

AN OPTIMISING MODEL FOR MONETARY POLICY ANALYSIS: CAN HABIT FORMATION HELP?

Jeffrey C. Fuhrer*

Research Discussion Paper
9812

September 1998

Reserve Bank of Australia

*Vice President and Economist, Research Department, Federal Reserve Bank of Boston, Boston, MA, 02106; Visiting Economist, Economic Research Department, Reserve Bank of Australia, September 1998; e-mail: Jeff.Fuhrer@bos.frb.org. I thank Michael Klein, Richard Kopcke, Scott Schuh, and Geoff Tootell, and seminar participants at the Reserve Bank of Australia, Johns Hopkins University and Columbia University for their helpful comments, and Hoyt Bleakley for his 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, or the Reserve Bank of Australia.

Abstract

In earlier work (Fuhrer 1997a), I document what I view as the failure of standard models of representative consumer and firm behaviour to replicate the dynamics that we observe in the aggregate data. In essence, these models fail because they imply that both inflation and real variables must ‘jump’ in response to monetary policy (and other) shocks, in contrast to identified VAR evidence that shows a gradual, ‘hump shaped’ response. This paper discusses a rigorous empirical standard for monetary policy models. The motivation for this discussion is that, if one wishes to conduct welfare analysis, one must be reasonably confident that the model provides a good approximation to underlying consumer and firm behaviour over the monetary policy horizon, i.e. in the short run. The paper examines a specific alternative to the standard consumption model in which consumers’ utility depends in part on current consumption relative to past consumption. This formulation of habit formation allows one to nest habit formation, life-cycle consumption, and Campbell and Mankiw’s ‘rule of thumb’ consumers within a more general model. The empirical tests developed in the paper show that one can reject the hypothesis of no habit formation with tremendous confidence. This result suggests that models that are unable to produce a hump-shaped response will be strongly rejected empirically.

JEL Classification Numbers: E43, E52

Keywords: consumer behaviour, habit formation, monetary policy

Table of Contents

1.	Introduction	1
2.	Problems with Standard Models	4
2.1	A Non-behavioural Solution to the Problem	5
3.	A Simple Habit Formation Model	6
3.1	A Linearised Consumption Function	8
3.2	Estimating and Testing the Consumption Function	9
4.	Empirical Results	10
4.1	Nested Tests of Habit Formation and Rule-of-thumb Behaviour	15
4.2	Adding Restrictions to the Model	17
5.	Monetary Policy Implications of the Model	20
5.1	Disinflation in the Models	22
6.	Accuracy of the Linear Approximation	24
7.	Welfare Considerations	26
8.	Conclusions	27
	Appendix A: First-order Conditions for the Nonlinear Model	29
	Appendix B: Deriving an Approximate Linear Consumption Function	31
	References	33

AN OPTIMISING MODEL FOR MONETARY POLICY ANALYSIS: CAN HABIT FORMATION HELP?

Jeffrey C. Fuhrer

1. Introduction

With the resurgence of interest in the effects of monetary policy on the macroeconomy, led by the work of the Romers (1989), Bernanke and Blinder (1992), Christiano, Eichenbaum and Evans (1994), 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 behaviour, so that changes in policy (or private) behaviour could cause shifts in reduced-form parameters (i.e. the critique of Lucas 1976). Others incorporate expectations, but derive key relationships from *ad hoc* behavioural assumptions, rather than from explicit optimising problems for consumers and firms.

Explicit expectations and optimising behaviour are both desirable, other things equal, for a model of monetary analysis. First, analysing the 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 modelling of welfare. This of course implies that the model be closely tied to the presumed objectives of consumers and firms, hence the emphasis on optimisation-based models.

A number of recent papers (see, for example, King and Wolman (1996); McCallum and Nelson (1997, 1998); Rotemberg and Woodford (1997)) have begun to develop models that incorporate explicit expectations, optimising behaviour, and an

economy in which monetary policy has real effects. However, to date, I would argue that most of these efforts have not been very successful. In essence, their failure arises from inadequate empirical validation of the restrictions imposed by the model.

It is certainly not the case that *any* model based on an optimisation problem with rational expectations will be a good candidate for use in monetary policy analysis. In particular, in earlier work (Fuhrer 1997a), I document what I view as the failure of simple standard optimising models to adequately mimic the dynamics found in the data for key variables. If a model fails significantly at matching these dynamics, it becomes much harder to claim that the model represents the underlying behaviour of consumers and firms. One therefore cannot trust the model's welfare rankings across alternative monetary policy strategies.

In order to identify models whose underlying assumptions reasonably approximate the objectives and decisions of consumers and firms, one must carefully test the model's implications for the dynamic evolution of key variables against the behaviour of these variables in the data. For monetary policy analysis, it is not enough to match first and second unconditional moments, or a subset of conditional moments implied by the model, to those implied by the data, as in any number of early equilibrium business cycle studies (see, for example, King and Plosser (1984); King, Plosser and Rebelo (1988); Christiano and Eichenbaum (1990)). The working assumption among most economists is that monetary policy has only short-run effects on real variables. If so, it would be a major omission not to fully evaluate the short-run dynamic effects of monetary policy in a candidate model.

Similarly, it is not sufficient in general to match a model's impulse response for a single shock to that in a reduced-form model (see, for example, Rotemberg and Woodford (1997) and my comments in the same volume). While such a procedure might be used to generate consistent parameter estimates given the model's restrictions, it does not validate the model's ability to match the dynamic behaviour of its key variables to the data. In particular, relying on the model's response to a monetary policy shock (typically a federal funds rate shock) may be quite misleading. As Leeper, Sims and Zha (1996) have shown, the fraction of the variance of output, inflation, or interest rates accounted for by the unanticipated component of monetary policy is generally quite small. By using a single impulse response to assess the validity of the model, especially a response to a monetary

policy shock, one is restricting oneself to a subset of the full range of dynamic behaviours implied by the model. The use of such a metric can be quite misleading.

I advocate, and in this paper I implement, the use of likelihood-based evaluation criteria for distinguishing among models. The simple rationale is that the likelihood incorporates *all* of the dynamic covariances among observable variables, weighted according to their contribution to the likelihood. It should as a result be less subject to the criticisms levied above against less formal evaluation criteria. As a graphical tool, I often compare the data's and the model's vector autocovariance functions (ACF). To be more precise, I compare the ACF implied by an unconstrained vector autoregression with the ACF of the constrained structural models that are nested within the VAR. This often reveals important differences in model and data dynamics that underlie statistically significant differences in the likelihood. Thus, the ACF may highlight visually a behavioural deficiency in the model that is more difficult to interpret from statistical evidence.¹

In the next section, I review the evidence on the deficiencies of several standard sticky-price models in the literature. One striking failure is the inability of a life-cycle model of consumption – even if augmented with Campbell–Mankiw rule-of-thumb behaviour – to adequately capture the dynamic interaction of consumption, income, and interest rates. In particular, the life-cycle model does not capture the ‘hump shaped’ response of consumption to shocks that appears to characterise the aggregate data. In Section 3, I develop a model of habit formation in consumer behaviour, based on the work of Carroll, Overland and Weil (1995), and related in spirit to the pioneering work of Duesenberry (1967). In Section 4, I examine the extent to which habit formation – one form of non-time separability in utility – can improve the dynamic behaviour of the simple model. I will argue 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 dynamic response of consumption to shocks. Section 5 examines the response of the model during a disinflation, Section 6 examines the quality of the linear approximations used, Section 7 discusses some welfare considerations, and Section 8 concludes.

¹ In this regard, this paper supports the conclusions of Fair (1992), i.e. a return to the Cowles Commission approach (perhaps somewhat modernised) to specification and testing of empirical macroeconometric models.

2. Problems with Standard Models

In Fuhrer (1997a) I document the inability of standard optimising models of consumer and investment behaviour to replicate the dynamic interactions among real and nominal variables found in the data. Here I briefly summarise these results, and motivate the exploration of a less standard description of consumer behaviour that involves habit formation.

The key results uncovered in my earlier work include both perverse parameter estimates and empirically contradicted dynamic behaviour. For consumption, I use a standard life-cycle model of consumption, augmented to include rule-of-thumb behaviour by a fraction of consumers that is empirically determined.² The problems with this specification include extremely significant unexplained serial correlation in the consumption-income ratio, parameter estimates that indicate very little or no forward-looking behaviour, and excessive sensitivity of consumption to current income arising from the rule-of-thumb behaviour.

For investment, I employed a model that allows for both time-to-build and costs of adjustment. In this sector, the problems included very significant unexplained serial correlation in the investment-capital ratio, extreme sensitivity of model stability and uniqueness to small perturbations in parameter estimates, and a *negative* estimate of the capital share in income.

Equally important for the purposes of monetary policy analysis are the dynamic implications of these specifications when embedded in a model with sticky prices and an explicit federal funds rate policy rule. The dynamic correlations implied by the model, summarised by the ACF, are seriously at odds with those from an unconstrained vector autoregression that nests the restricted model. A set of disinflation simulations identify an important source of the discrepancy: both consumption and investment act like ‘jump variables’, completely front-loading or pulling forward in time their responses to shocks. This stands in contrast to exercises with identified VARs (e.g. Christiano, Eichenbaum and Evans (1994); Leeper, Sims and Zha (1996)), in which these variables demonstrate a gradual response over several years, with the peak response at one year or so. Thus the

² Strictly speaking, the rule-of-thumb model suggests that a fraction of total income accrues to consumers who follow a rule of thumb.

problem, broadly speaking, is that the standard models imply a strongly counterfactual immediate response of consumption and investment to all shocks, but particularly to monetary policy shocks.

This paper attempts a solution to the problems for the consumer sector, developing and econometrically testing a habit formation model. The intuition behind this approach is simple: If the standard life-cycle model implies a too-rapid or ‘jump’ response to shocks, then a model is required in which the utility function implies a smoother, more hump-shaped response to shocks. The specification explored below achieves this goal by employing a utility function that implies a smoothing motive for the change in consumption as well as its level.

