Habit Formation in Consumption and Its Implications for Monetary Policy Models

Jeffrey C. Fuhrer*
Federal Reserve Bank of Boston
April, 2000

Abstract

This paper explores a monetary policy model with habit formation for consumers, in which consumers' utility depends in part on current consumption relative to past consumption. The empirical tests developed in the paper show that one can reject the hypothesis of no habit formation with tremendous confidence, largely because the habit formation model captures the gradual hump-shaped response of real spending to various shocks. The paper then embeds the habit consumption specification in a monetary policy model and finds that the responses of both spending and inflation to monetary policy actions are significantly improved by this modification. (JEL D12, E52, E43)

Forthcoming, *American Economic Review*, June 2000.

With the resurgence of interest in the effects of monetary policy on the macroeconomy, led by the work of the Christina D. and David H. Romer (1989), Ben S. Bernanke and Alan S. Blinder (1992), Lawrence J. Christiano, Martin S. Eichenbaum, and Charles L. Evans (1996), and others, the need for a structural model that could plausibly be used for monetary policy analysis has become evident. Of course, many extant models have been used for monetary policy analysis, but many of these are perceived as having critical shortcomings. First, some models do not incorporate explicit expectations behavior, so that changes in policy (or private) behavior could cause shifts in reduced-form parameters (i.e., the critique of Robert E. Lucas 1976). Others incorporate expectations, but derive key relationships from *ad hoc* behavioral assumptions, rather than from explicit optimizing problems for consumers and firms (Fuhrer and George R. Moore 1995b is an example).

Explicit expectations and optimizing behavior are both desirable, other things equal, for a model of monetary analysis. First, analyzing potential improvements to monetary policy relative to historical policies requires a model that is stable across alternative policy regimes. This underlines the importance of explicit expectations formation. Second, the "optimal" in optimal monetary policy must ultimately refer to social welfare. Many have approximated social welfare with weighted averages of output and inflation variances, but one cannot know how good these approximations are without more explicit modeling of welfare. This implies that the model be closely tied to the underlying objectives of consumers and firms, hence the emphasis on optimization-based models. A critical test for whether a model reflects underlying objectives is its ability to accurately reflect the dominant dynamic interactions in the data.

A number of recent papers (see, for example, Robert G. King and Alexander L. Wolman (1996), Bennett T. McCallum and Edward Nelson (1999a, 1999b); Julio R. Rotemberg and Michael Woodford (1997)) have developed models that incorporate explicit expectations, optimizing behavior, and frictions that allow monetary policy to have real effects. This paper continues in that line of research by documenting the empirical importance of a key feature of aggregate data: the "hump-shaped," gradual response of spending and inflation to shocks. It then develops a monetary policy model that can capture this feature, as well as all of the features (e.g. the real effects of monetary policy, the persistence of inflation and output) embodied in earlier models.

The key to the model's success on the spending side is the inclusion of habit formation in the consumer's utility function. This modification significantly improves the short-run dynamic behavior of the model, both qualitatively and statistically. Such improvements in the model's ability to accurately reflect significant short-run dynamic properties may be quite important, especially given the working assumption among most economists that monetary policy has only short-run effects on real variables.

The improvements afforded by habit formation arise in two ways. First, the data on real consumption spending exhibit a significant delay and hump-shaped response to monetary policy and other shocks. Habit formation allows the model to match the response of real spending to monetary policy shocks. In addition, given the link in most monetary policy models from real spending to inflation [see, for example, the price specification in John B. Taylor (1980), as well as the optimizing models of Rotemberg and Woodford (1997) and McCallum and Nelson (1999a)), a jump response in real spending can cause a corresponding jump response in inflation (see Figure 7 and section 5.1 below for a full discussion of this point). Conversely, then, a more gradual real spending response to monetary shocks implies a more gradual response of inflation to policy shocks. In particular, the model with habit formation can much more accurately replicate the gradual decline of inflation during a disinflation.

The next section reviews the reason that the simple permanent income model is unable to replicate the "hump-shaped" response of consumption to shocks that characterizes the aggregate data. It then motivates the use of habit formation as an *a priori* desirable modification to the model, and demonstrates how it can yield a hump-shaped impulse response of consumption to shocks. Section 2 more fully develops the model of habit formation in consumer behavior, based on the utility specifications used in Andrew B. Abel (1990) and Christopher D. Carroll, Jody R. Overland, and David N. Weil (1995), and related in spirit to the pioneering work of James S. Duesenberry (1949). Section 3 conducts a number of empirical tests to determine the extent to which habit formation can improve the dynamic behavior of the simple model. The results show that, because habit formation imparts a utility-based smoothing motive for both changes and levels of consumption, it significantly improves the ability of the model to match the hump-shaped response of consumption to shocks. Sections 4 and 5 incorporate the consumption specification into a simple model for monetary policy analysis, and examine the resulting improvement in the overall dynamic behavior of the model, and section 6 concludes.

1 The Hump-Shaped Response, Excess Smoothness, and Habit Formation

A key feature of the monetary policy transmission mechanism is that both spending and inflation variables demonstrate a gradual response to policy actions over several years, with the peak response at about one year, and the entire effect lasting two years or more. These empirical regularities have been emphasized recently in identified VAR work by Christiano, Eichenbaum, and Evans (1996) and Eric M. Leeper, Christopher A. Sims, and Tao Zha (1996), and immortalized in Milton Friedman's depiction of the long and variable lags of monetary policy.

Yet many standard specifications imply instead that spending and inflation act like "jump variables," completely front-loading or pulling forward in time their responses to shocks. This is a well-known feature of the permanent income hypothesis (PIH) model with rational expectations. Consumption jumps immediately in response to current "news" about lifetime resources, a direct implication of the random walk property of PIH consumption derived by Robert E. Hall (1978). John Y. Campbell and Angus S. Deaton (1989) showed that in fact consumption does not respond immediately to news, and that as a result, consumption exhibits "excess smoothness" that cannot be reconciled with the PIH model. Resolutions of the excess smoothness puzzle have proven elusive, as emphasized in recent work by Sydney C. Ludvigson and Alexander Michaelides (1998), which shows that models that include a fuller treatment of uncertainty do not solve the puzzle. One contribution that this paper makes is an empirically successful solution to the excess smoothness puzzle, a solution that relies on the presence of habit formation in the utility function.

The inability of theoretical models to match the dynamics in the data extends to price specifications as well. Fuhrer and Moore (1995a) and John M. Roberts (1997) discuss the jump problem for staggered wage and price contract models of inflation. For monetary policy models, then, the challenge is to build models that imply a gradual and hump-shaped response of spending and inflation to all shocks, and in particular to monetary policy shocks.[1]

This paper begins by focusing on consumption expenditures for nondurable goods and services. Interestingly, as I show below, the hump-shaped response to interest rate or income shocks is not linked exclusively to durable goods. Such a response is evident and statistically significant in the data for nondurables and services, which accounts for almost 60 percent of GDP.

But this observation raises a challenge in uncovering the source of the gradual response to shocks. A number of the modifications to the standard consumption and investment models that have been proposed, including costs of adjustment, durability, and time-to-build lags, simply don't make much sense for nondurables and services consumption.[2] The costs of adjusting the real quantity of haircuts per period or food consumed at home are likely quite small and not a material impediment to rapid adjustment to important changes in the macroeconomic environment. Thus, on *a priori* grounds, one must look elsewhere for justification of the hump-shaped response in nondurables and services consumption.

If the source of gradual responses is unlikely to be found in costs of adjustment or in a fuller accounting of uncertainty, a natural alternative is to reexamine the specification of the utility function. Thus, this paper explores the implications of a utility function that should be expected to produce more sluggish responses, because it allows for consumers who form slowly-changing habits.

