare can be approximated by a function that depends negatively upon the variances of inflation and the output gap, as well as the steady-state inflation rate. Notably, welfare is reduced only by fluctuations in output that are not caused by movements in the natural rate of output. In contrast, we show that the presence of habit formation leads to an approximation of welfare that does depend upon variability in output itself, as well as the output gap, with the relative weights on these terms changing with the size of the habit.

Regarding optimal policy, one result we obtain is that the variance of output increases dramatically with the size of the habit, even though it acquires a higher implicit weight in the welfare function. One way of interpreting this is through the effect that habit formation has on the behavior of the real interest rate in a hypothetical flexible price equilibrium. In particular, we show that the variance of the Wicksellian natural rate of interest is increasing in the habit. In order to maintain a reasonable degree of variability in interest rates (e.g., due to the zero lower bound on nominal interest rates), part of this greater volatility, *ceteris paribus*, is realized in a higher variance of output. As it turns out, even a higher weight on output in the policymaker's objective is not enough to reduce output variability under optimal policy in the light of more pronounced fluctuations in the natural rate of interest.

Our second set of results on optimal policy concerns the properties of simple interest rate rules. We show that rules restricted to include responses to current inflation and output, and the lagged interest rate, nearly achieve the welfare obtained under the optimal plan. This result confirms the similar conclusion reached by others using a wide variety of models and policymaker objectives (e.g. Taylor (1999)). Furthermore, we show that "super-inertial" policy — as represented by a coefficient greater than one on the lagged interest rate (Woodford (1999b)) — is optimal in the presence of a habit. However, habit formation has the effect of reducing the size of the lagged interest rate response because, given a particular path of persistently high interest rates, the additional inertia in other variables caused by habit formation forces deviations of the output gap and inflation from their steady-states to be larger and more persistent in the face of shocks. By toning-down the threat to keep interest rates high after an inflationary shock, less excessive fluctuations in all of the variables can be achieved.

The rest of the paper is structured as follows. In the next section, we introduce a model with habit formation in which agents make optimal choices about consumption and output supply, and we discuss our baseline calibration of the model. In section 3, we present an approximation to the welfare function that monetary policy aims to maximize, and characterize optimal policy by solving for the optimal plan. In section 4, we explore how closely simple interest rate rules come to mimicking optimal policy. In section 5, we conduct some sensitivity analysis by investigating how alternative calibrations impact upon the results provided in section 3. Some conclusions are offered in the last section. Appendix A provides details on the derivation of the log-linear approximations to the structural equations of the model and the second-order approximation of the welfare function. Appendix B presents the equations characterizing the optimal plan.

2 A Structural Model with Habit Formation

To analyze the consequences for optimal monetary policy of habit formation in consumption, we use a small structural model derived from optimizing behaviour of households and imperfectly competitive suppliers. Except for the feature of habit formation, our model is identical to that of Woodford (1999b). Specifically, we assume that the economy consists of a continuum of households, each of which is the monopolistic supplier of one differentiated product. Because households derive utility from consuming an aggregate of the differentiated products, suppliers face a downward-sloping demand schedule for their product. To keep the model as simple as possible, the economy is assumed closed, and there is no capital accumulation, so that goods market clearing requires that all output is being consumed each period.

In this section, we derive the implications of habit formation in consumption for the relationship between expected future real interest rates and aggregate demand for output as well as for the relationship between output and inflation. We then choose parameter values for the structural parameters and shock processes, and discuss the model's empirical performance.

2.1 Aggregate Demand and Supply

There is a continuum of households uniformly distributed on the unit interval, each of them indexed by the product of which it is the monopolistic supplier. Household i maximizes

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[u(C_t^i, C_{t-1}^i; \xi_t) - v(y_t(i); \xi_t) \right] \tag{1}$$

C_t^i is household i's consumption in period t of the usual Dixit-Stiglitz aggregate of differentiated products, with the elasticity of substitution between products described by θ. We assume that the felicity function u is given by

$$u(C_t, C_{t-1}; \xi_t) = \frac{1}{1-\sigma} \left(\frac{C_t}{C_{t-1}^h} \right)^{1-\sigma} \exp(a'\xi_t), \;\; 0 \leq h \leq 1 \tag{2}$$

which is a special case of the functional forms analyzed by Abel (1990) and Carroll et al. (2000). The parameter h measures the strength of habit persistence. In the limiting case that $h = 1$, the instantaneous utility derived from consumption depends only on the ratio of current to previous period's consumption. The opposite limiting case $h = 0$ is the standard case of time-separable utility. In the latter case, the parameter $\sigma > 0$ measuring the curvature of u corresponds to the inverse of the intertemporal elasticity of substitution. The vector-valued random disturbance ξ in (1) represents taste shifters affecting the utility of consumption $u(\cdot, \cdot; \xi)$ and the disutility of supply $v(\cdot; \xi)$ respectively. The effects of ξ on u and v are not required to be mutually orthogonal.

The household maximizes (1) subject to its budget constraint and the constraint that it satisfies demand for its product at the price $p_t(i)$ that household i charges per unit of its product, i.e.

$$y_t(i) = \left(\frac{p_t(i)}{P_t}\right)^{-\theta} Y_t \tag{3}$$

Here, P_t and Y_t denote the utility-based price index and the aggregate of individual households' output respectively, corresponding to the aggregate C_t. We assume that financial markets are complete, and that households insure themselves against all idiosyncratic risk, such that their wealth is identical *ex ante*. Therefore all households choose the same consumption path, which we denote by $\{C_t\}$. This path is characterized by the first-order condition

$$C_t^{-\sigma} C_{t-1}^{-h(1-\sigma)} e^{a'\xi_t} - \beta h E_t \left[C_{t+1}^{1-\sigma} C_t^{-(1+h(1-\sigma))} e^{a'\xi_{t+1}}\right] - \phi_t = 0 \tag{4}$$

where ϕ_t denotes the marginal utility of consumption in period t. To capture the idea that an increase in past consumption raises the marginal utility of current consumption, it is necessary to assume that $\sigma > 1$, so that the exponent on C_{t-1} is positive. While the first term in (4) measures the effect that past consumption has on the marginal utility of current consumption, the second term reflects the effect of a marginal change in current consumption on expected utility from next period's consumption.

Goods market clearing requires that $C_t = Y_t \; \forall t$. As shown in appendix A, taking a log-linear approximation of the condition (4) leads to an intertemporal IS equation of the form

$$\frac{\eta_2}{\sigma} E_t \Delta y_{t+1} = -\frac{\eta_1}{\sigma}[\Delta y_t + \beta E_t \Delta y_{t+2}] - \beta h E_t \Delta g_{t+2} + E_t \Delta g_{t+1} + \frac{1-\beta h}{\sigma}[r_t - E_t \pi_{t+1}] \tag{5}$$

Here, y_t denotes the percent deviation of aggregate output from its steady-state level, π_t is the growth rate of the aggregate price index, i.e. $\pi_t \equiv \log(P_t/P_{t-1})$, and r_t is the nominal interest rate on a one-period riskless bond.[2] The term g_t represents variation in spending that is not caused by changes in the real interest rate, such as disturbances to the marginal utility of consumption caused by fluctuations in ξ_t. The parameters η_1 and η_2 are defined as $h(1-\sigma)$ and $\sigma - \beta h(1 + h(1-\sigma))$ respectively. Assuming that $\sigma > 1$ and $h \in [0, 1]$ (and hence $\eta_1 < 0$) implies that $\eta_2 > 0$. The presence of habit formation implies a positive correlation between the change in output expected

[2] More specifically, all log-linearizations are taken around a steady state with zero inflation. Hence, π_t is by definition the percent deviation from its steady-state value, while r_t denotes the percent deviation of the interest rate from its steady state value associated with zero inflation.

from the current to the next period, and the change in output in the adjacent two periods. Without habit formation in consumption, $\eta_1 = 0$ and $\eta_2 = \sigma$, and (5) reduces to the standard intertemporal IS equation $E_t\Delta y_{t+1} = E_t\Delta g_{t+1} + \sigma^{-1}[r_t - E_t\pi_{t+1}]$ that characterizes the demand side in the model of Woodford (1999b) and other recent studies of the effects of monetary policy.

Real effects of monetary policy arise in this model from the assumption that within a given period not all suppliers are able to adjust their prices in response to fluctuations in demand. Specifically, we follow Calvo (1983) in assuming that each period a fraction $1 - \alpha$ of suppliers is offered the opportunity to choose a new price, while the remaining suppliers have to maintain whichever price they charged before. Moreover, suppliers are drawn randomly and independent of their own history, in particular independent of the time they were last offered the opportunity to adjust their price.

Since every supplier faces the same demand function (3), all suppliers chosen in period t to adjust their price will choose the same price, which we denote by p_t, and which maximizes

$$E_t\left\{\sum_{j=0}^{\infty}(\alpha\beta)^j\left[\phi_{t+j}\left(\frac{p_t}{P_{t+j}}\right)^{-\theta}Y_{t+j}\frac{p_t}{P_{t+j}} - v\left(\left(\frac{p_t}{P_{t+j}}\right)^{-\theta}Y_{t+j};\xi_{t+j}\right)\right]\right\} \tag{6}$$

The first term in brackets represents the household's utility from consumption in period $t+j$ if it chooses price p_t in the current period. It is the product of its revenue (expressed in consumption units) in period $t+j$ conditional on its price being p_t, and the marginal utility of consumption in period $t+j$. The second term represents the disutility incurred from supplying the amount of its product demanded in period $t+j$ if its price is still p_t. Since the price chosen in period t will still be in effect j periods later with probability α^j, the household discounts the stream of future utilities conditional on its choice of price today by the factor $\alpha\beta$.

The assumption that suppliers are offered the opportunity to reset their price independent of their history leads to a simple law of motion for the aggregate price index. Combining this law of motion with the expression for ϕ_t implied by (4) then leads to an aggregate supply equation (derived in appendix A)

$$\pi_t = \kappa\left[\left(\frac{\eta_2}{1-\beta h} + \omega\right)x_t + \frac{\eta_1}{1-\beta h}(x_{t-1} + \beta E_t x_{t+1})\right] + \beta E_t\pi_{t+1} \tag{7}$$

where $x_t \equiv y_t - y_t^n$ is the deviation of output from the natural rate of output y_t^n, the level of output (measured as percent deviation from steady-state output) that would obtain if prices were flexible.

In Appendix A this is shown to be

$$y_t^n \equiv [\beta\eta_1 L^{-1} + \eta_3 + \eta_1 L]^{-1} [\sigma(g_t - \beta h E_t g_{t+1}) + \omega(1-\beta h) z_t] \tag{8}$$

The natural rate of output is a composite of the shocks g_t and z_t that are measuring the effects of the disturbance ξ_t on the marginal utility of consumption and the marginal disutility from output supply respectively. $L^{-1}x_t$ is here defined as $E_t x_{t+1}$, and $\eta_3 \equiv \eta_2 + \omega(1-\beta h)$. The coefficient κ is a function of the structural parameters β, θ, α, and the parameter ω measuring the elasticity of the disutility function v. Without habit formation, $\eta_1 = 0$ and hence (7) reduces to the New-Keynesian Phillips curve $\pi_t = \kappa(\sigma + \omega)x_t + \beta E_t \pi_{t+1}$. Habit formation causes inflation to depend on the lagged and expected future output gap (beyond the dependence on expected future conditions captured by $E_t \pi_{t+1}$) because expected marginal revenues are valued by the marginal utility of consumption, which does no longer depend on current output only. This effect of habit formation on the aggregate supply side is absent from the model of McCallum and Nelson (1999) who assume inelastic labour supply. Similarly, it is absent from Fuhrer's (2000) study, who models the supply side by using a reduced-form VAR equation.

The definition of the output gap x_t can be used to rewrite the IS equation (5) as

$$\frac{\eta_2}{\sigma} E_t \Delta x_{t+1} = -\frac{\eta_1}{\sigma}[\Delta x_t + \beta E_t \Delta x_{t+2}] + \frac{(1-\beta h)}{\sigma}[r_t - r_t^n - E_t \pi_{t+1}] \tag{9}$$

where

$$r_t^n \equiv \frac{\sigma}{1-\beta h} E_t \left[\frac{\beta\eta_1}{\sigma} \Delta y_{t+2}^n + \frac{\eta_2}{\sigma} \Delta y_{t+1}^n + \frac{\eta_1}{\sigma} \Delta y_t^n + \beta h \Delta g_{t+2} - \Delta g_{t+1} \right] \tag{10}$$

is the Wicksellian natural rate of interest, the real interest rate that would obtain if all prices were flexible, and that would correspond to the equilibrium nominal interest rate in the case of price stability. Insofar as under optimal monetary policy the real interest rate is made to follow the natural rate (10), the effect of habit formation is that the real rate has to respond to higher-order distributed leads and lags of the exogenous shocks than it would have to absent habit formation.

2.2 Calibration of the Model

For the characterization of optimal policy and the analysis of interest rate rules presented in the following sections, the model's structural parameters as well as the shock processes have to be calibrated. The model's structural parameters are $\beta, \kappa, \sigma, \omega$ and h. Because the focus of this article is on exploring the effects of habit formation, below results are reported for various values of h. Specifically, we consider the case of time-separable utility ($h = 0$), the case corresponding to Fuhrer's (2000) estimate of this parameter ($h = 0.8$), and an intermediate case ($h = 0.4$).

The model is calibrated to simulate data at a quarterly frequency, and hence β is set to the conventional value of 0.99. The parameter κ is set to 0.031, a value consistent with an average lifetime of price contracts of three quarters ($\alpha = 0.66$) and an average markup in goods markets of 15% ($\theta = 7.88$, or $\frac{\theta}{\theta-1} = 1.15$). This value of κ is based on Rotemberg and Woodford's (1997) estimate of this parameter.[3] They argue that an average duration of price contracts of three quarters is consistent with survey evidence on firms' price setting, and that a 15% markup is neither implausibly high, nor too low to be consistent with firms engaging in staggered price adjustment.

