Working Paper no. 273

‘Real-world’ mortgages, consumption volatility and the low inflation environment

Sebastian Barnes and Gregory Thwaites

September 2005

Bank of England

‘Real-world’ mortgages, consumption volatility and the low inflation environment

*Sebastian Barnes**

and

*Gregory Thwaites***

Working Paper no. 273

* Structural Economic Analysis Division.
Email: sebastian.barnes@bankofengland.co.uk

** International Finance Division.
Email: gregory.thwaites@bankofengland.co.uk

We are grateful to John Campbell and João Cocco for making available the code from their 2003 NBER working paper. We are also grateful for comments and advice from Peter Andrews, Andrew Blake, Andrew Brigden, Jenni Greenslade, Andreas Lehnert, David Miles, Alasdair Scott, Chris Shadforth, James Talbot, seminar participants at the Bank of England and two anonymous referees. This paper was finalised on 15 August 2005.

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

Information on the Bank’s working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

ISSN 1368-5562

Contents

Abstract

This paper considers the interaction between the microeconomic decisions facing households and the macroeconomic environment in a setting where households have ‘real-world’ mortgage contracts. In particular, we consider the possible consequences of the important changes in the framework for setting monetary policy in the United Kingdom in recent decades that have coincided with a more stable and low inflation environment. We set a model of households with ‘real-world’ mortgages in a partial equilibrium overlapping generations framework calibrated to UK data. We find that the welfare gains of the change of regime would have been considerable. However, the baseline calibration of the model implies that the volatility of aggregate consumption growth would actually be higher in the steady state in the more stable environment of the 1990s regime. This is due to greater synchrony in the response of households to shocks, offsetting the smaller magnitude of macroeconomic shocks. This effect is amplified by higher levels of debt in the 1990s. The result that aggregate consumption volatility could be higher in the current regime suggests that the observed fall in aggregate consumption volatility cannot necessarily be attributed to the more stable macroeconomic environment and the role of mortgage debt. If this result applies, this would suggest that the observed fall in volatility may be due either to other factors or may be a transitional phenomenon rather than a feature of the new steady state.

Key words: Household consumption, household portfolios, mortgage choice, housing, monetary policy regimes, overlapping generations.

JEL classification: D11, D14, E21, E42.

Summary

This paper considers the interaction between the microeconomic decisions facing households and the macroeconomic environment in a setting where households have 'real-world' mortgage contracts. In particular, we consider the possible consequences of the important changes in the framework for setting monetary policy in the United Kingdom in recent decades.

The change in monetary policy regime from the 1980s to the 1990s has been associated with greater stability of the macroeconomic environment. 'Real-world' mortgages may provide an explanation of how more stable economic conditions have contributed to reducing the volatility of aggregate consumption through effects not captured in elementary textbook models of consumption with debt. In these models, it is typically assumed that household borrowing takes the form of successive one-period debt contracts, denominated in units of consumption. Actual mortgage contracts - the biggest financial commitments that most households ever make - look very different to this: they are denominated in nominal terms with repayments over many periods, sometimes with fixed nominal interest rates. This paper is concerned with the role of such real-world mortgage contracts in consumption volatility.

We use a model of real-world mortgages to show the effects at household level of the change in monetary policy regime under adjustable-rate and fixed-rate mortgages. We use this to model aggregate consumption uncertainty in a partial equilibrium overlapping generations framework.

At household level, we find that non-housing consumption would be smoother over the life cycle in the more stable 1990s regime. The change of regime generates substantial welfare gains for mortgage holders. Even though households now have more mortgage debt than in the past, we find that households could still enjoy similar levels of utility from non-housing consumption in the 1990s as in the 1980s regime. This suggests that households may have increased their demand for housing in response to the lower cost and greater certainty of mortgage borrowing in the 1990s.

The main parameterisation of the model suggests, counterintuitively, that aggregate consumption volatility under the 1990s regime would actually be higher in the steady state than in the 1980s regime, other things being equal. Although macroeconomic shocks have become less pronounced, this result suggests that households' responses to shocks have become more synchronised.