Habit formation may be modeled by assuming that consumers' current utility is determined by current consumption relative to a reference level of consumption. The notion that consumers form habits in their expenditure patterns certainly has intuitive appeal. And an examination of a simple form of a utility function with habit formation reveals how habit formation produces hump-shaped responses. Consider a specific form of the utility function explored in more detail below:

$$U_t = \frac{1}{(1-\sigma)} \left[\frac{C_t}{C_{t-1}^{\gamma}} \right]^{(1-\sigma)}$$

Here current utility, U_t, depends on current consumption, C_t, relative to lagged consumption, the habit reference level. The parameter γ indexes the importance of the reference level relative to current consumption. This utility function may be rewritten as

$$U_t = \frac{1}{(1-\sigma)} \left[\frac{C_t}{C_{t-1}} C_{t-1}^{(1-\gamma)} \right]^{1-\sigma} \quad ,$$

which highlights the essence of habit formation: Consumers wish to smooth both the level and the change in consumption. Thus, in response to shocks to interest rates or income, both the level and the change in consumption will respond gradually, leading to a hump-shaped response. This implication of habit formation will be developed more fully below.

Interestingly, other literatures have developed considerable theoretical and empirical support for habit formation. The extensive literature on asset pricing anomalies, most notably the equity premium puzzle, lends credence to the presence of habit formation (see, for example, Abel (1990), George M. Constantinides (1990), Campbell and John H. Cochrane (1999), and Urban J. Jermann (1998)). As explained above, habit-forming consumers dislike large and rapid cuts in consumption. As a result, the premium that they will require to hold risky assets that might force a rapid cut in consumption will be large relative to that implied by the time-separable utility model. While habit formation may not explain all asset pricing anomalies, it is becoming widely agreed that it "fits the data" better than time-separable utility asset pricing models.

Similarly, the growth literature provides support for the habit formation model. Much of the recent growth literature has attempted to explain the finding that growth Granger-causes saving [see Carroll and David N. Weil (1994) for the original documentation of the correlation, and more recent corroborating work by Orazio P. Attanasio, L. Picci, and A. Scorcu (1998), Dani Rodrik (1998), and N. Loayza, Klaus Schmidt-Hebbel, and Luis Servén (1998)]. This finding constitutes a serious violation of the PIH, because a PIH consumer would save *less* today in the face of strong growth that augments lifetime resources. Carroll, Overland, and Weil (1995) suggest that, as shown in this paper, habits imply a sluggish response of consumption to income shocks. Thus, after an income shock, growth in income could temporarily exceed consumption growth, raising savings while consumers gradually respond to the increase in income.

Finally, Lars Ljungqvist and Harald Uhlig (1999) examine a "catching up with the Joneses" utility function, a near cousin to habit formation models, in a productivity shock-driven model. They demonstrate that optimal tax policy in such a model can be procyclical. Such a "Keynesian" tax policy is optimal because it damps booms that arise sub-optimally because of the failure of individuals to take account of the external effect of their own consumption on the consumption of others.

2 A Simple Habit Formation Model

Following Abel (1990) and Carroll, Overland, and Weil (1995), consumers' t-period utility may be expressed as:

$$U_t = \frac{1}{(1-\sigma)} \left(\frac{C_t}{Z_t^{\gamma}} \right)^{(1-\sigma)} . \tag{1}$$

where Z_t is the habit-formation reference consumption level, defined as

$$Z_t = \rho Z_{t-1} + (1-\rho) C_{t-1} . \tag{2}$$

Note that utility is no longer time-separable, because the consumption choice today influences the future habit reference level in next period's and all future periods' utility. One advantage of this simple habit formation specification is that it conveniently parameterizes two features of habit formation:

1. The parameter γ indexes the importance of habit formation in the utility function. If $\gamma = 0$, then the standard model applies. If $\gamma = 1$, then *only* consumption relative to previous consumption matters. $\gamma > 1$ is not admissible, because it implies that steady-state utility is falling in consumption.

2. The parameter ρ indexes the persistence or "memory" in the habit formation reference level. If $\rho = 0$, then only last period's consumption is important. For $0 < \rho \leq 1$, the larger is ρ, the further back in time is the reference level determined (or, more accurately, the longer is the "mean lag" of the habit reference level).

2.1 The Euler Equation

Using the definition of period utility

$$U_t = \frac{1}{(1-\sigma)}\left(\frac{C_t}{Z_t^{\gamma}}\right)^{(1-\sigma)} , \tag{3}$$

the overall utility function

$$U = U_t + \beta U_{t+1} + \dots , \tag{4}$$

and the habit-formation reference consumption level

$$Z_t = \rho Z_{t-1} + (1-\rho)C_{t-1} , \tag{5}$$

we can compute the derivative of U with respect to C_t

$$\frac{\partial U}{\partial C_t} = \frac{\partial U_t}{\partial C_t} + \beta \frac{\partial U_{t+1}}{\partial Z_{t+1}}\frac{\partial Z_{t+1}}{\partial C_t} + \beta^2 \frac{\partial U_{t+2}}{\partial Z_{t+2}}\frac{\partial Z_{t+2}}{\partial C_t} \dots . \tag{6}$$

Noting that $\frac{\partial U_t}{\partial C_t} = \frac{1-\sigma}{C_t}U_t$, that $\frac{\partial U_t}{\partial Z_t} = \frac{-\gamma(1-\sigma)}{Z_t}U_t$, and that $\frac{\partial Z_{t+i}}{\partial C_t} = \rho^{i-1}(1-\rho)$, one can express the derivative of total utility with respect to consumption in period t as

$$\frac{\partial U}{\partial C_t} = \frac{1-\sigma}{C_t}U_t - \beta\frac{\gamma(1-\sigma)}{Z_{t+1}}U_{t+1}(1-\rho) - \beta^2\frac{\gamma(1-\sigma)}{Z_{t+2}}U_{t+2}\rho(1-\rho) - \dots \tag{7}$$

which in turn collapses to a more compact discounted summation

$$\frac{\partial U}{\partial C_t} = \frac{(1-\sigma)}{C_t}U_t - \gamma(1-\sigma)(1-\rho)\sum_{i=1}^{\infty}\beta^i\rho^{i-1}\frac{U_{t+i}(C_{t+i})}{Z_{t+i}} . \tag{8}$$

Defining

$$P_t \equiv \beta\rho E_t P_{t+1} + \frac{U_{t+1}(C_{t+1})}{\rho Z_{t+1}} \tag{9}$$

the derivatives of utility U with respect to C_{t+i} may be written

$$\frac{\partial U}{\partial C_{t+i}} = \frac{(1-\sigma)}{C_{t+i}}U_{t+i} - \gamma(1-\sigma)(1-\rho)P_{t+i} . \tag{10}$$

To derive the Euler equation, consider the effect on utility of shifting a unit of consumption from period t to period $t+1$. The optimal path of consumption should be such that an "epsilon" shift of consumption from one period to the next produces no change in utility. The decline in utility in period t, $-\frac{\partial U}{\partial C_t}$, must be equal to the discounted increase in utility in period $t+1$, $\beta\frac{\partial U}{\partial C_{t+1}}$, plus the real interest that would accrue on the income saved until period $t+1$ at rate r_{t+1}. This logic and the expression for the derivatives of utility with respect to C_{t+i} yields the Euler equation

$$(11)\quad \frac{1}{C_t}\left(\frac{C_t}{Z_t^\gamma}\right)^{(1-\sigma)} - \gamma(1-\sigma)(1-\rho)P_t =$$
$$\beta E_t\left[(1+r_{t+1})\frac{1}{C_{t+1}}\left(\frac{C_{t+1}}{Z_{t+1}^\gamma}\right)^{(1-\sigma)}\right] - \beta\gamma(1-\sigma)(1-\rho)E_t[(1+r_{t+1})P_{t+1}]$$

Note that when $\gamma = 0$ or $\rho = 1$, this Euler equation collapses to the familiar time-separable utility consumer's problem without habit formation. In the first case, the reference level of consumption no longer enters the utility function, so the marginal condition reduces to the standard one. In the second case, the reference level is zero, and again the consumer maximizes discounted time-separable utility of current consumption.