Estimates of the parameter σ based on aggregate consumption data by Hall (1988) and Attanasio and Weber (1993) are on the order of 3. Rotemberg and Woodford (1997), on the other hand, provide an estimate of 0.16, which they obtain by matching impulse response functions implied by their model to those obtained from a VAR using U.S. data. As they point out, their low estimate of σ is related to the fact that it measures the interest-rate sensitivity of total output (they use GDP data), not just that of nondurable consumption. When $h = 0$, the parameter σ in our model has the same interpretation as in Rotemberg and Woodford, implying that a low value would be of greater relevance. As discussed above, however, a value exceeding 1 is necessary for the effect of past consumption on current marginal utility to be positive, which is the essence of habit formation. We therefore choose σ to be 1.1.

The parameter ω in our model is a combination of the Frisch elasticity of labour supply and the elasticity of output with respect to hours. We set ω to 0.6, which is consistent with a Frisch elasticity of 5 and a Cobb-Douglas production technology with a coefficient on labour of 0.75 (which, together with a steady-state markup of 15%, implies a labour share of about 2/3). Our choice of ω is motivated by the impulse responses to an interest rate innovation displayed in Figure 1. We compute these impulse response functions by combining the structural equations (7) and (9) with an interest rate rule

$$r_t = 0.69\ r_{t-1} + 0.67\ \pi_t + 0.15\ y_t + \epsilon_t \tag{11}$$

This rule is a simplified version of the rule that characterizes U.S. monetary policy according to

[3]Their parameter κ corresponds to $\kappa(\sigma + \omega)$ in the present model.

Rotemberg and Woodford's (1997) estimates.[4] The responses displayed in Figure 1 are those to a 1% innovation in ϵ_t.[5] Given our choice of σ, a value for ω of 0.6 generates approximately the same response of inflation to an interest rate innovation as estimated by Rotemberg and Woodford.[6] As is evident from Figure 1, the impulse response of inflation is almost invariant with respect to changes in h, so that our choice of ω based on this impulse response remains valid across different h.

The definition of the natural rate of interest (10) shows that it is a composite of the natural rate of output y_t^n and the shock g_t. Likewise, the definition (8) implies that y_t^n is itself a composite of the shocks g_t and z_t measuring the effects of ξ_t on the marginal utility of consumption and the marginal disutility from output supply respectively. Both r_t^n and y_t^n are functions of the parameter h, while g_t and z_t are not. We therefore calibrate the latter two shocks, and construct the former two for any given value of h.

We calibrate these shocks using the method developed in Rotemberg and Woodford (1997). Specifically, since their model is nearly identical to ours when $h = 0$, we can construct shock processes such that, given our parameter values for σ and ω, their model together with these shock processes replicates exactly the law of motion of the endogenous variables as estimated by an unrestricted VAR. Throughout, the processes g_t, z_t and y_t^n are measured (like y_t and x_t) as percent of steady-state output, while r_t^n is measured as annual percent (like π_t and r_t). For $\sigma = 1.1$ and $\omega = 0.6$, the standard deviation of r_t^n implied by this procedure is 22.42. This standard deviation is much higher than 3.72, the value obtained under Rotemberg and Woodford's parametrization of σ and ω. Our choice of a higher value of σ in particular implies that considerably larger interest-rate movements are required to offset the effects of shocks g_t and y_t^n of a given magnitude on inflation, raising the variability of the natural rate of interest. We assume that g_t and z_t (and hence y_t^n and r_t^n) are serially uncorrelated. The standard deviations of g_t and z_t implied by Rotemberg and Woodford's model are 5.09 and 14.22 respectively, and the correlation coefficient between the two processes is -0.53. We scale the standard deviations by a factor 0.82 such that, for $h = 0$, the standard deviation of r_t^n is 22.42, equal to Rotemberg and Woodford's estimated value for this

[4]Rotemberg and Woodford's rule includes extra lags of the variables; the coefficients in (11) equal the sums of the coefficients for each variable in their rule.

[5]The responses of y_t are omitted because they are identical to the responses of x_t by construction.

[6]The main difference in Figure 1 from their impulse responses is the initial response of inflation and output, which is restricted to be 0 under their identification strategy. Nonetheless, the paths of responses are quite similar.

Table 1: Model Calibration

Structural Parameters		Shock Processes	
β	0.99/quarter	std dev (g_t)	4.20
κ	0.031	std dev (z_t)	11.75
σ	1.1	corr (g_t, z_t)	-0.53
ω	0.6		
h	0, 0.4, 0.8		

statistic.[7] The calibration of the model is summarized in Table 1.

Figure 1 suggests that changes in h affect primarily the dynamics of the output gap. As h increases, the response of the output gap to a monetary shock is initially more muted, but also more protracted. These two effects — the weaker initial response and the more gradual return of the output gap to zero — cause the impulse response of inflation to be almost unaffected by an increase in h, because a smaller value (in absolute terms) of x_t is offset by a larger value of $E_t x_{t+1}$ in (7).

Figure 2 shows the responses of the four endogenous variables to a one standard deviation innovation to the component of g_t that is orthogonal to z_t.[8] This innovation has the effect of raising the marginal utility of consumption, while leaving the marginal disutility of supply unaffected; its standard deviation is 3.56% of steady-state output. As in the previous figure, the different lines are indexed by h. In the case $h = 0$, the shock initially has an expansionary effect on output and the output gap since the interest rate increases less than r_t^n at first. Despite the increase in the output gap, inflation falls immediately in response to persistently negative future output gaps, which come about in part from smoothing of the interest rate (i.e. the positive dependence of r_t on r_{t-1} in (11)). As h increases, all of these dynamics are amplified. For example, under $h = 0.8$, the initial response of output is higher, even though the interest rate response is considerably larger as well. The large effect on the interest rate is partly the result of the initial impact of the shock on output, movements in which monetary policy responds to contemporaneously. Larger and more persistent

[7] Scaling is necessary because there is not an exact 1-to-1 correspondance between our processes g_t and z_t and their empirical counterparts in Rotemberg and Woodford's model.

[8] This shock is identified by a Cholesky decomposition by ordering z_t ahead of g_t. Along with the monetary shock ϵ_t, this leaves one remaining shock that is correlated with both g_t and z_t.

deviations of y_t below y_t^n causes similar fluctuations in π_t.

It is clear that different sizes of the habit have a much greater effect on these responses than on those to a monetary policy shock. The differences are mainly attributable to the change in the elasticity of x_t with respect to $(r_t - r_t^n)$ and the increase in the variance of r_t^n. While it is true that shocks to the natural rate of interest are dampened by the smaller elasticity, changes in interest rates are also less effective. However, a shock of any given size has a magnified impact on r_t^n, as evidenced by the solid line in Figure 4, and therefore all of the endogenous variables become more volatile.[9] This includes larger movements in the interest rate that partly, but not completely, offset larger movements in r_t^n. To a lesser extent, the appearance of lagged endogenous variables in (5) and (7) has an effect on all of the impulse responses that works in the same direction; the extra forward-looking terms, i.e. $E_t x_{t+2}$ in (9) and $E_t x_{t+1}$ in (7), have little effect on the responses.

A broader perspective of the effects of h on the dynamics of the model's endogenous variables is provided by the correlation functions displayed in Figure 3. For the computation of these correlation functions, we assume that the interest rate innovation ϵ_t in (11) has standard deviation 0.85, which is the standard deviation of the interest rate disturbance to Rotemberg and Woodford's estimated interest rate rule. As shown by the solid lines in the (1,1) and (2,2) panels of Figure 3, under our assumption of serially uncorrelated shocks, the model generates no positive serial correlation in either inflation or the output gap when utility is time-separable. This is a well-documented feature of models in which the structural equations are entirely forward-looking. Increasing h has the effect of inducing some positive serial correlation to the output gap and, quantitatively more important, to inflation. This is consistent with increasing persistence in the output gap being transmitted to inflation. The finding that the most pronounced effects of introducing habit formation are on the serial correlation of inflation points to the importance of modelling the effects of habit formation on the supply side of the model, as discussed in the introduction.

The largest changes in Figure 3 occur in the cross-correlations between inflation and the output gap on the one hand, and the interest rate on the other. A first point to note is that positive innovations to g_t and z_t cause the natural rates of both output and interest to rise. As long as the interest rate rule implies that r_t rises less than 1-for-1 with an increase in the natural rate (as is the case under the rule (11)), the effect of positive innovations to g_t and z_t on balance tends to be

[9] This occurs even though the shock to g_t is the same size regardless of the value of h.

positive on all three variables.[10] An interest rate innovation, by contrast, causes the interest rate to increase and both inflation and the output gap to fall. The positive contemporaneous correlations for $h = 0$ between the interest rate and inflation (0.3) and the interest rate and the output gap (0.5) indicate that the correlation patterns in Figure 3 are mostly driven by innovations to g_t and z_t. This is consistent with the fact that the standard deviation of the interest rate innovation is not even 4% of the standard deviation of the natural rate.

As h increases, the contemporaneous correlation between the interest rate and the output gap falls to 0.3, while that between the interest rate and inflation turns negative (-0.4). The intuition behind these changes in the contemporaneous correlations can be explained by distinguishing several effects of a change in h on equations (7) and (9). First, the effects of any innovations — whether to g_t, z_t, or the interest rate — on the output gap are affected by the increase in the coefficients on Δx_t and $E_t \Delta x_{t+2}$ as h increases. Second, the effect of innovations to either the interest rate or the natural rate on the output gap is weakened to the extent that $\frac{1-\beta h}{\eta_2 - \eta_1}$ decreases for higher values of h, as mentioned above. Moreover, as Figure 4 shows, the variance of r_t^n increases in h. Since the variance of the interest rate innovation is not affected by h, innovations to the natural rate are becoming an increasingly important source of fluctuations as h increases.

As the panel labelled x_t, r_{t-j} shows, the endogenous response of the interest rate to the various innovations has initially a weaker effect on x_t as h increases, but the inertia inherent in the interest rate rule (11) brings forth a series of negative output gaps. To explain the negative correlation between current inflation and interest rates up to lag 6, it is important to note that, although the factor $\frac{\eta_1}{1-\beta h}$ in (7) increases in absolute value, even for $h = 0.8$ the factor multiplying x_t is more than 7 times the size of the former. Because the latter factor is positive, the sequence of expected negative output gaps displayed in the x_t, r_{t-j} panel has the effect of reducing inflation. For $h = 0.8$, this sequence of expected output gaps is sufficiently large to cause the correlation between π_t and r_{t-j}, $0 \leq j \leq 6$ to turn negative.

[10]It was shown in Figure 2 that an orthogonalized innovation to g_t has a negative contemporaneous effect on π_t, but this is more than offset by the effect of the remaining innovation on both g_t and z_t.

3 Optimal Monetary Policy

In this section we present an approximation to the representative household's welfare when there is habit formation in consumption. We then characterize the policy that minimizes welfare losses according to this objective.

3.1 The Welfare Objective for Monetary Policy

The early literature on stabilization policy, such as Taylor (1979), postulates that the objective for monetary policy should be to minimize some convex combination of the variances of inflation and either output or the output gap. Within the context of an optimization-based model similar to the one above, but with time-separable utility, Rotemberg and Woodford (1997) and Woodford (1999a) show that maximization of the representative household's welfare indeed implies minimization of the variances of inflation and the output gap, where the relative weights on inflation and output gap stabilization are determined by the model's structural parameters. One effect of introducing habit formation in consumption is that it changes this objective, in addition to changing the model equations (5) and (7).

To describe how habit formation affects the objective for monetary policy, suppose that monetary policy chooses at some point $t = 0$ a plan that maximizes the representative household's welfare, defined by

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left[u(Y_t, Y_{t-1}; \xi_t) - \int_0^1 v(y_t(i); \xi_t) di \right] \right\} \tag{12}$$

The integral over $v(y_t(i); \xi_t)$ is taken in order to abstract from the effects on an individual household's supply of the particular date at which it last had the opportunity to adjust its price. We evaluate alternative policies under the assumption that subsidies for output are in place such that the steady-state level of output is efficient despite the presence of imperfect competition.[11] Therefore, monetary policy has no inflationary bias. Nevertheless, as pointed out by Woodford (1999b), the optimal response of monetary policy to shocks in a forward-looking model depends on whether it is conducted under commitment or under discretion. We assume throughout that monetary policy is able to act under commitment, and benefit from the associated stabilization gains. Finally, we assume that all state variables in the initial period are at their unconditional expectation of

[11] The efficiency condition is stated in Appendix A.

zero to ensure that the desirability of the chosen plan does not depend on any particular initial conditions at time 0.

In appendix A we show that a second-order Taylor series expansion of (12) around the same steady state that (5) and (7) have been derived can be expressed in the form of a loss function

$$W = \sum_{t=0}^{\infty} \beta^t E_0 L_t \tag{13}$$

where the period-loss function L_t is defined as

$$L_t = \pi_t^2 + \lambda_x \left[\eta_3 x_t^2 + \eta_1(x_{t-1}^2 + \beta x_{t+1}^2) + \eta_1(y_{t-1}\Delta y_t - \beta y_{t+1}\Delta y_{t+1})\right] \tag{14}$$

and the weight λ_x is defined as $\frac{\kappa}{\theta(1-\beta h)}$.[12] After some further simplifications (described in appendix A) and taking the unconditional expectation of (13) to abstract from initial conditions (as discussed above), the loss function becomes

$$\hat{W} = V[\pi_t] + \lambda_x\{(\eta_3 + (1+\beta)\eta_1)V[x_t] + \eta_1 \left(V[y_t + y_{t-1}] - 2(1+\beta)V[y_t]\right)\} \tag{15}$$

where the measure of variability for any variable z that is used here is defined by

$$V[z] \equiv E\left[(1-\beta)\sum_{t=0}^{\infty} \beta^t E_0 z_t^2\right] \tag{16}$$

Except for the discounting, this measure corresponds to the unconditional variance of z_t.