Furthermore, higher indebtedness in the 1990s has also tended to make aggregate consumption less stable. This result shows how the more stable economic environment associated with the 1990s regime would not necessarily translate into greater stability of aggregate consumption given real-world mortgages. If the assumptions necessary for this result hold, the observed fall in aggregate consumption volatility in the 1990s would either have to be explained by other offsetting factors or because the economy was in a period of transition between two regimes rather than the new steady state; households in the 1990s actually benefited simultaneously from the more stable macroeconomic environment, and the lower levels of indebtedness inherited from the past.

1 Introduction

Models of household consumption typically assume that household borrowing takes the form of successive one-period debt contracts, denominated in units of consumption. Mortgage contracts - the biggest financial commitments that most households ever make - look very different to this: they are denominated in nominal terms with repayments over many periods, sometimes with fixed nominal interest rates. This paper is concerned with such 'real-world' mortgage contracts, how changes in the macroeconomic environment affect their characteristics, and what this means for household consumption decisions. In particular, we consider the possible consequences of the important changes in the framework for setting monetary policy in the United Kingdom in recent decades. By altering the macroeconomic environment in which households make consumption and portfolio decisions, these changes in policymaking are likely to have changed the way households behave.

This paper considers the implications of these changes for both the consumption behaviour of individual households and their consumption in aggregate. We model and simulate the consumption decisions of real-world mortgage debtors. We use the results to assess the welfare implications of the change in monetary policy making and the choices households would make between two types of mortgages, adjustable-rate and fixed-rate. We also assess the implications for aggregate consumption volatility.

Understanding the implications of these household-level changes for *aggregate* consumption is an important issue for policymakers:

> 'So price stability appears to have led to a more stable macroeconomic environment, with fewer surprises to inflation, less inflation uncertainty, and a lower level of real interest rates. Has this improved nominal performance led to greater stability of the real economy?' King (2002)

Empirically, as we show below, simple measures of aggregate consumption variability appear to have been lower in the current monetary policy regime than its predecessor. From a theoretical perspective, there are many market imperfections that affect household consumption behaviour in ways that may ultimately lead it to depend on the monetary policy regime.

The literature on household consumption and portfolio behaviour emphasises a number of respects in which market imperfections lead behaviour to differ from what the basic buffer stock hypothesis would predict. Most notably, a substantial proportion of households may be credit constrained and consequently their consumption may be better explained by buffer stock saving models than by the permanent income hypothesis (Carroll (2001)), and consumption will be more sensitive to cash flows as a result.

Typical household portfolios also differ significantly from those predicted by elementary models of finance, with households holding portfolios that are highly concentrated in a single asset (housing) and with positions that are highly leveraged with respect to income and net wealth (see Brueckner (1997)). Strong incentives towards homeownership, arising from the favourable tax treatment of owner-occupation or the role of housing as a hedge against housing costs (Cocco (2001)), rationalise the existence of such portfolios (Flavin and Yamashita (2002)). These market imperfections affect the sensitivity of household behaviour to changes in the economic environment that alter the returns, risk and timing of different decisions in their lives.

This paper concentrates on the possible role that real-world mortgages, through their effect on household-level consumption decisions, may have in linking nominal stability to the real economy. Within consumption, this paper concentrates specifically on mortgage holders. Mortgage debt is likely to be an important factor for aggregate consumption in the United Kingdom: it stands at over 100% of post-tax household sector annual income, and around 50% of households have a mortgage (Hamilton (2003)).

The literature on mortgages is dominated by empirical studies of mortgage prepayment behaviour (see Lacour-Little (1999) for an overview or Fu, Lacour-Little and Vandell (2000)). There is also an extensive literature on the value of options under different mortgage contracts (Kau and Keenan (1995)). A subset of papers in this literature deals explicitly with the difference between actual real-world mortgage contracts and contracts that are optimal from the perspective of finance theory (see Taewon (1987), Sanders and Slawson (2001)). Fratantoni (2001) uses a calibrated model to consider the role of 'committed expenditure risk', where the commitment to make mortgage payments affects households' liquidity. There have also been a number of papers examining the choice between different, typically adjustable-rate and fixed-rate, mortgages either theoretically (Alm and Follain (1984) and Lehnert and Statman (2003)) or empirically (Shilling, Dhillon and Sirmans (1987)).