2.2 Deriving an Approximate Linear Consumption Function

In order to derive an explicit consumption function, I linearize the first-order conditions given in equation 11 and substitute into the linearized budget constraint. I approximate the first-order condition with its first-order expansion about steady-state values for consumption and the habit reference level, C_0 and Z_0:

$$(12)\ f(C,Z) \approx f(C_0,Z_0) + f_C(C_0,Z_0)(C-C_0) + f_Z(C_0,Z_0)(Z-Z_0) + O(n)\,,$$

where $O(n)$ represents higher-order terms. In the steady state, $Z = C$, simplifying the linearized first-order condition:

$$\text{(13)} \qquad a_1(C_t - C_0) + a_2 P_t - a_3(Z_t - C_0) + k_0 = (1+\bar{r})\beta[a_1(E_t C_{t+1} - C_0) + a_2 E_t P_{t+1} + a_3(E_t Z_{t+1} - C_0)] - \delta(1 + E_t r_{t+1})$$

where the coefficients a_i and δ are defined as $a_1 = \sigma C_0^{(1-\gamma)(1-\sigma)}$, $a_2 = \gamma(1-\sigma)(1-\rho)$, $a_3 = (1-\sigma)\gamma C_0^{(1-\gamma)(1-\sigma)}$, $\delta = \beta[-\gamma(1-\sigma)(1-\rho)\bar{P} + C_0^{(1-\gamma)(1-\sigma)-1}]$, $k_0 = C_0^{-\gamma(1-\sigma)}$.

Approximate the summation defined in P_t as

$$\text{(14)} \qquad P_t \approx C_0^{(1-\gamma)(1-\sigma)} + \frac{C_0}{Z_0^{\gamma}}^{(1-\sigma)}(C_{t+1} - C_0) - \frac{C_0}{Z_0^{\gamma}}^{(1-\sigma)} \frac{(\gamma(\sigma-1)-1)(\sigma-1)}{C_0^2}(Z_{t+1} - C_0) + \beta\rho P_{t+1}\,.$$

Utilizing the approximation in Campbell and N. Gregory Mankiw (1991), one can write the log-linearized budget constraint in consumption and income, with time-varying real interest rate r_{t+j}, as

$$\text{(15)} \qquad \begin{aligned} c_t - y_t &= \sum_{j=1}^{\infty} \mu^j (r_{t+j} - \Delta c_{t+j}) + \mu\kappa/(1-\mu) \\ &\quad + E_t \sum_{j=1}^{\infty} \mu^j (\Delta y_{t+j} - r_{t+j}) - \mu\kappa/(1-\mu) \\ &= \sum_{j=1}^{\infty} \mu^j (\Delta y_{t+j} - \Delta c_{t+j}) \end{aligned}$$

where lowercase letters denote logs. The parameter μ is the discount rate for future income (as distinguished from the real interest rate; see Campbell and Mankiw (1991)), and thus indexes the extent to which consumers look forward.

If one uses the approximation $(1+\bar{r})\beta \approx 1$ in the Euler equation, then the expected change in consumption is

$$E_t C_{t+1} - C_t = \frac{b_1}{a_1}(P_t - E_t P_{t+1}) + \frac{c_1}{a_1}(Z_t - E_t Z_{t+1}) + \frac{\delta}{a_1}(1 + E_t r_{t+1}) \tag{16}$$

Using the approximation that the change in the level of C will be proportional to the log change in C for a non-trending series (consumption is defined as per capita, less a segmented linear trend), and substituting this expression into the budget constraint, yields the approximate log-linear consumption function

$$c_t - y_t = E_t \sum_{j=1}^{\infty} \mu^j \left[\Delta y_{t+j} + a_1^*(p_{t+j+1} - p_{t+j}) + a_2^*(z_{t+j+1} - z_{t+j}) - \delta^* r_{t+j+1} \right] . \tag{17}$$

with p_t defined as

$$p_t \equiv \beta \rho E_t p_{t+1} + b_1 c_t - b_2 z_t \tag{18}$$

The parameters a_1^*, a_2^*, δ^* in equation 17 correspond to $\frac{b_1}{a_1}$, $\frac{c_1}{a_1}$, and $\frac{\delta}{a_1}$. With the steady-state value for C_0 (and thus for Z_0) set arbitrarily to unity, the values of these parameters are:

$$\begin{aligned}
a_1^* &= ((\gamma(1-\sigma)(1-\mu))/\sigma) \\
a_2^* &= ((1-\sigma)\gamma)/\sigma \\
\delta^* &= \beta[(-\gamma(1-\sigma)(1-\rho) - 1)/\sigma] \\
b_1 &= (\mu - \sigma)/(1-\sigma) \\
b_2 &= (\gamma(1-\sigma) - 1)(1-\sigma)
\end{aligned}$$

In the estimation step, δ^* is not fully constrained, that is, not all of the restrictions implied by the Euler equation are imposed. The final consumption function used in the empirical work is this equation with the addition of some rule-of-thumb consumers, as described below.

The expression for p_t captures the dependence of the marginal utility of current consumption on future consumption. This is the fundamental source of non-time-separability in the model, and it arises because the choice of a consumption level today affects the reference level of consumption in the future. With the form of the utility function employed here, this dependence takes the convenient recursive form of equation 18.

The *ex ante* real interest rate is defined as the discounted weighted average of model-consistent forecasts of short-term real interest rates, $f_t - \pi_{t+1}$, or

$$r_t \equiv (1-d) \sum_{i=0}^{\infty} d^i E_t(f_{t+i} - \pi_{t+i+1}) \tag{19}$$

where $d = \frac{D}{1+D}$, and D is the duration of the (implied) long-term real bond, which is set to ten years for this paper.

The consumption function implies that the log consumption-income ratio will be higher when (1) expected discounted income growth is higher, as in the standard PIH model, or (2) expected discounted real interest rates are lower (δ is positive for all plausible values of the underlying utility parameters). The effect of the p and z terms depends on whether the estimate of σ, the parameter that indexes the curvature of the utility function, is less than or greater than one. For values of σ greater than one (as in the estimates presented below), the higher is the expected growth in the reference level, the lower is the log consumption-income ratio. Higher reference levels lower marginal utility, because they "raise the bar" over which incremental consumption must rise to increase utility. Through the intertemporal link between current and future utility, a higher *expected* reference level lowers *current* marginal utility relative to future marginal utility, yielding more saving today (a lower current consumption-income ratio), holding income growth and real rates constant.

One key feature of the effect of expected real interest rates on the current log consumption-income ratio is worth noting. Section 1 above discusses the hump-shaped response of consumption to income in the habit formation model: the more important is habit formation, the more emphasis on smooth changes in consumption, and thus the more hump-shaped the response of consumption to income. But this smoothing motive operates with regard to real interest rates as well. Inspecting the definition of δ, one can verify that the derivative of δ with respect to γ is negative for $\sigma > 1$. Thus, for sufficient curvature of the utility function, increasing habit formation implies a more muted response of the consumption-income ratio to real interest rates.

2.3 "Rule-of-Thumb" Consumers

Campbell and Mankiw (1989, 1990, 1991) provide compelling evidence for the existence of so-called "rule-of-thumb" consumers, i.e., consumers whose current consumption equals current income. To be more precise, they provide empirical evidence that the predictable component of current income is correlated with current consumption. This constitutes a strong violation of the permanent income theory. A permanent income consumer would consume in period $t-1$ the annuity value of the component of current income that was predictable in period $t-1$. I allow for the possibility of "rule-of-thumb" consumers in the log-linear consumption function by modifying it as

$$c_t - y_t = (1-\lambda)E_t\left(\sum_{j=1}^{\infty}\mu^j[\Delta y_{t+j} + a_1^*(p_{t+j+1} - p_{t+j}) + a_2^*(z_{t+j+1} - z_{t+j}) - \delta^* r_{t+j+1}]\right) + \epsilon_{ct} \tag{20}$$

where λ represents the fraction of total income accruing to rule-of-thumb consumers (who follow the rule $c_t = y_t$), and ϵ_{ct} is the structural innovation in the consumption equation. With the income process explicitly modeled, this innovation represents transitory shifts in preference parameters.