Habit formation in consumption has three distinct effects compared to the case of time-separable utility, in which the period-loss function is given by $\pi_t^2 + \frac{\kappa(\sigma+\omega)}{\theta} x_t^2$. With habit formation, variability in output as well as the output gap is welfare reducing — while $\eta_1 < 0$, the product $\eta_1(V[y_t + y_{t-1}] - 2(1+\beta)V[y_t])$ is positive for $h > 0$. In our model, in which consumption equals output, output variability reduces welfare because of the link from past output to the current marginal utility of consumption. Second, the current period-loss function (14) depends on past and expected future, as well as current, output and output gaps. As shown in (15), this has the effect of altering the weights on output and output gap variability relative to inflation variability. Finally, the weight λ_x is increasing in h. When inflation is measured as annual percent, while y_t and x_t are measured as percent of steady-state output, the parameter values presented in Table 1 imply λ_x is 0.063 for $h = 0$. As h increases to 0.8, λ_x increases to 0.30. The combined effect is that the weight

[12] This formulation of the loss function, while useful for gaining intuition, omits one initial product $y_0 y_{-1}$. The formulation in (15), on which the computations are based, does not.

$\lambda_x(\eta_3 + (1+\beta)\eta_1)$ on output gap stabilization decreases from 0.108 to 0.102 as h rises from 0 to 0.8, while $\lambda_x \eta_1$ increases in absolute value from 0 to 0.024.

Apart from a suboptimally low steady-state level of output due to imperfect competition, the only departure of the model presented above from the efficient allocation is caused by distortions in relative prices due to price stickiness. This distortion can be eliminated by a policy that implements price stability, in which the real interest rate equals the natural rate of interest in every period. As Rotemberg and Woodford (1997) point out, however, for a sufficiently variable process $\{r_t^n\}$, such a policy requires a positive steady-state rate of inflation to avoid the zero lower bound on the nominal interest rate. The second-order approximation of welfare given by (15) assumes that both π_t and x_t have steady-state values of 0. To account for the possibility that steady-state inflation $\bar{\pi}$ may be non-zero under the optimal policy, (15) is augmented by a term measuring the consequences of interest rate variability under a given interest rate policy for steady-state inflation.[13] We follow Rotemberg and Woodford (1997) and penalize interest rate variability by assuming that, for any level of interest rate variability, all realizations of the interest rate will be distributed within the interval $[E[r] - kV[r]^{1/2}, E[r] + kV[r]^{1/2}]$. Hence, only interest rate policies such that $kV[r]^{1/2}$ exceeds the steady-state net real interest rate $\beta^{-1} - 1$ cause positive steady-state inflation rates. Thus, the criterion for evaluating alternative policies is based on a modification of (15) given by

$$\hat{W}_1 = \hat{W} + \left(\max\{kV[r]^{1/2} - (\beta^{-1} - 1), 0\}\right)^2 \tag{17}$$

In the simulations below, we use $k = 2.26$ and a steady-state net real interest rate of 3% per annum, which are based on the VAR estimated in Rotemberg and Woodford. This implies that a tradeoff between perfect stabilization and zero steady-state inflation exists whenever the standard deviation of the natural rate of interest exceeds 1.32.

Under the calibration of the shock processes discussed above (like under the estimated shock processes of Rotemberg and Woodford (1997)), the zero lower bound is binding because $V[r^n]^{1/2}$ is found to exceed 1.32 for any value of h. In this case the modification of the loss function (15) can alternatively be achieved by adding a term $\lambda_r r_t^2$ to the period-loss function (14), as discussed in Woodford (1999b). The loss function (15) is then modified as

$$\hat{W}_2 = \hat{W} + \lambda_r V[r] \tag{18}$$

[13]In addition, the loss function contains a similar term related to the steady-state level of the output gap $\bar{x}$ equal to $\lambda_x(\eta_3 + \eta_1(1+\beta))\bar{x}^2$. By equation (7), $\bar{x} = [\kappa(\eta_3 + \eta_1(1+\beta))]^{-1}(1-\beta)(1-\beta h)\bar{\pi}$.

Table 2: Optimal Plan, Base Case Calibration

	$h = 0$	$h = 0.4$	$h = 0.8$
λ_r	0.15	0.18	0.28
std(r^n)	22.42	25.77	36.29
$V[r]$	1.86	1.88	1.97
$V[\pi]$	0.50	0.66	1.46
$V[x]$	17.49	19.84	22.38
$V[y]$	13.47	32.41	146
$\bar{\pi}$	0.091	0.10	0.18
$\hat{W}_1$	2.39	3.03	10.26

The weight λ_r is chosen such that the policy that minimizes the loss function (18) is the same that minimizes (17), although the actual values of the two loss functions may be different.

3.2 The Optimal Plan

We begin to evaluate the consequences of habit formation for optimal monetary policy by computing, for different values of the parameter h, the optimal plan, i.e. the stochastic processes $\{\pi_t, y_t, r_t\}$ which minimize the loss function (17) subject to the constraint that the model's structural equations (7) and (9) hold at any date $t \geq 0$. The problem, and the first-order conditions characterizing the optimal plan, are formally stated in appendix B. The first-order conditions are taken using the loss function (14), augmented by a term $\lambda_r r_t^2$, and numerical optimization is performed to determine the value of λ_r that leads to the lowest welfare loss according to the original loss function (17). For each of the three values of h, we characterize the optimal plan by the variability statistics for the endogenous variables and the implied steady-state level of inflation that jointly determine the welfare loss $\hat{W}_1$. These statistics, together with λ_r and the standard deviation of r_t^n, are reported in Table 2.

In the standard case with time-separable utility, the standard deviation of r_t^n is 22.42, the value discussed in the previous section. By contrast, the variability of the interest rate, $V[r]$, is not even 1/2% of the variance of the natural rate, implying that the stabilization gains from matching the natural rate more closely are outweighed by the losses from increased steady-state inflation. Steady-state inflation is non-zero, though still very low (0.1% at annual rate), and therefore the marginal effect of reducing $V[r]$ would be to reduce the welfare loss $\hat{W}_2$ by a fraction 0.15 of the reduction in $V[r]$, other things equal. Despite the fact that stabilization is not nearly complete, $V[\pi]$ is quite low at 0.5% annual rate. The consequences of not matching r_t^n perfectly are most visible in the variability measures for output and the output gap.

As h increases, so does the standard deviation of the natural rate. There are two separate effects of h on r_t^n working in opposite directions. As the dashed line in Figure 4 shows, the variance of the expectational term in (10), i.e. of $\frac{1-\beta h}{\sigma} r_t^n$, is decreasing in h, because the polynomial in y_t^n inside the brackets is becoming more similar to a moving average.[14] The dominant effect, however, is that the factor in front of the expectations operator increases in h. This rise in the standard deviation of r_t^n causes all the statistics in Table 2 to rise as well. It therefore amounts to an outward shift of the "optimal policy frontier", the locus of feasible minimum variances of all of the endogenous variables.

The second effect of increasing h is to change the relative weights on inflation, output, and output gap variability in the objective $\hat{W}_1$. Despite the fivefold increase in λ_x as h rises from 0 to 0.8, the weight on $V[x]$ falls slightly, from 0.108 to 0.103. By contrast, the combined contribution of the terms related to output variability increases from 0 to a fraction 0.044 of $V[y]$.

However, while the relative concern for output fluctuations increases, the rise in output variability is the most pronounced among the endogenous variables because interest rate variability increases only marginally, compared to the precipitous increase in the variability of the natural rate. The small change in interest rate variability reflects the substantial increase in λ_r as h, and with it the variability of the natural rate, increases. Hence, the increase in $V[r]$ is not nearly sufficient to fully offset the effects of a more variable natural rate on the variability of inflation, output, and the output gap.

Figure 5 shows impulse responses of the four endogenous variables to a shock that raises the marginal utility of consumption, while leaving the marginal disutility of supply unaffected. (It is the same shock analysed in Figure 2.) The solid lines in the four panels show that for $h = 0$ the initial effect of the innovation is positive on all four variables. Under the optimal plan, the interest

[14] In our case, in which g_t and z_t are serially uncorrelated, $\frac{1-\beta h}{\sigma} r_t^n$ takes the form $g_t - \sigma^{-1}\left[(\eta_2 - \eta_1) y_t^n + \eta_1 y_{t-1}^n\right]$. Notice that the sum of the weights on the y^n-terms is η_2, which is a decreasing function of h under our parameterization. Thus, not only is a moving average of the y^n-terms taken for h greater than zero, a scaling reduction is also applied.

rate returns to zero only very gradually. Expectations of persistently high interest rates have the effect of reducing both inflation and the output gap below zero in the period following the shock, from where both variables return to their steady-state levels. The concern for minimizing inflation variability is evident by the deflationary periods following the initial inflation. By reducing inflation below zero for some periods, monetary policy succeeds in reducing the discounted sum of squared deviations from the steady-state.

As the dashed and the dashed-dotted lines in Figure 5 show, the effect of increasing h is to magnify the responses of all four variables. The initial response of all variables is again positive. Inflation falls for the next two periods, the output gap for the next three, and output for the next four, while the interest rate peaks only in the quarter following the innovation. The gradual declines in the impulse responses of x_t and y_t are partly driven by the increasing importance of lagged endogenous variables in the structural equations, as well as the interaction of a larger impact of the shock on r_t^n with a reduced elasticity of the output gap with respect to the interest rate gap. These are the factors that were evident in Figure 2 (see section 2.3). However, notice that for larger values of h, both output and the output gap are allowed to rise substantially at first. This occurs because further dampening of the impact of the shock on these two variables would involve a more than offsetting rise in welfare losses due to the higher interest rates necessary, i.e. higher interest rate variability and steady-state inflation.

Comparing Figure 5 with Figure 2 helps clarify the nature of optimal policy relative to alternatives, such as our characterization of historical policy in (11). Optimal policy has a strong influence on how the endogenous variables respond to this type of shock. This can be explained by two differences of optimal policy versus (11). First, optimal policy entails much smaller responses to output, which means that the interest rate hardly increases in response to the substantial rise in output following this type of shock. Second, optimal policy is much more inertial. Thus, the rise in the interest rate is smaller, but it is more drawn out. The combination of these two factors allows output to increase more initially (i.e. policy is more expansionary to begin with), but subsequent fluctuations around the steady state are greatly reduced because large negative values of the output gap are avoided by a smaller rise in the interest rate in the first few periods. These effects become stronger as h increases. The consequences for inflation are that it actually increases initially, and then it falls below the steady-state in a less pronounced way. One conclusion from this analysis is that a rule such as (11) does not respond optimally to a shock of this type.

Table 3: Simple Rules, Base Case Calibration

	$h = 0$	$h = 0.4$	$h = 0.8$
a	1.72	1.66	1.32
b	2.37	1.97	0.76
c	0	-0.001	0.007
$V[r]$	1.86	1.87	1.96
$V[\pi]$	0.54	0.70	1.58
$V[x]$	17.47	19.86	22.33
$V[y]$	13.35	32.62	146
$\bar{\pi}$	0.090	0.10	0.18
$\hat{W}_1$	2.42	3.07	10.34

4 Simple Interest Rate Rules

The results presented in the previous section characterize optimal monetary policy, yet they do not provide any instruction of how such a policy would be implemented. In this section we analyze feedback rules for the interest rate as a mechanism to implement alternative policies. In principle, the first-order conditions presented in appendix B can be solved for the feedback rule that implements the optimal plan. However, such a feedback rule would be quite complicated, as it would depend on the entire history of inflation, output and the interest rate, and involve several different viewpoints of expectations. Moreover, it would involve feedback from the output gap. Orphanides (1998) argues that rules that respond to the output gap may perform poorly because of the problems associated with accurate measurement of the natural rate of output. Instead, we consider the performance of an interest rate rule of the form

$$r_t = ar_{t-1} + b\pi_t + cy_t \tag{19}$$

This rule is of the form proposed by Taylor (1993), augmented to allow feedback from the lagged interest rate.[15] It does not require the central bank to observe any of the model's shock processes. The response to the lagged interest rate generates history dependence of policy, which is an important feature of optimal policy in forward-looking models, as Woodford (1999b) emphasizes.

Table 3 reports the coefficients a, b and c that minimize the loss function $\hat{W}_1$ in (17), along with

[15]Notice that this rule is of the same form as (11), which will help to highlight the differences between an optimized rule and our description of historical policy.

the variability measures for the endogenous variables, steady-state inflation, and the welfare loss for the same three values of h as in the previous section. Comparing the entries in the last six rows of Table 3 to those in Table 2 reveals how remarkably close the simple rules come to replicating the optimal plan. The measure $V[r]$, and hence steady-state inflation, are almost identical under the simple rules compared to their counterparts in Table 2. The same holds for output and output gap variability. The only marked difference in the statistics occurs for $V[\pi]$, the values of which under the simple rules exceed those obtained under the optimal plan by up to 8% of the latter. The higher values for $V[\pi]$ explain why the welfare losses under the simple rules are about 1% higher than those under the optimal plan. Variation in h causes each of the six statistics to change exactly like under the optimal plan, and the interpretation presented in the previous section applies.

