Estimating real interest rates for the United Kingdom

Jens Larsen
Ben May
and
*James Talbot**

Working Paper no. 200

E-mail: jens.larsen@bankofengland.co.uk
ben.may@bankofengland.co.uk
james.talbot@bankofengland.co.uk

* *Corresponding author.*

The views expressed are those of the authors and do not necessarily reflect those of the Bank of England. This paper builds on earlier work at the Bank, conducted by Laurent Clerc, Richard Harrison and Katharine Neiss, to whom we owe a large debt. The authors would also like to thank Nicola Anderson, Peter Andrews, Alec Chrystal, Jan Groen, Martin Lettau, Jack McKeown, Craig Shute, Claire Williamson and especially Charlie Bean, for many helpful comments and discussions. All remaining errors are our own.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone 020 7601 4030, fax 020 7601 3298, e-mail mapublications@bankofengland.co.uk

Working papers are also available at www.bankofengland.co.uk/wp/index.html

The Bank of England's working paper series is externally refereed.

ISSN 1368-5562

Contents

Abstract

Any monetary policy maker using a short-term nominal interest rate as the primary policy tool will have an interest in understanding developments in *ex-ante* real interest rates. In this paper, several methods for calculating real interest rates for the United Kingdom are explored. These include: yields on index-linked bonds; yields on nominal bonds minus an appropriate measure of inflation expectations; and a 'consumption-based' measure – derived from manipulating the first-order condition of a standard household intertemporal optimisation problem. It is found that the basic (power utility) version of the consumption-based model suffers from the standard problems outlined in the literature, so the basic framework is augmented to allow for (external) habit formation in consumption, and a general *k*-period real interest rate is derived. Interestingly, although the different approaches outlined above can sometimes yield very different estimates of real interest rates, all the measures move more closely together during the post-1992 inflation-targeting period than before. Before 1992, uncertainty about the monetary regime, coupled with persistent expectational errors, may have made it more difficult for agents to forecast real interest rates and inflation.

Key words: Real interest rates, asset pricing, consumption.

JEL classification: E21, E43, G12.

Summary

The *ex-ante* real interest rate is a key variable in the transmission mechanism of monetary policy. Any change in the short-term nominal interest rate set by the monetary authority will – if prices are sluggish – lead to a change in real interest rates, which will affect demand, and subsequently inflation, via the consumption and savings or investment decisions of households and firms.

In general, there are few *direct* measures of the *ex-ante* real interest rate because almost all debt contracts are specified in nominal terms. So this paper explores a number of methods for calculating UK real interest rates. The pros and cons of each approach are evaluated carefully: after constructing a long and consistent time series of each measure, a rigorous sensitivity analysis is conducted and, where appropriate, error bands are constructed around the estimates in order to assess their accuracy.

The United Kingdom has a well-developed market for government bonds (gilts) that are indexed to the retail prices index (RPI), so the first approach considers real interest rates derived from these bond prices. But more recently, estimation has been complicated by the combined effect of limited supply and artificially price-inelastic demand on yields in this market.

A second approach uses yields on nominal gilts minus an appropriate measure of inflation expectations. But this method is also subject to several important problems. First, nominal bond yields will be subject to the same distortions identified above. Second, estimating inflation expectations is not an easy task. The paper adopts two approaches to devise inflation expectations: one is based on surveys and another one uses forecasts from a vector autoregressive model of inflation, unemployment and interest rates. Third, such estimates will include a measure of the inflation risk premium, and so are not directly comparable with those from index-linked gilts.

The third approach uses a ‘consumption-based’ measure – derived from manipulating the first-order condition of the standard household intertemporal optimisation problem. The basic (power utility) version of this model suffers from the standard problems outlined in the literature: the so-called ‘risk-free rate’ and ‘equity premium’ puzzles. So the basic framework is augmented to allow for (external) habit formation in consumption, and extended to estimate general *k*-period real interest rates.

Real interest rates at one, three and ten-year maturities derived using this approach look reasonably plausible: real interest rates peak during the recession of the early 1980s and fall during the economic expansions of the late 1980s and late 1990s. But because the model is based on a relatively simple process for consumption growth (a random walk), the term structure of interest rates contains less information, remaining relatively flat throughout the sample period.

Interestingly, although the different approaches outlined above can sometimes yield different estimates of real interest rates, all the measures move more closely together during the post-1992 inflation-targeting period than before. Before 1992, uncertainty about the monetary regime, coupled with persistent expectational errors, may have made it more difficult for agents to forecast real interest rates and inflation.

Another question is whether the fall in long real yields observed in the index-linked gilt market post 1997 is based on movements in real fundamentals? Evidence from the model with habit formation suggests that there has been some fall in the ten-year real interest rate since the mid-1990s. But it would appear that at least some of the decline observed in the index-linked market has been driven by the institutional factors described above, underlining the value of taking an eclectic approach when assessing movements in real interest rates.

1. Introduction

In theory, the *ex-ante* real interest rate[1] should be a key determinant of consumption decisions because it provides a measure of the relative price of consumption today against consumption tomorrow, or equivalently, the return to deferring consumption. It is also an important component of the user cost of capital, and therefore investment decisions. For these reasons alone, a monetary policy maker using a short-term nominal interest rate as the primary policy tool will have an interest in understanding developments in real interest rates. For example, a tightening in the policy rate will – if prices are sluggish – result in higher real interest rates, which will affect demand and subsequently inflation via consumption and savings decisions.

In general, the real rate of interest cannot be measured directly for two main reasons. First, almost all debt contracts are written in nominal terms, and second, expected inflation is not easily observable.[2] For the United Kingdom, the first problem is mitigated somewhat by the existence of a well-established market for government bonds that are indexed to the retail price index (RPI). Subject to certain caveats explored in more detail below, the yields on these index-linked gilts (IGs) should provide one directly observable measure of the risk-free real rate of interest over time. And combining these yields with a consumption-based asset pricing equation (known as the consumption-CAPM) will give an implicit forecast of consumption growth. Consumption will, of course, have other determinants, but under some restrictive assumptions, the real interest rate is a *sufficient* statistic for consumption growth in the sense that it summarises all the information that is relevant for producing a forecast.

More recently, prices in both the nominal and IG markets may have been driven out of line with economic fundamentals by the combined effect of limited supply and artificially price-inelastic demand. So it is attractive to use the consumption-based theory outlined above in reverse, in a similar style to Ireland (1996), in order to calculate alternative estimates of real interest rates: any given forecast of consumption growth will imply a unique real interest rate. This is the essence of the exercise undertaken in this paper – using a set of restrictive assumptions about financial markets and the preferences of the agents within them; we can derive real interest rates using a selection of empirical models.

The problems of the C-CAPM are well known, and are discussed in some detail in this paper.[3] As such, the estimates of real interest rates derived from this technique are likely to be of little interest on their own. It is possible for example that divergences between real rate estimates from IG yields and the C-CAPM are simply a function of the inadequate theoretical concepts

(1) Throughout this paper, the real interest rate is *ex ante* unless otherwise indicated.
(2) We discuss the advantages and disadvantages of survey-based inflation expectations later.
(3) Readers seeking an even more exhaustive review should see the excellent survey paper by Campbell (1999).

underlying the latter. But consumption-based estimates may become a useful comparator when studied alongside a range of other estimates; particularly if we believe that there are market-specific or technical factors that may make real rates derived from IG yields a less reliable estimate of a theoretical risk-free rate. And quarter-to-quarter changes in theoretically derived real interest rates may be useful in calibrating the impact of a shock.

A first necessary step to understanding the properties of each of the estimates is to construct a long and consistent time series. The second step is to conduct a rigorous sensitivity analysis of these measures, which are often dependent on the nature of the calibration exercise; and the third is to construct error bands around our estimates in order to assess their accuracy. This eclectic approach not only allows the policy-maker to choose from a wide range of estimates, but also to evaluate the pros and cons of these measures carefully.

To this end, the outline of the paper is as follows. First, in Section 2 we develop a set of 'reference' measures of real interest rates to which the consumption-based real rate measures calculated later in the paper can be compared. Measures based on IGs are obviously key components of this reference set, but we also draw on other information. In doing so, we present measures that are based on a manipulation of the Fisher identity to calculate real interest rates by subtracting a measure of inflation expectations from nominal interest rates on government bonds.(4) Such measures are, of course not 'perfect' – interest rates implied by yields on conventional gilts may be affected by market-specific factors too, and our measures of inflation expectations may be imprecise or subject to distortions. But by having a range of different measures, we hope to gain insights not only into the underlying risk-free real interest rate, but also on the factors that may cause divergence between the different measures.

In Section 3, we outline the theory behind the C-CAPM that relates the price of any asset to the product of its expected return and a *stochastic discount factor*, or *pricing kernel* that is common to all assets. We base our description on Campbell (1999) and Cochrane (2001) although our focus is slightly different. These two authors are primarily interested in constructing and assessing theories of asset prices, where we focus on providing the associated estimates of the risk-free real interest rate.

The forecast for consumption is a key ingredient in our measure, and we devote Section 4 to a discussion of our forecast models and their properties. This includes a short section on how to derive variance bounds around our consumption forecasts and associated real interest rates.

(4) In sum, the Fisher identity says that the nominal interest rate is the sum of the real interest rate and inflation expectations.

Section 5 is given over to the derivation of a basic measure of real interest rate, using the theoretical model outlined in Section 3 and the consumption forecasts calculated in Section 4. We also examine the sensitivity of these models to their calibration as well as the underlying theory. On the latter we propose a new application of the current consumption-based asset-pricing literature by deriving and estimating a *k*-period real interest rate with habit formation.

Section 6 provides some conclusions and suggestions for future research.

2. Reference measures of the real interest rate

Before calculating the consumption-based real rate, we first construct 'reference' estimates that are independent of consumption data. As mentioned in the introduction, the purpose of this reference set is not an evaluation of the performance of the consumption-based measure; instead, we are interested in understanding the divergences between the different measures in order to evaluate the marginal informational content of consumption data.

One way to estimate the risk-free real interest rate is by measuring returns on index-linked gilts (IGs). Although we can obtain estimates of real rates at long maturities from yields on IGs, at shorter horizons it is not always possible to extract reliable estimates because of a shortage of bonds and relatively thin trading conditions.(5) And even where the appropriate securities exist, there may be reasons to think that IG yields may not always be completely reliable because of institutional features of the market. So in order to broaden the range of maturities we can look at, and to ensure that any comparison is not distorted by temporary gilt-market specific factors, we complement the IG-based measures with estimates obtained by subtracting either survey or model-based econometric forecasts of inflation from a risk-free nominal interest rate. Taken together, these measures should be able to provide a menu of risk-free real rates over time and across different maturities, in keeping with our eclectic approach. The main drawback is that all the measures reported here are based on information extracted from the gilt market in some way.

2.1 IG-based real rate measures

Since 1981, the UK government has been issuing bonds whose returns are linked to the retail prices index (RPI). The market is well established and redemption dates are reasonably numerous, so a zero-coupon yield curve can be estimated with reasonable precision.(6) In the absence of any distortions, the estimates of real interest rates extracted from the IG yield curve should be equal to the 'true' risk-free rate of interest.

---

(5) For a comprehensive summary of these problems see Anderson and Sleath (2001), and previously, Deacon and Derry (1994a,b).

(6) The UK market is second to the US market in terms of absolute size, but has a larger number of different bonds.

There may, however, be a number of reasons why real rates derived from an estimated yield curve can deviate from the true, underlying risk-free real interest rate. These are normally grouped into two broad categories: technical and institutional factors. Here we provide only a brief discussion of these, but Bank of England (1995) and Scholtes (2002) discuss these issues in more depth.

The main technical factors that may cause a wedge between IG-based and 'true' real rates are *indexation lags* and *price index mismeasurement*. An indexation lag – returns are linked to RPI inflation eight months ago – means that a proportion of returns at the end of the bond's life has nominal rather than real certainty and hence may be subject to inflation risk. In addition, a more generic problem is that the price index may not reflect the true rate of inflation because of *mismeasurement*, either because prices are measured with error, or because the index is not representative of the typical consumption bundle of the investor. But as a first approximation, it is probably safe to assume that neither of these factors should bias real rates too much, at least at maturities beyond the short end of the yield curve.

Institutional factors, such as regulatory measures, may lead to more important distortions. For the United Kingdom, probably the most significant of these has been the Minimum Funding Requirement (MFR) legislation, announced as part of the 1997 Pensions Act. The MFR was designed to ensure that – in the event of an employer becoming insolvent – pension schemes would have sufficient assets to ensure that current liabilities were covered at all times. Although the legislation does not directly compel pension funds to hold large quantities of IGs, the use of certain type and maturity gilts as discount factors for valuing liabilities creates strong incentives for pension funds to hold similar gilts on the asset side of their balance sheet. This may have contributed to the substantial fall in real yields seen since 1997.