It is important to note here a logical distinction between rule-of-thumb behavior and habit formation in consumption. Rule-of-thumb consumers respond immediately and one-for-one to the shock in current income, as well as to the predictable component of current income. Consumers with a habit formation utility function will delay some of the response to an income shock, smoothing the change in consumption. Thus, these two consumption motives are both logically and, as will be shown below, empirically distinct.

Thus specified, the model nests a number of interesting alternatives, including: the standard PIH model ($\lambda = 0, \gamma = 0$), the PIH with some rule-of-thumbers ($\gamma = 0$), a forward-looking habit formation model ($\gamma \neq 0$), as well as other combinations. In addition, the parameter μ, which is the discount factor applied to future income and the future marginal effects of current consumption decisions through habit formation, indexes the degree of forward-lookingness in the model.

2.4 Habit-Formation and the Hump-Shaped Response

It is straightforward to demonstrate the ability of the habit formation model to produce hump-shaped responses to shocks. Figure 1 examines the impulse response of consumption to a transitory but persistent income shock. The response is computed for the simple habit formation model for several values of γ, and compared with the same impulse response from a VAR. Setting $\gamma = 0$ and $\lambda = 0$ yields the simple permanent income model. For the purposes of this exercise, I use a first-order autoregressive process for income with coefficient 0.95. The fraction of income accruing to rule-of-thumb consumers, λ, is set to zero. The utility curvature parameter $\sigma = 2.0$, $\mu = 0.98$, $\rho = 0.5$, and $\beta = 0.9875$. The real interest rate term is set to zero. None of the qualitative results is sensitive to the precise values chosen for any of these parameters.

The VAR is estimated from 1966:1 to 1995:4 on quarterly data for the effective federal funds rate (quarterly average of monthly observations), log per capita chain-weighted nondurable goods and services consumption, log per capita real disposable income, log per capita non-consumption chain-weighted GDP, the log change in the chain-type price index for consumption, and the Journal of Commerce industrial materials commodity price index for all items. The consumption, income, and GDP data are detrended using a segmented linear trend with a break in 1974. The ordering used allows the funds rate to react contemporaneously to commodity prices and inflation but not to the output gap or income. The beginning of the sample is motivated by the earliest time at which the federal funds rate consistently traded above the discount rate, indicating the use of the funds rate as the primary instrument of monetary policy.[3]

The figure illustrates several important points. First, the VAR shows a clear hump in the response of nondurables and services consumption to income, with a peak at about one year. Second, the model without habit formation ($\gamma = 0$) produces no hump in the response to the income shock, as expected. The permanent income model implies an immediate jump at the time of the income shock to the maximal response, decaying geometrically thereafter. As the figure indicates, the greater the importance of habit formation (the larger the value taken by γ), the more hump-shaped the response to an income shock (shown in the dashed lines). While I will conduct more formal estimation and testing below, this pictures highlights the reason behind the empirical success of the habit formation model.

3 Empirical Results

An estimated consumption function is required that explicitly links consumption, income, and interest rates, in order to examine the monetary policy model that I develop in section 4 below. To estimate the parameters in the habit formation consumption function, I use a method akin to that used in Campbell and Robert J. Shiller (1987) for present value models. In essence, the linearized consumption function is a present value model, somewhat complicated by non-time-separability. In a manner similar to Campbell and

Shiller, I employ an unconstrained vector autoregression to generate the forecasts of the future changes in consumption, income, and real interest rates that enter on the right-hand side of the consumption function.[4]

More precisely, the linear rational expectations model that comprises the consumption function (equation 21), the definitions of the reference level variables z_t and p_t (equations 2 and 18), the *ex ante* real interest rate (equation 19), and the VAR equations for income, inflation, and interest rates may be expressed as a set of stochastic linear difference equations

$$\sum_{i=-\tau}^{0} H_i x_{t+i} + \sum_{i=1}^{\theta} H_i E_t(x_{t+i}) = \epsilon_t \,, \tag{21}$$

The maximum lag and lead in the model are denoted τ and θ, respectively, and the structural coefficients are collected in the matrices H_i. The vector of structural errors, ϵ_t, is assumed to be iid. The procedure of Gary S. Anderson and Moore (1985) allows us to solve for the expectations in terms of current and lagged variables,

$$E_t(x_{t+k}) = \sum_{i=-\tau}^{-1} B_i E_t(x_{t+k+i}), \qquad k > 0 \,. \tag{22}$$

The reduced-form solution coefficients B_i may be used to substitute out the expectations from the original structural model, equation 21, yielding a set of constrained decision rules in observable variables, with the original structural shocks identified

$$\sum_{i=-\tau}^{0} \mathbf{S_i} x_{t+i} = \epsilon_t \,. \tag{23}$$

The likelihood for this system is

$$\mathcal{L} = T(\log |\mathbf{J}| - .5 \log |\hat{\Omega}|) \tag{24}$$

where T is the sample size, $\mathcal{J}$ is the Jacobian of transformation (which is time-invariant by assumption), and Ω is the variance-covariance matrix of the structural residuals ϵ_t. The parameters of the structural models in this paper are estimated by numerical maximization of the likelihood function in equation 24. More details of the procedure are available upon request from the author.

One advantage of this approach is that it allows estimation to proceed from an unrestricted linear vector autoregression, which nests all of the models considered here, to successively more-restricted linear models. Each succeeding restriction is nested within the preceding less-restricted model and within the VAR. Once I have estimated the parameters of the consumption function, I can proceed to impose additional restriction on the VAR equations as I incorporate the consumption specification into the monetary policy model considered below.[5]

The ultimate goal of this paper will be to embed the estimated consumption function in a monetary policy model with sticky prices and sticky inflation, in order to determine to what extent the modifications to consumption entertained here alleviate the problems identified in earlier work. Thus, I begin with an unconstrained vector autoregression that includes the set of variables necessary to nest the final monetary policy model. These are nondurables and services consumption, disposable personal income, the federal funds rate, the price level, and GDP excluding nondurables and services consumption, as described in section 2.4 above.

In the first stage of estimation, I estimate only the parameters of the log-linear consumption function. The processes for income, the funds rate, prices, and other GDP are unconstrained equations from the VAR. The definitions of z_t, p_t, and *ex ante* real rates r_t are as defined in equations 2, 18, and 19 above.

3.1 Estimating and Testing the Consumption Function

Using the data described above, and estimating via maximum likelihood over the sample 1966:1 to 1995:4, I obtain the parameter estimates shown in the second column of Table 1. At the estimated parameter values, I find that: (1) habit formation is an economically important determinant in the utility function ($\gamma \neq 0$); (2) the habit formation reference level is essentially last period's consumption level ($\rho \approx 0$); (3) rule-of-thumb behavior is important, with about one-fourth of income accruing to rule-of-thumb consumers ($\lambda =$ .26); (4) the parameter indexing the curvature of the utility function, σ, is much larger than one; (5) for those who look forward, the horizon is long; the parameter μ takes the estimated value .996 on a quarterly basis, .984 on an annual basis; (6) the effect of expected real interest rates on the consumption-income ratio, δ, is negative and significant (recall that the positive coefficient is preceded by a negative sign in the linear consumption model). The sign of this not-fully-constrained real rate effect is the same as the sign of the fully-constrained coefficient at the estimates for σ, γ, and β, although the magnitude is somewhat larger; and (7) the model explains most, but not all, of the autocorrelation in the consumption data, as evidenced by the low $p-$value for the Ljung-Box test for serial correlation in the first 12 residual autocorrelations in the consumption equation.[6]