One key paper considers mortgage choice in the United States using a richer framework of household behaviour and considers the effect of a subperiod ('Greenspan-Volcker') where the monetary policy was different from over the sample as a whole (Campbell and Cocco (2003)). To our knowledge, that is the only previous paper that explicitly considers how optimal mortgage choice would vary with the macroeconomic environment, and parts of our analysis and method follow it closely. In the United Kingdom, many of these issues were considered in the analysis of supply-side and demand-side factors limiting the development of a fixed-rate mortgage market recently undertaken on behalf of the government (Miles (2004)).

The rest of the paper is structured as follows. Section 2 presents stylised facts about consumption and the UK macroeconomic environment, including simple measures of the change in consumption volatility. Section 3 explains the theoretical implications of a change in the monetary policy regime on the macroeconomic environment facing households. It also discusses how the costs of real-world mortgages would vary with both real and nominal features of the macroeconomic environment. This section also discusses how buffer stock consumption theory translates these changes of economic circumstance into household consumption, portfolio decisions and the risks that people face. Section 4 presents a model of individual consumption with real-world mortgages under adjustable-rate and fixed-rate mortgages, closely following Campbell and Cocco (2003). We extend this to a partial equilibrium overlapping generations (OLG) model of aggregate consumption. Section 5 discusses calibrating the model to UK data. Section 6 presents results for the volatility of household-level consumption and assesses the welfare implications of the change in regime. This section also presents results for the volatility of aggregate consumption, which may be important for understanding the macroeconomy under the new regime. Results are presented for several types of mortgage contract under two stylised monetary regimes, corresponding to the 1980s and 1990s. Section 7 concludes.

2 Consumption and the UK macroeconomy

This section sets out key stylised facts about consumption and the macroeconomic environment in the United Kingdom under two different broadly-defined monetary policy regimes. These two regimes are the 1980s regime - from 1980 Q4 to 1992 Q3 - and the 1990s - from 1992 Q4 to 2003

Q1. The periods correspond to changes in the monetary policy framework, such as the introduction of inflation targeting in 1992, and are consistent both with statistical tests and the existing empirical literature (Benati (2004)).[(1)]

The change in regime appears to coincide with changes in the behaviour of key macroeconomic variables. The unconditional standard deviation of annual GDP growth in the 1990s regime is less than half that in the 1980s (Table A). The standard deviation of aggregate consumption growth is also materially lower in the 1990s. Consumption remains more volatile than GDP, albeit to a lesser extent, in the more stable monetary regime. Furthermore, the contemporaneous correlation between annual GDP growth and consumption growth has become weaker, falling from 0.88 in the 1980s to 0.23 in the 1990s regime.

Table A: Variability of consumption and GDP in different regimes

	GDP	Consumption	Correlation
1980s regime	0.02	0.028	0.88
1990s regime	0.009	0.011	0.23

Standard deviation of annual growth rates using quarterly data.

The change in the monetary policy regime can be characterised in terms of inflation and real interest rates (Table B). Compared to the 1980s regime, the 1990s has experienced lower mean inflation, a lower unconditional standard deviation of inflation and less persistent inflation. Furthermore, the mean and the standard deviation of the real interest rate have been lower.

Table B: Inflation and real interest rates in different regimes

	1980s regime	1990s regime
Mean inflation	0.056	0.025
Standard deviation of inflation	0.02	0.008
Persistence of inflation	0.89	0.75
Mean real interest rate	0.061	0.034
Standard deviation of real interest rate	0.026	0.013

Annual rates estimated using quarterly data. Real *ex-ante* interest rates estimated using the method set out in Section 5.