It is thought that UK life assurance and pension funds hold more than half of the outstanding stock of IGs. So the MFR legislation may have created inelastic demand for certain types of bond. At the same time, there was a substantial reduction in the supply of UK government debt: net gilt issuance was negative for fiscal years 1998/99-2000/01. So these factors, taken together, may have substantially distorted the shape of the real yield curve. But because both technical and institutional distortions will tend to be both time-varying and maturity dependent, it is not straightforward to devise a simple adjustment that allows more accurate estimation of real interest rates.

So how do we assess the real yields derived from the IG market in practice? According to the Fisher hypothesis, inflation expectations should be equal to the wedge between the real and nominal yield curves. At the same time, under rational expectations, the *ex-ante* real interest rate on conventional bonds should be an unbiased predictor of the *ex-post* real rate of interest on the same bonds. One 'test' could therefore be to compare yields on IGs with *ex-post* real interest rates. This is shown in Charts 2.1 and 2.2. Over the limited period for which we have an appropriate comparison, *ex-ante* real yields appear to have persistently underestimated actual real interest rates, particularly at longer maturities, although as pointed out by Breedon (1995), this may reflect the short sample period we consider here.

Chart 2.1

IG-based measures of the real interest rate against *ex-post* real rate

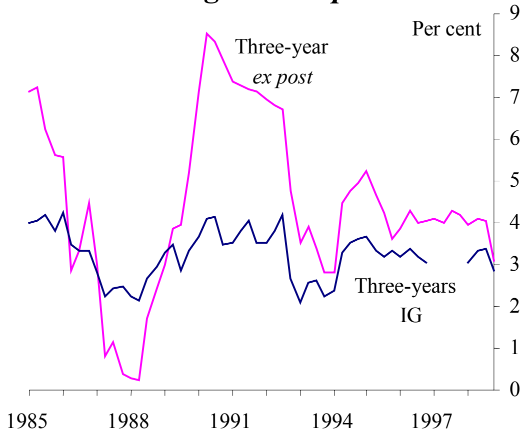

Chart 2.2

IG-based measures of the real interest rate against *ex-post* real rate

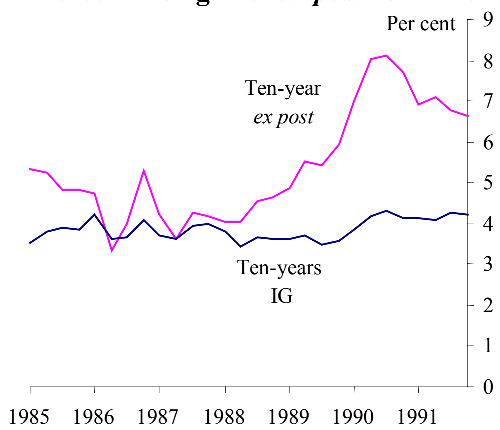

But there are several problems with this comparison. First of all it is a joint test with the implicit assumption of 'rationality' and second, the Fisher identity does not hold in practice because of the existence of risk premia.

On the first point we know that the monetary regime has changed several times over the period, which may have changed inflation dynamics and hence caused significant and persistent expectational errors, particularly at longer horizons. Scholtes (2002) tests for this, and finds that even at horizons as short as two-year, forecast rationality can be rejected – therefore, the importance of recognising that this is at best a joint test cannot be underestimated.

On the second point, in practice, nominal yields on government debt can be decomposed by the following equation:

$$(1 + y_{kt}) = (1 + r_{kt}) \cdot (1 + E_t \pi_{t,t+k}) \cdot (1 + \rho_{t,t+k}) \qquad \textbf{(2.1)}$$

Where y_{kt} is the nominal yield on a *k*-period bond at time *t*; r_{kt} is the corresponding real interest rate; $E_t\pi_{t,t+k}$ is the expected inflation over the remainder of the life of the bond; and $\rho_{t,t+k}$ is the inflation risk premium over the same period.

The index-linked yield curve provides a measure of the 'pure' real interest rate r_{kt}, whereas policy-makers may be most interested in the *ex-ante* 'real cost of borrowing' derived by some simple manipulation of **(2.1)**:

$$\frac{(1+y_{kt})}{(1+E_t\pi_{t,t+k})} = (1+r_{kt})\cdot(1+\rho_{t,t+k}) \qquad \textbf{(2.2)}$$

This measure – represented by nominal yields adjusted by agents' inflation expectations – includes an inflation risk premium, which is not separately identified by our yield curve estimation techniques. It is precisely this measure of real interest rates that can be derived by subtracting inflation expectations measured by surveys (or from an econometric model) from nominal yields. And this is discussed in more detail in the next two subsections.

2.2 Survey-based measures

As an alternative to IG-based measures, a real rate estimate can be constructed by subtracting survey-based measures of inflation expectations from a nominal interest rate. Such survey-based measures have a number of advantages: forecasts can be obtained over a range of short and medium-term maturities – horizons where IG-based measures may not be available. These surveys are also diverse in coverage (for example, price indices other than RPI) and participation (estimates can be constructed for different respondents). The first point will be particularly important if the RPI is not the correct measure on which to base inflation expectations – arguably investors may think in terms of RPIX – particularly since the adoption of an explicit inflation target in the United Kingdom in 1992. But the second could be a potential problem if those who are part of the survey have different preferences to those who invest in conventional gilts.

Survey-based measures do, of course, have a number of drawbacks. Unlike market-based measures, surveys may only be available infrequently, say monthly or quarterly. Often these surveys, particularly at longer horizons, provide estimates of the average rate over a calendar year, which is less useful than an estimate of the entire term structure. In addition, surveys are generally taken over a period of time, and the survey window is often not published, so it can be difficult to provide a precise match with nominal interest rate data. And indeed, survey respondents may not have updated their forecasts during the survey period to reflect the latest economic data. Updating is clearly a costly process, and there may be few incentives for the respondents to spend much time on this, so responses may not be consistent with actual behaviour.

An equivalent test to that done in Section 2.1 above, is to compare the quarterly Barclays Basix forecast[7] since 1987 Q1 with inflation outturns (see Chart 2.3: where the Basix forecast is pushed forward one year). Since the second half of 1991, Basix respondents have consistently forecast inflation above actual outturns. Unlike breakeven inflation forecasts, such an upward bias cannot be explained by institutional factors or inflation risk premia, although this pattern would be consistent with the effect of a change in the policy regime after 1992. But this cannot explain the previous pattern of forecast errors, and, using formal tests, Bakhshi and Yates (1998) found sufficient evidence to reject the null hypothesis of rationality for these data anyway.

Chart 2.3

Barclays Basix RPI forecast[(1)] one year ahead against inflation

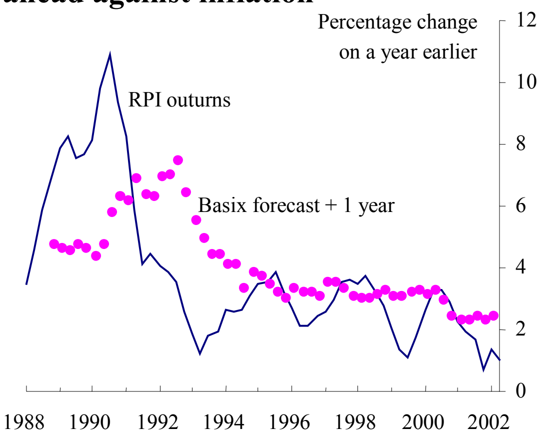

[(1)] Excluding responses by the general public.

Source: Barclays Capital.

2.3 Forecasting inflation using econometric models

The final option that we consider is to use econometric models to provide estimates of expected inflation, and then to subtract these from nominal yields. Like survey-based measures, such real rate estimates will include an inflation risk premium, and will also be affected by other factors, technical or otherwise, that are associated with the nominal gilt market. But, unlike survey-based estimates, we can assess the informational content of econometric-based forecasts more directly: we know the model structure, so we can estimate the model to exploit the information in the data as efficiently as possible. We can also construct error bands around our forecasts to provide some assessment of how 'accurate' our estimates are likely to be based on historical experience.

[(7)] Excluding responses from the general public for annual RPI inflation one year ahead.

The econometric models we consider here are reduced-form vector autoregressions (VARs) and autoregressive (AR) models – where the log difference of the RPI index (RPI) is the inflation variable (to ensure consistency with the estimates obtained from IGs). We base our choice of VAR on the models used by Dotsey and Scholl (2000) – who use such models to forecast the US real interest rate – and by Neiss and Nelson (2001) – who estimate a VAR for the United Kingdom. As well as the inflation measure, we include the quarterly average of the three-month Treasury bill rate[8] (TB) as our interest rate variable in all our models. The first VAR includes the change in the quarterly average claimant count unemployment rate (UR) as an additional variable,[9] while the second uses the annualised log differences of (non-durable) consumption per capita (C). The final VAR includes all four variables. All models are estimated on quarterly data from 1975 Q1 to 2001 Q4.

In addition to this, we use dummy variables to take into account the seasonal pattern of RPI, and the short-run effect of indirect tax changes on inflation such as the near-doubling of the VAT rate in 1979 Q3, and the 1990 Q2 introduction of the community charge. We also introduce a step dummy for the introduction of inflation targeting in 1992 Q4. Although this is sufficient to obtain a well-specified VAR, it may be insufficient to account for the implications of the structural break for the inflation forecasts in the period immediately after the adoption of inflation targeting (see below).

Given that we are interested in constructing and comparing *ex-ante* measures of the real interest rate, we restrict ourselves to using only the data available at the time of the forecast; which implies recursive estimation of the VAR model. Subject to ensuring a rigorous econometric specification, the choice of model for our reference set is driven by forecast performance measured by root mean square errors (RMSEs) and mean average errors (MAEs) for out-of-sample forecasts.[10]

A summary of the performance of the different types of model is shown in Table 2.A below. It is difficult to choose a specification that performs well across all criteria, but on balance, the VAR(3) model in inflation, unemployment and the three-month Treasury bill rate seems to be best out of the class of VAR models that we consider. Although the AR models have lower

[8] We use the T-bill rate to extend the sample period. Although the secondary market was relatively illiquid for much of the period and the correlation between its daily change and the daily change in the three-month GC repo rate is very low, the correlation between the quarterly average changes is close to one.
[9] There was insufficient evidence to reject the null hypothesis that the unemployment rate was I(1) over the sample period.
[10] We crosscheck these results by estimating rolling regressions, and these are consistent with the best recursive specifications.

MSEs than the VARs,[11] we choose the latter because AR models seem to be even worse at predicting turning points in inflation than their VAR counterparts (which is something that will be of obvious interest to policy-makers).

Table 2.A

RPI forecasts: RMSE and MAE for one-year and three-year forecasts

Model	RMSE results(1)		MAE results(1)		Full sample info.criteria		Criteria by
	Recursive estimation 1975 Q1 – 2001 Q4(2)				1975Q1 – 2001Q4		which the model
	1yr	3yr	1yr	3yr	Akaike	Schwarz	is 'best'
	"R1"	"R3"	"M1"	"M3"	"A"	"S"	
RPI, TB, C (3)	2.30	2.10	1.68	1.58	2.02	2.43	R1,R3,M1,M3,S
RPI, TB, C (7)	2.46	2.61	1.88	2.05	1.86	2.59	A
RPI, TB UR (3)	2.24	2.03	1.79	1.58	2.05	2.46	R1,R3,M1,M3,S
RPI, TB UR (7)	2.95	2.73	2.14	2.18	1.97	2.69	A
RPI, TB, C, UR (3)	2.38	2.17	1.79	1.64	2.05	2.53	R1,A,S
RPI, TB, C, UR (5)	2.39	1.99	1.73	1.57	2.12	2.81	R3, M1, M3
Memo AR models:							
AR (6)	1.97	2.37	1.61	1.80	2.08	2.41	S
AR (10)	1.93	2.17	1.60	1.70	2.05	2.49	R1, R3, M1, M3,A

(1) Results based on forecast growth rates.

(2) Data are used from 1975 Q1 to make forecasts beginning in 1985 Q1.

Chart 2.4 shows the one-year RPI forecast from our chosen VAR, compared with actual inflation outturns; the dashed lines represent a band of one standard error around the point forecast. By comparison with the survey and IG-based measures of inflation expectations, the point estimate from the econometric model shows less sign of a consistent upward bias in the latter part of the 1990s.[12] On the whole – but especially in the latter half of the sample – inflation outturns have remained well within one standard error of our best point estimate confidence band.

[11] This finding is consistent with, for example, Canova (2002) who finds that VAR models with fixed coefficients may not be much better than simple AR models at forecasting G7 inflation.

[12] Although it appears from Chart 2.4 that forecast errors are serially correlated; this pattern may be driven by the method for calculating one-year inflation forecasts. This is done by computing forecasts for each of the next four quarters at each point in time, which are then combined to give a one-year inflation forecast – and it is this methodology that imparts the pattern seen in the chart.