2.1 A Non-behavioural Solution to the Problem

In this paper and in my earlier work, I begin with the assumption that the structural innovations in the econometric model are uncorrelated across time, although they may be correlated across equations. The rationale behind the first assumption is that we should first attempt to model the dynamic behaviour evident in the data as the outcome of the behaviour of consumers and firms. While some of the correlations in the data may arise from somewhat correlated shocks, it would be unsatisfying to attribute most or all of the fluctuations in key variables to shocks; this would in essence be admitting that consumption and investment fluctuate for reasons that we do not understand and cannot model as economic processes. This seems to take all the fun out of dynamic macroeconomic modelling.

However, in principle, one can augment a structural model that is dynamically deficient with an arbitrary error structure so as to exactly replicate the dynamic structure in the data. This approach is taken in Rotemberg and Woodford (1997), for example. A key problem with this approach, however, is that the error processes so identified cannot be considered ‘structural’ in any meaningful sense. We have no idea in what way the errors are linked to underlying behaviour, and thus we can have no more confidence about their policy invariance than we have for reduced-form VARs or 1960s structural models *sans* explicit expectations. In essence, by

putting the dynamic structure in the errors, the model becomes vulnerable to a Lucas critique of its errors.³

For these reasons, I find the augmentation of the error structure an unappealing solution to the problem of finding a dynamically satisfying monetary policy model. The next section describes an attempt to modify the behavioural assumptions underlying the consumer sector to better capture the dynamics in the data.

3. A Simple Habit Formation Model

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

$$U_t = \frac{1}{(1-s)} \left(\frac{C_t}{Z_t^g} \right)^{(1-s)} \quad (1)$$

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

$$Z_t = r_z Z_{t-1} + (1 - r_z) C_{t-1} \quad (2)$$

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 parameterises two features of habit formation:

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

³ For more on this point, see my discussion of Rotemberg and Woodford in the 1997 *NBER Macroeconomics Annual* (Fuhrer 1997b).

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

Employing a standard budget constraint with time-varying interest rate, one can express the (nonlinear) first-order conditions for the consumer’s problem as:

$$\begin{aligned} & \frac{1}{C_t} \left(\frac{C_t}{Z_t^g} \right)^{(1-s)} - \mathbf{g}(1-s)(1-\mathbf{r}_z)P_t = \\ & \mathbf{b}E_t \left[(1+r_{t+1}) \frac{1}{C_{t+1}} \left(\frac{C_{t+1}}{Z_{t+1}^g} \right)^{(1-s)} \right] - \mathbf{b}_g(1-s)(1-\mathbf{r}_z)E_t[(1+r_{t+1})P_{t+1}] \end{aligned} \quad (3)$$

where Z_t is defined above and P_t is defined as:

$$P_t \equiv \mathbf{b}\mathbf{r}_z E_t P_{t+1} + \frac{U_t}{Z_t} \quad (4)$$

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

$$r_t \equiv (1-d) \sum_{i=0}^{\infty} d^i E_t \left(i_{t+i} - \mathbf{p}_{t+i+1} \right) \quad (5)$$

where $d = \frac{D}{1+D}$, and D is the duration of the (implied) long-term real bond, which is set to 10 years for this paper. The parameter \mathbf{r} 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. (See Appendix A for a full derivation of the first-order conditions and the linear approximations used below.)

3.1 A Linearised Consumption Function

In order to derive an explicit consumption function, I linearise the first-order conditions and substitute into the linearised budget constraint, obtaining the approximate log-linear consumption function (see Appendix A for details):

$$c_t - y_t = \sum_{j=1}^{\infty} \mathbf{r}^j E_t \left[\Delta y_{t+j} + a_1 (p_{t+j+1} - p_{t+j}) + a_2 (z_{t+j+1} - z_{t+j}) - \delta r_{t+j+1} \right] \quad (6)$$

with P_t defined as:

$$P_t \equiv \mathbf{b} \mathbf{r}_z E_t p_{t+1} + \frac{U_t}{Z_t} \quad (7)$$

where lower-case letters denote logarithms, and the parameters a_1 , a_2 , and \mathbf{d} are nonlinear functions of the underlying parameters \mathbf{g} , \mathbf{s} , \mathbf{r}_z , \mathbf{b} , and I impose these nonlinear constraints on the parameters.⁴

Campbell and Mankiw (1989, 1990, 1991) provide compelling evidence for the existence of rule-of-thumb consumers, i.e. consumers whose current consumption equals current income. This constitutes a strong violation of the permanent income theory, because a significant fraction of this period's income is predictable as of last period. A permanent income consumer would consume beginning in last period the annuity value of the component of current income that was predictable

⁴ In particular, the coefficients are defined as:

$$\begin{aligned} a_1 &= ((\mathbf{g}^* (1 - \mathbf{s}) * (1 - \mathbf{r})) / \mathbf{s}) \\ a_2 &= ((1 - \mathbf{s}) * \mathbf{g}) / \mathbf{s} \\ \mathbf{d} &= \mathbf{b} [(-\mathbf{g}(1 - \mathbf{s}) (1 - \mathbf{r}_z) - 1) / \mathbf{s}] \\ b_1 &= (\mathbf{r} - \mathbf{s}) / (1 - \mathbf{s}) \\ b_2 &= (\mathbf{g}^* (1 - \mathbf{s}) - 1) * (1 - \mathbf{s}) \end{aligned}$$

and steady-state detrended consumption is assumed to be zero. Note that the linearised consumption function retains the parametric features of the nonlinear model: when $\mathbf{g} = 0$, habit formation does not enter the consumer's problem and the consumption function reduces to a standard life-cycle/permanent-income specification.

last period. I allow for the possibility of rule-of-thumb consumers in the log-linear consumption function by modifying it as:

$$c_t - y_t = (1 - \mathbf{I}) \sum_{j=1}^{\infty} \mathbf{r}^j E_t[\Delta y_{t+j} + a_1(p_{t+j+1} - p_{t+j}) + a_2(z_{t+j+1} - z_{t+j}) - \mathbf{c}r_{t+j+1}] + \mathbf{e}_{ct} \quad (8)$$

where \mathbf{I} represents the fraction of total income accruing to rule-of-thumb consumers (who follow the rule $c_t = y_t$), and \mathbf{e}_{ct} is the structural innovation in the consumption equation, usually interpreted as the innovation to lifetime resources.

Thus specified, the model nests a number of interesting alternatives, including: the standard PIH model ($\mathbf{I} = 0$, $\mathbf{g} = 0$), the PIH with some rule-of-thumbers ($\mathbf{g} = 0$), a forward-looking habit formation model ($\mathbf{g} \neq 0$), as well as other combinations. In addition, the parameter \mathbf{r} , 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.

3.2 Estimating and Testing the Consumption Function

To estimate the underlying parameters, I employ a numerical maximum likelihood method which is documented in Fuhrer and Moore (1995a). The advantages of this system approach are that; (1) it allows estimation to proceed naturally from an unrestricted linear vector autoregression that nests all of the linear models considered to successively more-restricted linear models, with each succeeding restriction nested within the preceding less-restricted model and within the VAR; and (2) the finite sample properties of the estimator may be more desirable than method-of-moments estimators, as documented in Fuhrer, Moore and Schuh (1995), and West and Wilcox (1993). A drawback to the approach is that, to the extent that any equation in the system is mis-specified, estimates of all the parameters in the system will (in principle) be affected. However, I pursue an estimation strategy below that is designed to minimise the exposure to this risk.

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 minimum set of variables necessary to nest the final monetary policy model. These are log per capita nondurables and services consumption, log per capita disposable personal income, the federal funds rate, the price level, and log per capita GDP other than nondurables and services consumption. Their definitions are provided in Table 1.

Table 1: Data

Variable	Definition
Consumption	Chain-weighted expenditures on nondurables and services, per capita, detrended, trend segmented in 1974
Income	Chain-weighted personal disposable income per capita, detrended as above
Short-term interest rate	Quarterly average of the effective federal funds rate
Prices	Consumer price index, excluding food and energy
Non-consumption GDP	Chain-weighted per capita GDP, excluding nondurables and services consumption, detrended as 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 above.

4. Empirical Results

Using the data detailed in Table 1, and estimating over the sample 1966:Q1 to 1995:Q4, I obtain the parameter estimates shown in the first column of Table 2. At the estimated parameter values, we find that: (1) habit formation is an economically important determinant in the utility function; (2) the habit formation reference level is essentially last period's consumption level; (3) rule-of-thumb behaviour is important, with about one-fourth of income accruing to rule-of-thumb consumers; (4) the intertemporal elasticity of substitution is quite small; (5) for those who look forward, the horizon is long; the parameter \mathbf{r} takes the estimated value 0.996 on a quarterly basis, 0.984 on an annual basis; and (6) 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.