Table 1 also reports alternative estimates for the parameters of the linearized consumption function obtained from a nonlinear Generalized Method of Moments (GMM) estimator. As the table indicates, the GMM estimates are quite similar to those obtained via maximum likelihood, although the standard errors are uniformly larger (as expected given the diminished efficiency afforded by the GMM estimator). Here the habit formation parameter γ is estimated at 0.9, a bit higher than the Full Information Maximum Likelihood (FIML) estimate. The estimated fraction of income accruing to rule-of-thumb consumers, λ, is nearly identical to the FIML estimate. The parameter indexing the curvature of the utility function, σ, is large and not significantly different from the FIML estimate given the precision of the estimate. The habit "memory" parameter ρ is again estimated to be nearly zero (see footnote 7 for estimation details).[7]

In the maximum likelihood estimates, the structural consumption equation error has one significant autocorrelation of about .53. The standard errors reported in Table 1 are corrected for the estimated correlation in this error. However, all of the impulse response, likelihood ratio, autocorrelation function, and simulations reported below assume the errors to be white. That is, *none* of the dynamics in the impulse responses or other results reported below may be attributed to across-time correlation in the error terms.

The low estimated value of the parameter that indexes the "memory" in the habit reference level, ρ, suggests that the operative reference level is last quarter's consumption. One presumes that habits are formed over horizons longer than one quarter, so this estimate of ρ is perhaps lower than expected. However, the estimate can be justified on several grounds. First, if the level of (detrended) consumption exhibits significant autocorrelation, then last period's consumption contains information about consumption in previous periods. So the lagged level of consumption may have considerable "memory" itself.

Second, note that rewriting the period utility function as

$$U_t = \left(\frac{C_t}{Z_t} Z_t^{(1-\gamma)} \right)^{1-\sigma}$$

and setting $\rho = 0$ yields the special case discussed above

$$U_t = \left(\frac{C_t}{C_{t-1}} C_{t-1}^{(1-\gamma)} \right)^{1-\sigma}$$

This form of the utility function distills the essence of habit formation. Habit formation mixes utility from the level of consumption with utility from the *change* in consumption. That is, the habit formation model with any normally shaped utility function will imply smoothing of both the level of consumption and its changes (provided γ is not zero). Larger values of ρ simply define the changes relative to a longer distributed lag of past consumption.[8]

Seen in this light, it becomes clear that a single lag of consumption in the reference level may be sufficient to impart the smoothness to changes in consumption expenditures that is absent in the standard permanent income model. In addition, note that the linearized consumption function with $\rho = 0$ is

$$c_t - y_t = E_t \sum_{j=1}^{\infty} \mu^j \left[\Delta y_{t+j} + a_1^*(p_{t+j+1} - p_{t+j}) + a_2^*(c_{t+j} - c_{t+j-1}) - \delta^* r_{t+j+1} \right] .$$

The third term on the right-hand side, the weighted sum of expected future changes in consumption, will differ relatively little from the weighted sum of expected future deviations of consumption from a moving average of past consumption (the corresponding term in the consumption function with $\rho \neq 0$). The difference will manifest itself for the most part in a small difference in the weights on future consumption changes. In essence, this specification of the habit formation model builds enough linkage between current consumption and future changes in consumption with or without a long memory in the reference level.[9]

Obtaining sensible parameter estimates is a necessary but not a sufficient condition for obtaining a reliable model for monetary policy analysis. Figures 2 and 3 provide a more complete picture of the dynamic interactions implied by the habit formation model. Figure 2 examines the extent to which the habit formation model with time-varying *ex ante* real interest rates can match the hump-shaped impulse responses in the VAR. Figure 3 displays the autocovariances for the unconstrained VAR (the solid lines) and the constrained consumption function (the dashed lines). The autocovariance function simply summarizes the covariances among variables across time that are implied by a model. An advantage of the autocovariances is that they do not require the same identifying assumptions and orthogonalization across models for comparability.

As the figures show, the model reflects the VAR's dynamics for consumption expenditures quite well, without doing damage to the interactions among other variables. The impulse responses of consumption to all four shocks match quite well. The hump-shaped responses of consumption to income and interest rate shocks (as well as the preference shock) match those exhibited by the VAR. However, the caveat in footnote 3 about strict comparability of these two impulse responses certainly applies.[10]

The autocorrelation function in figure 3 shows that the model implies sensible and data-consistent dynamic correlations, capturing the persistence in the own correlation of consumption, as well as the persistent dynamic correlations between consumption and income, interest rates, and inflation. The lighter dotted lines in Figure 3 display the 90% confidence intervals around the VAR's vector autocovariance function. As the plot shows, the differences between the two autocorrelation functions are generally insignificant at the 10% level. Thus the correlations that the structural model cannot match perfectly are generally not precisely determined in the data. In the following section, I perform a series of likelihood ratio tests to determine the statistical significance of a variety of restrictions on the model.[11]

3.2 Nested Tests of Habit Formation and Rule-of-Thumb Behavior

The hypothesis that habit formation is unimportant in this model—that the exponent γ on the reference level of consumption is zero—is overwhelmingly rejected. The χ^2 likelihood ratio test for this single restriction takes the value 21.4, with $p-$value $4X10^{-6}$. Similarly, the hypothesis that rule-of-thumb behavior is unimportant is strongly rejected. The χ^2 likelihood ratio test for the restriction $\lambda = 0$ takes the value 12.6, with $p-$value $= 4X10^{-4}$. It is interesting to note that the likelihood ratio test for the constrained baseline model, which incorporates the many zero restrictions and cross-equation restrictions implied by the structure of the consumption model and by rational expectations, takes the value 32.8, not significant at even the 10 percent level. This is one of relatively few cases in which the joint restrictions imposed by an optimization-based model with rational expectations cannot be rejected relative to the unconstrained model in which the constrained model is nested.

The vector autocovariance function illustrates the importance of habit formation and rule-of-thumb behavior in replicating the dynamic interactions among consumption, income, interest rates, and inflation. As Figure 4 shows, the primary consumption dynamics in the model with neither habit formation nor rule-of thumb behavior (setting γ and λ to zero) are almost totally missing. Note also in the bottom row of figure 4 that the absence of habit formation and rule-of-thumb behavior has caused a deterioration in the relationship between inflation and consumption or income. This autocovariance comparison reinforces and extends the insight gained from the illustrative impulse responses in figure 1. The simple PIH model cannot replicate the dynamics in the spending and inflation data. Both rule-of-thumb behavior and habit formation are statistically significant modifications to add to the model, and they represent both logically and empirically distinct additions.

3.3 Caveats

A few caveats about data and methodology are in order (in addition to the measurement difficulties noted above and emphasized by David W. Wilcox (1992)). First, note that the use of detrended income and consumption data, while in keeping with convention for most business cycle models, does not adhere to the letter of the theory. Consumers' utility depends on the level of consumption, not the level less a deterministic trend. In addition, the break in the trend in 1974 would not have been known immediately by consumers forecasting detrended income.

Second, I have glossed over a potentially serious aggregation issue. Even if individual consumers exhibit habit formation, they can easily differ in their rates of time preference, degree of habit formation, or degree of risk aversion. Similarly, they presumably are endowed with different levels of wealth, and are subject to different shocks to their streams of resources. All of this unaccounted-for heterogeneity renders the aggregation implicit in the empirical model imperfect at best.

Of additional concern in this vein is the contrast between the relatively strong evidence in favor of habit formation in the aggregate data and the failure to develop support for habit formation in Karen E. Dynan's (1999) microdata study. Her analysis of food consumption in the PSID finds no evidence of habit formation. Whether this discrepancy arises because food is not representative of other consumption, or because aggregation over individuals imparts to aggregate consumption the appearance of habit formation, is unclear. Serious consideration of these issues lies outside the scope of this paper.