These changes are in line with what one might expect to be the result of a more effective, credible and transparent monetary policy framework (see Section 3.1, below). Such a change in the policy

(1) These regimes are broadly defined, partly to allow for longer sample periods, and so do not capture all changes to the monetary policy framework. In particular, we do not consider separately the period of UK membership of the Exchange Rate Mechanism (ERM) or the granting of independence to the Bank of England.

environment would also affect how macroeconomic variables covary, so it may be misleading to consider each one in isolation. Consistent with this idea, Table C shows that the contemporaneous correlation of key macroeconomic variables, as well as their unconditional moments, has also changed materially between the two regimes.

Table C: Cross-correlation of variables in different regimes

1980s regime	Inflation	Real interest rate	Aggregate real labour income growth
Inflation	1	0.13	-0.35
Real interest rate		1	-0.60
Aggregate labour income growth			1
1990s regime	Inflation	Real interest rate	Aggregate real labour income growth
Inflation	1	0.60	0.00
Real interest rate		1	-0.33
Aggregate labour income growth			1

Contemporaneous correlations using annual data at quarterly frequency.

3 Theory

This section outlines the theoretical connections between (i) the change in the monetary framework and changes in the macroeconomic environment, (ii) the change in the economic environment and changes to the cost and time-profile of mortgages for households with real-world mortgages, and (iii) changes in the cost of mortgages and the risks households face in taking consumption or portfolio decisions.

3.1 Changes in the economic environment

The change in the behaviour of key macroeconomic variables between the two periods (see Table B) is broadly consistent with a change in the preferences of policymakers, under standard New Keynesian assumptions.[2] Rotemberg and Woodford (1999) assume that policy is set according to a standard Taylor feedback rule of the form:[3]

$$r_t - r^* = \sum_{j=0}^{m_\pi} a_j(\pi_t - \pi^*) + \sum_{j=0}^{m_y} b_j(y_t - y^*) + \sum_{j=1}^{m_r} c_j(r_t - r^*)$$

(2) The evidence and theory here are somewhat ambiguous in that the implied relationship between the variables is sensitive to the exact specification adopted. The results presented in Table B hold for contemporaneous correlations but may not hold for other relevant specifications of the relationship between the variables. Equally, different theoretical models have different implications for the correlation between variables. In some cases, real interest rates would be expected to covary positively with inflation (where real interest rates are increased to reduce inflation) whereas in other cases a negative relationship would be expected as high inflation results from low real rates.
(3) It should be noted that the Taylor rule is a statistical representation and is only used here for heuristic purposes.

Under this rule, they show how the unconditional variance of π, r, y and the price level depend on the coefficients a, b and c of the policy reaction function. In their calibrated model, the relationship between the three policy-rule coefficients is complex. In some cases, there is a trade-off between (say) higher weight on inflation stabilisation, which reduces inflation variability, and lower output variability. However, there are some policy settings that dominate others: these dominant rules imply lower variability of inflation, output and real interest rates than some other determinate rules. The shift in the United Kingdom from the 1980s to the 1990s regime may be such a case, where a more efficient rule is introduced that generates lower variance of all variables. The persistence of inflation shocks also depends on how much policy responds to current deviations from the target, as well as the other parameters. Stronger policy responses to deviations of current inflation from target could also explain the lower observed persistence of UK inflation in the 1990s regime. The steady-state rate of inflation π^* in this type of model is determined independently of the dynamics, but can also be interpreted as set by an inflation target or policymakers. UK experience suggests that the implied target rate of inflation is also now lower than in the past. The framework outlined later in our paper differs from that used by Rotemberg and Woodford (1999) in that it uses a simpler dynamic structure and, more significantly, allows the steady-state level of inflation to influence the dynamic properties of the model. However, the New Keynesian framework does provide some motivation for why the level of inflation and the variability of key macroeconomic variables could all be lower in the current monetary policy regime than in the past.