Table 2: Estimation Results

Coefficient	Baseline model	
	Estimate	Standard error
g	0.80	0.19
r_z	0.0015	0.0039
l	0.26	0.13
s	6.11	1.81
r	0.99	0.01
d	28.49	5.17
Ljung-Box Q(12) (p-value)		
Consumption	83.7 (0.00)	
Income	6.1 (0.91)	
Federal funds rate	13.5 (0.33)	
Inflation	8.2 (0.77)	
Log-likelihood	2366.43	
Detail on error correlations for consumption equation (standard errors in parentheses)		
Lag	Autocorrelation	Partial correlation
1	0.55(0.14)	0.55(0.09)
2	0.39(0.14)	0.12(0.09)
3	0.42(0.14)	0.25(0.09)
4	0.15(0.14)	-0.27(0.09)

The structural consumption equation error shows one significant autocorrelation of about 0.55. This correlation might be the manifestation of time-averaging in the data (Ermini 1989), or might reflect the durability of some components of consumption that are nonetheless classified as nondurables and services (Mankiw 1982). The standard errors reported in Table 1 are corrected for the estimated correlation in this error. However, all of the autocorrelation function computations, likelihood ratio tests, and simulations reported below assume the errors to be white. That is, *none* of the dynamics in the 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, r_z , 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 r_z 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:

$$U_t = \left(\frac{C_t}{Z_t} Z_t^{1-g} \right)^{1-s} \quad (9)$$

and setting $r_z = 0$ yields the special case:

$$U_t = \left(\frac{C_t}{C_{t-1}} C_{t-1}^{1-g} \right)^{1-s} \quad (10)$$

which shows that the essence of habit formation is that it 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 g is not zero). Larger values of r_z simply define the changes relative to a longer distributed lag of past consumption.⁵

Seen in this light, it becomes clear that a single lag of consumption in the reference level is sufficient to impart the smoothness to changes in consumption expenditures that is absent in the standard life-cycle model. In addition, note that the linearised consumption function with $r_z = 0$ is:

⁵ In this sense, the habit formation may provide a reasonable approximation to a model with a standard utility function and costs of adjustment in DC_t

$$c_t - y_t = (1 - I) \sum_{j=1}^{\infty} p^j E_t [\mathbf{D}y_{t+j} + a_1(p_{t+j+1} - p_{t+j}) + a_2(c_{t+j+1} - c_{t+j}) - \mathbf{d}r_{t+j+1}] + \mathbf{e}_{ct} \quad (11)$$

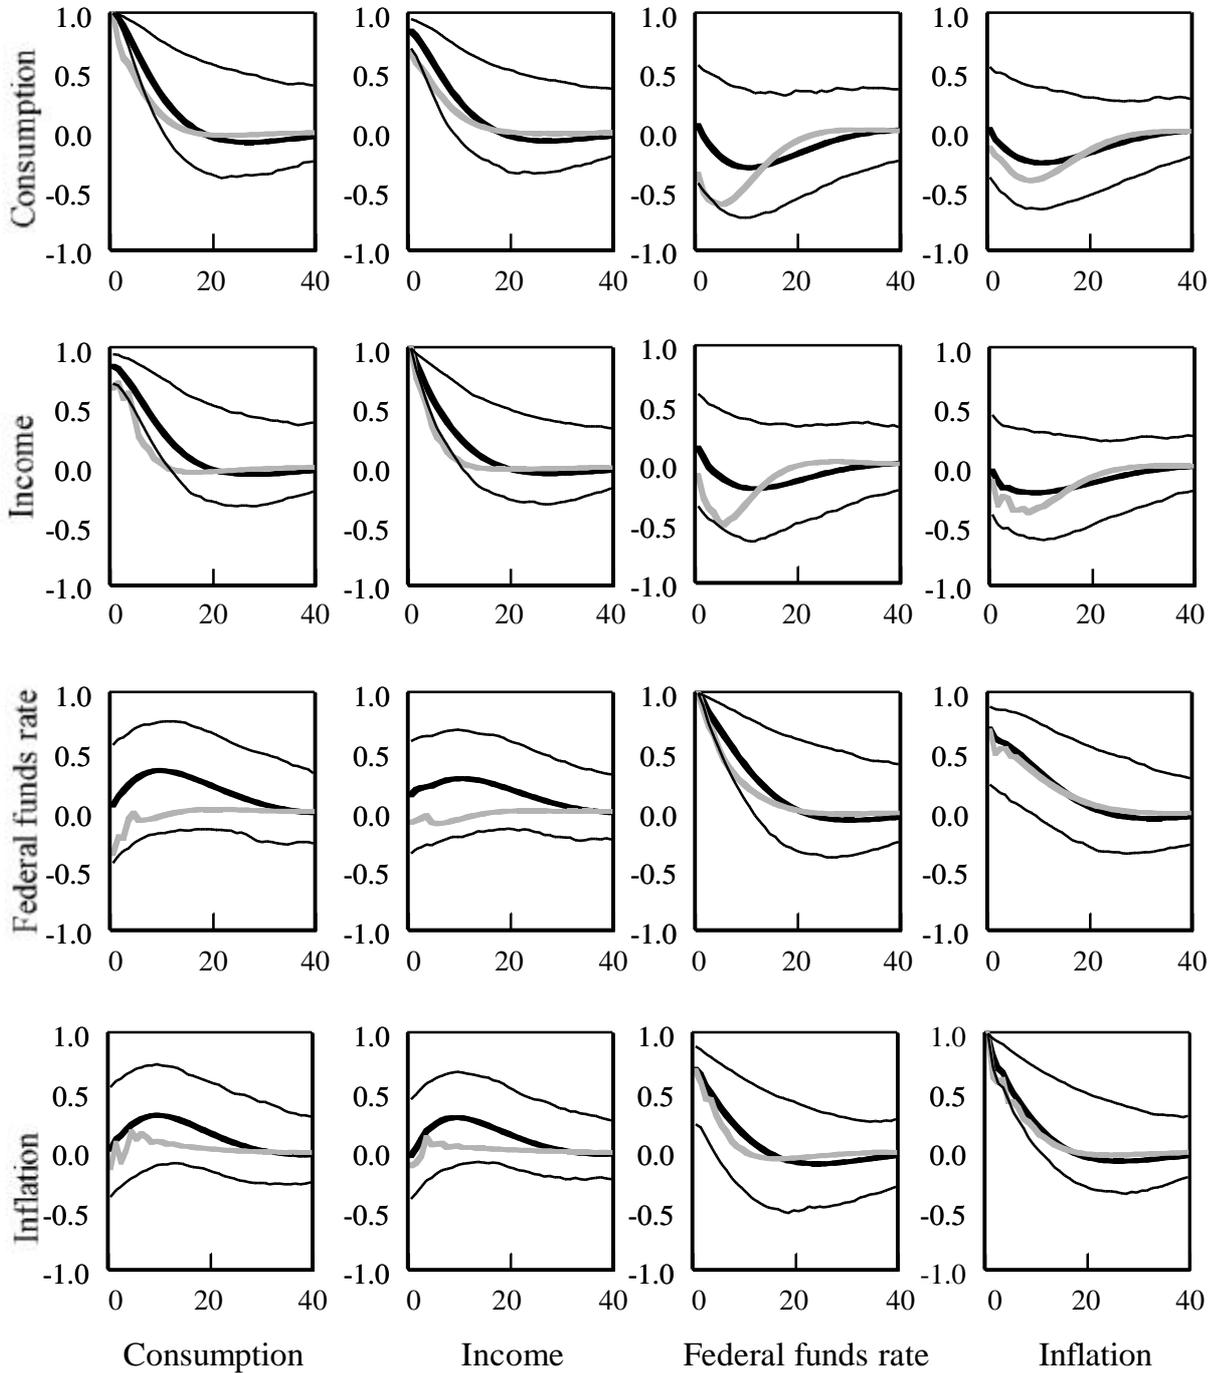
Note that the third term on the right-hand side, the weighted sum of expected future changes in consumption, will differ relatively little from the weighted sum of expected future deviations of consumption from a moving average of past consumption (the corresponding term in the consumption function with $r_z \neq 0$). 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.⁶

As I argue above, obtaining sensible parameter estimates is a necessary but not a sufficient condition for obtaining a reliable model for monetary policy analysis. Figure 1 displays the full set of autocovariances for the unconstrained VAR (the thick black lines) and the constrained consumption function (the grey lines). As the figure shows, the model has recovered the dynamic covariances for consumption expenditures quite well, capturing the persistence in the autocorrelation, as well as the persistent dynamic correlations between consumption and income, interest rates, and inflation.

⁶ The simulation presented in Figure 6 shows the effect of a longer-memory reference level on a standard disinflation simulation.

Figure 1: Comparison of Autocorrelation Functions

VAR versus habit-formation model



Notes: Ordinate: lags in quarters. Co-ordinate: correlation function. Thick black lines, VAR; grey lines, constrained habit-formation model; thin black lines, VAR standard error bands.

The thin black lines in Figure 1 display the 90 per cent 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 per cent level. Thus the correlations that the structural model cannot match are generally not precisely determined in the data. In Section 4.1 below, I perform a series of likelihood ratio tests to determine the statistical significance of a variety of restrictions on the model.⁷