4 Towards a Model for Monetary Policy Analysis

I now progressively add restrictions to the unconstrained VAR equations in the model (the reduced-form equations for inflation, short-term interest rates, and income) in order to identify the systematic component of monetary policy and the pricing decisions of firms. I begin with the monetary policy function, imposing zero restrictions to the reduced-form funds rate equation so that it takes the form of a simple Taylor rule (1993).

$$(25)\ f_t = (1-\sum \alpha_{f,i})(\bar{r}+\bar{\pi})+\sum_{i=1}^{2} \alpha_{f,i} f_{t-i}+\sum_{j=0}^{2} \alpha_{\pi,j}(\pi_{t-j}-\bar{\pi})+\sum_{k=0}^{2} \alpha_{y,k} y_{t-k}+\epsilon_{ft}$$

where $\bar{r}$ and $\bar{\pi}$ are the equilibrium real interest rate and the inflation target, respectively. These simple restrictions yield insignificant deterioration of the likelihood from its baseline model value, and the vector autocovariance function shows little sign that the imposition of the Taylor rule on the model has constrained the dynamics in an economically significant way.

The reaction function is estimated from 1966:1-1995:4. The estimated parameters are

$$\begin{aligned} f_t \;=\; & 1.1 f_{t-1} - 0.18 f_{t-2} + 0.26\pi_t - 0.40\pi_{t-1} + 0.25\pi_{t-2} \\ & + 0.31(y_t^o + c_t) + 0.04(y_{t-1}^o + c_{t-1}) - 0.26(y_{t-2}^o + c_{t-2})\,. \end{aligned}$$

Over this sample, the fit in a static simulation is quite good, showing no obvious signs of instability. However, numerous authors have suggested and documented at least one break in the reaction function after October 1979. As a robustness check, I compute the vector autocovariance function and impulse responses for the same model, but substituting a reaction function estimated from 1980 to 1998. The differences are quantitatively small and statistically insignificant. Relative to the rather large differences between spending equations that capture a hump-shaped response as opposed to a jump response, the shifts in reaction functions across time are not important.

The second step is to constrain the price process. I begin by using a very simple version of a Fuhrer-Moore contracting model, which can be shown to be equivalent to a two-sided inflation specification (see Fuhrer and Moore

(1992, 1995a), Roberts (1997)):

$$\pi_t = (1/2)\sum_{i=1}^{k}(\pi_{t-i} + \pi_{t+i}) + \alpha y_t + \epsilon_{pt}\,. \tag{26}$$

k is set to 3 to correspond to the optimal lag lengths chosen in the unconstrained VAR. This additional set of restrictions does not significantly deteriorate the likelihood from the baseline model's likelihood value. In addition, further constraining the price dynamics exactly as in Fuhrer and Moore (1995a), with explicit nominal price contracts, does not cause a statistically significant deterioration in the likelihood.

Finally, I allow the non-consumption components of GDP to enter the model. The importance of this addition is that the funds rate in the policy reaction function can now respond to the total GDP gap, rather than just consumption of nondurable goods and services. In addition, the overall GDP gap enters as the excess demand term in the contract price specification. Other GDP is entered as in the earlier "I–S" specification of Fuhrer and Moore (1995b). That is, the gap between non-consumption GDP and its trend depends positively on its own lag and negatively on the difference between the *ex ante* long-term (model-consistent) real rate and its equilibrium:

$$\tilde{y}^o_t = \omega \tilde{y}^o_{t-1} - y_r(r_{t-1} - \bar{r}) + \epsilon_{yt}\,. \tag{27}$$

The addition of this equation and of the feedback of total GDP into interest rate and price determination does not significantly deteriorate the likelihood.

Figure 5 compares the vector autocovariance function for this more fully constrained (and identified) model with the unconstrained VAR autocovariance function. As the figure indicates, the constrained model largely replicates the dynamic behavior of the unconstrained VAR. Of course, the constrained model cannot perfectly replicate unconstrained dynamic behavior. For example, while the correlation between consumption and the lagged funds rate or lagged inflation is negative, it is too strongly so. In addition, the correlation between the funds rate and lagged consumption is negative, while the VAR says it should be mildly positive. Recall from figure 3, however, that these dynamic correlations are not so precisely determined in the VAR that the differences between the constrained model and the VAR are significant. These autocorrelation comparisons provide graphical verification of the likelihood ratio tests conducted above.

5 Monetary Policy Implications of the Model

An alternative interpretation of the results of this paper is that the restrictions imposed on the price specification and the funds rate reaction function are invalid, and are interfering with the real-side dynamics of consumption and output. To test this possibility, I estimate a model with reduced-form processes for consumption and income, so that only the restrictions from the price and interest rate specifications constrain the model. This model allows us to isolate the effects of these restrictions.[12]

A comparison of the autocovariance function for this model (Figure 6) with the unconstrained autocovariance function suggests that this interpretation is invalid. The Fuhrer-Moore price specification and the simple reaction function capture the dynamics in these variables without distorting their dynamic interactions with consumption and income (or vice versa).

It is the case, however, that improper specification of the real side of the model can seriously distort the dynamics of inflation and nominal interest rates. This should not come as a surprise, given the structural links between real output and inflation in many price specifications, and given the assumed response of nominal interest rates to real output in the policy reaction function. Figure 7 below provides an example of such a case.

5.1 Disinflation in the Models

In Fuhrer (1997), a disinflation simulation highlighted the "front-loaded" responses of real variables to a monetary shock. This information is, of course, contained in the vector autocovariance function, but the simple disinflation simulation makes a key problem of the specification quite clear. Note that the specification *including* rule-of-thumb consumers still exhibits rapid response to shocks; I wish to determine whether the addition of habit formation improves this counterfactual behavior in the model.

The simulation is straightforward. Starting from a steady state, I decrease the long-run inflation target from 3% to 0%. The decrease is unanticipated. It is informative to compare the response of the model without habit formation to the model that includes it. Figure 7 displays the results of the simulation.

In the model with habit formation, shown in the dashed lines, inflation (the top panel) falls gradually from its old steady state to the new, lower equilibrium. Interestingly, consumption, in the bottom panel, also responds gradually, with its peak response at a year or so; the full response takes three to four years. This response contrasts markedly with that of the model excluding habit formation–although the model includes rule-of-thumb consumers–shown in the solid lines in the figure (from Figure 4 in Fuhrer (1997)). Thus the solid lines demonstrate that the PIH model with rule-of-thumb consumers not only misrepresents consumption dynamics, it also compromises the behavior of inflation. The persistence of inflation in this model, as indicated in the solid line in the top panel of figure 7, is significantly decreased by the rapid (and counterfactual) response of real variables to a disinflationary shock.

Figure 8 assesses the impact of the length of "memory" in the habit reference level on the model's dynamics. In this figure, the same disinflation simulation is performed, substituting a value of .9 for ρ (recall that the estimated value is .01). As a comparison of figures 7 and 8 shows, the model's behavior is altered only slightly by the change from one-quarter memory to much more persistent memory in the reference level. This simulation confirms the intuition developed in section 3.

Taking all of the empirical results together, one can conclude the following:

1. Habit formation is a statistically significant and economically important feature of aggregate macroeconomic data. It captures an element of consumer behavior that is logically and empirically distinct from rule of thumb behavior.

2. The inclusion of habit formation in a monetary policy model improves the dynamics of *both* real spending and inflation, without worsening any other dynamic interactions in the model. The improvements are significant both statistically and economically.

3. Excluding habit formation significantly worsens both real spending and inflation dynamics.

6 Conclusions

A model for monetary policy analysis should be closely related to the underlying objectives of consumers and firms, should explicitly model expectations, and should capture the dynamic interactions among variables that are exhibited in the data. While many recently developed models explicitly model expectations, and build close ties to underlying agents' objectives, many simple optimization-based macroeconomic models fail to replicate economically important and statistically significant dynamic correlations in the data. A direct implication of these models' failure to replicate key dynamic correlations is that the models are unlikely to represent agents' dynamic behavioral decisions. As a result, such models may not be suitable for monetary policy analysis.