The interest rate rules that achieve this close approximation to the optimal plan are characterized by coefficients a larger than one. Such "super-inertial" behaviour of interest rates is consistent with a stationary equilibrium of the system of equations because the private sector's anticipation of this behaviour leads to responses of output and inflation that prevent interest rates from following an explosive path. In the standard case of time-separable utility, Rotemberg and Woodford (1997) and Woodford (1999b) discuss the optimality of super-inertial interest-rate rules. As shown in Table 3, the response to the lagged interest rate weakens as h increases, because the higher inertia in both output and inflation under habit formation implies more drawn-out responses of these variables to the interest rate. A higher degree of habit formation therefore acts to some extent as a substitute for interest rate inertia. Nevertheless, the coefficient a is well above 1 even when $h = 0.8$.

The increasing inertia of inflation and output for higher values of h has a stronger effect on the response to inflation. For $h = 0.8$, the coefficient b is only one third of its size in the time-separable case. The response to output is weak regardless of the value of h, despite the fact that the weight on output stabilization rises as h increases from 0 to 0.8 (see above). While the specific values for the coefficients a, b, and c are likely to be sensitive to the structure of the model as well as its calibration, our findings suggest that an optimal value of a exceeding 1 is robust with respect to the introduction of habit formation in consumption.

Finally, comparing the optimal rules to the rule (11) illustrates the differences between optimal and historical policy discussed in the previous section. The response to the lagged interest rate is much larger in the optimal rules across all values of h. Even in the case $h = 0.8$, this coefficient is nearly twice as large. Moreover, the response to output in (11) is at least 20 times the size of the response in any of the optimal rules. These two properties explain the differences in the responses to shocks presented in Figures 2 and 5.

5 Sensitivity Analysis

In view of the uncertainty surrounding the appropriate values for σ and ω, in this section we assess the effects on our results from choosing alternative values for these two parameters. As mentioned before, Hall (1988) and Attanasio and Weber (1993) arrive at estimates of σ considerably higher than 1.1. The interpretation of σ in the present study as the interest-rate sensitivity of output, instead of nondurable consumption, implies that a value lower than found in those studies is appropriate. Nevertheless, a value between 1.1, and the preferred estimates of Hall and Attanasio and Weber of around 3, might be interesting. We therefore consider in this section a value $\sigma = 2$.

Similarly, while estimates of the elasticity of labour supply reported by Mulligan (1998) are mostly between 0.5 and 2, those surveyed in Pencavel (1986) do not exceed 1. This suggests exploring the consequences of an elasticity of labour supply much lower than the value 5 implied by our calibration of $\omega = 0.6$. Combined with a coefficient on labour in a Cobb-Douglas production function of 0.75, a labour supply elasticity of 0.8 implies $\omega = 2$. We therefore choose ω to be 2. Setting both σ and ω to 2 affects the values of κ and λ_x.[16] The former is now 0.01, while $\lambda_x = 0.022$ for $h = 0$, and $\lambda_x = 0.10$ for $h = 0.8$.

The calibration of the shock processes is again based on the properties of the processes g_t, z_t, and r_t^n implied by Rotemberg and Woodford's model under our alternative parameter values, as described in section 2.2. For the sake of comparison to our earlier results, we first scale the standard deviations of g_t and z_t such that the standard deviation of r_t^n remains at 22.42, given the correlation between g_t and z_t of -0.53. This calibration is labelled "Case 1" in Table 4. Alternatively, the standard deviation of r_t^n implied by Rotemberg and Woodford's model when $\sigma = \omega = 2$ is 40.54, while the standard deviations of g_t and z_t are 3.45 and 8.31, and the correlation between the two is -0.13. We again scale the standard deviations by a factor 1.08 such that, for $h = 0$, the standard deviation of r_t^n is 40.54. This calibration is labelled "Case 2" in Table 4.

The two cases of shock processes presented in Table 4 provide two directions in which to inves-

[16]This is seen from the form for κ given in Appendix A, and recalling that λ_x is a function of κ (section 3.1). We keep α and θ constant at their previous values of 0.66 and 7.88, respectively.

Table 4: Alternative Model Calibration

Structural Parameters		Shock Processes		
β	0.99/quarter		Case 1	Case 2
κ	0.010	std dev (g_t)	1.63	3.72
σ	2	std dev (z_t)	4.56	8.96
ω	2	corr (g_t, z_t)	-0.53	-0.13
h	0, 0.4, 0.8			

tigate the sensitivity of our results. The first is the changing of σ and ω to 2, while keeping the variance of r_t^n constant (for $h = 0$). This corresponds to comparing the Case 1 alternative to our base case calibration. The appeal of this exercise is as follows. When $h = 0$, the model reduces to the standard case. Moreover, none of the elements of the welfare function involve y_t itself, so the only barrier to complete stabilization is the zero lower bound on nominal interest rates. Thus, we can focus directly on fluctuations in r_t^n as the source of welfare reductions. Keeping the variance of r_t^n constant, for $h = 0$, allows us to isolate the effects of changing the structural parameters in question. Of course, the variance of r_t^n may not be equal across parameterizations of σ and ω as h changes, but this is precisely one of the effects of habit formation we wish to explore. The second exercise we consider is changing the variances of the processes g_t and z_t, and hence r_t^n, while holding the structural parameters constant. This involves comparing our two alternative calibrations, Cases 1 and 2, against each other. We investigate these changes in turn, respectively.

5.1 Implications of Changing σ and ω

Statistics characterizing the equilibrium under the optimal plan for our alternative Case 1 calibration are presented in Table 5. The first observation is that all of the statistics are lower under Case 1. This is best explained by examining Figure 6, which plots the components of the variance of r_t^n, analogous to Figure 4. While the values of the standard deviation of r_t^n are identical in the two graphs for $h = 0$, the portion of this statistic corresponding to the expectation of the shocks (see (10)) is much lower for $\sigma = \omega = 2$ (as shown by the dashed line). The larger value of σ reduces the size of the elasticity of output with respect to the one-period expected real interest rate. The scaling factor in r_t^n (which is the inverse of this elasticity) thus becomes relatively large, meaning that the variance of the expectational component must be smaller to achieve the same variance for

Table 5: Optimal Plan, Alternative Calibrations

	Case 1			Case 2		
	$h = 0$	$h = 0.4$	$h = 0.8$	$h = 0$	$h = 0.4$	$h = 0.8$
λ_r	0.040	0.043	0.051	0.093	0.10	0.13
std(r^n)	22.42	28.38	36.86	40.54	51.55	67.91
$V[r]$	1.78	1.78	1.79	1.82	1.82	1.85
$V[\pi]$	0.10	0.12	0.18	0.40	0.48	0.74
$V[x]$	5.61	4.35	1.76	20.30	15.72	6.85
$V[y]$	2.03	3.05	5.77	12.64	18.62	31.92
$\overline{\pi}$	0.024	0.023	0.030	0.055	0.055	0.086
$\hat{W}_1$	0.59	0.53	0.84	2.16	2.12	4.27

r_t^n as in our base case. Thus, the overall impact of shocks in the Case 1 calibration is smaller. At the same time, the potency of policy is reduced by a larger σ. On balance, the variability of all the endogenous variables, and hence welfare losses, are reduced under the alternative calibration.

Figure 7 shows the impulse responses of the endogenous variables to a marginal utility of consumption shock under the alternative calibrations. To focus on the effects that the alternative calibrations have, Figure 7 presents the impulse responses for $h = 0$. The impulse responses under our base case calibration are identical to the solid lines in Figure 5, and are reproduced here for convenience. The size of the shock to r_t^n is identical in both the base case (solid) and Case 1 (dashed). Although the initial response of the interest rate is slightly smaller in the latter case, the output gap does not expand as much in the period of the shock.[17] Moreover, the output gap returns and stays closer to the steady-state, which results in lower inflation variability. These results are driven by slightly higher interest rate inertia under the optimal policy in Case 1, which succeeds in reducing fluctuations in inflation and output.

In the alternative Case 1 calibration, as in our base case, the variance of r_t^n increases with h, while the variance of the term $\frac{1-\beta h}{\sigma} r_t^n$ decreases. Similarly, as h increases, the overall impact of shocks is lessened, and so is the potency of policy through a lower elasticity in the IS equation and a smaller κ in the AS equation. The decline in the variance of the output gap as h increases is the most noticeable difference between Case 1 and the base case. The reason for such different performance

[17] Nonetheless, the output gap is positive in both instances since policy does not respond 1-to-1 to the shock in r_t^n in either case.

is that the trade-off between output gap and interest rate variability improves substantially in favor of the output gap as σ is raised from 1.1 to 2. Large reductions in $V[x]$ are obtained with virtually no increase in $V[r]$. Moreover, the shadow value on interest rate variability hardly increases in Case 1, from 0.04 for $h = 0$ to 0.05 for $h = 0.8$, compared to the increase in λ_r from 0.15 to 0.28 in the base case. Similarly, the increases in $V[y]$ are relatively small compared to Table 2. Finally, one interesting feature of the outcomes under Case 1 is the non-monotonicity in the level of welfare achieved for different values of h. In particular, $\hat{W}_1$ is smaller when $h = 0.4$ than when h is 0 or 0.8. Apparently, the improvement in the trade-off between interest rate and output gap variability as h increases more than offsets the welfare-reducing effects of larger variation in r_t^n.

Comparison of the solid and dashed lines in Figure 8 to their counterparts in Figure 7 shows that the differences between the impulse responses due to the alternative calibrations are becoming more pronounced as h increases. In particular, the initial response of the interest rate shown in Figure 8 is much lower in Case 1 (the dashed line), and it reaches its peak after 3 quarters, compared to 1 quarter in the base case (the solid line). The response of the output gap and inflation are relatively more muted compared to the base case; however, the overshooting of the output gap to below its steady-state is delayed, and a similar course is followed by inflation, which does not occur to the same extent in the base case.

One notable effect of the combination of larger σ and h is that the optimal responses of the variables are more hump-shaped, something that does not occur when $h = 0$ for $\sigma = \omega = 2$. The same behavior is reflected in the impulse responses in the base case to a much lesser extent. This is a result of having at the same time smaller elasticities in the IS and AS equations and a muted transmission of interest rate changes due to the appearance of extra output gap terms in the structural equations.

5.2 Implications of Changing g_t, z_t, and $var(r_t^n)$

In our second robustness check, we change the variances of the shock processes g_t and z_t, and hence, r_t^n and y_t^n, while keeping the structural parameters at the same values. A comparison between our alternative Cases 1 and 2 is relevant here. The statistics describing the equilibrium under the optimal plan for our Case 2 calibration are also presented in Table 5. As can be seen in the second row of the table and Figure 9, the variance of r_t^n is almost double the value in Case 1 across the range of h. With the structural parameters being identical across these two cases, this naturally leads to higher variances for the endogenous variables, as well as higher welfare losses.

Since the variance of the shock g_t is larger, the standard deviation of the component of r_t^n that is orthogonal to z_t is 2.67 times larger than in Case 1. However, as shown in Figure 7, for $h = 0$ the increase in the initial response of the interest rate in Case 2 (the dashed-dotted line) compared to Case 1 (the dashed line) is much lower than this factor. Instead, the greater magnitude of the shock to r_t^n is reflected by interest rates remaining higher for a prolonged period. Nevertheless, the response of the interest rate does not prevent the other endogenous variables to increase substantially at time 0 as well.

As shown by the impulse responses in Figure 8, the initial response of the interest rate is relatively small for $h = 0.8$ as well, compared to the increase in the standard deviation of the shock in Case 2. Relatively larger upward movements in the interest rate are now required in later quarters compared to when $h = 0$. The magnitude of this dynamic response is due in large part to the higher variance of r_t^n for the higher value of h. However, the relatively persistent larger movements under Case 2 compared to Case 1 are not due to changes in the variance of r_t^n alone, but also the fact that the variance of y_t^n is larger in Case 2 and variation in output, as well as the output gap, now matters for welfare.[18] As a consequence of the delayed, yet persistent rise in the interest rate, the output gap remains positive longer and takes longer to converge back to the steady-state compared to when $h = 0$, with similar consequences for inflation.

As can be seen in Table 5, one major difference between Cases 1 and 2 is the percentage increase in the shadow value λ_r. Since the structural parameters as well as the weights on the other terms in the welfare objective are the same under the two cases, the change in λ_r means that the indifference curves implied by the welfare objective are shifting in favour of interest rate stabilization. Nevertheless, the variance of the interest rate increases under the optimal plan.[19]

[18] In fact, the variance of r_t^n increases by a greater percentage under Case 1 in going from $h = 0$ to $h = 0.8$.

[19] Note that, despite the overall increase in variability in Case 2, the result described in the previous sub-section about the variance of the output gap declining with h continues to hold since this is entirely driven by the larger values for σ and ω compared to the base case.

6 Conclusions

In this paper, we developed a model for the analysis of optimal monetary policy that includes households making optimal choices over consumption and output supply, and features habit formation in consumption. The log-linearized version of our model consists of generalizations of the standard forward-looking IS curve and the New-Keynesian Phillips curve, to include extra terms involving past output and expected future output. The presence of habit formation introduces endogenous persistence into the model, and under our base case calibration, we showed that a large habit in consumption can lead to a substantial increase in the serial correlation of inflation.

We then used the model to analyze the implications of habit formation for optimal policy. We found large increases in the variance of output under optimal policy, despite the fact it acquires a higher implicit weight in the welfare function. Similarly, interest rate variability also increases even with a higher weight in the objective, although much less than in the case of output. We obtained these results because the variance of the natural rate of interest increases precipitously with the size of the habit. We also investigated the properties of a simple interest rate rule, and found that it performs nearly as well as the optimal plan. In particular, we found that rules with a coefficient greater than one on the lagged interest rate are preferred across the range of values of the habit that we consider, and that the optimal value of this coefficient, as well as the coefficient on the current inflation term, decline as the habit increases. Our characterization of optimal policy in the presence of habit formation is thus supportive of two findings that have emerged from the recent literature, namely, Woodford's (1999b) result that super-inertial policy is desirable and the approximately optimal nature of simple interst rate rules for conducting monetary policy. Our specific results show exactly how the character of policy should change depending upon the degree of habit formation in consumption.