4.1 Nested Tests of Habit Formation and Rule-of-thumb Behaviour

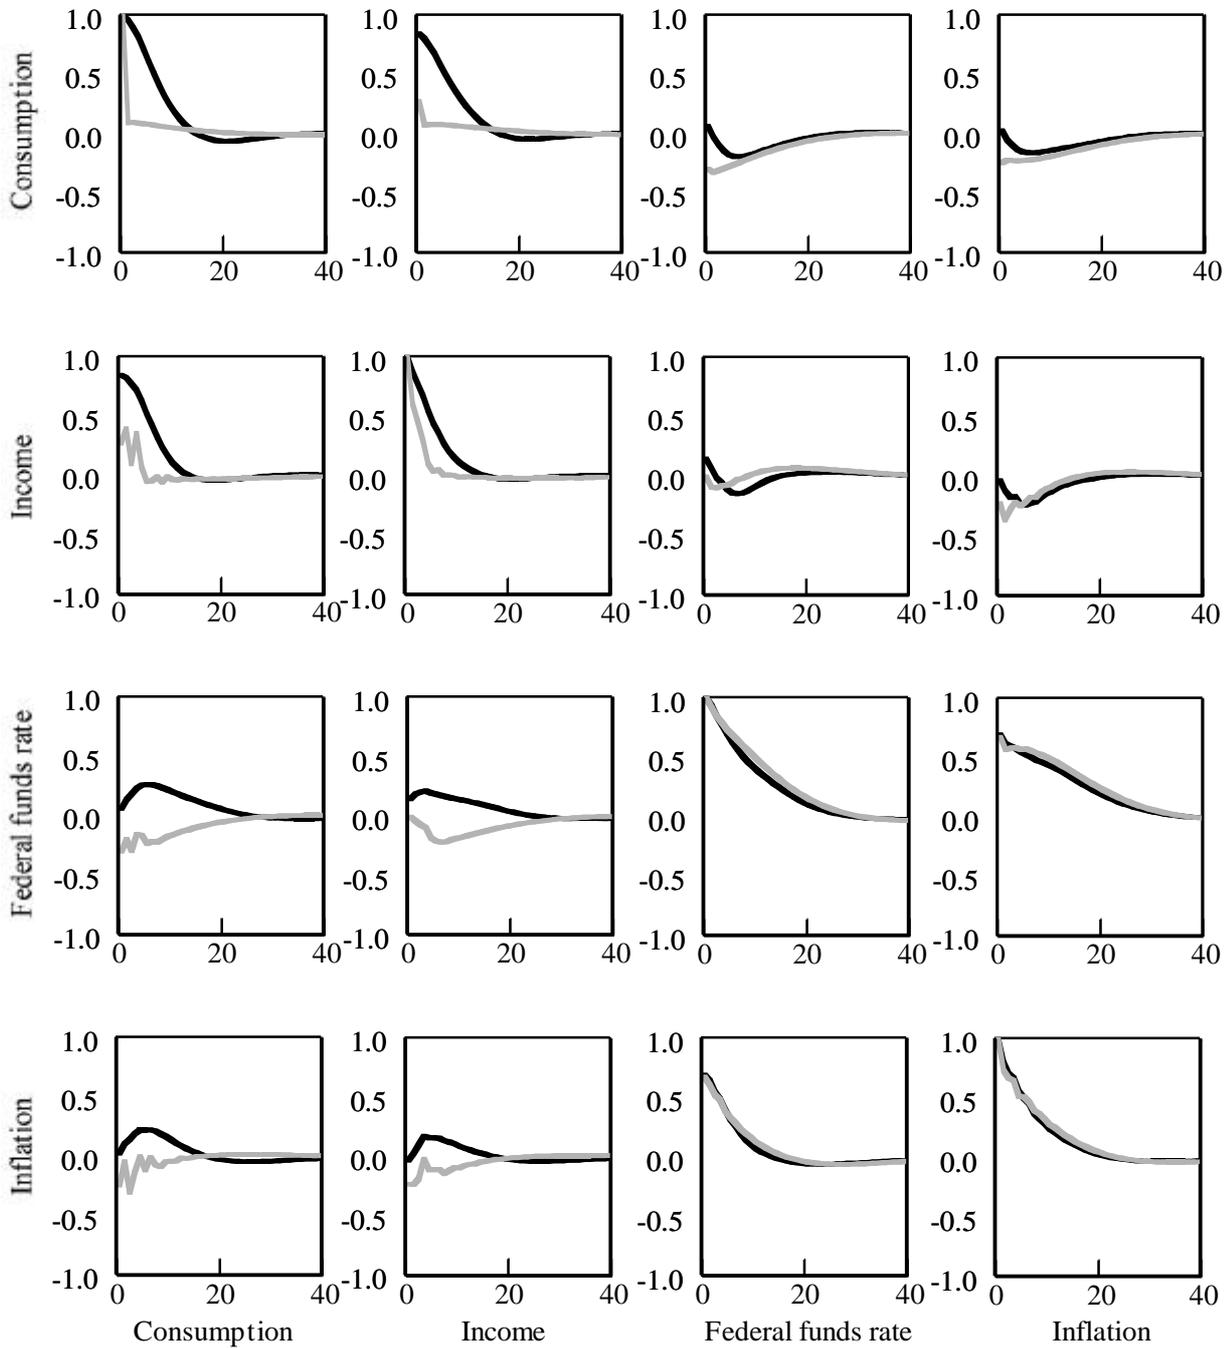
The hypothesis that habit formation is unimportant in this model – that the exponent \mathbf{g} on the reference level of consumption is zero – is overwhelmingly rejected. The \mathbf{c}^2 likelihood ratio test for this single restriction takes the value 21.4, with p-value of 4×10^{-6} . Similarly, the hypothesis that rule-of-thumb behaviour is unimportant is strongly rejected. The \mathbf{c}^2 likelihood ratio test for the restriction $\mathbf{I} = 0$ takes the value 12.6, with p-value = 4×10^{-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 per cent level. This is one of relatively few cases in which the joint restrictions imposed by an optimisation-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 behaviour in replicating the dynamic interactions among consumption, income, interest rates and inflation. As Figure 2 shows, the primary consumption dynamics in the model that sets \mathbf{g} and \mathbf{I} to zero are almost totally missing. The simple PIH model cannot replicate the dynamics in the data. Both rule-of-thumb behaviour and habit formation are statistically significant modifications to add to the model.

⁷ The confidence intervals are computed as follows. I assume the distribution of coefficient estimates to be asymptotically normal. I follow a Monte Carlo technique that draws a vector of coefficient estimates from the multivariate normal distribution centred on the sample estimates, with covariance matrix as estimated from the sample. For each vector of estimates, I compute the corresponding vector autocovariance function, holding the residual covariance matrix fixed. The 90 per cent confidence intervals are bounded by the 5th and 95th percentiles of the ranked autocovariance functions.

Figure 2: Comparison of Autocorrelation Functions

VAR versus pure life-cycle model



Notes: Ordinate: lags in quarters. Co-ordinate: correlation function. Black lines, VAR; grey lines, constrained model.

4.2 Adding Restrictions to the Model

We now progressively add restrictions to the unconstrained portions of the model, 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):

$$f_t = (1 - \sum_{i=1}^2 \mathbf{a}_{f,i})(\bar{r} + \bar{p}) + \sum_{i=1}^2 \mathbf{a}_{f,i} f_{t-i} + \sum_{j=0}^2 \mathbf{a}_{p,j} (p_{t-j} - \bar{p}) + \sum_{k=0}^2 \mathbf{a}_{y,k} y_{t-k} + \mathbf{e}_{ft} \quad (12)$$

where \bar{r} and \bar{p} are the equilibrium real interest rate and the inflation target, respectively. These simple restrictions do not significantly deteriorate 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 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) and Roberts (1997)):

$$\mathbf{p}_t = \left(\frac{1}{2} \right) \sum_{i=1}^k [(\mathbf{p}_{t-i} + \mathbf{p}_{t+i}) + \mathbf{a}y_t] + \mathbf{e}_{pt} \quad (13)$$

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, we 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 can drive 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:

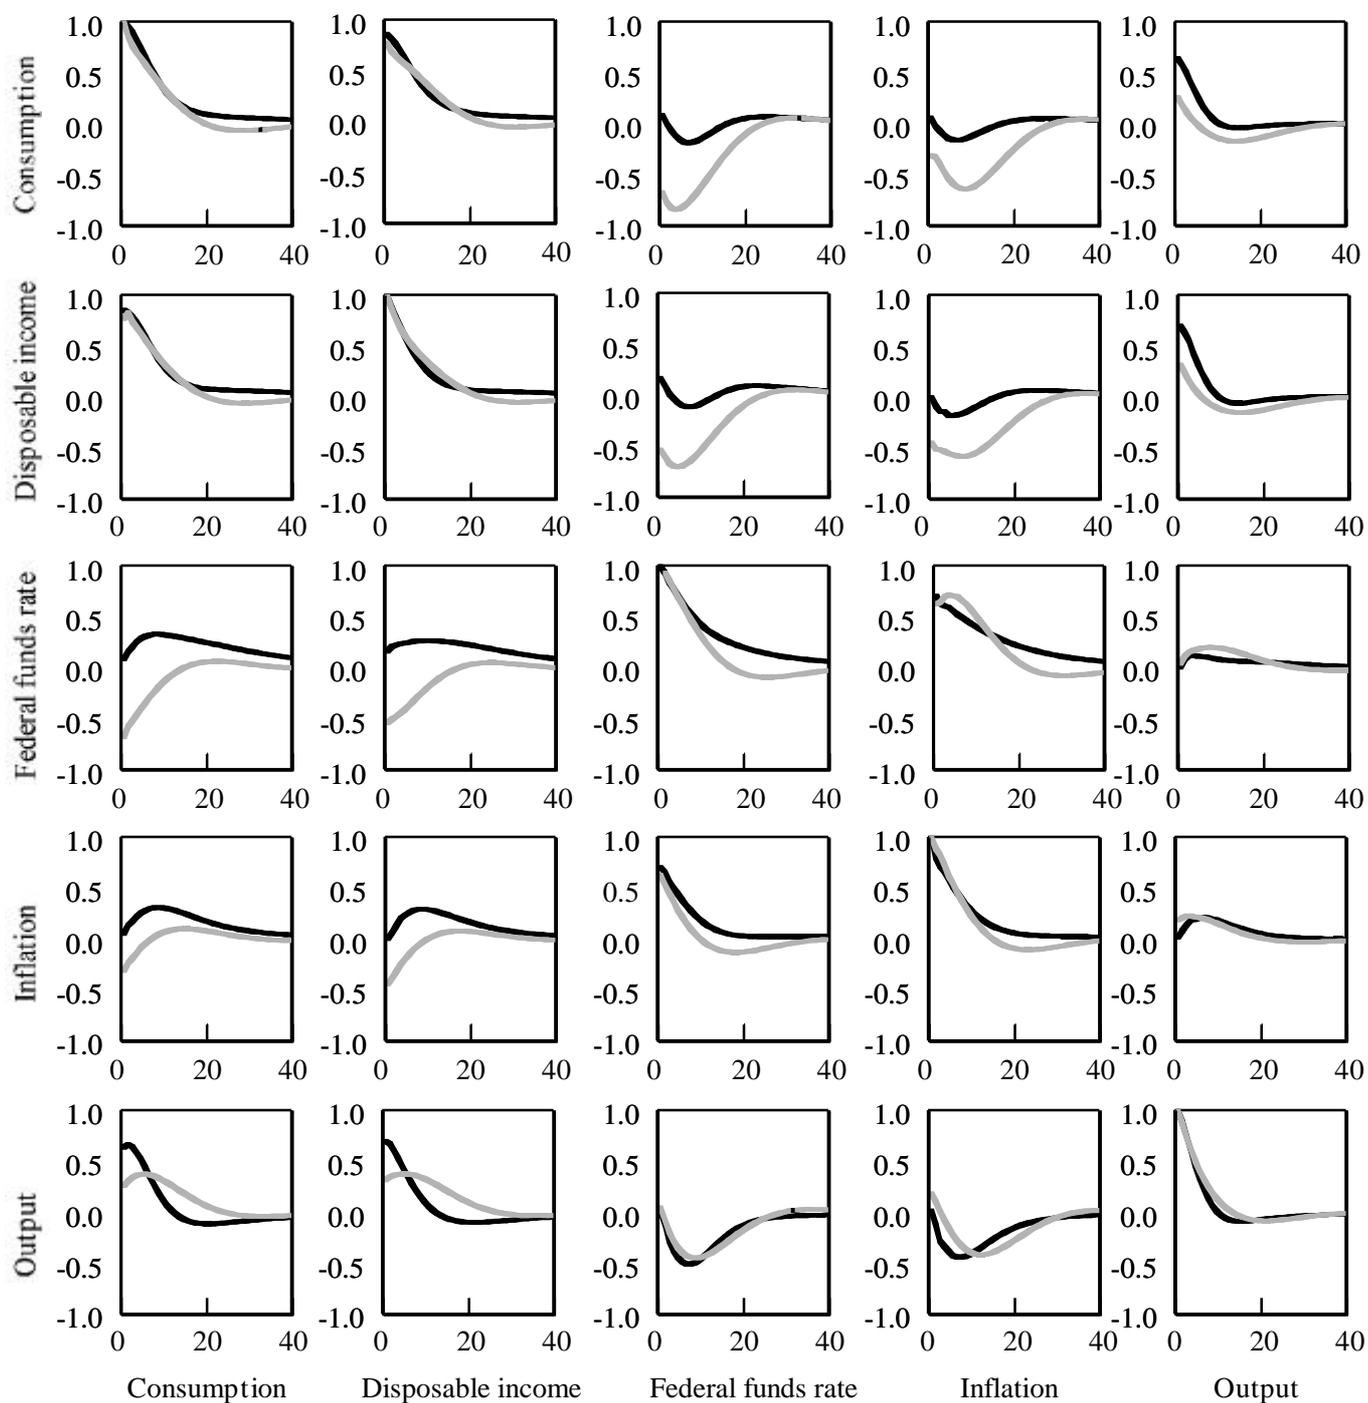
$$\tilde{y}_t^\circ = \mathbf{w} \tilde{y}_{t-1}^\circ - y_r(\mathbf{r}_{t-1} - \bar{\mathbf{r}}) + \mathbf{e}_{yt} \quad (14)$$