This paper makes some progress towards a model that meets the standards itemized above. It does so by including a particular form of non-time-separability in the utility function, namely "habit formation," or the assessment by consumers of utility relative to a habit level of consumption. The paper develops evidence that shows that augmenting the model in this way allows the model to replicate key dynamic correlations among consumption, output, interest rates, and inflation to a degree that standard models cannot. In particular, the model can match the hump-shaped response of consumption to income, interest rate, and inflation shocks. The habit formation specification improves upon the standard specification because it imparts a motive for consumers to smooth the *change*, as well as the level of consumption. Another improvement afforded by the smooth response of consumption is that the model implies a more realistic and data-consistent gradual decline in inflation during a disinflation.

Other specifications may also afford improvements in the empirical performance of the standard model. However, specifications that rely on costs of adjustment do not apply to nondurables and services consumption on theoretical grounds. Recent work has found that models that incorporate uncertainty do not produce the required smoothness in the response of consumption to income shocks. Finally, work in asset pricing and growth has found corroborating evidence for the importance of habit formation in consumer behavior. The gradual or hump-shaped response of consumption to shocks that is found in reduced-form and other empirical studies is a statistically significant feature of the data. This feature must be incorporated in monetary policy models that wish to accurately reflect the gradual responses of spending and inflation to monetary policy actions.

Appendix
Accuracy of the Linear Approximation

For computational tractability, all of the computations reported above depend on the linearized approximate consumption function. An important question is how well the linear approximation reflects the underlying nonlinear model from which it is derived.

I examine several measures of the approximation's accuracy. First, I solve the nonlinear model (substituting equation 11 for equation 17), using the parameters estimated from the linear model, for the standard disinflation simulation of the previous section. I obtain nearly identical results.[13]

In addition, substituting the linear model's solutions for consumption, income, and real interest rates into the nonlinear first-order conditions, I find that they hold quite well. The maximum absolute error in the nonlinear Euler equations is about .01, compared to steady-state marginal utility of about -1. Finally, the estimate of lifetime utility for the disinflation simulation is very similar whether computed using the solution paths from the linear model or from the nonlinear model.

Overall, then, it appears that the linear model provides a very good approximation to the behavior implied by the nonlinear model.

References

Abel, Andrew B. "Asset prices under habit formation and catching up with the Joneses." *American Economic Review*, May 1990, vol. **80**(2), 38–42.

Anderson, Gary S. and Moore, George R., "A linear algebraic procedure for solving linear perfect foresight models." *Economics Letters*, 1985, vol. **17**, 247–252.

Attanasio, Orazio P. and Picci, L. and Scorcu, A. "Savings, Growth, and Investment: A Macroeconomic Analysis Using a Panel of Countries." Forthcoming, *World Bank Policy Research Working Papers*, 1998.

Bernanke, Ben S. and Blinder, Alan S. "The federal funds rate and the channels of monetary transmission." *American Economic Review*, September 1992, vol. **82(4)**, 901–921.

Campbell, John Y. and Cochrane, John H. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, April 1999, vol. **107(2)**, 205–51.

Campbell, John Y. and Deaton, Angus S., "Why Is Consumption So Smooth?" *Review of Economic Studies*, July 1989, vol. **56(3)**, 357–73.

Campbell, John Y. and Mankiw, N. Gregory, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." O. Blanchard and S. Fischer, *NBER Macroeconomics Annual*. Cambridge: MIT Press, 1989, 185–216.

Campbell, John Y. and Mankiw, N. Gregory. "Permanent Income, Current Income, and Consumption." *Journal of Business and Economic Statistics*, July 1990, vol. **8**(3), 265–279.

Campbell, John Y. and Mankiw, N. Gregory, "The response of consumption to income: a cross-country investigation." *European Economic Review*, May 1991, vol. **35**(4), 723–756.

Campbell, John Y. and Shiller, Robert J., "Cointegration and tests of present value models." *Journal of Political Economy*, October 1987, vol. **95(5)**, 1062–1088.

Carroll, Christopher D. and Weil, David N. "Saving and Growth: A Reinterpretation." *Carnegie-Rochester Conference Series on Public Policy*, June 1994, vol. **40**, 133–92.

Carroll, Christopher D., Overland, Jody R., and Weil, David N., "Saving and Growth with Habit Formation." FEDS Working paper #95–42, September 1995.

Christiano, Lawrence J., Eichenbaum, Martin S., and Evans, Charles L., "Identification and the Effects of Monetary Policy Shocks." Blejer, Mario I, and Eckstein, Evi and Leiderman, Leonardo, eds., *Financial Factors in Economics Stabilization and Growth.* Cambridge: Cambridge University Press, 1996, 36–74.

Constantinides, George M. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy*, June 1990, vol. **98(3)**, 519–43.

Duesenberry, James S., *Income, Saving, and the Theory of Consumer Behavior.* New York: Oxford University Press, 1949.

Dynan, Karen E., "Habit Formation in Consumer Preferences: Evidence from Panel Data." forthcoming, *American Economic Review*, July 1999.

Estrella, Arturo and Fuhrer, Jeffrey C., "Dynamic inconsistencies: Counterfactual implications of a class of rational expectations models." Prepared for the NBER Universities Research Conference, December 4-5, Cambridge, MA, December 1998.

Fuhrer, Jeffrey C., "Towards a Compact, Empirically-Verified Rational Expectations Model for Monetary Policy Analysis." *Carnegie-Rochester Conference Series on Public Policy*, November 1997, vol. **47**, 197–230.

Fuhrer, Jeffrey C. and Moore, George R., "Monetary Policy Rules and the Indicator Properties of Asset Prices." *Journal of Monetary Economics*, April 1992, vol. **29**(2), 303–336.

Fuhrer, Jeffrey C. and Moore, George R., "Inflation Persistence." *Quarterly Journal of Economics*, February 1995a, vol. **110**(1), 127–159.

Fuhrer, Jeffrey C. and Moore, George R., "Monetary Policy Trade-offs and the Correlation Between Nominal Interest Rates and Real Output." *American Economic Review*, March 1995b, vol. **85**(1), 219–239.

Fuhrer, Jeffrey C., Moore, George R., and Schuh, Scott D., "Estimating the Linear-Quadratic Inventory Model: Maximum Likelihood versus Generalized Method of Moments." *Journal of Monetary Economics*, February 1995, vol. **35**(1), 115–157.

Hall, Robert E. "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, December 1978, vol. **86(6)**, 971–87.

Jermann, Urban J. "Asset Pricing in Production Economies." *Journal of Monetary Economics*, April 1998, vol. **42(2)**, 257–75.

King, Robert G. and Wolman, Alexander L., "Inflation Targeting in a St. Louis Model of the 21st Century." NBER Working Paper #5507, March 1996.

Leeper, Eric M., Sims, Christopher A., and Zha, Tao, "What Does Monetary Policy Do?" *Brookings Papers on Economic Activity* 1996, vol. **2**, 1–63.

Ljungqvist, Lars and Harald Uhlig, "Tax Policy and Aggregate Demand Management under Catching Up with the Joneses." forthcoming, *American Economic Review*, July 1999.

Loayza, N. and Schmidt-Hebbel, Klaus and Servén, Luis "What Drives Saving Across the World?" Forthcoming, *World Bank Policy Research Working Papers*, 1998.

Lucas, Robert E. Jr., "Econometric policy evaluation: a critique." *Carnegie-Rochester Conference Series on Public Policy*, 1976, vol. **1**, 7–33.

Ludvigson, Sydney C. and Alexander Michaelides, "Can Buffer Stock Saving Explain the 'Consumption Excesses'?" Working Paper, Federal Reserve Bank of New York, 1998.