Chart 2.4

One-year econometric inflation forecast with error bands

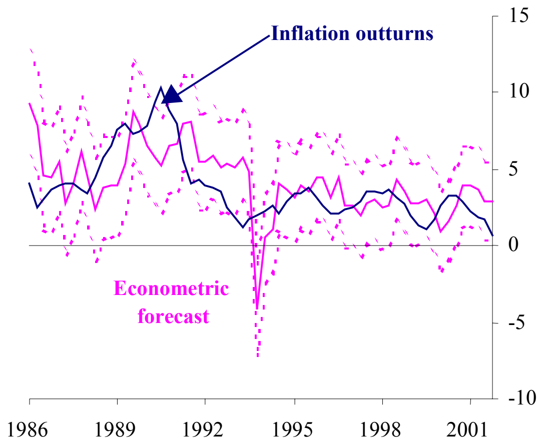

As noted above, we introduce a step dummy for the inflation-targeting period. There is a clear downward shift in the mean of the inflation forecast during this period; but there is also a sharp downward spike from 1992 Q4 -1993 Q3 before the forecast begins to fluctuate around its new, lower mean.[13] Although we use a dummy variable to account for this in the VAR specification, the recursive quarterly inflation forecast for 1992 Q4 will still be affected – and therefore so will the annual forecasts for the next three quarters. Where we calculate a real interest rate in the charts shown below, we shade out the period from 1992 Q4-1993 Q3 and concentrate on pre and post-inflation targeting periods for most of our analysis.

2.4 Reference measures of the real interest rate

Drawing together these measures, we can construct a reference set of real interest rate measures at most horizons. For ease of exposition we consider only one, three and ten-year rates, with two different measures at each maturity, shown in Charts 2.4, 2.5 and 2.6 respectively. At short horizons, where no IG-based measure is available, we use the survey and the model-based estimates.[14] The two co-move substantially[15] – the correlation coefficient is 0.76 – and the maximum deviation between the series is 1.5 percentage points during the inflation-targeting period. The survey-based measures are also generally well within one standard error of the model-based estimates, so statistically we would not reject the null hypothesis of their being the same.[16]

---

(13) This is shown as 1993 Q4-1994 Q3 in Chart 2.3 as we are comparing one-year inflation forecasts with outturns.

(14) In order to compare the real rate derived from the Basix survey with the econometrically based real rate, we treat the Basix forecast of inflation twelve months ahead as a spot inflation forecast.

(15) Excludes the 1992 Q4-1993 Q3 period.

(16) Although, perhaps unsurprisingly the standard error bands are fairly wide – with a sample average of 2.9 percentage points.

Chart 2.5

Econometric and survey-based measures of the one-year spot real rate

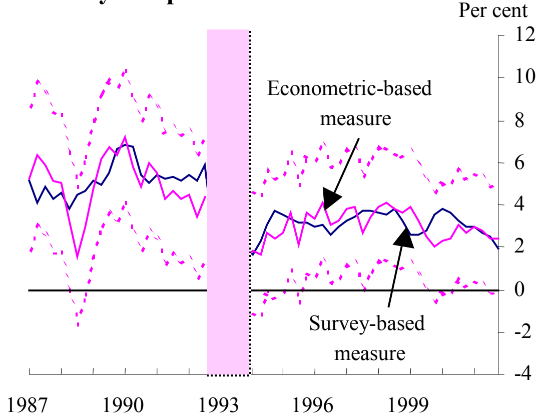

Dotted line indicates plus or minus one standard error around the forecast.
In 2001 Q3 no survey was carried out, line interpolated over this point.

At longer maturities, IG-based measures are available. At three and ten-year horizons, the average real yields based on the econometric measures are very close to their IG counterparts. And there is also a high degree of comovement between the two series, which is particularly noticeable during the period of inflation targeting for which the correlation coefficient is around 0.8 at both maturities.

Chart 2.6

Econometric and IG-based measures of the three-year spot real rate

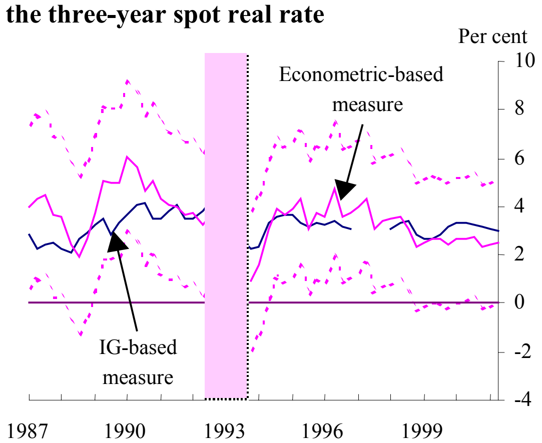

Dotted line plus or minus one standard error around the forecast.
IG based measure unavailabe between 1997 Q1 and 1997 Q4.

Chart 2.7

Econometric and IG-based measures of the ten-year spot real rate

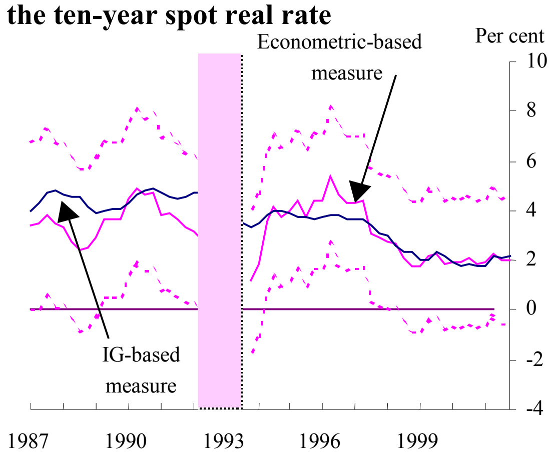

Dotted line plus or minus one standard error around the forecast.

At the ten-year horizon, both measures fall in the period after 1997: this (possibly MFR-related) fall is something that we will want to cross-check with the consumption-based measures of Section 5. At all horizons, current real interest rates appear to be at or around their historical lows.

3. Theoretical concepts

The consumption capital asset pricing model (C-CAPM) was first proposed by Breeden (1979), extending previous work by Sharpe (1964) and Lintner (1965) on the CAPM. Although the earlier work assumed that investors were concerned primarily with the mean and variance of portfolio returns, Breeden (1979) brought consumption into the model, unifying the fundamental theory of asset pricing with the utility-maximisation problem of the representative agent. Although the limitations of the C-CAPM are widely discussed, the model remains the workhorse of the modern finance literature – see, for example, Cochrane (2001) – and much of New Keynesian macroeconomics is based on the utility-maximisation problem outlined below.

To facilitate our exposition, we first outline the conceptual issues at a general level, following the approach in Campbell (1999) where possible. We begin with the standard intertemporal optimisation problem where an investor, who can trade freely in a set of assets indexed by *i* with prices P_{t+j}^i, maximises the expectation of a time-separable utility function in each time period by choosing consumption[17] C_{t+j} and a portfolio allocation $\{X_{t+j}^i\}$:

$$MaxE_t\left[\sum_{j=0}^{\infty}\delta^j U(C_{t+j})\right]$$
$$\text{s.t. } C_{t+j} + \sum_i P_{t+j}^i X_{t+j}^i = \sum_i \left(1+R_{t+j-1}^i\right)X_{t+j-1}^i \qquad \textbf{(3.1)}$$

Here δ is the time discount factor and $U(C_{t+j})$ is the period utility of future consumption at time *t+j*. The constraint implies that the real value of the investor's assets carried forward from the previous time period, X_{t+j-1}^i, with a return R_{t+j-1}^i, must be equal to consumption in period *t+j* added to the real value of the assets remaining at time *t+j*.

The first-order condition of this maximisation problem is given by:

$$P_t^i = \delta E_t\left[(1+R_t^i)\frac{U'(C_{t+1})}{U'(C_t)}\right] \qquad \textbf{(3.2)}$$

And for a one-period risk-free asset with a price of 1:

[17] As is standard in the C-CAPM literature, we assume that consumption follows a deterministic process in the analysis that follows.

$$1 = \delta E_t \left[(1 + R_t^f) \frac{U'(C_{t+1})}{U'(C_t)} \right] \quad \textbf{(3.3)}$$

Next, we define a term known as the one-period 'stochastic discount factor', M_{t+1} (sometimes referred to as a pricing kernel) and equal to the discounted ratio of marginal utilities of consumption:

$$M_{t+1} = \delta U'(C_{t+1}) / U'(C_t) \quad \textbf{(3.4)}$$

Then, by substituting **(3.4)** into **(3.3)** we get an equation often referred to as the fundamental asset pricing condition:

$$1 = E_t[(1 + R_t^i) M_{t+1}] \quad \textbf{(3.5)}$$

M_{t+1} is defined – in this specific example – as the discounted ratio of the marginal utilities of consumption in the current period and next, and is known as the 'one-period stochastic discount factor'.[18]

So in this paper, we assume that agents maximise a utility function of the form specified in equation **(3.1)**. This implies that the level of the real interest rate can be determined by aggregate non-durable consumption data.[19] Using the further (restrictive) assumption of power utility, the utility function can be expressed as:

$$U(C_t) = \frac{(C_t^{1-\gamma} - 1)}{(1 - \gamma)} \quad \textbf{(3.6)}$$

And the stochastic discount factor under power utility can be expressed in terms of the *k*-period-ahead consumption growth rate:

$$M_{t+k} = \delta (C_{t+k} / C_{t+k-1})^{-\gamma} \quad \textbf{(3.7)}$$

where γ is the coefficient of relative risk aversion.

(18) The general condition given by equation **(3.5)** follows merely from the assumption of no arbitrage. It does not require that investors maximise well-behaved utility functions, or that utility depends solely on consumption; this is just the way in which we have chosen to specify it. For a more detailed explanation see Cochrane (2001, pages 69-74).

(19) We use non-durable consumption data, because the relevant concept for this sort of analysis is the flow of consumption services the investor receives in any given quarter.

From now on – for ease of notation – we only consider risk-free assets, or those which have zero correlation with the stochastic discount factor (this enables us to drop the superscript on R). But we add another subscript to define the remaining maturity of the asset, so $R_{k,t+j}$ refers to the real holding period return[20] on a k-period risk-free bond at time $t+j$, and $P_{k,t+j}$ is its price. The relationship between price and yield is given by the following equation:

$$1 + R_{k,t+1} = \frac{P_{k-1,t+1}}{P_{kt}} \qquad \textbf{(3.8)}$$

In other words, the change in the price of a k-period bond purchased at time t (P_{kt}) over the period t to $t+1$, at the end of which the bond's price is given as $P_{k-1,t+1}$, is equal to its return (or one-period yield) over that period $R_{k,t+1}$. Substituting **(3.8)** into **(3.5)**, we get the following relationship:

$$P_{kt} = E_t[P_{k-1,t+1} \cdot M_{t+1}] \qquad \textbf{(3.9)}$$

And by forward-substitution, we can write **(3.9)** as the product of k stochastic discount factors:

$$P_{kt} = E_t[M_{t+1} \cdot M_{t+2} \cdot .... \cdot M_{t+k}] \qquad \textbf{(3.10)}$$

So, by substituting **(3.7)** into **(3.10)** we derive an expression for the price of a k-period bond in terms of the expected k-period consumption growth:

$$P_{kt} = \delta^k E_t \left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \qquad \textbf{(3.11)}$$

Once bond prices are known, we can calculate their associated yields and therefore a term structure of (real) interest rates. And, working with continuously compounded yields for ease, the log price of a k-period bond is related to the log of its yield to maturity by:

$$y_{kt} = -\frac{1}{k} p_{kt} \qquad \textbf{(3.12)}$$

(20) Where the holding period is defined as a quarter.

So, assuming that the stochastic discount factor is conditionally lognormal,[(21)] we take logs of equation **(3.9)**[(22)] and substitute for the *k*-period bond price from **(3.10)**, to get (after some rearrangement) an expression for the real yield to maturity on a *k*-period bond:

$$y_{kt} = -\log\delta + \frac{\gamma}{k} E_t \Delta c_{t+k} - \frac{\gamma^2}{2k}\sigma_c^2 \qquad \textbf{(3.13)}$$

where Δc_{t+k} is the change in consumption from period *t* to *t*+*k* and $\sigma_c^2 = Var(\Delta c_{t+k} - E_t \Delta c_{t+k})$. The equation implies that the real yield is linear in expected *k*-period-ahead consumption growth, with a slope coefficient equal to the coefficient of relative risk aversion divided by the maturity of the bond. The conditional variance of consumption growth has a negative effect on the riskless rate, which is often interpreted as a precautionary savings effect. In other words, investors are risk averse, and when their future consumption stream is more uncertain, they will tend to save more, which means that the real interest rate will – *ceteris paribus* – be lower than otherwise.

So far, our statements are still reasonably generic: in addition to the (strong) assumptions about either complete markets or representative agents, we have imposed that utility is time-separable and described by the power-utility form. In Section 5, we will discuss the implications of preference specifications that are non-separable over time and different forms of the utility function.