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 3 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 behaviour of the unconstrained VAR. However, the model cannot perfectly replicate unconstrained dynamic behaviour. 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, 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. In a sense, these autocorrelation comparisons provide graphical verification of the likelihood ratio tests conducted above.

Figure 3: Comparison of Autocorrelation Functions

VAR versus habit-formation model with reaction function and price specification



Notes: Ordinate: lags in quarters. Co-ordinate: correlation function. Black lines, VAR; grey lines, habit-formation.

5. Monetary Policy Implications of the Model

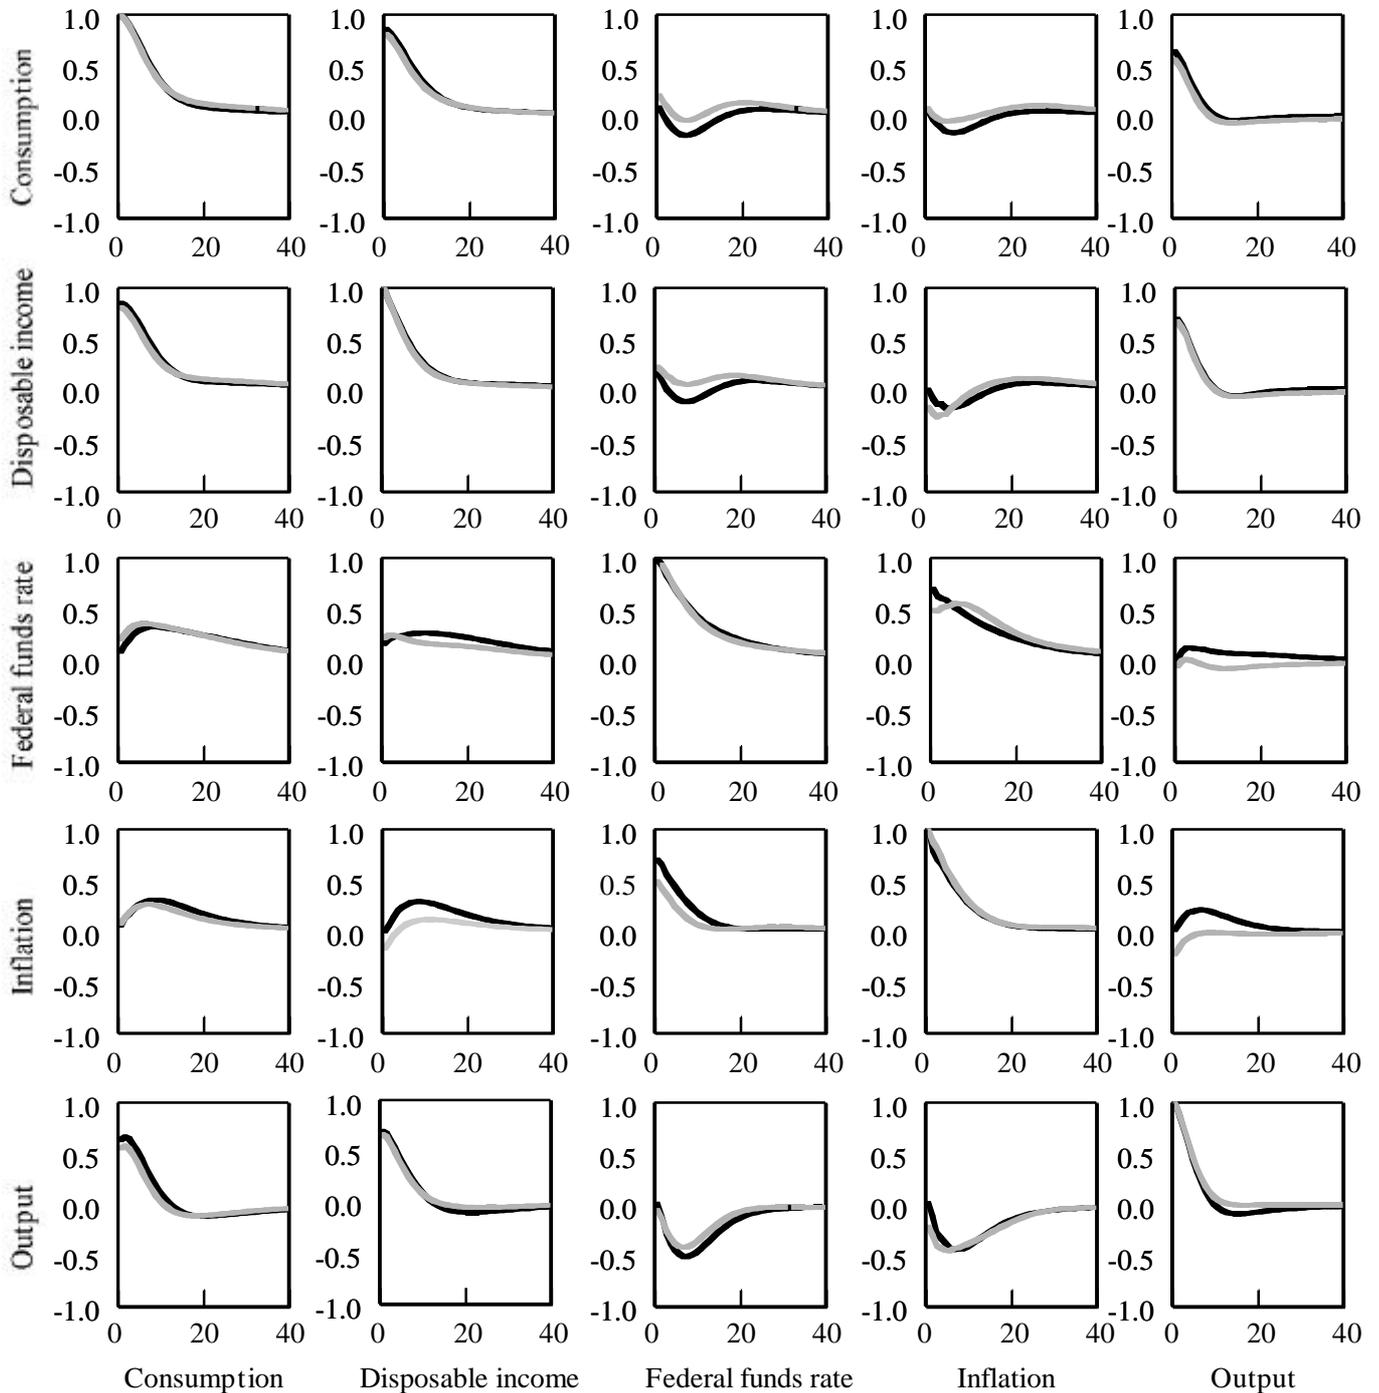
An alternative interpretation of the results of this paper and Fuhrer (1997a) is that the restrictions imposed on the price specification and the funds rate reaction function are invalid, and are interfering with the real-side dynamics of consumption and output. To test this possibility, I estimate a 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.⁸

A comparison of the autocovariance function for this model (Figure 4) 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).