A Log-linear Approximations

A.1 Derivation of Equations (5) and (7)

In deriving (5) and (7) we take first-order Taylor-series expansions of the exact non-linear equilibrium conditions. If $\|\xi\|$ is a measure of the magnitude of fluctuations of the process $\{\xi_t\}$, the expansions in this subsection are accurate up to a remainder of size $O(\|\xi\|^2)$, which we omit from the equations. For ease of exposition, we assume that the steady-state level of output is at its efficient level Y^*, characterized in our model by

$$(1-\beta h){Y^*}^{-\nu} = v_y(Y^*;0) \tag{20}$$

where $\nu \equiv \sigma + h(1-\sigma)$, and all Taylor series expansions are taken around Y^*. For a discussion of the case in which the steady-state output level is inefficient, see Woodford (1999a). Moreover, all expansions are developed around a zero inflation steady state, and we assume that the steady state value of ξ_t is zero as well.

The Lagrangian of household i at time 0 is

$$E_0\left[\sum_{t=0}^{\infty}\beta_t\left\{\frac{1}{1-\sigma}\left(\frac{C_t}{C_{t-1}^h}\right)^{1-\sigma}e^{a'\xi_t} - v(y_t(i);\xi_t) - \phi_t\left[\frac{B_t}{P_t} - \frac{R_{t-1}B_{t-1}}{P_t} + C_t - \frac{p_t(i)}{P_t}y_t(i)\right]\right\}\right] \tag{21}$$

where B_t denotes the amount of nominal riskless one-period bonds bought by the household at the end of period t that pay a gross return of R_t units of the numeraire in every state in period $t+1$. The first-order conditions with respect to C_t and B_t are given by (4) and

$$\frac{\phi_t}{P_t} = \beta R_t E_t\left[\frac{\phi_{t+1}}{P_{t+1}}\right] \tag{22}$$

The first-order condition (4) implies that $\bar{\phi} = (1-\beta h){Y^*}^{-\nu}$. A log-linear approximation of (4) can then be written as

$$\frac{1}{1-\beta h}\left[-\sigma c_t - h(1-\sigma)c_{t-1} + \sigma g_t - \beta h E_t[(1-\sigma)c_{t+1} - (1+h(1-\sigma))c_t + \sigma g_{t+1}]\right] = \hat{\phi}_t \tag{23}$$

where

$$c_t \equiv \log\left(\frac{C_t}{Y^*}\right),\ \hat{\phi}_t \equiv \log\left(\frac{\phi_t}{\bar{\phi}}\right),\ g_t \equiv \sigma^{-1}a'\xi_t$$

(22) can be approximated as

$$\hat{\phi}_t = E_t\hat{\phi}_{t+1} + r_t - E_t\pi_{t+1} \tag{24}$$

where $r_t \equiv \log(\beta R_t)$. Combining (23) and (24) and substituting $y_t \equiv \log(Y_t/Y^*)$ for c_t yields (5).

The first step in deriving (7) is to take the partial derivative of (6) with respect to p_t and approximate it around the same steady state as before:

$$E_t \sum_{j=0}^{\infty} (\alpha\beta)^j \left\{ \hat{\phi}_{t+j} + \hat{p}_t - \sum_{k=1}^{j} \pi_{t+k} - \omega[y_{t+j} - z_{t+j} - \theta(\hat{p}_t - \sum_{k=1}^{j} \pi_{t+k})] \right\} = 0 \tag{25}$$

where

$$\hat{p}_t \equiv \log\left(\frac{p_t}{P_t}\right), \ \omega \equiv \frac{v_{yy}(Y^*;0)Y^*}{v_y(Y^*;0)}, \text{ and } z_t \equiv -\frac{v_{y\xi}(Y^*;0)}{v_{yy}(Y^*;0)Y^*}\xi_t$$

In the special case in which prices are flexible, $\alpha = 0$, and since all firms choose the same price, $\hat{p}_t = 0$ as well. Therefore (25) simplifies to the condition that

$$\hat{\phi}_t = \omega(y_t - z_t) \tag{26}$$

By combining again (23) and (24), substituting y_t for c_t, then substituting for $\hat{\phi}_t$ in (26) and solving for y_t we obtain the "natural rate of output" defined in (8), i.e. the level of output that would obtain if prices were flexible. In the general case in which some firms cannot adjust their price, i.e. $\alpha > 0$, substituting from (23) for $\hat{\phi}_{t+j}$ in (25) and solving for $\hat{p}_t$ yields

$$\hat{p}_t = (1-\alpha\beta)E_t \sum_{j=0}^{\infty} (\alpha\beta)^j \left[\frac{1}{(1-\beta h)(1+\omega\theta)} (\eta_3 x_{t+j} + \eta_1(x_{t+j-1} + \beta x_{t+j+1})) + \sum_{k=1}^{j} \pi_{t+k} \right] \tag{27}$$

where $x_t \equiv y_t - y_t^n$ is the output gap.

Because all firms that are offered the opportunity to choose a new price in period t will choose the same price, the aggregate price index evolves according to

$$P_t = [\alpha P_{t-1}^{1-\theta} + (1-\alpha)p_t^{1-\theta}]^{\frac{1}{1-\theta}} \tag{28}$$

Log-linearizing this equation yields

$$\hat{p}_t = \frac{\alpha}{1-\alpha}\pi_t \tag{29}$$

Furthermore, the double sum in (27) can be simplified as

$$\sum_{j=0}^{\infty} (\alpha\beta)^j \sum_{k=1}^{j} \pi_{t+k} = (1-\alpha\beta)^{-1} \left[\sum_{j=0}^{\infty} (\alpha\beta)^j \pi_{t+j} - \pi_t \right] \tag{30}$$

Substituting (29) and (30) into (27), rearranging, and quasi-differencing the resulting expression, we obtain (7), where

$$\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \cdot \frac{1}{1+\omega\theta}$$

A.2 The Representative Household's Welfare

The second-order approximation (13) to the representative household's welfare (12) is derived in this section, using methods discussed in more detail in Woodford (1999a). The second-order Taylor series expansion of (12) is formed around the steady state characterized by the efficient output level Y^* defined in (20) and zero inflation, i.e. the same steady state around which the model's exact equilibrium conditions have been log-linearized.

Let $U_t \equiv u(Y_t, Y_{t-1}; \xi_t) - \int_0^1 v(y_t(i); \xi_t)di$, where the market clearing condition $C_t = Y_t \; \forall t$ has been substituted in. The first term in U_t can be approximated as

$$\begin{aligned} u(Y_t, Y_{t-1}; \xi_t) = {Y^*}^{1-\nu} \Big[& y_t + \frac{1-\sigma}{2} y_t^2 - h y_{t-1} \\ & + \frac{h^2(1-\sigma)}{2} y_{t-1}^2 - h(1-\sigma) y_t y_{t-1} + \sigma y_t g_t - h\sigma y_{t-1} g_t \Big] + t.i.p. + \mathcal{O}(\|\xi\|^3) \end{aligned} \tag{31}$$

where the notation $t.i.p.$ stands for terms that are independent of monetary policy.

Following the same steps as e.g. in Rotemberg and Woodford (1997), the second term in U_t can be approximated as

$$\begin{aligned} & \int_0^1 v(y_t(i); \xi_t) di \\ = \; & v_y Y^* \left[y_t + \frac{1}{2}(1+\omega) y_t^2 + \frac{1}{2}(\theta^{-1} + \omega) var_i(\hat{y}_t(i)) - \omega y_t z_t \right] + t.i.p. + \mathcal{O}(\|\xi\|^3) \end{aligned} \tag{32}$$

where z_t and ω are defined as in (23), $\hat{y}_t(i) \equiv \log(y_t(i)/Y^*)$, and var_i measures the variance of $\hat{y}_t(i)$ across i. Subtracting (32) from (31), and using from (20) that $v_y Y^* = (1-\beta h){Y^*}^{1-\nu}$, and from (3) that $var_i(\hat{y}_t(i)) = \theta^2 var_i(\log p_t(i))$, we obtain

$$\begin{aligned} U_t \;=\; & -\frac{{Y^*}^{1-\nu}}{2} \left\{ 2h \left[y_{t-1} + \frac{1}{2} y_{t-1}^2 - \beta(y_t + \frac{1}{2} y_t^2) \right] + (\sigma + (1-\beta h)\omega) y_t^2 \right. \\ & + 2h(1-\sigma) y_t y_{t-1} - h(1 + h(1-\sigma)) y_{t-1}^2 - 2(1-\beta h)\omega z_t y_t \\ & \left. - 2\sigma g_t (y_t - h y_{t-1}) + (1-\beta h)\theta(1+\omega\theta) var_i(\log p_t(i)) \right\} + t.i.p. + \mathcal{O}(\|\xi\|^3) \end{aligned} \tag{33}$$

Forming the sum $\sum_{t=0}^{\infty} \beta^t U_t$, rearranging terms across summands, and using the definition (8) of y_t^n, the objective (12) can be expressed as

$$\begin{aligned} E_0 \sum_{t=0}^{\infty} \beta^t U_t \;=\; & -\frac{{Y^*}^{1-\nu}}{2} E_0 \left[2h(y_{-1} + \frac{1}{2} y_{-1}^2) - h(1 + h(1-\sigma)) y_{-1}^2 + h(1-\sigma) y_0 y_{-1} \right. \\ & + \sum_{t=0}^{\infty} \beta^t \left\{ \eta_3 x_t^2 + \eta_1 (x_{t-1}^2 + \beta x_{t+1}^2) + \eta_1 (y_{t-1} \Delta y_t - \beta y_{t+1} \Delta y_{t+1}) \right. \\ & \left. + (1-\beta h)\theta(1+\omega\theta) var_i(\log p_t(i)) \} \right] + t.i.p. + \mathcal{O}(\|\xi\|^3) \end{aligned} \tag{34}$$

Finally, using the approximation that

$$\sum_{t=0}^{\infty}\beta^t var_i(\log p_t(i)) = \frac{\alpha}{(1-\alpha)(1-\alpha\beta)}\sum_{t=0}^{\infty}\beta^t\pi_t^2 + t.i.p. + \mathcal{O}(\|\xi\|^3) \tag{35}$$

(see Woodford (1999a)) and suppressing a constant and terms independent of policy, we obtain the loss function (14), where $\lambda_x \equiv \frac{\kappa}{(1-\beta h)\theta}$.

Further rearranging of terms across summands in (34) yields

$$\begin{aligned} E_0\sum_{t=0}^{\infty}\beta^t U_t \;=\; & -\Omega E_0\Big[\lambda_x\Big\{2h(y_{-1}+\frac{1}{2}y_{-1}^2) - h(1+h(1-\sigma))y_{-1}^2 + \eta_1(x_{-1}^2 - y_{-1}^2)\Big\} \\ & + \sum_{t=0}^{\infty}\beta^t\Big\{\pi_t^2 + \lambda_x\Big[(\eta_3+\beta\eta_1)x_t^2 + \beta\eta_1 x_{t+1}^2 + 2\eta_1 y_t y_{t-1} - \beta\eta_1(y_t^2 + y_{t+1}^2)\Big]\Big\}\Big] \\ & + t.i.p. + \mathcal{O}(\|\xi\|^3) \end{aligned} \tag{36}$$

where $\Omega \equiv \frac{{Y^*}^{1-\nu}}{2}\frac{(1-\beta h)\theta}{\kappa}$. Note that all right-hand side terms on the first line are independent of the policy adopted at time 0. Replacing the cross-product $2y_t y_{t-1}$ by $(y_t + y_{t-1})^2 - y_t^2 - y_{t-1}^2$, applying the operator (16), and noting that for any variable z_t, $V[z_t] = \beta V[z_{t+1}]$, we obtain (15).