4. The consumption forecast

As shown in Section 2, under certain restrictive assumptions, a forecast for aggregate consumption implies a unique stochastic discount factor, and in turn a risk-free real interest rate. So in order to calculate a *k*-period real rate, we need a *k*-period ahead consumption forecast, and in this section, we describe the methods used to obtain such forecasts.

The 'quality' of such an estimate will depend on the accuracy of the forecast and how close it is to agents' expectations of future consumption growth. To fix ideas, assume that there is a representative, rational agent. Under these assumptions, (log) consumption growth in the next

---

(21) This follows from the assumption that aggregate consumption is conditionally lognormal – see Hansen and Singleton (1983). When a random variable *X* is conditionally lognormally distributed, the following relationship also holds: $\log E_t X = E_t \log X + \frac{1}{2} Var_t \log X$.

(22) If, in addition, the r.v. X is conditionally homoscedastic then: $Var_t \log X = E[(\log X - E_t \log X)^2]$ $= Var(\log X - E_t \log X)$.

period (Δc_{t+1}) is given by the sum of the representative agent's forecasts of consumption growth ($E_t \Delta c_{t+1}$) and the associated one-period ahead forecast error u_{t+1}:

$$\Delta c_{t+1} = E_t \Delta c_{t+1} + u_{t+1} \qquad \textbf{(4.1)}$$

The forecast error in period t+1 should be orthogonal to all past information, although it may have time-varying variance σ^2_{ut}. We next assume that the forecast can be explained by a set of variables (X) – assumed to be observable to both the agent and the econometrician – and a component v_{t+1}, that is observed by the former, but not the latter. In other words, the econometrician needs to model the evolution of the rational expectation consumption forecast, and estimates one-step-ahead consumption growth as:

$$E_t \Delta c_{t+1} = X_t \beta + v_{t+1} \qquad \textbf{(4.2)}$$

where X_t is the current observation of the explanatory variables and β is estimated by linear regression over the full sample. Because we allow for the possibility of time-varying variance for the unobservable component (σ^2_{vt}), the econometric estimate of future consumption growth $\hat{E}_t \Delta c_{t+1}$ is given by the generalised least squares estimator:

$$\hat{E}_t \Delta c_{t+1} = X_t (X' \Omega^{-1} X \ )^{-1} X' \Omega^{-1} \Delta c \qquad \textbf{(4.3)}$$

where $X' = \{ X'_1, X'_2, ..., X'_t \}$ and Ω is a diagonal matrix with the sum of the (time-varying) variances of u and v $\{\sigma_{ui} + \sigma_{vi}\}_{i=1,..,T}$ as its diagonal elements.[23]

The difference between the econometrician and representative agent's consumption growth forecast errors are given by (see Appendix 1 for the full derivation):

$$\begin{aligned} Var(\hat{E}_t \Delta c_{t+1} - E_t \Delta c_{t+1}) &= E[\{X_t (X' \Omega^{-1} X \ )^{-1} X' \Omega^{-1} (u \ + v \ ) - v_t\}^2] \\ &= (1 - 2\rho_t) X_t (X' \Omega^{-1} X \ )^{-1} X_t' + \rho_t (\sigma^2_{ut} + \sigma^2_{vt}) \qquad \textbf{(4.4)} \end{aligned}$$

where $u' = \{u'_1, u'_2, ..., u'_t\}$, $v' = \{v'_1, v'_2, ..., v'_t\}$ and:

[23] Both X and Ω are time-varying matrices that grow in dimension with t. For ease of notation, we have left out time-subscripts on these matrices, but the variance bounds are calculated using time-varying rather than full-sample matrices.

$$\rho_t = \frac{\sigma_{vt}^2}{\sigma_{ut}^2 + \sigma_{vt}^2} \tag{4.5}$$

Given that the individual variances σ_{ut}^2 and σ_{vt}^2 are unknown and unobservable, the variance of the difference between the econometric estimate and the agent's estimate cannot be computed directly. But under the assumption that *v* is orthogonal to *u,* the individual variance terms sum to the diagonal elements of Ω. And, under this assumption, ρ_t is given by equation **(4.5)**. This implies that ρ_t will approach its lower bound of zero when agents' expectational errors are relatively large, and its upper bound of one when the econometrician's model is a relatively poor fit. So formally, the variance bounds on the consumption growth forecast are given by:

$$X_t(X'\Omega^{-1}X\ )^{-1}X'_t \leq Var(\hat{E}_t\Delta c_{t+1} - E_t\Delta c_{t+1}) \leq (\sigma_{ut}^2 + \sigma_{vt}^2) - X_t(X'\Omega^{-1}X\ )^{-1}X'_t \tag{4.6}$$

The width of the bounds obviously depends on the performance of the econometrician's model, so it is therefore important to carefully consider the choice of model for our consumption forecast. If the econometric model perfectly represents the information used by the rational agent, then the variance bounds will be tight whereas the bounds will be a lot wider if the agent were to place no weight on the econometric model. In practice though, it is impossible to know the position of the 'true' bounds between these two extremes – and, as shown below, the maximum variance bounds tend to be quite wide.

We consider a number of simple linear models including an AR model of consumption and some of the VARs estimated in the previous section. The models all tend to produce large variance bounds for the real interest rate estimates. But we stress that these models are just examples – in principle, any consumption forecast could be used. Based on the available back-run of data,[24] and subject to achieving a robust econometric specification, the VAR(3) consumption model gave forecasts with the lowest RMSE and MAE across the class of models considered – at least for one and three-year forecasts (see Table 4.A below).[25]

---

[24] We wanted to estimate real interest rates across as long a time frame as possible.

[25] We also looked at a VAR including Treasury bill yields. But these data were available over a more restricted time period, and the forecasts from such a VAR had a higher MSE anyway.

Table 4.A

Consumption forecasts: RMSE and MAE for one, three and ten-year forecasts

	RMSE results (1) Recursive estimation 1963Q1 -2001Q4 (2)			MAE results (1) Recursive estimation 1963Q1 -2001Q4 (2)		
	1yr	3yr	10yr	1yr	3yr	10yr
AR (3)	35.60	86.78	164.17	28.29	70.81	143.92
VAR (3): C, UR	35.35	86.65	171.48	28.31	70.28	150.29

(1) Results based on non-durable consumption per capita in levels implied by forecast consumption growth rates.

(2) RMSE and MAE for one-year forecast based on 1974 Q3-2001 Q4, three-year 1976 Q3-2001 Q4 and ten-year 1983 Q3-2001 Q4.

5. The consumption models and estimates

This section uses the general conceptual framework outlined in Section 2 to calculate a *k*-period real interest rate using aggregate actual and forecast non-durable consumption data. After a short digression on some of the shortcomings of the basic consumption-based approach, we explore the issue of calibration and the sensitivity of our basic estimates to these assumptions. As a cross-check, we also compare our estimates with the reference measures presented in Section 3.

We then consider a variant of the model, suggested by Campbell and Cochrane (1995, 1999) which allows for habit formation in consumption. This should solve some of the problems associated with the basic model highlighted in the next section. Finally, we attempt to derive and estimate a *k*-period real interest rate with habit formation.

5.1 Sensitivity to the basic framework

The limitations of the model specified in equation **(2.13)** are well known – see Campbell (1999) for an excellent summary. The first is known as 'the equity premium puzzle', first outlined by Mehra and Prescott (1985). Within the general asset-pricing framework described in equation **(2.3)** above, the equity risk premium is measured as the covariance between the stochastic discount factor and the return on equities. As discussed above, in the more general consumption-based framework, the stochastic discount factor is replaced by expected log consumption growth, and this covariance can be estimated subject to calibrating the parameter γ. The logarithm of the equity risk premium is therefore given by the right-hand side of the following equation:

$$E_t[r_{i,t+1} - r_{f,t+1}] + \frac{\sigma_i^2}{2} = \gamma\sigma_{ic} \qquad \textbf{(5.1)}$$

That is to say the excess return on a risky asset (adjusted for the normal Jensen's inequality term σ_i^2), represented by $r_{i,t+1}$, minus the riskless rate $r_{f,t+1}$,[26] is equal to the product of the coefficient of relative risk aversion and the covariance between the return on equities and log consumption growth. But when this model is used to estimate risk premia for international equity returns data – eg Campbell (1999) – the implied estimates for the coefficient of relative risk aversion are huge. For example, for the United Kingdom, γ is around 40 for the post-1970 period – almost four times the maximum plausible value considered by Mehra and Prescott (1985).

But even if we were prepared to accept these implausibly high values for γ in order to solve the 'equity premium puzzle', this creates a second problem – known as 'the risk-free rate puzzle'. When using equation **(2.11)** to estimate real interest rates, 'sensible' estimates can only be achieved with an implausibly high level for the discount factor. In other words, a value of δ close to (or sometimes even larger than) one, implying that agents are virtually indifferent between (or actually prefer) consuming tomorrow rather than today.

One well-known disadvantage of the power utility specification is that it constrains the elasticity of inter-temporal substitution (ψ below) – or the inverse of the slope of equation **(2.13)** – to be equal to the coefficient of relative risk aversion. This is arguably an inappropriate restriction because the former links the willingness of an investor to substitute consumption across different time periods, whereas the latter refers to the willingness to substitute between different states of the world – see Hall (1988) for a discussion.

Epstein and Zin (1989, 1991) develop a more flexible version of the power utility model by removing this restriction, and their utility function is defined recursively as:

$$U_t = \left\{ (1-\delta) C_t^{\frac{1-\gamma}{\theta}} + \delta \left(E_t [U_{t+1}^{1-\gamma}] \right)^{\frac{1}{\theta}} \right\}^{\frac{\theta}{1-\gamma}} \qquad \textbf{(5.2)}$$

where: $\theta = (1-\gamma)/(1-1/\psi)$ and, when θ=1 this simplifies to the power utility function described above. And using the standard assumptions that asset returns and consumption are both homoscedastic and jointly lognormal, the equation for the one-period riskless real rate is given by:

$$y_{1t} = -\log\delta + \frac{1}{\psi} E_t \Delta c_{t+1} + \frac{\theta - 1}{2}\sigma_w^2 - \frac{\theta}{2\psi^2}\sigma_c^2 \qquad \textbf{(5.3)}$$

---

(26) Where $r_{i,t+1} = \log(1+R_{i,t+1})$ and $r_{f,t+1} = \log(1+R_{f,t+1})$.

where σ_w^2 is the variance of wealth. Within this framework, a high risk-aversion coefficient does not does not necessarily imply a low average risk-free rate because the slope coefficient of **(5.3)** does not depend on γ. But given that there is direct empirical evidence pointing to a low elasticity of inter-temporal substitution anyway, this leaves another unresolved puzzle.

It is also much easier to explain the equity premium puzzle within this framework because the excess return on a risky asset depends on its covariance with the aggregate wealth portfolio as well as aggregate consumption. But again, this creates a further problem. Even if we assume the aggregate wealth portfolio to be much more volatile than consumption, the two are linked via the inter-temporal budget constraint. So volatile wealth is another puzzle; which must be explained in the light of the observation of 'smooth' consumption data.

One possible explanation is that risk aversion varies over time. And in response to this, a class of models exhibiting time non-separability in consumption were developed. These 'habit formation' models have the property that consumers only derive utility from that part of consumption which is in excess of some subsistence or 'habit' level.

Some discussion exists in the literature with regard to how the level of habit should be determined. 'Internal' habit models such as Constantinides (1990) and Sundaresan (1989) suggest that habit should be related to the agent's own historic consumption level. But other papers suggest an 'external' habit model; see for example, Abel (1990, 1999) Campbell and Cochrane (1995, 1999). In these models habit is determined by aggregate consumption, that is to say agents are interested in 'catching up with the Jones' when deciding on their habit level of consumption.

We focus on external habit formation models, using a 'difference' rather than 'ratio' specification. Abel (1990) proposes the latter, which implies that the utility function is written as:

$$U_t = E_t \sum_{j=0}^{\infty} \delta^j \frac{(C_{t+j} / X_{t+j})^{1-\gamma} - 1}{1-\gamma} \qquad \textbf{(5.4)}$$

where X_t is the level of habit at time t, and is a function of last period's aggregate consumption $(X_t = C_{t-1}^{\kappa})$. Equation **(2.11)** is therefore augmented by an additional term in current consumption:

$$y_{1t} = -\log \delta + \gamma E_t \Delta c_{t+1} - \gamma^2 \sigma_c^2 / 2 - \kappa(\gamma - 1)\Delta c_t \qquad \textbf{(5.5)}$$

when κ is large, the coefficient of risk aversion can be increased to solve the equity premium puzzle without creating a 'risk-free rate puzzle'. But the addition of this final term also makes the risk-free rate more volatile – a general problem with habit formation models. So Campbell and Cochrane (1995, 1999) – hereafter known as CC – suggest a 'difference model', where the agents' utility function is specified as:

$$E_t \sum_{j=0}^{\infty} \delta^j \frac{(C_{t+j} - X_{t+j})^{1-\gamma} - 1}{1-\gamma} \qquad \textbf{(5.6)}$$

This tends to reduce the variability of real interest rates, while at the same time allowing for time-varying risk aversion, which is intuitively appealing. We explore this sort of model in more detail in Section 5.3 below.