⁸ 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 behaviour to match the moments in the data.

Figure 4: Comparison of Autocorrelation Functions

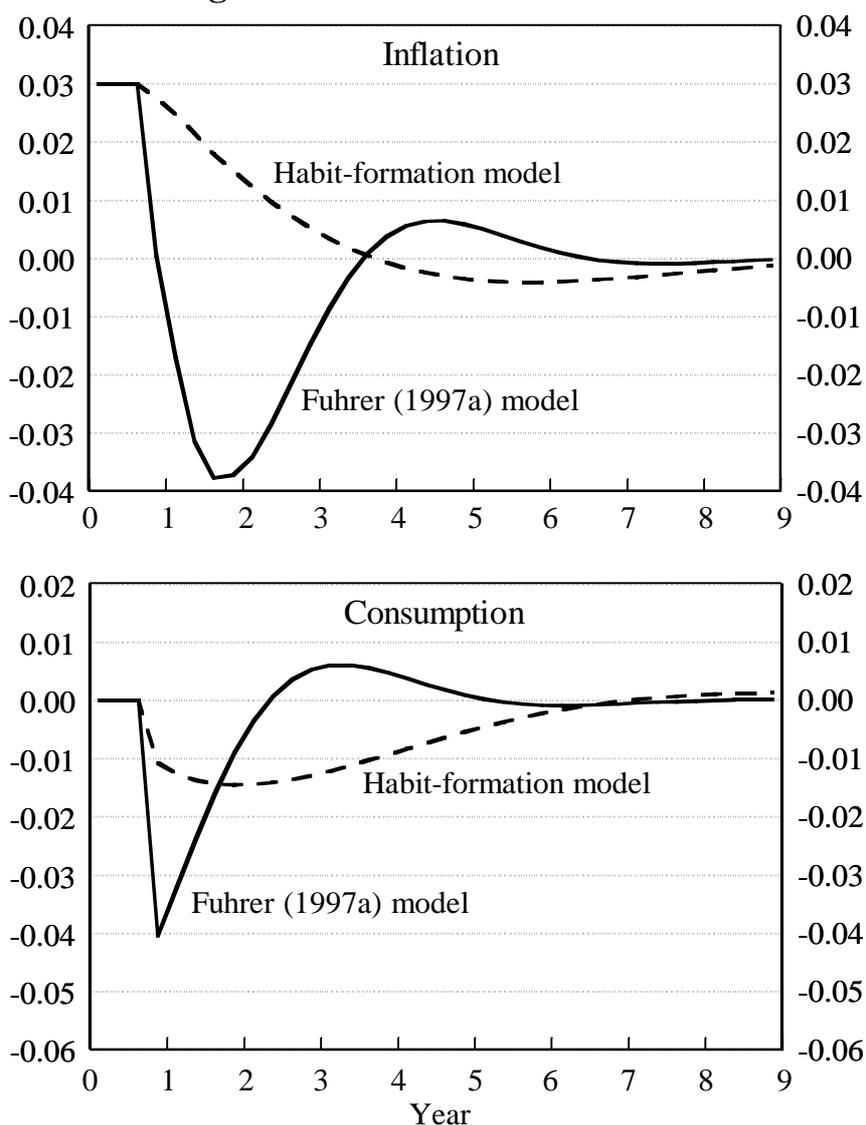
Reaction function/price specification restrictions imposed



Notes: Ordinate: lags in quarters. Co-ordinate: correlation function. Black lines, unrestricted; grey lines, restricted.

It is the case, however, that improper specification of the real side of the model can 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 almost any price specification, and given the assumed response of nominal interest rates to real output in the policy reaction function. Figure 5 below provides an example of such a case.

Figure 5: Disinflation Simulation



5.1 Disinflation in the Models

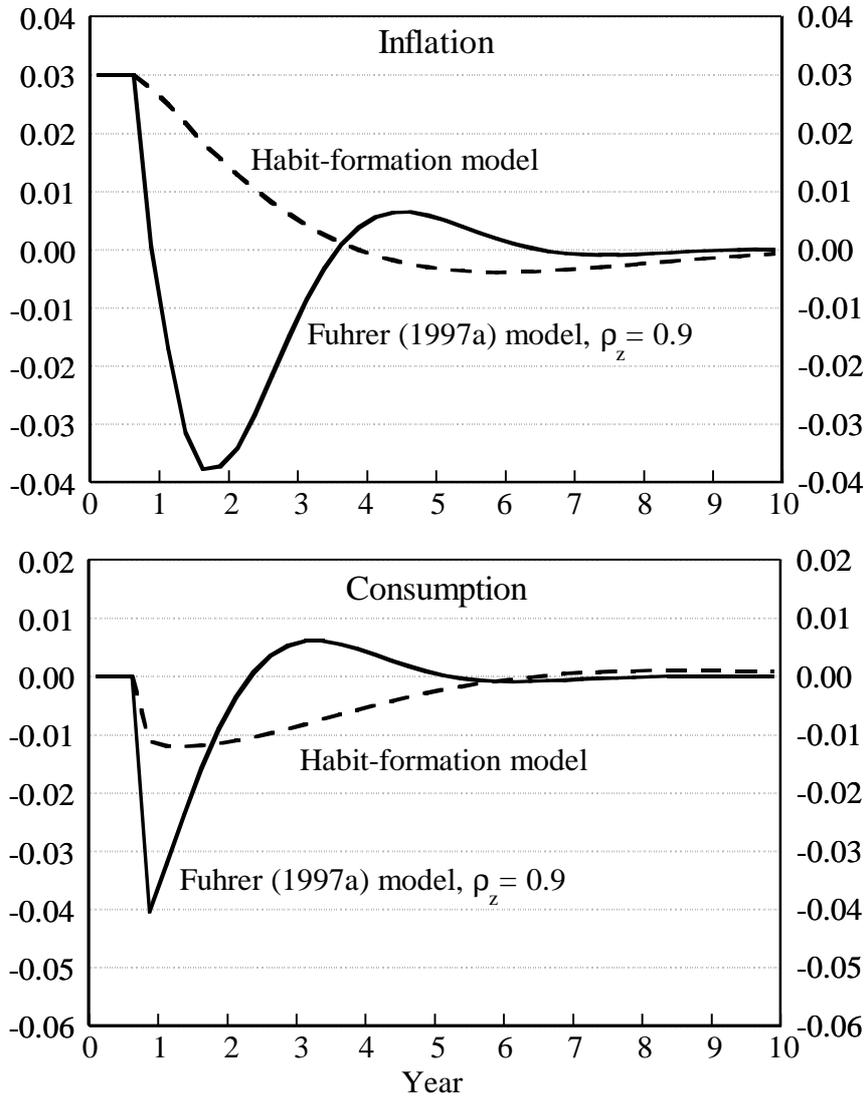
In Fuhrer (1997a), a disinflation simulation highlighted the ‘front loaded’ responses of real variables to a monetary shock. This information is, of course, contained in the vector autocovariance function, but the simple disinflation simulation makes a

key problem of the specification quite clear. Note that the specification *including* rule-of-thumb consumers still exhibited rapid response to shocks; we wish to determine whether the addition of habit formation improves this counterfactual behaviour in the model.

The simulation is straightforward. Starting from a steady state, I decrease the long-run inflation target from about 3 per cent to 0 per cent. The decrease is unanticipated. It is informative to compare the response of the model without habit formation to the model that includes it. Figure 5 displays the results of the simulation.

In the model with habit formation, inflation falls gradually from its old steady state to the new, lower equilibrium. Interestingly, consumption 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 (from Figure 4 in Fuhrer (1997a)), shown in the solid lines in the figure. Note that the solid lines exhibit an example of a model in which the mis-specification of the real side compromises the behaviour of the nominal side. The persistence of inflation in this model, as indicated in the solid line in the top panel of Figure 5, is significantly decreased by the rapid (and counterfactual) response of real variables to a disinflationary shock.

Figure 6 assesses the impact of the length of ‘memory’ in the habit reference level on the model’s dynamics. In this figure, I perform the same disinflation simulation, substituting a value of 0.9 for r_z (recall that the estimated value is 0.001). As a comparison of Figures 5 and 6 shows, the model’s behaviour is altered only slightly by the change from one-quarter memory to more persistent memory in the reference level.