B Conditions Characterizing the Optimal Plan

From (36) we obtain a loss function of the form (13), where the period-loss function is

$$L_t = \pi_t^2 + \lambda_x\Big[(\eta_3+\beta\eta_1)x_t^2 + \beta\eta_1 x_{t+1}^2 + 2\eta_1 y_t y_{t-1} - \beta\eta_1(y_t^2 + y_{t+1}^2)\Big]$$

Appending the term $\lambda_r r_t^2$ to L_t for the reasons discussed in section 2, replacing y_t by $x_t + y_t^n$, and omitting terms independent of policy, the problem of choosing at time 0 an optimal plan for the endogenous variables $\{x_t, \pi_t, r_t\}$ can be stated as a Lagrangian

$$\begin{aligned} E_0\sum_{t=0}^{\infty}\beta^t\Big\{ & \pi_t^2 + \lambda_x\Big[\eta_3 x_t^2 + 2\eta_1(x_t x_{t-1} + x_t y_{t-1}^n + x_{t-1}y_t^n) - 2\beta\eta_1(x_t y_t^n + x_{t+1}y_{t+1}^n)\Big] + \lambda_r r_t^2 \\ + \; & 2\phi_{1,t}\Big[\frac{\eta_2-\eta_1}{\sigma}x_t - \frac{\beta\eta_1}{\sigma}x_{t+2} - \frac{\eta_2-\beta\eta_1}{\sigma}x_{t+1} + \frac{\eta_1}{\sigma}x_{t-1} + \frac{1-\beta h}{\sigma}(r_t - r_t^n - \pi_{t+1})\Big] \\ + \; & 2\phi_{2,t}\Big[\pi_t - \frac{\kappa\eta_3}{1-\beta h}x_t - \frac{\kappa\eta_1}{1-\beta h}x_{t-1} - \frac{\kappa\beta\eta_1}{1-\beta h}x_{t+1} - \beta\pi_{t+1}\Big]\Big\} \end{aligned} \tag{37}$$

where (9) and (7) are the constraints, and $\{\phi_{1,t}, \phi_{2,t}\}$ are sequences of Lagrange multipliers associated with the constraints. The first-order conditions with respect to x_t, π_t, and r_t are respectively

$$0 \;=\; \lambda_x\{\eta_3 x_t + \eta_1(x_{t-1} + y_{t-1}^n) + \beta\eta_1 E_t[x_{t+1} + y_{t+1}^n] - (1+\beta)\eta_1 y_t^n\}$$

$$
\begin{aligned}
& +\frac{1}{\sigma}\left[(\eta_2-\eta_1)\phi_{1,t}-\frac{\eta_1}{\beta}\phi_{1,t-2}-\frac{\eta_2-\beta\eta_1}{\beta}\phi_{1,t-1}+\beta\eta_1 E_t\phi_{1,t+1}\right] \\
& -\frac{\kappa}{1-\beta h}[\eta_3\phi_{2,t}+\eta_1\phi_{2,t-1}+\beta\eta_1 E_t\phi_{2,t+1}] \qquad (38)
\end{aligned}
$$

$$
0 = \pi_t - \frac{1-\beta h}{\sigma\beta}\phi_{1,t-1}+\phi_{2,t}-\phi_{2,t-1} \qquad (39)
$$

$$
0 = \lambda_r r_t + \frac{1-\beta h}{\sigma}\phi_{1,t} \qquad (40)
$$

An optimal plan is defined as a bounded solution $\{\pi_t, x_t, r_t, \phi_{1,t}, \phi_{2,t}\}_{t=0}^{\infty}$ to the system of equations (5), (7), and (38)-(40) together with the initial condition that $\phi_{1,-1}=\phi_{1,-2}=\phi_{2,-1}=0$.

References

[1] Abel, A.B., 1990. Asset Prices under Habit Formation and Catching up with the Joneses. American Economic Review 80, Papers and Proceedings, 38-42.

[2] Attanasio, O., Weber, G., 1993. Consumption Growth, the Interest Rate, and Aggregation. Review of Economic Studies 60, 631-649.

[3] Blinder, A.S., 1994. On Sticky Prices: Academic Theories Meet the Real World. In: Mankiw, N.G. (Ed.), Monetary Policy. The University of Chicago Press, Chicago, pp. 117-150.

[4] Calvo, G.A., 1983. Staggered Prices in a Utility-Maximizing Framework. Journal of Monetary Economics 12, 383-398.

[5] Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton.

[6] Carroll, C.D., Overland, J., Weil, D.N., 2000. Saving and Growth with Habit Formation. American Economic Review 90, 341-355.

[7] Coenen, G., Wieland, V., 2000. A Small Estimated Euro Area Model with Rational Expectations and Nominal Rigidities. ECB Working Paper No. 30, European Central Bank.

[8] Dynan, K.E., 2000. Habit Formation in Consumer Preferences: Evidence from Panel Data. American Economic Review 90, 391-406.

[9] Edge, R., 2000. Time-to-Build, Time-to-Plan, Habit-Persistence, and the Liquidity Effect. International Finance Discussion Papers, No. 673, Federal Reserve Board.

[10] Erceg, C.J., Henderson, D.W., Levin, A.T., 2000. Optimal Monetary Policy with Staggered Wage and Price Contracts. Journal of Monetary Economics 46, 281-313.

[11] Fuhrer, J.C., 2000. Habit Formation in Consumption and Its Implications for Monetary-Policy Models. American Economic Review 90, 367-390.

[12] Fuhrer, J.C., Moore, G., 1995. Inflation Persistence. Quarterly Journal of Economics 110, 127-159.

[13] Hall, R.E., 1988. Intertemporal Substitution in Consumption. Journal of Political Economy 96, 339-357.

[14] McCallum, B.T., Nelson, E., 1999. Nominal Income Targeting in an Open-Economy Optimizing Model. Journal of Monetary Economics 43, 553-578.

[15] Mulligan, C.B., 1998. Substitution over Time: Another Look at Life-Cycle Labor Supply. In: Bernanke, B.S., Rotemberg, J.J. (Eds.), NBER Macroeconomics Annual 1998. MIT Press, Cambridge, pp. 75-134.

[16] Orphanides, A., 1998. Monetary Policy Evaluation with Noisy Indicators. Finance and Economics Discussion Series, No. 1998-50, Federal Reserve Board.

[17] Pencavel, J.H., 1986. Labor Supply of Men: A Survey. In: Ashenfelter, O., Layard, R. (Eds.), Handbook of Labor Economics Vol. 1. North Holland, Amsterdam.

[18] Rotemberg, J.J., Woodford, M., 1997. An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy. In: Bernanke, B.S., Rotemberg, J.J. (Eds.), NBER Macroeconomics Annual 1997. MIT Press, Cambridge, pp. 297-346

[19] Rotemberg, J.J., Woodford, M., 1999. Interest-Rate Rules in an Estimated Sticky Price Model. In Taylor, J.B. (Ed.), Monetary Policy Rules. The University of Chicago Press, Chicago, pp. 57-119.

[20] Rudebusch, G.D., Svensson, L.E.O., 1999. Policy Rules for Inflation Targeting. In Taylor, J.B. (Ed.), Monetary Policy Rules. The University of Chicago Press, Chicago, pp. 203-246.

[21] Taylor, J.B., 1979. Estimation and Control of a Macroeconomic Model with Rational Expectations. Econometrica 47, 1267-1286.

[22] Taylor, J.B., 1993. Discretion Versus Policy Rules in Practice. Carnegie-Rochester Conference Series on Public Policy 39, 195-214.

[23] Taylor, J.B., 1999. Monetary Policy Rules (Ed.) The University of Chicago Press, Chicago.

[24] Woodford, M., 1999a. Inflation Stabilization and Welfare. Mimeo, Princeton University.

[25] Woodford, M., 1999b. Optimal Monetary Policy Inertia. NBER Working Paper No. 7261.

Figure 1: Model Impulse Responses to a Monetary Policy Shock

under Estimated Interest Rate Rule

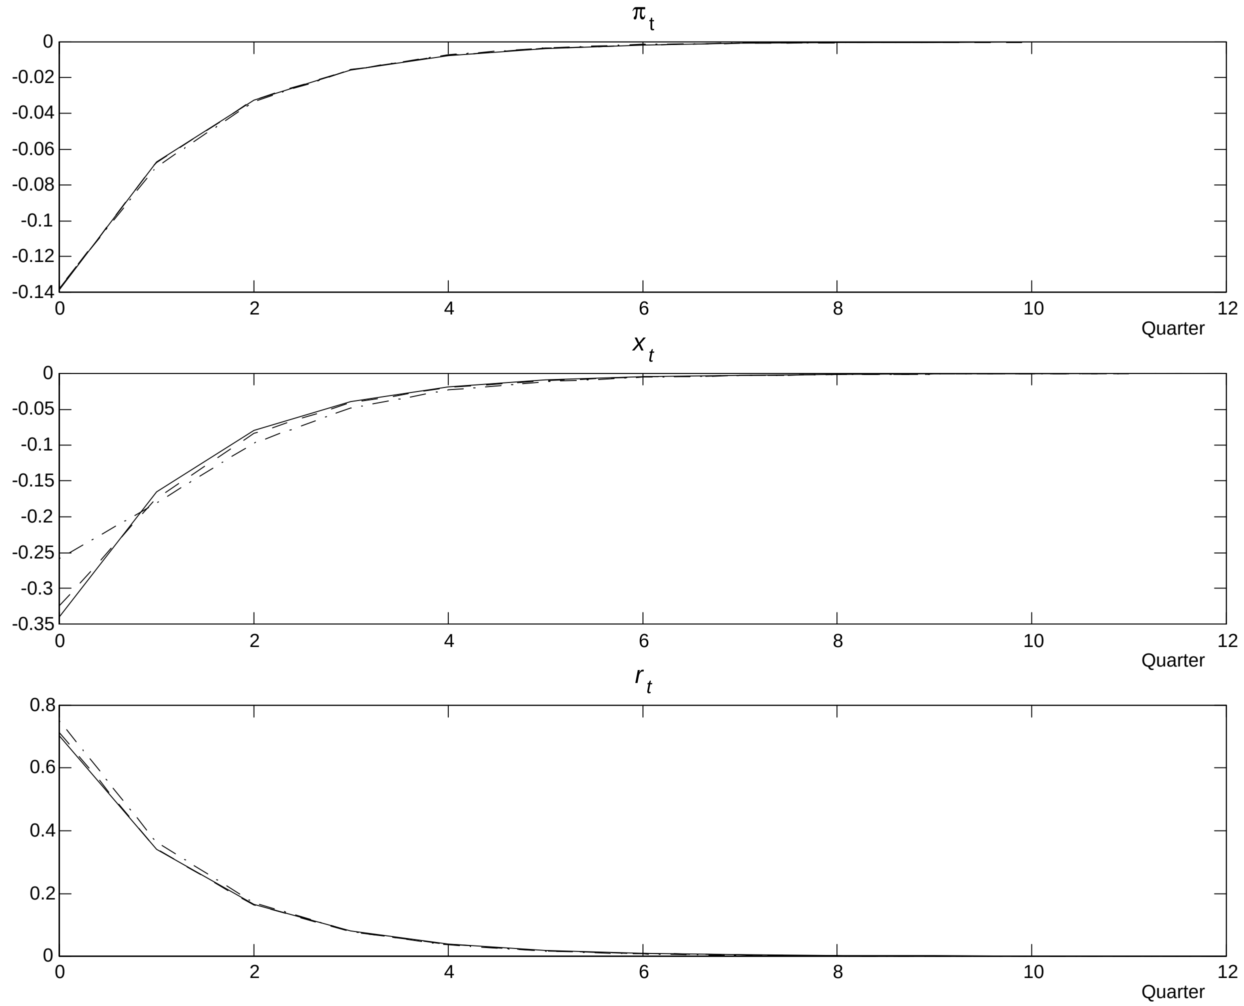

Notes: The figure shows the impulse responses of some of the endogenous variables to a one-unit monetary policy shock at time 0, using the model equations (7) and (9) and the interest rate rule given in (11). The different lines correspond to different values of h: 0 (solid), 0.4 (dash) and 0.8 (dash-dot). One period is equal to a quarter. Inflation and the interest rate are expressed in annualized percentages, the output gap in percentages.

Figure 2: Model Impulse Responses to a Marginal Utility of Consumption Shock under Estimated Interest Rate Rule

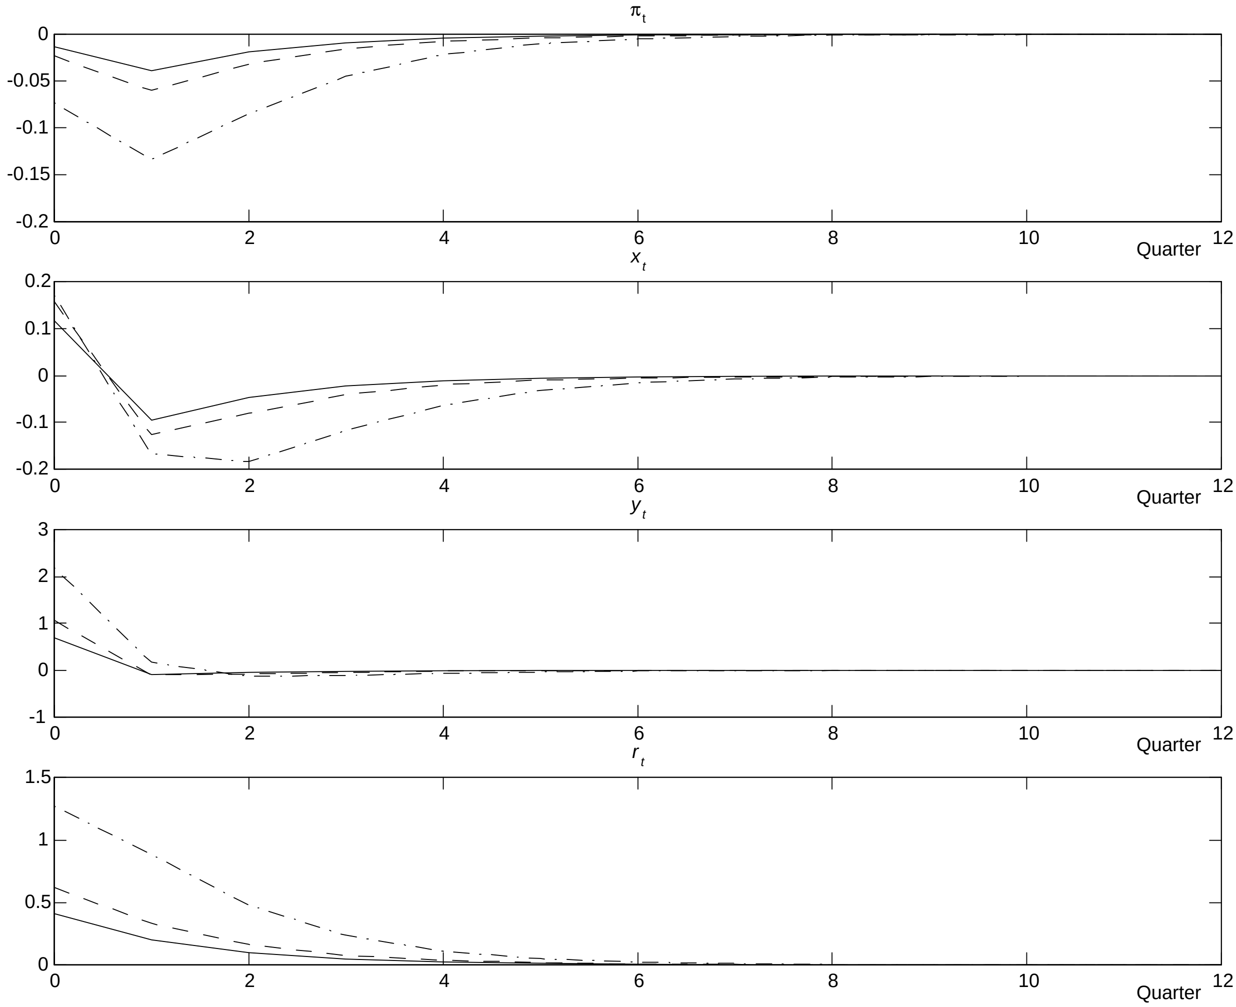

Notes: The figure shows the impulse responses of the endogenous variables to a one-standard deviation shock to the orthogonalized innovation in the process g_t at time 0. This has the interpretation of a shock to the marginal utility of consumption that does not also directly affect the disutility of labour supply. These are based on the model equations (7) and (9) and the interest rate rule given in (11). The different lines correspond to different values of h: 0 (solid), 0.4 (dash) and 0.8 (dash-dot). One period is equal to a quarter. Inflation and the interest rate are expressed in annualized percentages, output and the output gap in percentages.