5.2 Calibration of the consumption-based measure with power utility

First we consider a *k*-period real rate calculated by the basic consumption-based model with power utility. There are several steps needed to make this real interest rate operational. First we need a *k*-period consumption forecast, and the variance of the associated consumption forecast error. In principle any consumption forecast can be used to determine the real rate, although the variance bounds around the central projection will vary according to the standard error of consumption in the model. We use the VAR(3) model specified in Section 4 above.

Second we must calibrate the 'deep' parameters δ (the time discount factor) and γ (the coefficient of risk aversion). These are set so that the average *ex-ante* three-month real rate calculated by equation **(2.13)** is equal to the three-month *ex-post* real rate over the same horizon. Between 1975 and 2001, the average *ex-post* three-month real rate is equal to approximately 2.6% on an annualised basis.(27) But in order to set two parameters, we need at least one more constraint.

This is provided by theory and reference to the work of others. Chart 5.1 below shows a stylised representation of the calibration problem. The solid line in the chart is a locus of points where *ex-ante* three-month real yields are equal to their *ex-post* counterparts on average over the chosen sample period. The dotted lines show restrictions on γ and δ provided by theory: Mehra and Prescott (1985) suggest that γ should not be larger than ten; and δ must be less than one, otherwise people will prefer to consume tomorrow rather than today. This provides a local region in which to search for suitable parameters.

(27) Based on the consumption deflator price series.

Chart 5.1

Range of values for γ and δ

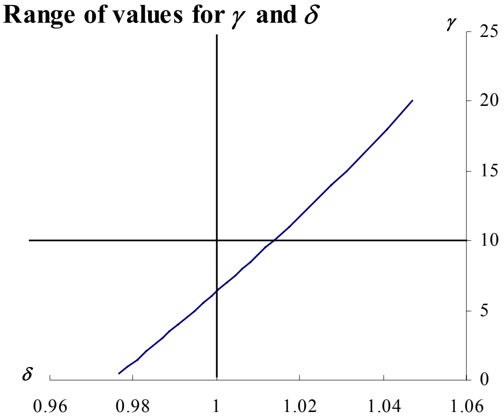

But the chart does illustrate the standard calibration problem in such models (discussed in more detail below). Although γ can be set in a plausible range, the associated value of δ is extremely high – implying that agents are almost indifferent between consumption today and tomorrow. We pick a combination of $\delta = 0.99$ and $\gamma = 3.8$, and explore the sensitivity of our measures to these parameters in more detail below. This is the standard value for δ chosen in most of the literature – see for example, McCallum and Nelson (1999) for the United States, or Neiss and Nelson (2001) for the United Kingdom.

Chart 5.2 shows the associated real rate measure at one, three and ten-year maturities. As we move further out along the yield curve, the real interest rate becomes less volatile – which is consistent with our prior that longer-run real interest rates should be determined more by real fundamentals rather than the current stance of monetary policy. In the model, as maturity increases, the average value of the second term in equation **(2.13)** falls while the third term becomes less negative. Although the changes in these two terms broadly offset each other – implying that mean real rates are broadly similar across maturities (at around the calibrated level), the variance of the expected average growth rate falls, resulting in lower real rate volatility as k increases.

Chart 5.2

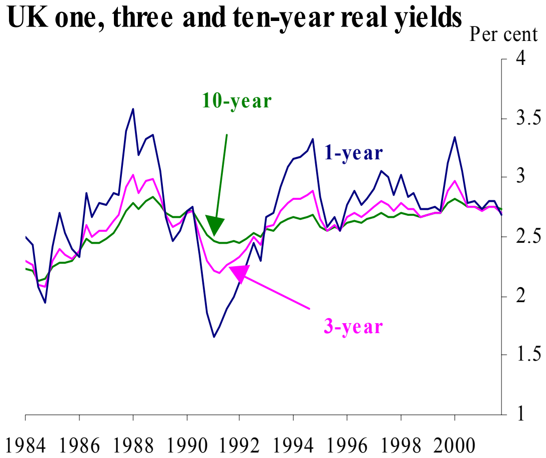

But how sensitive are these estimates to the calibration we choose? Assuming for the moment that we still want to match *ex-ante* and *ex-post* real yields, and restricting γ and δ to be in the 'theoretically' plausible zone shown by the dashed lines in Chart 5.1, Chart 5.3 below shows some different possible calibrations. As γ falls towards zero, the real rate becomes less volatile (reflecting lower risk aversion). Again, δ always takes a high value if we assume that γ should be bounded between zero and ten.

On the other hand, Chart 5.4 provides an illustration of the 'risk-free rate' puzzle. If we wish to reduce δ, we need to increase γ by a lot. But by maintaining the constraint that *ex-post* yields should be equal to *ex-ante* yields, the real interest rate series becomes extremely volatile (and implausible).

Chart 5.3

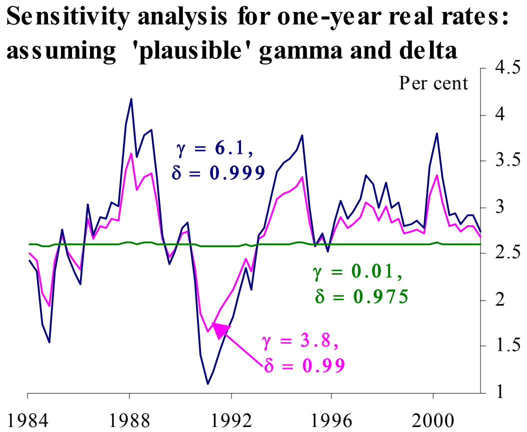

Chart 5.4

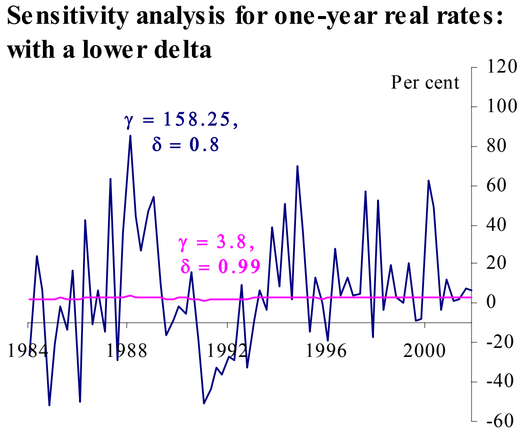

So, using the 'sensible' calibration of the model suggested above, together with the analysis in Section 4, we estimate variance bounds on our real rates. In Charts 5.5 and 5.6 the bold line represents the central estimate of one and three-year real interest rates and the lighter lines show the maximum (solid) and minimum (dotted) possible error band around the (unobservable) rational agents' forecast. Recall from Section 4 above that the dotted lines correspond to the situation where the econometrician's estimate is exactly equal to the rational expectations forecast and the solid lines where her model has zero explanatory power. Both the central estimate and the bounds become progressively smoother across time as k increases, but the average width of the bands remains virtually unchanged.

Chart 5.5

UK one-year risk free real interest rate

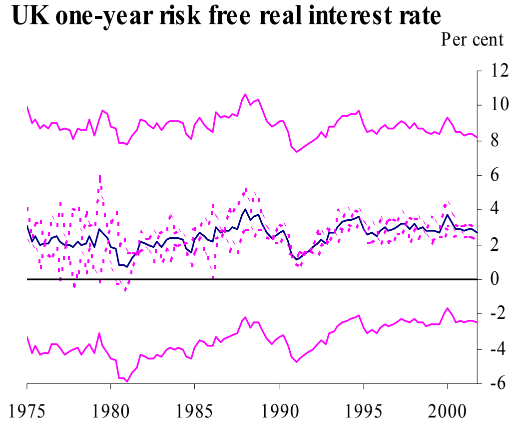

Chart 5.6

UK three-year annualised risk free real interest rate

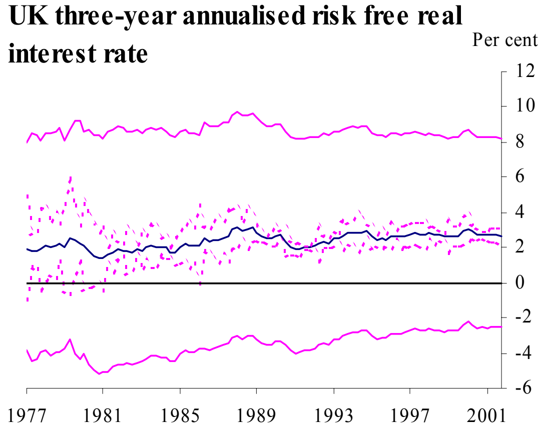

So how do these estimates compare with the reference measures we estimated in Section 3? Charts 5.7 - 5.9 show a comparison.[28] The consumption-based real rates are a lot smoother – possibly reflecting the fact that the reference measures are all based on gilt market yields in some way, which tend to be more volatile.[29][30] Again, the fit seems to be much better after the introduction of inflation targeting in late 1992, possibly suggesting that real interest rates are easier for investors to forecast. Interestingly, Chart 5.9 shows that the fall in ten-year IG yields starting in 1997 is not replicated in our simple consumption-based measure, suggesting that the

[28] Within the reference measures, the four years after the ERM exit are excluded from the 'econometric measure' – for reasons discussed above – and the three-year IG measure is incomplete because of data unavailability.
[29] As noted above, the models are calibrated so that the mean of the estimated one-quarter consumption-based real rate is set equal to the sample average of the *ex-post* real interest rate. Because we attempt to match the mean of the data more closely, it is possible that the volatility of the consumption-based measures may therefore be more different. In addition – measures that allow for habit formation in consumption allow for more volatility and appear to match the data more closely – see Section 5.3 below.
[30] In addition, it may actually be inappropriate to compare our consumption-based real interest rate measures with those derived from 'risk-free' bond prices. In practice, households are credit-constrained, and to the extent that aggregate consumption growth reflects these constraints, it may be appropriate to compare our measure with 'actual' borrowing rates faced by households – eg 'real' mortgage rates. But in practice, such rates are not readily available and are difficult to calculate. Another factor not explicitly considered here is the effect of different tax rates on the real interest rates faced by different individuals/investors.

recent falls in the former may have been driven largely by institutional factors rather than 'real fundamentals'.

Chart 5.7

One-year consumption-based and 'reference' real rates

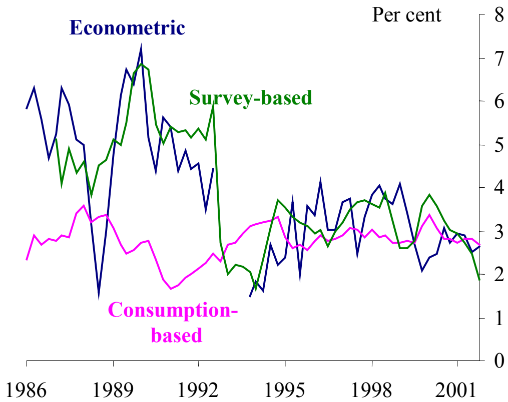

Chart 5.8

Three-year consumption-based and 'reference' real rates

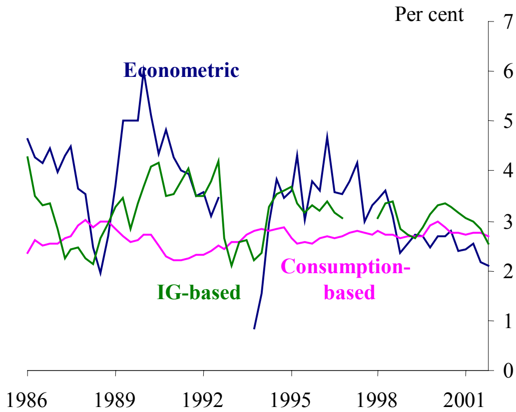

Chart 5.9

Ten-year consumption-based real rates and IG yields

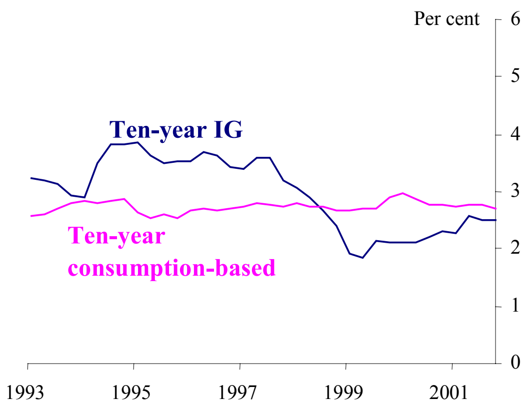

5.3 A consumption-based measure with habit formation

As noted above, we concentrate on external habit formation models. So specifying the utility function as in equation **(5.6)** above:

$$E_t \sum_{j=0}^{\infty} \delta^j \frac{(C_{t+j} - X_{t+j})^{1-\gamma} - 1}{1-\gamma} \qquad \textbf{(5.6)}$$

where X_{t+j} is the level of habit at time $t+j$. The model is written in terms of the surplus consumption ratio, because consumption only yields utility to the extent that it exceeds habit. The surplus consumption ratio is therefore calculated as:

$$S_t \equiv \frac{C_t - X_t}{C_t} \qquad \textbf{(5.7)}$$

Consumption is assumed to follow a random walk – which is not inconsistent with the empirical evidence. Campbell and Cochrane (CC) also model the log surplus consumption ratio (s_t)[31] as an AR(1) process which evolves over time according to the following equation:

$$s_{t+1} = (1-\phi)\bar{s} + \varphi s_t + \lambda(s_t)\varepsilon_{c,t+1} \qquad \textbf{(5.8)}$$

The parameter φ governs the persistence of this ratio and habit is a non-linear function of current and past consumption. In contrast to other types of habit formation model described above, the lambda function (described in more detail in Appendix 2) allows us to control the sensitivity of the real interest rate to innovations in consumption growth $\varepsilon_{c,t+1}$. This can eliminate the problem of real interest rates that are too variable – for example, in CC, the real interest rate is calibrated to be constant over time (although for our purposes, this is not very appealing).