Figure 6: Disinflation Simulation

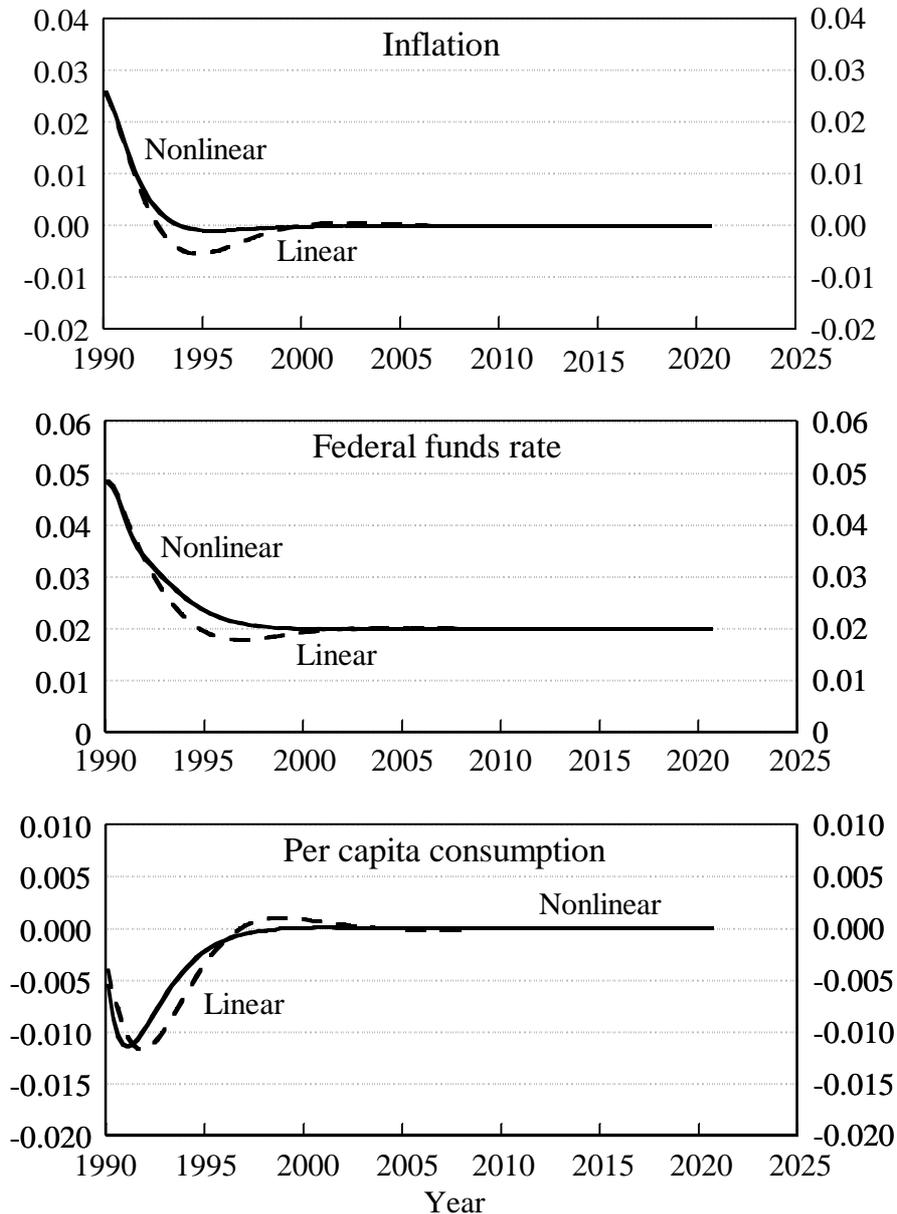
6. Accuracy of the Linear Approximation

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

I present several measures of the approximation's accuracy. First, I solve the nonlinear model (substituting Equation 3 for Equation 6), using the parameters

estimated from the linear model, for the standard disinflation simulation of the previous section. As Figure 7 shows, I obtain nearly identical results.⁹

Figure 7: Comparison of Linear and Nonlinear Simulation Paths
Disinflation from 3 to 0 per cent



⁹ 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.

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 0.01, compared with 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 behaviour implied by the nonlinear model.

7. Welfare Considerations

Having identified the utility parameters in this simple monetary policy model, it is tempting to use the model to compute the welfare implications of alternative monetary policy strategies. That is, after all, the ultimate goal in developing models of this type.

However, I would hesitate to do so with this model, for several reasons. First, given the empirical significance of consumers who appear to follow a rule of thumb, it is difficult to know whose utility we should maximise in evaluating policy alternatives. One could maximise the utility of the (forward-looking, rational) habit-formation consumers, but one could not know the welfare implications for the rule-of-thumbers.

Second, the model as it stands includes no explicit cost of inflation! The agents in the model know that the Fed cares about deviations of inflation from target (insofar as this motive is reflected in the reaction function). They know that, as a result, the Fed will cause real disruptions in order to move inflation back towards its target when it deviates, and these real disruptions will cause them to suffer welfare losses. However, it is only through these indirect effects that inflation affects consumers. Without any direct cost of inflation, the optimal policy from the consumers' point of view is one that minimises fluctuations in consumption. This is not a satisfying or interesting policy conclusion.

Finally, the representative agent nature of this model makes welfare analysis somewhat suspect. Because the bulk of the welfare cost arguably arises from

discrete shifts in employment status for a small fraction of the population, the representative agent model may not provide an accurate measure of the relative welfare costs of pursuing different monetary policies.

A commonly used alternative is to posit an indirect utility function or approximate loss function that depends on a weighted average of output deviations (around potential) and inflation deviations (around the Fed's target for inflation). Although the mapping between this loss function and utility cannot be known *a priori*, this may be a reasonable approximation to use until the thornier issues of explicitly modelling inflation losses and characterising welfare in a nonrepresentative agent framework are tackled. I leave the computation of 'optimal' policy responses in this model for future work.

8. Conclusions

A model to be used 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 purport to build close ties to underlying agents' objectives, most simple optimisation-based macroeconomic models fail to replicate important 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 behavioural decisions. As a result, such models are not suitable for monetary policy analysis. In many cases, the model's empirical failings are not widely understood, because the authors have not attempted rigorous empirical testing of the model.

This paper suggests a reasonably rigorous empirical standard for dynamic econometric models, and makes some progress towards a model that meets the standards itemised 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.

Other specifications may also afford improvements in the empirical performance of the standard model. This paper suggests, however, that only specifications that impose some smoothness on the change in consumption will be successful empirically. 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.

The specification set forth in this paper might not be robust across shifts in monetary or other policy regimes. But only through rigorous econometric testing of this and alternative specifications across regime shifts can observational equivalence (or empirical dominance) of alternative specifications and stability of any one specification across policy shifts be determined. I believe that this paper takes a small step in the direction of developing a rigorous standard of empirical validation for macroeconometric models, and in the implementation of that standard to provide a modest improvement in optimising models for monetary policy.

Appendix A: First-order Conditions for the Nonlinear Model

Beginning with the definition of period utility:

$$U_t = \frac{1}{(1-s)} \left(\frac{C_t}{Z_t^g} \right)^{(1-s)} \quad (\text{A1})$$

the overall utility function:

$$U = U_t + bU_{t+1} + \dots \quad (\text{A2})$$

and the habit-formation reference consumption level:

$$Z_t = r_z Z_{t-1} + (1-r_z)C_{t-1} \quad (\text{A3})$$

The derivative of U with respect to C_t is:

$$\frac{\mathcal{J}U}{\mathcal{J}C_t} = \frac{\mathcal{J}U_t}{\mathcal{J}C_t} + \frac{\mathcal{J}U_t}{\mathcal{J}Z_t} \frac{\mathcal{J}Z_t}{\mathcal{J}C_t} + b \frac{\mathcal{J}U_{t+1}}{\mathcal{J}Z_{t+1}} \frac{\mathcal{J}Z_{t+1}}{\mathcal{J}C_t} + b^2 \frac{\mathcal{J}U_{t+2}}{\mathcal{J}Z_{t+2}} \frac{\mathcal{J}Z_{t+2}}{\mathcal{J}C_t} + \dots \quad (\text{A4})$$

Noting that $\frac{\mathcal{J}U_t}{\mathcal{J}C_t} = \frac{(1-s)}{C_t} U_t$ that $\frac{\mathcal{J}U_t}{\mathcal{J}Z_t} = \frac{-g(1-s)}{Z_t} U_t$ and that $\frac{\mathcal{J}Z_{t+i}}{\mathcal{J}C_t} = r_z^i (1-r_z)$

then we can express:

$$\begin{aligned} \frac{\mathcal{J}U}{\mathcal{J}C_t} = & \frac{(1-s)}{C_t} U_t - \frac{g(1-s)}{Z_t} U_t (1-r_z) - b \frac{g(1-s)}{Z_{t+1}} U_{t+1} r_z (1-r_z) - \\ & b^2 \frac{g(1-s)}{Z_{t+2}} U_{t+2} r_z^2 (1-r_z) - \dots \end{aligned} \quad (\text{A5})$$

which collapses to a more compact discounted summation:

$$\frac{\mathcal{J}U}{\mathcal{J}C_t} = \frac{(1-s)}{C_t} U_t - g(1-s)(1-r_z) \sum_{i=0}^{\infty} b^i r_z^i \frac{U_{t+i}(C_{t+i})}{Z_{t+i}} \quad (\text{A6})$$

Defining:

$$P_t \equiv \mathbf{b}r_z E_t P_{t+1} + \frac{U_t(C_t)}{Z_t} \quad (\text{A7})$$

then we have the derivatives of utility U with respect to C_t and C_{t+1}

$$\frac{\mathcal{J}U}{\mathcal{J}C_t} = \frac{(1-s)}{C_t} U_t - \mathbf{g}(1-s)(1-r_z)P_t \quad (\text{A8})$$

$$\frac{\mathcal{J}U}{\mathcal{J}C_{t+1}} = \frac{(1-s)}{C_{t+1}} U_{t+1} - \mathbf{g}(1-s)(1-r_z)P_{t+1} \quad (\text{A9})$$

Combining these with a standard budget constraint (with time-varying real interest rate r_t) in a Lagrangian, we obtain the first-order conditions:

$$\frac{\mathcal{J}L}{\mathcal{J}C_t} = \frac{1}{C_t} \left(\frac{C_t}{Z_t^{\mathbf{g}}} \right)^{(1-s)} - \mathbf{g}(1-s)(1-r_z)P_t - L \quad (\text{A10})$$

$$\frac{\mathcal{J}L}{\mathcal{J}C_{t+1}} = \frac{1}{C_{t+1}} \left(\frac{C_{t+1}}{Z_{t+1}^{\mathbf{g}}} \right)^{(1-s)} - \mathbf{g}(1-s)(1-r_z)E_t P_{t+1} - E_t \left(\frac{L}{1+r_{t+1}} \right) \quad (\text{A11})$$