Figure 3: Model Autocorrelation Functions
under Estimated Interest Rate Rule

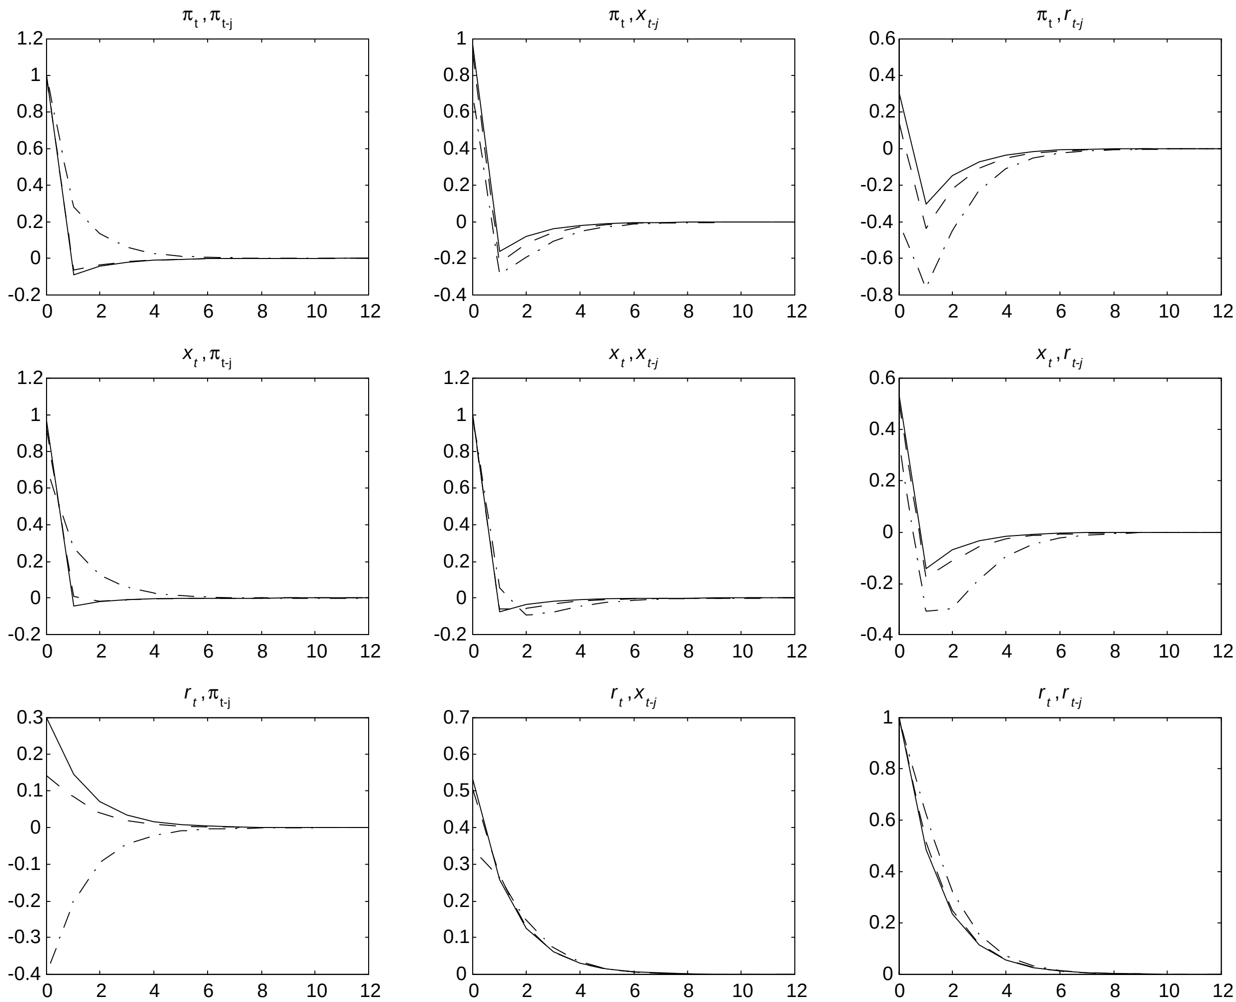

Notes: The figure shows the autocorrelation functions of some of the endogenous variables based on the model equations (7) and (9) and the interest rate rule given in (11). The different lines correspond to different values of h: 0 (solid), 0.4 (dash) and 0.8 (dash-dot). One period is equal to a quarter.

Figure 4: Standard Deviation of the Natural Rate of Interest
Base Case Calibration

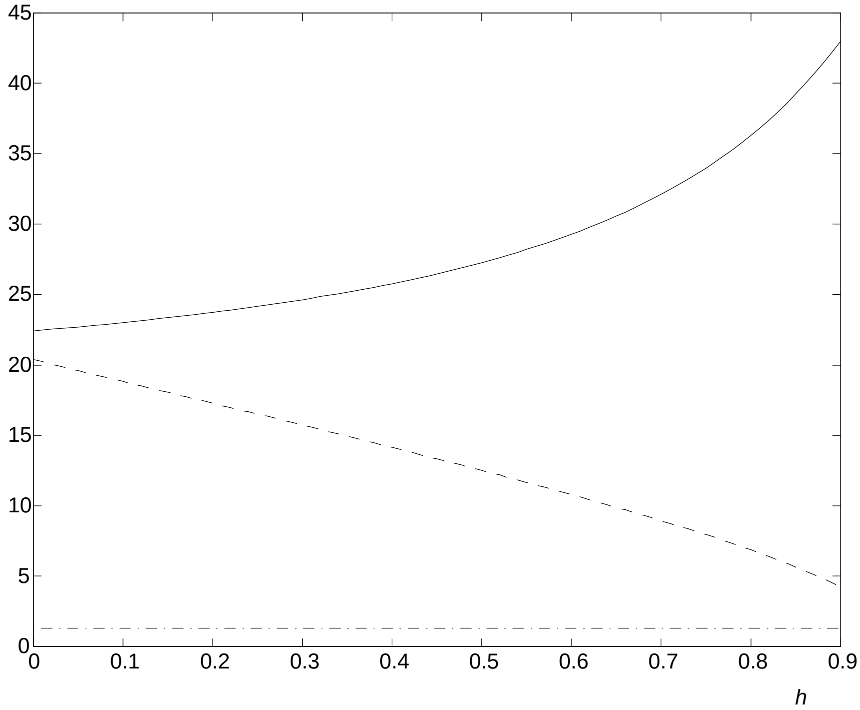

Notes: The figure shows the unconditional standard deviation of the natural rate of interest r_t^n (solid), given in equation (10), across different values of h. Also shown is the component of the standard deviation related to the conditional expectation of the shock processes, i.e. $std.dev.(r_t^n)(1-\beta h)/\sigma$ (dash), and the threshold at which the zero lower bound on nominal interest rates becomes binding, i.e. $std.dev.(r_t^n) = 1.32$ (dash-dot). The natural rate of interest is expressed in annualized percentages.

Figure 5: Impulse Responses to a Marginal Utility of Consumption Shock
Optimal Plan, Base Case Calibration

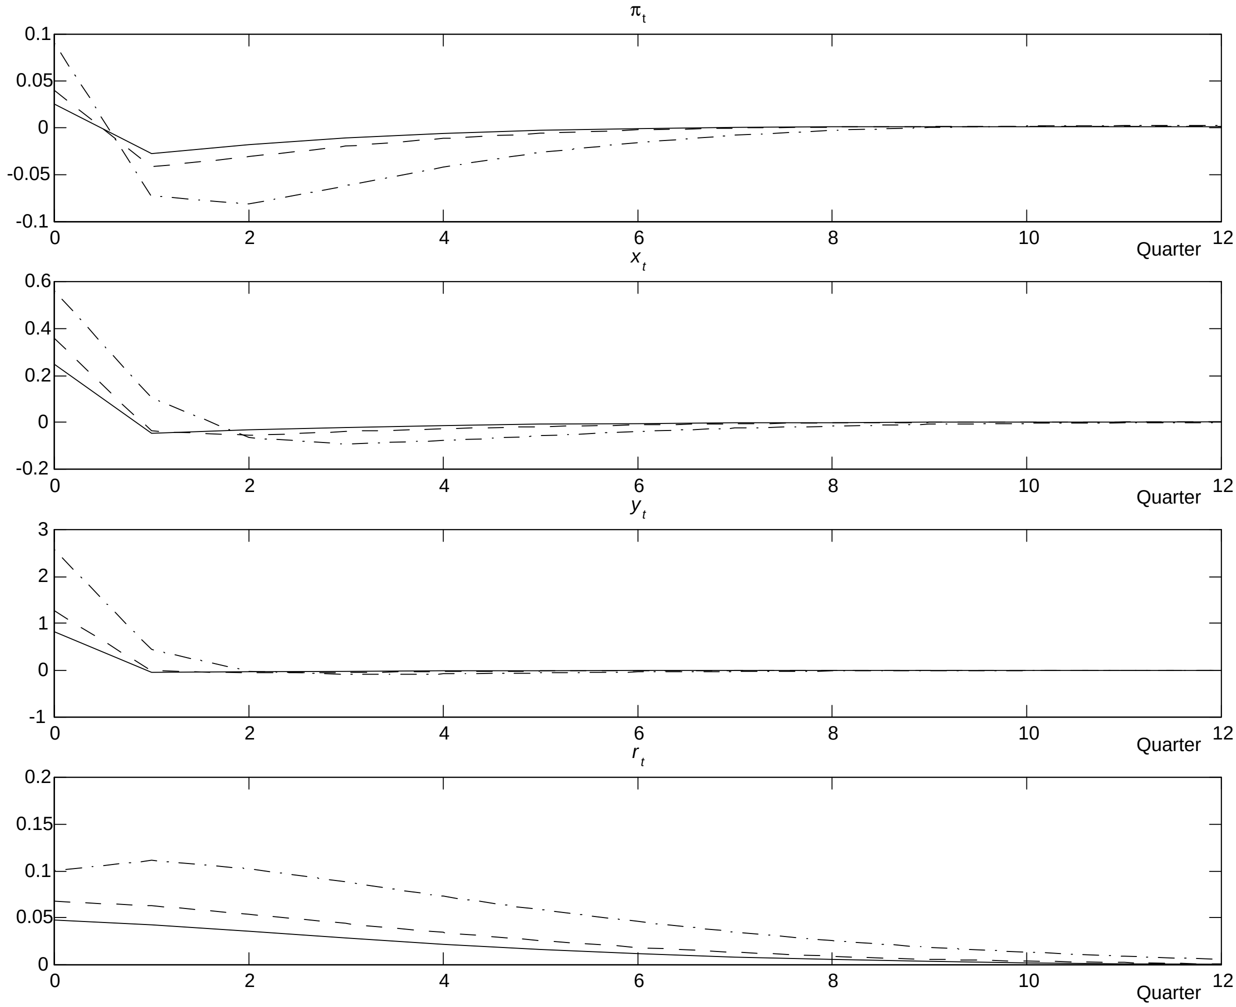

Notes: The figure shows the impulse responses of the endogenous variables to a one-standard deviation shock to the orthogonalized innovation in the process g_t at time 0. This has the interpretation of a shock to the marginal utility of consumption that does not also directly affect the disutility of labour supply. The different lines correspond to different values of h: 0 (solid), 0.4 (dash) and 0.8 (dash-dot). One period is equal to a quarter. Inflation and the interest rate are expressed in annualized percentages, output and the output gap in percentages.