The one-period real interest rate is therefore given by:[32]

$$y_{1,t+1} = -\log\delta + \gamma g - \gamma(1-\varphi)(s_t - \bar{s}) - \tfrac{1}{2}\gamma^2\sigma_c^2(\lambda(s_t)+1)^2 \qquad \textbf{(5.9)}$$

which has both similarities and differences to the power utility model outlined in equation **(3.13)**. The first two terms are similar – although the sample average growth rate has replaced the one period ahead expectation of consumption growth in the second term, reflecting the assumption that consumption follows a random walk. But the third and fourth terms are new – the former reflecting intertemporal substitution and the latter a different precautionary savings term.

For the third term, if the surplus consumption ratio is below its steady state $\bar{s}$, marginal utility is expected to fall in the future, which implies that consumers would like to borrow rather than save, driving up real interest rates. Intertemporal substitution is driven by an additional channel: mean reversion in marginal utility rather than consumption – which is modelled as a random walk.

The fourth term implies that as uncertainty increases, consumers become more willing to save which acts to depress the real interest rate *ceteris paribus*. Although uncertainty about consumption is constant in this model (because consumption is homoscedastic), habit formation makes a given level of consumption uncertainty more serious for marginal utility when consumption is low relative to habit through the lambda function.

(31) The log specification means that consumption cannot fall below habit.
(32) See Appendix 2 for a more thorough derivation.

Next we turn to the estimation of a one-period real interest rate with habit formation using equation **(5.9)** above. There are several issues to be resolved. First, the surplus consumption ratio must be estimated as a recursion, with the initial level of habit set as the average aggregate consumption level over the previous year. It is then possible to calculate the lambda function at time *t*. Using the calibration of $\delta = 0.99$ and $\gamma = 3.8$ suggested above for the power utility model, we still need to calibrate the persistence of the log surplus consumption ratio (φ). CC suggest that φ should be close to (but less than) 1: we choose this parameter by setting the sample average of the one-period real interest rate equal to its *ex-post* value, as above. In doing so, this yields $\varphi = 0.984$.

A time series of the associated real interest rate is given by Chart 5.10 below.[33] The chart shows three-month (annualised) real interest rates. Real interest rates reach a peak of over 5% in 1982, and then fall for the remainder of the 1980s, before rising back to just over 2% in mid-1992. After this real interest rates fell to their current level of 1% – close to historical lows.

The 'UK calibration' of our habit formation model is compared with the 'CC95' and 'CC99' calibrations applied to UK data. These refer to the calibrations used in Campbell and Cochrane's original *NBER Working Paper* (where the possibility of time-varying real interest rates is allowed) and their 1999 *JPE* article (where equation **(5.9)** is calibrated to ensure a constant real interest rate for US data) respectively. The series are similar, although the former is more volatile, and the latter less volatile than our calibration, which is what we would expect. We prefer our calibration since it matches average *ex-post* to average *ex-ante* yields, but in practice, the parameters we derive are quite similar to CC's as noted above; and the important point is that we could derive a 'plausible' series for the real interest rate without recourse to an implausibly high (>1) discount factor.

[33]We do not show either the one-period real rate under power utility, or the 'reference' measures (based on inflation forecasts) because both are too volatile for a meaningful comparison to be made. On this criteria, the one-period real rate with habit formation looks to be a reasonably good fit, particularly since one of the well-known drawbacks of many habit-based measures is that the real interest rate is too volatile. Another possible comparison from 1975 onwards might be with *ex-post* real three-month T-bill rates. Again these are volatile: the series with habit formation lies above real interest rates based on T-bill yields from the mid-1970s - mid-1980s and below the series for the remainder of the period.

Chart 5.10

One-period real rate with habit formation

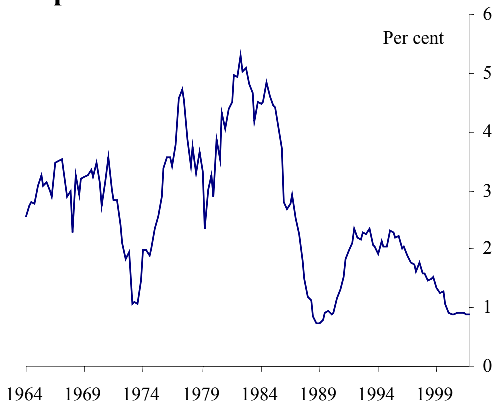

Chart 5.11

One-period real rate with habit formation - UK vs CC calibration

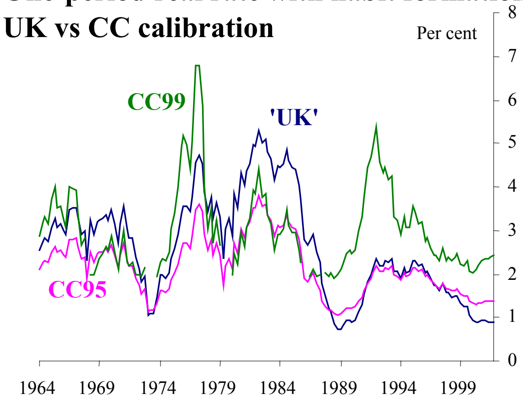

But we also extend CC's habit formation model to derive a *k*-period real interest rate. In the spirit of eclecticism, this will enable us to construct another comparator real interest rate, and – as far as we are aware – this is the first paper that currently performs such an exercise. The derivation is complicated, and is given in more detail in Appendix 2. The formula for a *k*-period real interest rate with habit formation is therefore given as:

$$y_{kt} = -\log\delta - \frac{\gamma}{k}(1-\varphi^k)(s_t - \bar{s}) + \gamma g - \frac{\gamma^2\sigma_c^2}{2k}\sum_{i=1}^{k}\left[\varphi^{2(k-i)} Var\lambda(s_{t+i-1})\right] - \frac{\gamma^2\sigma_c^2}{2}$$
$$-\frac{\gamma^2\sigma_c^2}{k}\left[E_t\sum_{i=1}^{k}\varphi^{k-i}\lambda(s_{t+i-1})\right] \quad \textbf{(5.10)}$$

In order to make this equation operational we need to compute expressions for the conditional expectation and variance of $\lambda(s_{t+i-1})$. Analytically, this cannot be done for $i > 1$, so we apply numerical techniques to obtain approximating functions for these, using a simple variant of the parameterised expectations algorithm – see Den Haan and Marcet (1990). We discuss this algorithm in further detail in Appendix 3.

Next we need to calibrate the three main parameters δ, γ and ϕ. δ determines the intercept of the real rate series, and in practice we have to set this at a high value, 0.999, to avoid excessively high real interest rates (although we are careful not to let it exceed 1 – see discussion in Section 5.1 above).[34] We set γ to 2 in accordance with CC: although this is low, steady-state risk

[34] We recognise that a discount factor of 0.999 is quite high, but for the current specification of the model, lower discount factors led to less desirable real interest rate series.

aversion is given by $\gamma / \bar{S}$, which will be much larger and therefore able to fit the observed equity risk premium. As noted above, φ should be less than, but close to 1, and we set it to be 0.99.[(35)]

Chart 5.12 shows a time series of one, three and ten-year real yields for the model with habit formation. There are two immediately striking features of these series. The first is that the term structure of real interest rates is quite flat, the second is that movements at different maturities are highly correlated over the sample period. Neither of these are features of the simple measures shown in Chart 5.2.

The first feature is caused by the calibration of γ. When γ is low, agents are less risk averse and so demand lower real interest rates to save over longer horizons, which tends to flatten the yield curve. So raising γ makes the yield curve more upward sloping. The second is caused by the assumption that consumption is a random walk. This implies that the surplus consumption ratio at time *t* is a sufficient statistic for the same ratio at all periods in the future (discounted by the persistence parameter). In other words, movements in real interest rates at different maturities are driven by the same underlying factor. Wachter (2002) is able to generate a more plausibly sloped yield curve by relaxing the assumption that consumption follows a random walk.

Chart 5.12

One, three and ten-year real yields with habit formation

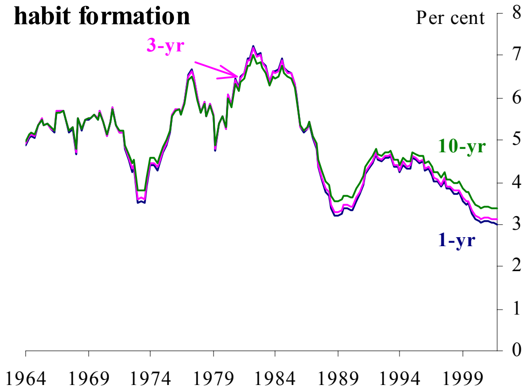

Chart 5.13

A comparison of ten-year real yields

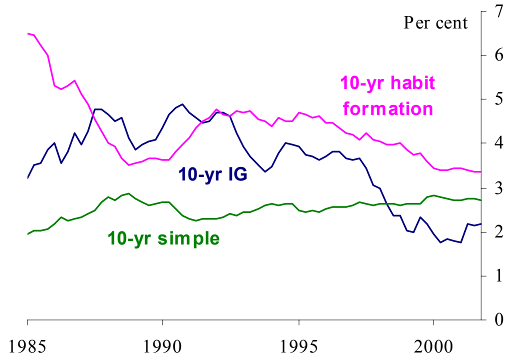

The third notable feature is the striking fall in real yields from their peak of around 7% in the mid-1980s to current – historically low – levels of around 3%. So interestingly, as shown in Chart 5.13, the series with habit formation suggest that a 'fundamentals-based' may explain at least part of the fall in real yields since 1997. This is because the rapid consumption growth of the late 1990s has driven the surplus consumption ratio well above its steady state, such that the second term in equation **(5.10)** exerts a significant downward impact on the real interest rate

(35) In practice, the model is quite sensitive to the calibration of these parameters - especially δ and φ, and there is not much room for variation before the model begins to give implausible estimates.

during this period. Indeed the habit-formation measure tracks the ten-year IG yield reasonably well over the period since 1988.

6. Conclusions and suggestions for further research

In this paper we have sought to estimate a broad set of *ex-ante* real interest rate measures at a selection of maturities, focusing particularly on a ‘theoretical’ consumption-based measure derived from the standard representative agent’s intertemporal optimisation problem. Although there were some persistent deviations across all these measures during the 1970s (where calculated) and 1980s, the various measures have moved together more closely since the introduction of inflation targeting in 1992. This could reflect increased macroeconomic stability over this period, which may have made both inflation and consumption easier to foreacst.

For the ‘theoretical’ real interest rates, we show that models based on the basic power utility specification of the C-CAPM are subject to the standard problems discussed in the literature. So we use a model with habit formation in consumption to generate *ex-ante* real interest rates at various maturities. All the real interest rates generated by this method exhibit substantially more volatility over time. This is a well-known feature of such models, and it is probably likely that the volatility of the actual real interest rate will be somewhere between these two.

We believe that calculating real interest rates using the C-CAPM with habit formation is a worthwhile avenue for future research. This version of the model appears to represent the best attempt so far to solve most of the standard problems within the asset-pricing literature, although we recognise that there are more issues to be resolved. That is why we place weight on developing a ‘menu’ of measures.

References

Abel, A B (1990), 'Asset prices under habit formation and catching up with the Joneses', *American Economic Review*, Papers and Proceedings, Vol. 80, pages 38-42.

Abel, A B (1999), 'Risk premia and term premia in general equilibrium', *Journal of Monetary Economics,* Vol. 43, pages 3-33.

Anderson, N L and Sleath, J (2001), 'New estimates of the United Kingdom real and nominal yield curves', *Bank of England Working Paper no. 126*.

Bakhshi, H and Yates, A (1998), 'Are United Kingdom inflation expectations rational?', *Bank of England Working Paper no. 81*.