McCallum, Bennett T. and Nelson, Edward, "An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis." *Journal of Money, Credit, and Banking*, August 1999a, vol. **31 (3)**, Part 1, 296–316.

McCallum, Bennett T. and Nelson, Edward, "Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model," in Taylor, John B., Monetary Policy Rules, Chicago: University of Chicago Press, 1999b, 15–45.

Roberts, John M., "Is Inflation Sticky?" *Journal of Monetary Economics*, July 1997, vol. **39**(2), 173–196.

Rodrik, Dani, "A Study of Transitions from Low to High Saving Rates." Forthcoming, *World Bank Policy Research Working Papers*, 1998.

Romer, Christina D. and Romer, David H., "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." Olivier J. Blanchard and Stanley Fischer, *NBER Macroeconomics Annual*. Cambridge: MIT Press, 1989, 121–170.

Rotemberg, Julio R. and Woodford, Michael, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." Ben S. Bernanke and Julio J. Rotemberg, *NBER Macroeconomics Annual*. Cambridge: MIT Press, 1997, 297–346.

Taylor, John B., "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy*, February 1980, vol. **88(1)**, 1–23.

Taylor, John B., "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, December 1993, vol. **39**, 195–214.

West, Kenneth D. and Wilcox, David W., "Some Evidence on Finite Sample Behavior of an Instrumental Variables Estimator of the Linear Quadratic Inventory Model." Finance and Economics Discussion Paper #93–29, August 1993.

Wilcox, David W. "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work," *American Economic Review*, September 1992, vol. **82(4)**, 922–41.

Table 1

Estimation Results

	FIML Estimates		GMM Estimates	
Coefficient	Estimate	Standard Error	Estimate	Standard Error
γ	0.80	0.19	0.90	1.83
ρ	0.0015	0.0039	0.052	7.41
λ	0.26	0.13	0.29	0.37
σ	6.11	1.81	13.02	7.85
μ	0.99	0.01	0.99	0.11
δ	28.49	5.17	94.89	8.00
Ljung-Box Q(12)	($p-$value)		Hansen's J-Test: 4.99 (df=11)	
Consumption	85.3 (.00)		p-value = 0.93	
Income	10.7 (.56)			
Funds rate	25.8 (.01)			
Inflation	8.2 (.77)			
Log-likelihood	2366.4			

Error correlations for consumption equation

Lag	Autocorrelation	Partial Correl.
1	0.53 (0.13)	0.53 (0.09)
2	0.37 (0.13)	0.12 (0.09)
3	0.40 (0.13)	0.23 (0.09)
4	0.14 (0.13)	-0.25 (0.09)

Asymptotic standard errors in parentheses

Footnotes

* Vice President and Economist, Research Department, Federal Reserve Bank of Boston, Boston, MA, 02106. e-mail: Jeff.Fuhrer@bos.frb.org. I thank Chris Carroll, Michael Klein, Richard Kopcke, Giovanni Olivei, Scott Schuh, Christopher Sims, Geoff Tootell, seminar participants at Johns Hopkins University and Columbia University, and anonymous referees and the editor for their helpful comments, and Hoyt Bleakley and Ann Ferris for excellent research assistance. The views expressed in this paper are the author's and do not necessarily reflect those of the Federal Reserve Bank of Boston.

1. Arturo E. Estrella and Fuhrer (1998) discuss the problem of jump variables in monetary policy models in more detail.

2. Of course, definitional problems with the NIPA data suggest some caution in this regard. Some nondurable goods are in fact somewhat durable, and some services are tied to the ownership of a home (e.g. the imputed service flow from housing and expenditures on household operation, including electricity and heating fuels). See Wilcox (1992) for a comprehensive treatment of problems with NIPA consumption data.

3. Lag lengths for the VAR are chosen according to conventional criteria. The results displayed in Figure 1 are not sensitive to the ordering chosen for the VAR. Note that the impulse responses from the VAR and the structural model are not strictly comparable. Here, the impulse responses are constructed by computing the approximate reduced form for the structural model from the vector autocovariance function, ordering and orthogonalizing the model's reduced form as in the VAR, and computing the resulting impulse responses. I thank Christopher Sims

for pointing out a computational method for obtaining the VAR representation. Note that this computation will only yield an approximate VAR representation of the structural model. For the monetary policy models presented later in the paper, the approximation will be less accurate, due to the expectational identities that define the contract price. As shown in Fuhrer and Moore (1995a), the observable "Phillips curve" representation of the real contract price model contains an infinite moving average of the excess demand term (see equation 25 in that paper). As a result, the structural model is not formally nested within a finite-lag unconstrained VAR.

4. In fact, the forecasts required to compute future p_t, z_t, and Δy_t are solved for using the method of Anderson and Moore (1985), and the likelihood computed using the numerical maximum likelihood method which is documented in Fuhrer and Moore (1995a). In this way, all of the rational expectations restrictions implied by the consumption model are imposed in estimation.

5. An additional advantage of this method is that its finite sample properties may be more desirable than method-of-moments estimators, as documented in Fuhrer, Moore, and Scott D. Schuh (1995) and Kenneth D. West and Wilcox (1993). Of course, a potential drawback to the approach is that, to the extent that any equation in the system is misspecified, estimates of all the parameters in the system will (in principle) be affected. However, I pursue an estimation strategy below that is designed to minimize the exposure to this risk.

6. Note that the real data are detrended using a segmented linear trend. Estrella and Fuhrer (1998) show that the impulse response and autocovariance properties in these data do not depend strongly on the detrending method employed, which include trends derived from the

Hodrick-Prescott filter, a band-pass filter, an hours-based trend, and the Congressional Budget Office's estimate of potential output.

7. The linearized consumption function is given in equation 21. I substitute for the definitions of p_t, z_t, and r_t to obtain a sixth-order expectational difference equation. I employ a standard nonlinear GMM estimator that minimizes the criterion $uZWZ'u'$, where u is the innovation in the linearized consumption function, Z is the instrument matrix, and W is a consistent estimate of the optimal weight matrix. Instruments include lags two through four of detrended disposable personal income, inflation (derived from the chain-weight consumption price index), and the federal funds rate, current and lagged relative oil prices, and four lags of chain-weighted government defense expenditures.

8. In this sense, habit formation may provide a reasonable approximation to a model with a standard utility function and costs of adjustment in ΔC_t.

9. The simulation presented in Figure 8 and in section 5 below shows the effect of a longer-memory reference level on a standard disinflation simulation.

10. Note the presence of a "price puzzle" in the estimated VAR impulse responses. This arises despite the presence of a commodity price index in the VAR. The puzzle arise for all orderings for the impulse responses, and regardless of the inclusion of either commodity price inflation or relative commodity prices. The 90% confidence region indicates that the positive response of inflation to a funds rate shock is significant for two quarters. Interestingly, substitution of the inflation rate in the GDP chain-type price index eliminates the problem. Estimation of the consumption parameters substituting GDP for CPI inflation yields

nearly identical results. Further investigation of this puzzle lies outside the scope of this paper.

11. The confidence intervals for the impulse response and the autocovariance functions are computed as follows. I assume the distribution of coefficient estimates to be asymptotically normal. I employ a Monte Carlo technique that draws a random vector of coefficient estimates from the multivariate normal distribution centered on the sample estimates, with covariance matrix as estimated from the sample. For each vector of estimates, I compute the corresponding impulse responses and vector autocovariance function, holding the residual covariance matrix fixed. I use 10,000 replications of this process. The 90% confidence intervals are bounded by the 5th and 95th percentiles of the ranked responses and autocovariance functions.

12. The converse of this test is performed above: The model with restrictions on consumption, but without restrictions on prices and interest rates, requires rule-of-thumb and habit formation behavior to match the moments in the data.

13. Note that for the nonlinear solution exercises presented here, I use a "certainty equivalence" solution technique that does not compute the stochastic distribution of the endogenous variables via value-function programming. The state dimension of the model would make the computation time for such a method prohibitive.