Figure 6: Standard Deviation of the Natural Rate of Interest
Alternative Calibration, Case 1

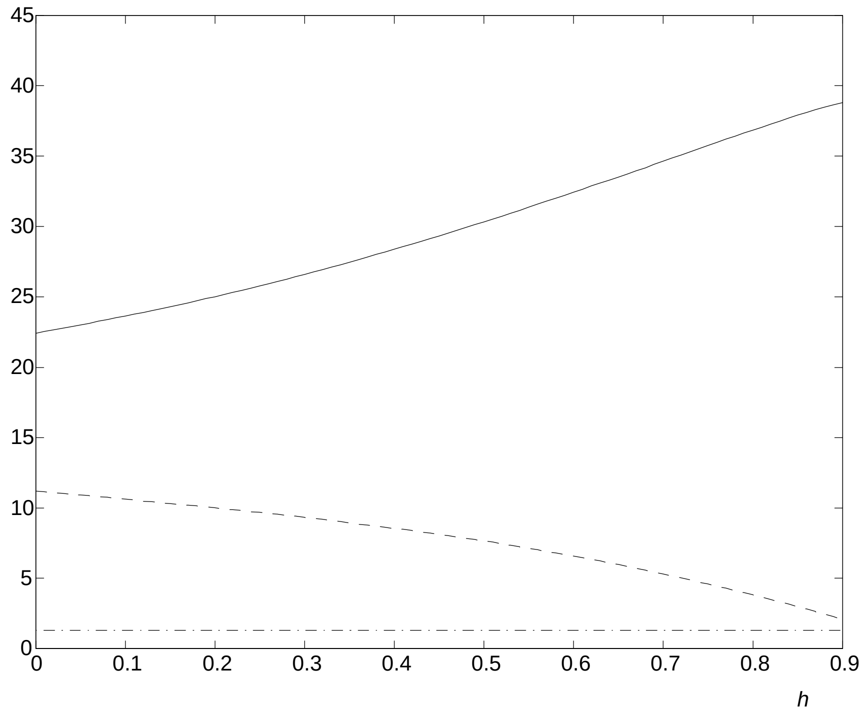

Notes: The figure shows the unconditional standard deviation of the natural rate of interest r_t^n (solid), given in equation (10), across different values of h. Also shown is the component of the standard deviation related to the conditional expectation of the shock processes, i.e. $std.dev.(r_t^n)(1-\beta h)/\sigma$ (dash), and the threshold at which the zero lower bound on nominal interest rates becomes binding, i.e. $std.dev.(r_t^n) = 1.32$ (dash-dot). The natural rate of interest is expressed in annualized percentages.

Figure 7: Impulse Responses to a Marginal Utility of Consumption Shock
Optimal Plan, Base Case vs. Alternative Calibrations, h = 0

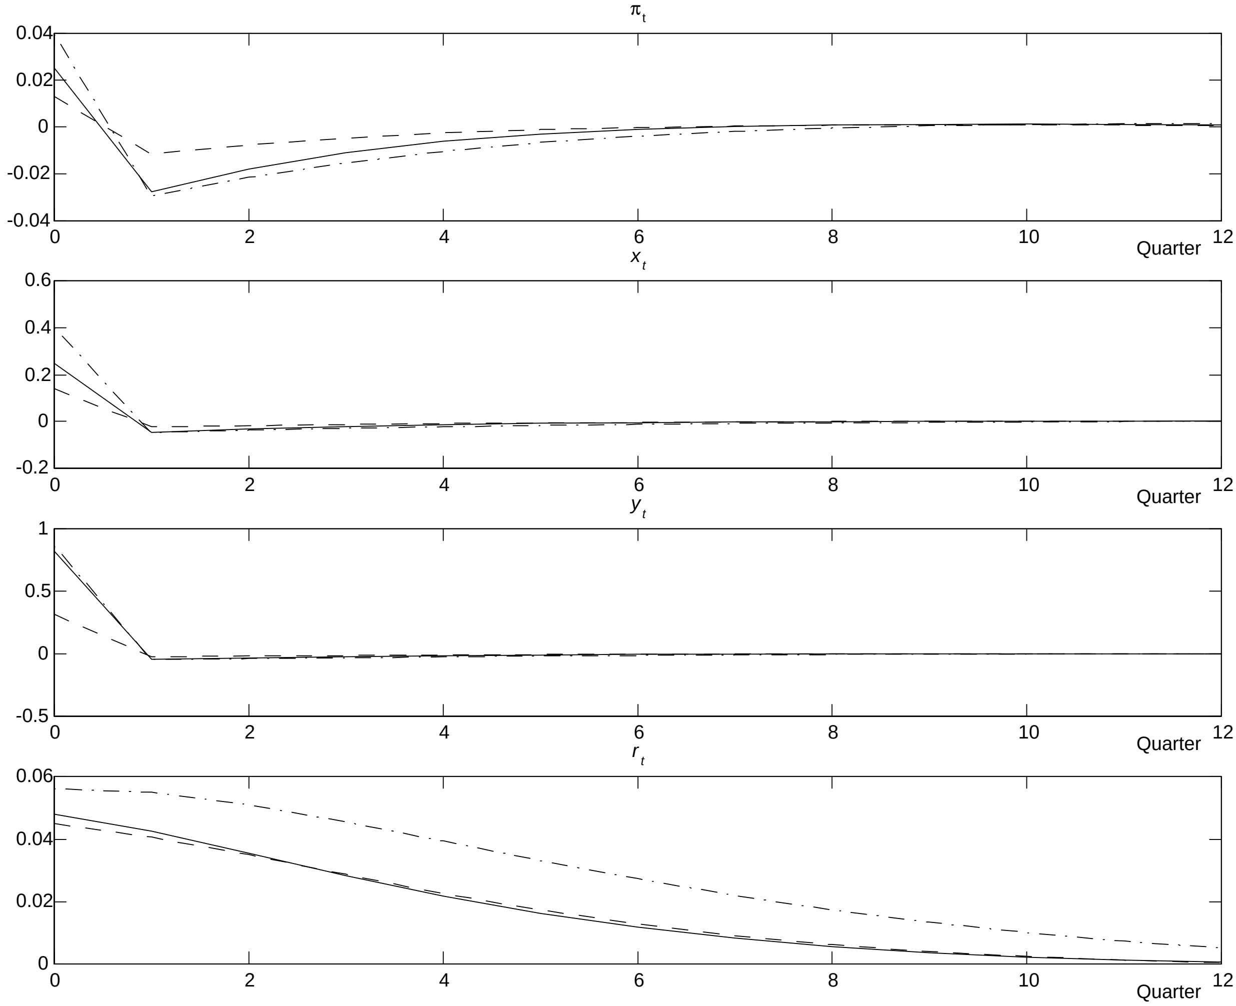

Notes: The figure shows the impulse responses of the endogenous variables to a one-standard deviation shock to the orthogonalized innovation in the process g_t at time 0. This has the interpretation of a shock to the marginal utility of consumption that does not also directly affect the disutility of labour supply. The different lines correspond to different cases for calibration of the parameters and shocks: base case (solid), alternative case 1 (dash) and alternative case 2 (dash-dot). The parameter h is set to 0. One period is equal to a quarter. Inflation and the interest rate are expressed in annualized percentages, output and the output gap in percentages.

Figure 8: Impulse Responses to a Marginal Utility of Consumption Shock
Optimal Plan, Base Case vs. Alternative Calibrations, h = 0.8

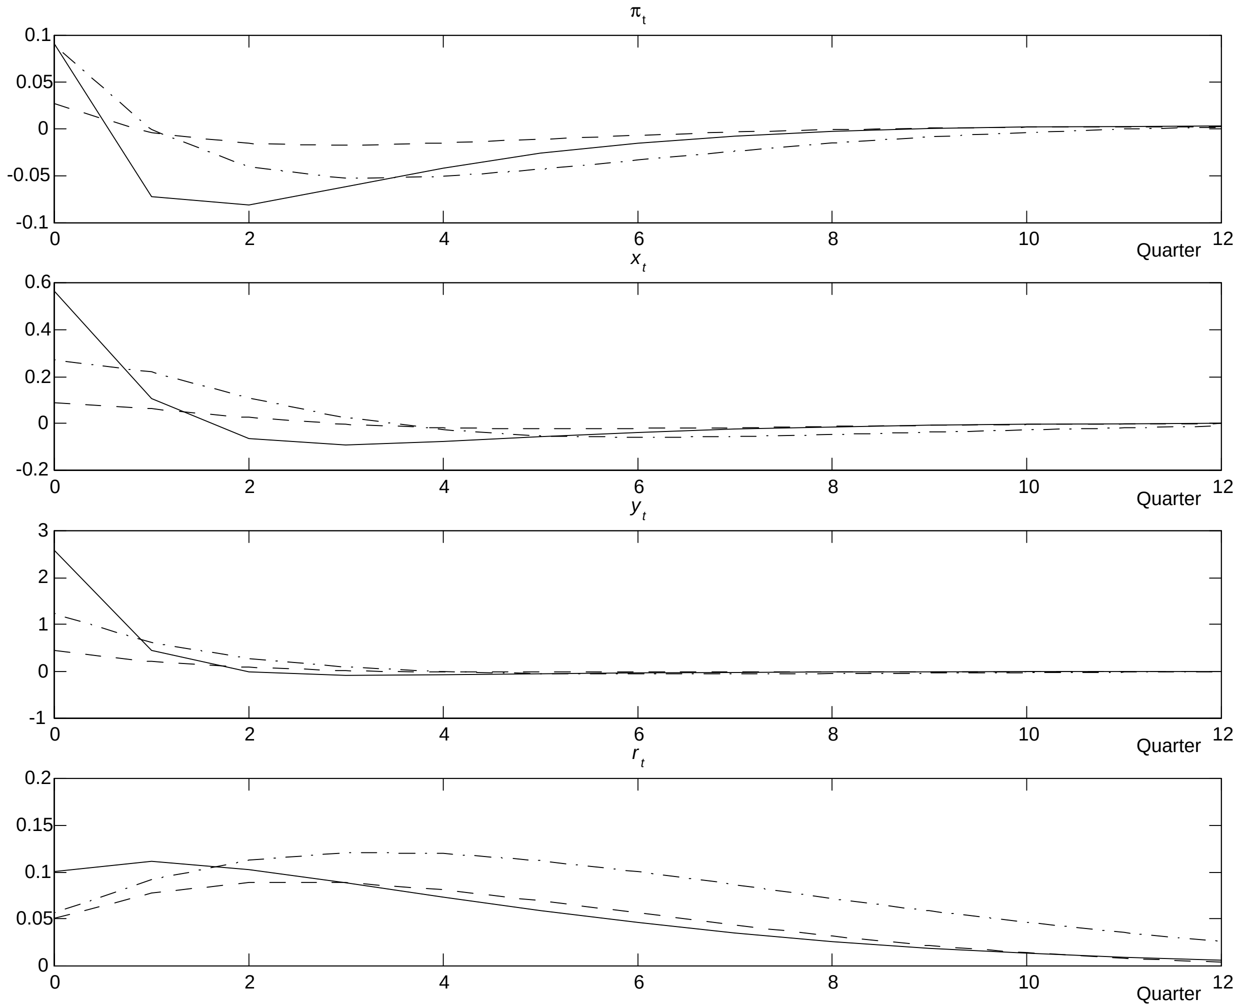

Notes: The figure shows the impulse responses of the endogenous variables to a one-standard deviation shock to the orthogonalized innovation in the process g_t at time 0. This has the interpretation of a shock to the marginal utility of consumption that does not also directly affect the disutility of labour supply. The different lines correspond to different cases for calibration of the parameters and shocks: base case (solid), alternative case 1 (dash) and alternative case 2 (dash-dot). The parameter h is set to 0.8. One period is equal to a quarter. Inflation and the interest rate are expressed in annualized percentages, output and the output gap in percentages.

Figure 9: Standard Deviation of the Natural Rate of Interest
Alternative Calibration, Case 2

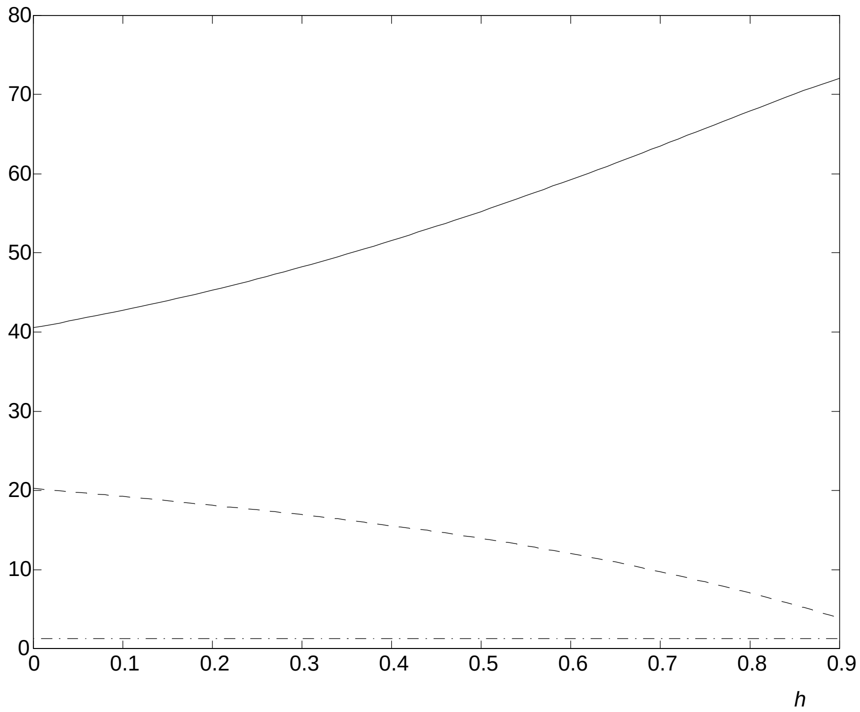

Notes: The figure shows the unconditional standard deviation of the natural rate of interest r_t^n (solid), given in equation (10), across different values of h. Also shown is the component of the standard deviation related to the conditional expectation of the shock processes, i.e. $std.dev.(r_t^n)(1-\beta h)/\sigma$ (dash), and the threshold at which the zero lower bound on nominal interest rates becomes binding, i.e. $std.dev.(r_t^n) = 1.32$ (dash-dot). The natural rate of interest is expressed in annualized percentages.