Bank of England (1995), *Index-linked debt: papers presented at the Bank of England conference September 1995.*

Breeden, D (1979), 'An inter-temporal asset-pricing model with stochastic consumption and investment opportunities', *Journal of Financial Economics,* Vol. 7, pages 265-96.

Breedon, F (1995), 'Bond prices and market expectations of inflation', *Bank of England Quarterly Bulletin,* May, pages 160-65.

Campbell, J Y (1999), 'Asset prices, consumption and the business cycle', in Taylor, J B and Woodford, M (eds), *Handbook of macroeconomics,* Vol. 1, Elsevier Science.

Campbell, J Y and Cochrane J H (1995), 'By force of habit: a consumption-based explanation of aggregate stock market behaviour', *NBER Working Paper,* No. 4995, January.

Campbell, J Y and Cochrane, J H (1999), 'By force of habit: a consumption-based explanation of aggregate stock market behaviour', *Journal of Political Economy,* Vol. 104, pages 205-51.

Canova, F (2002), 'G-7 inflation forecasts', *ECB Working Paper,* No. 151.

Cochrane, J H (2001), *Asset pricing,* Princeton University Press.

Constantinides, G M (1990), 'Habit formation: a resolution of the equity premium puzzle', *Journal of Political Economy*, Vol. 98, pages 519-43.

Deacon, M and Derry, A (1994a), 'Deriving estimates of inflation expectations from the prices of UK government bonds', *Bank of England Working Paper no. 23*.

Deacon, M and Derry, A (1994b), 'Estimating the term structure of interest rates', *Bank of England Working Paper no. 24*.

Den Haan, W J and Marcet, A (1990), 'Solving the stochastic growth model by parameterizing expectations', *Journal of Business and Economic Statistics*, Vol. 8(1), pages 31-34.

Dotsey, M and Scholl, B (2000), 'Behavior of the real rate of interest over the business cycle', *Federal Reserve Bank of Richmond Working Paper,* No. 00-9.

Epstein, L G and Zin, S E (1989), 'Substitution and risk aversion, and the temporal behaviour of consumption and asset returns: a theoretical framework', *Econometrica,* Vol. 57, pages 937-68.

Epstein, L G and Zin, S E (1991), 'Substitution and risk aversion, and the temporal behaviour of consumption and asset returns: an empirical investigation, *Journal of Political Economy,* Vol. 99, pages 263-86.

Hall, R E (1988), 'Inter-temporal substitution in consumption', *Journal of Political Economy,* Vol. 96, pages 221-73.

Hansen, L P and Singleton, J L (1983), 'Stochastic consumption, risk aversion, and the temporal behaviour of asset returns', *Journal of Political Economy*, Vol. 91, pages 249-68.

Ireland, P N (1996), 'Long-term interest rates and inflation: a Fisherian approach', *Federal Reserve Bank of Richmond Economic Quarterly,* Winter, pages 21-36.

Lintner, J (1965), 'The valuation of risky assets and the selection of risky instruments in stock portfolios and capital budgets', *Review of Economics and Statistics,* Vol. 47, pages 13-37.

McCallum, B T and Nelson, E (1999), 'Performance of operational policy rules in an estimated semi-classical structural model', in Taylor, J B (ed), *Monetary policy rules*, University of Chicago Press, pages 15-45.

Mehra, R and Prescott, E C (1985), 'The equity premium puzzle', *Journal of Monetary Economics*, Vol. 15, pages 145-61.

Neiss, K S and Nelson, E (2001), 'The real interest rate gap as an inflation indicator', *Bank of England Working Paper no. 130*.

Scholtes, C T J (2002), 'On market-based measures of inflation expectations', *Bank of England Quarterly Bulletin,* Spring, pages 67-77.

Sharpe, W (1964), 'Capital asset prices: a theory of market equilibrium under conditions of risk', *Journal of Finance*, Vol. 19, pages 425-42.

Sundaresan, S M (1989), 'Inter-temporally dependent preferences and the volatility of consumption and wealth', *Review of Financial Studies,* Vol. 2, pages 73-88.

Wachter, J (2002), 'Habit formation and returns on bonds and stocks', *unpublished working paper,* New York University.
Available at: http://pages.stern.nyu.edu/~jwachter/research/habit.pdf

Appendix 1: Derivation of the standard error bands of the econometric consumption forecast around the rational representative agent's forecast

The value of a random variable in the next period is equal to the sum of the representative agent's forecast of it ($E_t y_{t+1}$) and the associated one period ahead forecast error u_{t+1}:

$$y_{t+1} = E_t y_{t+1} + u_{t+1} \qquad \textbf{(A1.1)}$$

The forecast error in period *t*+1 will be orthogonal to all past information, although it may have time-varying variance σ_{ut}^2.

If the forecast can be explained by a set of variables (*X*) – assumed to be observable to both the agent and the econometrician – and a component *v*, that is observed by the former, but not the latter, then the general linear regression model is given as:

$$y = X\beta + \varepsilon \qquad \textbf{(A1.2)}$$

Where ε is white noise. But the dependent variable we wish to model is actually the representative agent's consumption forecast one period ahead ($E_t y_{t+1}$). This can be calculated as:

$$E_t y_{t+1} = X_t \beta + v_{t+1} \qquad \textbf{(A1.3)}$$

Where X_t is the observation of the independent variables in the current time period, and v_{t+1} is an unknown forecast error. Because we allow for the possibility of

time-varying variance for this unobservable component (σ_{vt}^2), the econometric estimate of y_{t+1} is given by the generalised least squares estimator:

$$\hat{E}_t y_{t+1} = X_t (X' \Omega^{-1} X\ )^{-1} X' \Omega^{-1} y \qquad \textbf{(A1.4)}$$

And when this is estimated recursively, at time *t*: $X = \{X'_1, X'_2, ..., X'_t\}$, and Ω is a diagonal matrix with the sum of the variances of u_t and v_t, $\{\sigma_{ui} + \sigma_{vi}\}_{i=1,..,t}$ as its diagonal elements. Because consumption is assumed to be homoscedastic, the diagonal elements are equal to a constant, σ_{ε}^2, although the individual variances of u_t and v_t may vary over time.

Estimation by GLS implies that the new variance-covariance matrices of u_t and v_t, will be white noise. So that we can write:

$E(u_{GLS}u_{GLS}')=\sigma_u^2\mathrm{I}$, $E(v_{GLS}v_{GLS}')=\sigma_v^2\mathrm{I}$, and: $E(u_i\ u_i)=0 \quad j \neq i \quad E(v_i\ v_i)=0 \quad j \neq i$

And given what we know about the matrix Ω, we can expand these expressions as:

$$E(u_{GLS}u_{GLS}')=\sigma_u^2\mathrm{I}=\frac{\sigma_u^2}{\sigma_u^2+\sigma_v^2}\Omega \qquad \textbf{(A1.5)}$$

$$E(v_{GLS}v_{GLS}')=\sigma_v^2\mathrm{I}=\frac{\sigma_v^2}{\sigma_u^2+\sigma_v^2}\Omega \qquad \textbf{(A1.6)}$$

The difference between the econometrician and representative agent's consumption forecast errors can be written as:

$$Var\left(\hat{E}_t y_{t+1}-E_t y_{t+1}\right)$$

$$=Var\left(X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}\mathrm{y}-\left(X_t\beta+v_t\right)\right)$$

$$=Var\left(X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}\left(X\beta+u+v\right)-X_t\beta-v_t\right)$$

$$=E\left(\left\{X_t\left(X'\Omega^{-1}\mathrm{X}\right)^{-1}X'\Omega^{-1}\left(u+v\right)-v_t\right\}^2\right)$$

$$=E\left(\left(X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}\left(u+v\right)\right)^2-X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}\left(u+v\right)v_t\right.$$
$$\left.-v'_t\left(u'+v'\right)\Omega^{-1}X\left(X'\Omega^{-1}X\right)^{-1}X'_t+v_t^2\right)$$

$$=E\left[X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}u\mathrm{u}'\Omega^{-1}X\left(X'\Omega^{-1}X\right)^{-1}X'_t+\right.$$
$$X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}vv'\Omega^{-1}X\left(X'\Omega^{-1}X\right)^{-1}X'_t-X_t\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}vv'_t \qquad \textbf{(A1.7)}$$
$$\left.-v_t v'\Omega^{-1}X\left(X'\Omega^{-1}X\right)^{-1}X'_t+v_t^2\right)$$

Using the conditions given in **(A1.5)** and **(A1.6),** the first two terms can be written as:

$$=\frac{\sigma_u^2}{\sigma_u^2+\sigma_u^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t+\frac{\sigma_v^2}{\sigma_u^2+\sigma_v^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t$$

Next we try to simplify the third and fourth terms of **(A1.7)**. We start by noting that these will be a scalar, but also note that the expression $v_t v'\Omega^{-1}X$ (or its transpose) is equal to a (1 x *k*) row (column) vector. Where the variance of the econometrician's error at time *t* is divided by the sum of the rational expectations error and the econometrician's error, and multiplied by the *t*-th observation of the independent variables X_t . This allows us to write:

$$v_t v'\Omega^{-1}X = vv'\Omega^{-1}X_t \qquad \textbf{(A1.8)}$$

And so the third and fourth terms of **(A1.7)** simplify to:

$$-2\frac{\sigma_v^2}{\sigma_u^2+\sigma_v^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t$$

Which allows us to re-write the whole expression as:

$$Var\left(\hat{E}_t y_{t+1}-E_t y_{t+1}\right)=\frac{\sigma_u^2}{\sigma_u^2+\sigma_v^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t+\frac{\sigma_v^2}{\sigma_u^2+\sigma_v^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t$$

$$-2\frac{\sigma_v^2}{\sigma_u^2+\sigma_v^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t+\sigma_v^2$$

$$=\frac{\sigma_u^2-\sigma_v^2}{\sigma_u^2+\sigma_v^2}X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t+\sigma_v^2$$

$$\textbf{(A1.9)}$$

$$=(1-2\rho)X_t\left(X'\Omega^{-1}X\right)^{-1}X'_t+\rho(\sigma_u^2+\sigma_V^2)$$

Where $\rho=\frac{\sigma_V^2}{\sigma_u^2+\sigma_V^2}$, the ratio of the variance of the econometrician's forecast error to the sum of the variances of the rational expectations forecast error and the econometrician's forecast error.

We know that the minimum and maximum values for ρ are 0 and 1 respectively (see main paper for an intuitive explanation), so that the variance bounds will be:

$$X_t(X'\Omega^{-1}X)^{-1}X'_t\leq Var(\hat{E}_t\Delta C_{t+1}-E_t\Delta C_{t+1})\leq(\sigma_{ut}^2+\sigma_{vt}^2)-X_t(X'\Omega^{-1}X)^{-1}X'_t$$

$$\textbf{(A1.10)}$$

Appendix 2: Derivation of a *k*-period real interest rate with habit formation

(i) Defining a one-period risk-free real interest rate with habit formation

Under habit formation it is assumed that the utility function is a power function of the difference between consumption and habit (X_t):

$$U(C_t) = \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma} \qquad \textbf{(A2.1)}$$

As in CC, we define the relationship between consumption and habit as the surplus consumption ratio:

$$S_t \equiv \frac{C_t - X_t}{C_t} \qquad \textbf{(A2.2)}$$

This implies that the marginal utility of consumption is:

$$U'(C_t) = (C_t - X_t)^{-\gamma} = S_t^{-\gamma} C_t^{-\gamma} \qquad \textbf{(A2.3)}$$

Therefore the one-period real yield is represented by the following equation:

$$(1 + R_{t+1}) = \frac{1}{E_t(M_{t+1})} = \left(\delta E_t \left(\frac{S_{t+1}}{S_t} \right)^{-\gamma} \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right)^{-1} \qquad \textbf{(A2.4)}$$

Taking logs of **(A2.4)** gives:

$$y_{1,t+1} = -\log\delta - \log\left(E_t \left(\frac{S_{t+1}}{S_t} \right)^{-\gamma} \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right) \qquad \textbf{(A2.5)}$$

Generally, if we assume conditional lognormality for two processes *X* and *Y*:

$$\ln E_t(XY) = E_t(x) + E_t(y) + \tfrac{1}{2}[Var_t(x) + Var_t(y) + 2Cov_t(x, y)] \qquad \textbf{(A2.6)}$$

So combining equation **(A2.5)** with the result from **(A2.6)** gives:

$$y_{1,t+1} = -\log\delta + \gamma E_t(\Delta c_{t+1}) + \gamma E_t(\Delta s_{t+1}) - \tfrac{1}{2}\gamma^2 Var_t(\Delta c_{t+1}) - \tfrac{1}{2}\gamma^2 Var_t(\Delta s_{t+1}) - \gamma^2 Cov_t(\Delta c_{t+1}, \Delta s_{t+1}) \quad \textbf{(A2.7)}$$

As in CC, we assume that log consumption follows a random walk, so:

$$\Delta c_t = g + \varepsilon_{c,t+1} \quad \textbf{(A2.8)}$$

where $\varepsilon_{c,t+1}$ is assumed to be normal and homoskedastic, with variance σ_c^2.