Both of these expressions must equal zero for an optimum, yielding an Euler equation:

$$\begin{aligned} & \frac{1}{C_t} \left(\frac{C_t}{Z_t^{\mathbf{g}}} \right)^{(1-s)} - \mathbf{g}(1-s)(1-r)P_t = \\ & \mathbf{b}E_t \left[(1+r_{t+1}) \frac{1}{C_{t+1}} \left(\frac{C_{t+1}}{Z_{t+1}^{\mathbf{g}}} \right)^{(1-s)} \right] - \mathbf{b}\mathbf{g}(1-s)(1-r_z)E_t((1+r_{t+1})P_{t+1}) \end{aligned} \quad (\text{A12})$$

Appendix B: Deriving an Approximate Linear Consumption Function

We approximate the first-order condition with its linear approximation about the steady-state values for C and Z :

$$f(C, Z) \cong f(C_0, Z_0) + f_c(C_0, Z_0)(C - C_0) + f_z(C_0, Z_0)(Z - Z_0) + \text{Higher order terms} \quad (\text{B1})$$

In the steady state, $Z = C$, simplifying the linearised first-order condition, and we obtain:

$$a_1(C_t - C_0) + a_2 P_t - a_3(Z_t - C_0) + k_0 = E_t \left[(1 + \bar{r}) \mathbf{b} (a_1(C_{t+1} - C_0) + a_2 P_{t+1} - a_3(Z_{t+1} - C_0)) - \mathbf{d}(1 + r_{t+1}) \right] \quad (\text{B2})$$

where the coefficients a_1 and \mathbf{d} are defined as:

$$\begin{aligned} a_1 &= \mathbf{s} C_0^{(1-\mathbf{g})(1-\mathbf{s})} & a_2 &= \mathbf{g}(1-\mathbf{s})(1-\mathbf{r}) \\ a_3 &= (1-\mathbf{s}) \mathbf{g} C_0^{(1-\mathbf{g})(1-\mathbf{s})} & \mathbf{d} &= \mathbf{b} \left[-\mathbf{g}(1-\mathbf{s})(1-\mathbf{r}_z) \bar{P} + C_0^{(1-\mathbf{g})(1-\mathbf{s})-1} \right] \\ k_0 &= C_0^{-\mathbf{g}(1-\mathbf{s})} \end{aligned} \quad (\text{B3})$$

We approximate the summation defined in P_t as:

$$P_t \cong C_0^{(1-\mathbf{g})(1-\mathbf{s})} + \frac{C_0^{(1-\mathbf{s})}}{Z_0^{\mathbf{g}}} (C_t - C_0) - \frac{C_0^{(1-\mathbf{s})}}{Z_0^{\mathbf{g}}} \frac{(\mathbf{g}(\mathbf{s}-1)-1)(\mathbf{s}-1)}{C_0^2} (Z_t - C_0) + \mathbf{b} r_z E_t P_{t+1} \quad (\text{B4})$$

Utilising the approximation in Campbell and Mankiw (1991), we can write the log-linearised budget constraint in consumption and income as:

$$c_t - y_t = \sum_{j=1}^{\infty} p^j (r_{t+j} - \mathbf{D}c_{t+j}) + rk / (1 - r) + E_t \sum_{j=1}^{\infty} r^j (\mathbf{D}y_{t+j} - r_{t+j}) - rk / (1 - r) = \sum_{j=1}^{\infty} p^j (\mathbf{D}y_{t+1} - \mathbf{D}c_{t+1}) \quad (\text{B5})$$

where lower-case letters denotes logs.

If we use the approximation $(1 + \bar{r}) \mathbf{b} \cong 1$ in the Euler equation, then the expected change in consumption is:

$$E_t C_{t+1} - C_t = \frac{b_1}{a_1} (P_t - E_t P_{t+1}) + \frac{c_1}{a_1} (Z_t - E_t Z_{t+1}) + \frac{\mathbf{d}}{a_1} (1 + E_t r_{t+1}) \quad (\text{B6})$$

Using the approximation that the changes in the level of C will be proportional to log changes 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:

$$c_t - y_t = \sum_{j=1}^{\infty} r^j E_t [\Delta y_{t+j} + \frac{b_1}{a_1} (p_{t+j+1} - p_{t+j}) + \frac{c_1}{a_1} (z_{t+j+1} - z_{t+j}) - \frac{\mathbf{d}}{a_1} r_{t+j+1}] + \mathbf{e}_{ct} \quad (\text{B7})$$

The parameters a_1 , a_2 , \mathbf{d} in Equation (6) correspond to b_1/a_1 , c_1/a_1 , and \mathbf{d}/a_1 ; the steady-state values for C_0 (and hence Z_0) are set arbitrarily to unity, and the steady-state value for P is determined accordingly. In the estimation step, I estimate \mathbf{d} as a parameter, not imposing all of the restrictions implied by the Euler equation. The final consumption function used in the empirical work is this equation with the addition of a fraction of income I accruing to rule-of-thumb consumers.

References

Bernanke, B.S. and Blinder, A.S. (1992), ‘The Federal Funds Rate and the Channels of Monetary Transmission’, *American Economic Review*, 82(4), pp. 901–921.

Campbell, J.Y. and Mankiw, N.G. (1989), ‘Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence’, *NBER Macroeconomics Annual*, pp. 185–216.

Campbell, J.Y. and Mankiw, N.G. (1990), ‘Permanent Income, Current Income, and Consumption’, *Journal of Business and Economic Statistics*, 8(3), pp. 265–279.

Campbell, J.Y. and Mankiw, N.G. (1991), ‘The Response of Consumption to Income: a Cross-country Investigation’, *European Economic Review*, 35(4), pp. 723–767.

Carroll, C.D., Overland, J. and Weil, D.N. (1995), ‘Saving and Growth with Habit Formation’, Federal Reserve System Finance and Economics Discussion Series No. 95-42.

Christiano, L.J. and Eichenbaum, M. (1990), ‘Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations’, Institute for Empirical Macroeconomics Discussion Paper No. 24.

Christiano, L.J., Eichenbaum, M. and Evans, C. (1994), ‘Identification and the Effects of Monetary Policy Shocks’, Federal Reserve Bank of Chicago Working Paper Series WP 94-7.

Duesenberry, J.S. (1967), *Income, Saving, and the Theory of Consumer Behavior*, Oxford University Press, New York.

Ermini, L. (1989), ‘Some New Evidence on the Timing of Consumption Decisions and on their Generating Process’, *Review of Economics and Statistics*, 71(4), pp. 643–650.

Fair, R.C. (1992), ‘The Cowles Commission Approach, Real Business Cycle Theories, and New Keynesian Economics’, NBER Working Paper No. 3990.

Fuhrer, J.C. (1997a), ‘Towards a Compact, Empirically-Verified Rational Expectations Model for Monetary Policy Analysis’, *Carnegie-Rochester Conference Series on Public Policy*, 47, pp. 197–230.

Fuhrer, J.C. (1997b), ‘An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy, Comment’, *NBER Macroeconomics Annual*, pp. 346–355.

Fuhrer, J.C. and Moore, G.R. (1992), ‘Monetary Policy Rules and the Indicator Properties of Asset Prices’, *Journal of Monetary Economics*, 2(2), pp. 303–336.

Fuhrer, J.C. and Moore, G.R. (1995a), ‘Inflation Persistence’, *Quarterly Journal of Economics*, 110 (1), pp. 127–159.

Fuhrer, J.C. and Moore, G.R. (1995b), ‘Monetary Policy Trade-offs and the Correlation Between Nominal Interest Rates and Real Output’, *American Economic Review*, 85(1), pp. 219–239.

Fuhrer, J.C. and Moore, G.R. and Schuh, S.D. (1995), ‘Estimating the Linear-quadratic Inventory Model: Maximum Likelihood Versus Generalized Methods of Moments’, *Journal of Monetary Economics*, 35(1), pp. 115–157.

King, R.G. and Plosser, C.I. (1984), ‘Money, Credit and Prices in a Real Business Cycle’, *American Economic Review*, 74(3), pp. 363–380.

King, R.G., Plosser, C.I. and Rebelo, S.T. (1988), ‘Production, Growth and Business Cycles I. The Basic Neoclassical Model’, *Journal of Monetary Economics*, 21(2/3), pp. 195–232.

King, R.G. and Wolman, A. (1996), ‘Inflation Targeting in a St. Louis Model of the 21st Century’, NBER Working Paper No. 5507.

Leeper, E., Sims, C. and Zha, T. (1996), ‘What Does Monetary Policy Do?’, *Brookings Papers on Economic Activity*, 2, pp. 1–78.

Lucas, R.E. (Jr) (1976), ‘Econometric Policy Evaluation: A Critique’, *Carnegie-Rochester Conference Series on Public Policy*, 1, pp. 7–33.

Mankiw, N.G. (1982), ‘Hall’s Consumption Hypothesis and Durable Goods’, *Journal of Monetary Economics*, 10(3), pp. 417–426.

McCallum, B.T. and Nelson, E. (1997), ‘An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis’, NBER Working Paper No. 5875.

McCallum, B.T. and Nelson, E. (1998), ‘Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model’, NBER Working Paper No. 6599.

Roberts, J.M. (1997), ‘Is Inflation Sticky?’, *Journal of Monetary Economics*, 39(2), pp. 173–196.

Romer, C. and Romer, D. (1989), ‘Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz’, *NBER Macroeconomics Annual*, pp. 121–170.

Rotemberg, J.R. and Woodford, M. (1997), ‘An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy’, *NBER Macroeconomics Annual*, pp. 297–346.

Taylor, J.B. (1993), ‘Discretion Versus Policy Rules in Practice’, *Carnegie-Rochester Conference Series on Public Policy*, 39, pp. 195–214.

West, K.D. and Wilcox, D.W. (1993), ‘Some Evidence on Finite Sample Behavior of an Instrumental Variables Estimator of the Linear Quadratic Inventory Model’, Federal Reserve System Finance and Economics Discussion Series No. 93-29.