The surplus consumption ratio evolves over time as:

$$s_{t+1} = (1-\varphi)\bar{s} + \varphi s_t + \lambda(s_t)\varepsilon_{c,t+1} \quad \textbf{(A2.9)}$$

Where φ determines the degree of persistence of the log surplus consumption ratio and $\lambda(s_t)$ controls the sensitivity of the consumption ratio to deviations of consumption growth away from the mean.

Combining equations **(A2.7)**, **(A2.8)** and **(A2.9)** gives:

$$r_{t+1}^f = -\log\delta + \gamma g - \gamma(1-\varphi)(s_t - \bar{s}) - \tfrac{1}{2}\gamma^2 Var_t(g + \varepsilon_{c,t+1}) - \tfrac{1}{2}\gamma^2 Var_t((1-\varphi)(\bar{s} - s_t) + \lambda(s_t)\varepsilon_{c,t+1}) - \gamma^2 Cov_t(g + \varepsilon_{c,t+1}, (1-\varphi)(\bar{s} - s_t) + \lambda(s_t)\varepsilon_{c,t+1}) \quad \textbf{(A2.10)}$$

If a variable X is conditionally lognormal and homoskedastic then:

$$Var_t \log X = E[(\log X - E_t \log X)^2] = Var(\log X - E_t \log X) \quad \textbf{(A2.11)}$$

Using the result in **(A2.11)**, **(A2.10)** can be re-written as:

$$\begin{aligned} y_{1,t+1} &= -\log\delta + \gamma g - \gamma(1-\varphi)(s_t - \bar{s}) - \tfrac{1}{2}\gamma^2\sigma_c^2 - \tfrac{1}{2}\gamma^2\lambda(s_t)^2\sigma_c^2 - \gamma^2\lambda(s_t)\sigma_c^2 \\ &= -\log\delta + \gamma g - \gamma(1-\varphi)(s_t - \bar{s}) - \tfrac{1}{2}\gamma^2\sigma_c^2(\lambda(s_t)+1)^2 \end{aligned} \quad \textbf{(A2.12)}$$

Which corresponds to equation **(70)** in CC (page 1,287).

(ii) A k-period real rate with habit formation

The price of a k-period bond bought at time t is the product of k stochastic discount factors, which (by cross-cancellation) simplifies to:

$$P_{kt} = E_t[M_{t+1} \cdot M_{t+2} \cdot \ldots\ldots\ldots \cdot M_{t+k}]$$
$$= \delta^k E_t\left[\left(\frac{S_{t+k}}{S_t}\right)^{-\gamma}\left(\frac{C_{t+k}}{C_t}\right)^{-\gamma}\right] \qquad \textbf{(A2.13)}$$

To solve **(A2.13)**, we need to define the expected change in consumption and the surplus consumption ratio between t and $t+k$. Under the assumption that consumption follows a random walk and extending to the k-period model:

$$c_{t+k} = c_t + kg + \sum_{i=1}^{k} \varepsilon_{c,t+i} \qquad \textbf{(A2.14)}$$

As $\varepsilon_{c,t+i}$ is white noise, the expected change in log consumption between t and $t+k$, $E_t \Delta c_{t+k}$, will be the average k-year growth rate, kg.

Using backward substitution we can use the one-period log consumption growth model to recursively define the k-period-ahead log surplus consumption ratio:

$$s_{t+k} = (1-\varphi)\bar{s} + \lambda(s_{t+k-1})\varepsilon_{c,t+k} + \varphi s_{t+k-1}$$
$$= (1-\varphi)\bar{s} + \lambda(s_{t+k-1})\varepsilon_{c,t+k} + \varphi[(1-\varphi)\bar{s} + \lambda(s_{t+k-2})\varepsilon_{c,t+k-1} + \varphi s_{t+k-2}]$$
$$= (1-\varphi^k)\bar{s} + \left[\sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1})\varepsilon_{c,t+i}\right] + \varphi^k s_t \qquad \textbf{(A2.15)}$$

So the change in the log surplus consumption ratio between t and $t+k$ will be:

$$\Delta s_{t+k} = s_{t+k} - s_t = (\varphi^k - 1)(s_t - \bar{s}) + \sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1})\varepsilon_{c,t+i} \qquad \textbf{(A2.16)}$$

And the expected change in the log surplus consumption ratio between t and $t+k$, $E_t \Delta s_{t+k}$, will therefore be equal to $(\varphi^k - 1)(s_t - \bar{s})$.

The log of the yield of a k-period bond at time t, y_{kt}, can be expressed as:

$$y_{kt} = -\, p_{kt} / k \qquad \textbf{(A2.17)}$$

Combining **(A2.17)** with **(A2.6), (A2.13)** can be written as:

$$-\, ky_{kt} = k \log \delta - \gamma E_t \Delta s_{t+k} - \gamma E_t \Delta c_{t+k} + \frac{\gamma^2}{2} Var_t(\Delta s_{t+k}) + \frac{\gamma^2}{2} Var_t(\Delta c_{t+k}) + \gamma^2 Cov_t(\Delta s_{t+k}, \Delta c_{t+k}) \qquad \textbf{(A2.18)}$$

Using **(A2.14)**, the variance at time *t* of the change in log consumption can be expressed as:

$$\begin{aligned} Var_t \log E_t \Delta c_{t+k} &= Var(\Delta c_{t+k} - E_t \Delta c_{t+k}) \\ &= Var(kg + \sum_{i=1}^{k} \varepsilon_{c,t+i} - kg) = k\sigma_c^2 \end{aligned} \qquad \textbf{(A2.19)}$$

Similarly, the variance at time *t* of the log surplus consumption ratio can be written as:

$$\begin{aligned} Var_t \log E_t \Delta s_{t+k} &= Var[\Delta s_{t+k} - E_t \Delta s_{t+k}] \\ &= Var\left[(1-\varphi)^k (\bar{s} - s_t) + \sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \varepsilon_{c,t+i} - (1-\varphi)^k (\bar{s} - s_t)\right] \\ &= Var\left[\sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \varepsilon_{c,t+i}\right] \\ &= \sigma_c^2 \sum_{i=1}^{k} \left[\varphi^{2(k-i)} Var \lambda(s_{t+i-1})\right] \end{aligned} \qquad \textbf{(A2.20)}$$

Where: $Var[\lambda(s_{t+k})] = E[\lambda(s_{t+k})^2] - [E(\lambda(s_{t+k}))]^2$, and we assume that the variance of lambda and the variance of consumption growth innovations are orthogonal.

It follows that the covariance term can be written as:

$$\begin{aligned} Cov_t(\Delta s_{t+k}, \Delta c_{t+k}) &= Cov[(\Delta s_{t+k} - E_t \Delta s_{t+k}), (\Delta c_{t+k} - E_t \Delta c_{t+k})] \\ &= Cov\left[\sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \varepsilon_{c,t+i}, \sum_{i=1}^{k} \varepsilon_{c,t+i}\right] \\ &= E_t \sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \varepsilon_{c,t+i}^2 - E_t \sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \varepsilon_{c,t+i} . E_t \sum_{i=1}^{k} \varepsilon_{c,t+i} \end{aligned}$$

$$= \sigma_c^2 \left[\sum_{i=1}^{k} \varphi^{k-i} E_t \lambda(s_{t+i-1}) \right] \quad \textbf{(A2.21)}$$

Combining these results with **(A2.18)** gives:

$$- ky_{kt} = k \log \delta - \gamma(\varphi^k - 1)(s_t - \bar{s}) - \gamma kg + \frac{\gamma^2 \sigma_c^2}{2} \sum_{i=1}^{k} \left[\varphi^{2(k-i)} Var \lambda(s_{t+i-1}) \right] + \frac{k\gamma^2 \sigma_c^2}{2}$$
$$+ \gamma^2 \sigma_c^2 \left[E_t \sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \right]$$

(A2.22)

Hence the *k*-period yield is equal to:

$$y_{kt} = -\log \delta - \frac{\gamma}{k}(1 - \varphi^k)(s_t - \bar{s}) + \gamma g - \frac{\gamma^2 \sigma_c^2}{2k} \sum_{i=1}^{k} \left[\varphi^{2(k-i)} Var \lambda(s_{t+i-1}) \right] - \frac{\gamma^2 \sigma_c^2}{2}$$
$$- \frac{\gamma^2 \sigma_c^2}{k} \left[E_t \sum_{i=1}^{k} \varphi^{k-i} \lambda(s_{t+i-1}) \right]$$

(A2.23)

Finally, we express the sensitivity of the log surplus consumption function to changes in $\varepsilon_{c,t}$ in the same way as CC:

$$\lambda(s_{t+k}) = \frac{1}{\bar{S}} \sqrt{1 - 2(s_{t+k} - \bar{s})} - 1 \quad \textbf{(A2.24)}$$

where $\bar{S} = \sigma \sqrt{\frac{\gamma}{1 - \varphi - B/\gamma}}$. In order to be able to solve **(A2.23)** we begin by assuming that in period *t*, habit is defined as average aggregate consumption over the previous year. This allows us to define the initial log surplus consumption ratio (s_t), and hence the sensitivity function, $\lambda(s_t)$.

Appendix 3: Approximating the *k*-period conditional expectation and variance

To make expression **(A2.23)**, the k-period yield with habit formation, operational, we need to compute expressions for the conditional expectation and conditional variance of $\lambda(s_{t+i-1})$. Analytically, this cannot be done for $i > 1$, so we apply numerical techniques to obtain approximating functions for these. We use a simple variant of the parameterised expectations algorithm, see Den Haan and Marcet (1990), which can be summarised as follows.

The assumption that consumption follows a random walk implies that s_t is a sufficient statistic for $\lambda\left(s_{t+j}\right)$. That is, rewriting **(A2.15),** we have that

$$\lambda\left(s_{t+j}\right) = F\left(s_t, \{\varepsilon_{t+i}\}_{i=1,..,j}; \varphi, \bar{s}, \bar{S}\right)$$

In words, $\lambda\left(s_{t+j}\right)$ is characterised by a function F of s_t, conditional on the parameters $\left\{\varphi, \bar{s}, \bar{S}\right\}$ and the shocks to consumption growth, $\{\varepsilon_{t+i}\}_{i=1,..,j}$, which are unforecastable at time t. The task at hand is to to approximate $E\left(\lambda\left(s_{t+j}\right)|s_t\right)$ and $Var\left(\lambda\left(s_{t+j}\right)|s_t\right)$. We do this in the following steps:

- Generate a sequence of normally distributed random shocks $\{\varepsilon_{t+n}\}_{n=1,..,N}$, with standard error σ_c. N should be a large number: we set N =100,000.
- Calculate the sequence $\{\lambda\left(s_{t+n}\right)\}_{n=1,..,N}$, conditional on some initial value s_0, which we set at $\bar{s}$. We discard the first 1,000 observations to avoid dependence on the initial value.
- Run the regression

$$\lambda\left(s_{t+n}\right) = G_j\left(s_{t+n-j}; \gamma^j\right) + e_t$$

where G_j is a polynomial with parameters γ and j is the forecast horizon, where we look at $j = 1, .., 40$. In practice we set

$$G_j\left(s_{t+n-j}; \gamma_j\right) = \gamma_{j0} + \gamma_{j1}s_{t+n-j} + \gamma_{j2}s_{t+n-j}^2 + \gamma_{j3}s_{t+n-j}^3$$

Under certain regularity assumptions

$$E\left(\lambda\left(s_{t+j}\right)|s_t\right) \approx G_j\left(s_t; \hat{\gamma}_j\right)$$

where $\hat{\gamma}_j$ is the least squares estimate of γ.

- Using this approximation, generate a sequence for the (actual) conditional variance, that is

$$\begin{aligned}\overline{Var}_j\left(\lambda\left(s_{t+n}\right)|s_{t+n-j}\right) &= \left(\lambda\left(s_{t+n}\right) - E\left(\lambda\left(s_{t+n-j}\right)|s_t\right)\right)^2 \\ &\approx \left(\lambda\left(s_{t+n}\right) - G_j\left(s_{t+n-j};\hat{\gamma}_j\right)\right)^2\end{aligned}$$

and run the regression

$$\overline{Var}_j\left(\lambda\left(s_{t+n}\right)|s_{t+n-j}\right) = H_j\left(s_{t+n-j};\eta_j\right) + e_t$$

where

$$H_j\left(s_{t+n-j};\eta_j\right) = \eta_{j0} + \eta_{j1}\left(s_{t+n-j} - \tilde{s}\right)^2$$

where $\tilde{s}$ is the sample mean. We have used a quadratic function to ensure that the function is monotic, and ensuring that the function is always positive—we are trying to approximate a variance—is simply a matter of checking that $\eta_{j0} > 0$ for all j. The functions $H_j\left(s_{t+n-j};\eta_j\right)$ are the approximations of the conditional variances, ie

$$Var_j\left(\lambda\left(s_{t+j}\right)|s_t\right) \approx H_j\left(s_t;\eta_j\right).$$