For the 1990s regime, real wages and inflation are essentially uncorrelated (see Table C), while for the 1980s there is a large negative correlation.(4) If nominal wages are set one period ahead, real wages will be negatively correlated with inflation as, for example, unexpected positive inflation shocks will lower the *ex-post* real value of the *ex-ante* nominal wage in force. The negative correlation in the 1980s regime between real wages and inflation under this hypothesis may then be explained by the importance of unanticipated 'inflation shocks' during the period.(5) It is also possible that wage-setting practices have changed in response to the shift in regime. By contrast, the absence of a correlation in the 1990s can be interpreted as consistent with the absence of important inflation surprises.

---

(4) The Rotemberg and Woodford (1999) model cannot be applied directly in this case, as it assumes that there is no wage-stickiness.

(5) In the Taylor rule framework, these might be interpreted as an additional monetary policy shock term to reflect uncertainty or variability in the policy reaction function.

3.2 'Real-world' mortgages

Household borrowing is often modelled in real terms as a succession of one-period loans, which evolve according to:

$$d_{t+1} = (1 + r_t)d_t - q_t \qquad \textbf{(1)}$$

where d indicates the real value of debt, r is the real rate of interest and q is the value of debt amortisation. It is typically assumed that, possibly subject to an upper limit on the amount of debt, the household can vary q as it pleases, giving flexibility over cash flow subject to an overall solvency constraint.

This specification ignores many important features of the actual menu of mortgage contracts typically available to households. First, contracts are usually specified in nominal, rather than real, terms. Second, mortgages typically last for many years. This is especially important for fixed-rate mortgages (FRMs), where the interest rate is also set for a number of periods, although it is also relevant for adjustable-rate or variable-rate mortgages (ARMs).[6] Third, nominal mortgage payments (M) (including both interest payments and debt repayments) are typically set so that they are equal in each period during the life of the mortgage, using the formula:

$$M = D_t \frac{i}{\left(1 - (1 + i)^{-(T')}\right)} \qquad \textbf{(2)}$$

where D is the nominal principal of the loan outstanding, i is the nominal interest rate, and T' is the remaining term of the mortgage.[7] In general, the standard model **(1)** and the real-world specification **(2)** will have very different characteristics.

Although the monthly *nominal* mortgage payment is held constant over the remaining life of the mortgage (for a given rate of interest), the schedule of *real* mortgage payments (m) will depend on the inflation rate. These are given by (lowercase variables are used to denote real variables):[8]

$$m_{t+k} = \frac{M_{t+k}}{\prod_{i=t}^{i=k} E_t\left[(1 + \pi_i)\right]} \qquad \textbf{(3)}$$

(6) Under some ARMs, monthly mortgage payments are reviewed only annually, although interest is adjusted on a monthly basis (the rate of debt repayment is varied to offset any change in interest payments). This practice, known as 'annual rest', has become less widespread in the United Kingdom in recent years.

(7) This is the standard formula for calculating nominal mortgage payments, where interest falls due on the principal at the start of each period. This formula can be obtained by writing an expression for the balance at the end of the mortgage in terms of the original balance, nominal interest payments, the number of periods of the mortgage and the (unknown) equal monthly payment. Obtain the monthly payment by solving this equation with the terminal balance set to zero, which gives the equal monthly nominal payments necessary to pay off the mortgage.

(8) In terms of period t prices.

A positive rate of (anticipated) inflation causes the real value of mortgage payments to decline over the life of the mortgage; a higher inflation rate will cause them to decline more quickly. Combining equations **(2)** and **(3)**, a higher steady-state inflation rate gives rise to higher initial nominal (and real) payments, but causes them to decline more quickly. This phenomenon, whereby real mortgage payments are concentrated earlier in the mortgage under higher inflation is known as 'front-end loading' (or 'mortgage tilt'). The timing of mortgage payments also has an effect on the timing of *repayments* of capital (amortisation). Although cash repayments are initially lower under high inflation, this is more than offset by the faster rate at which higher inflation erodes the real value of the debt. As a result, the debt is initially amortised at a faster rate in real terms under higher inflation, but is repaid more slowly later in the life of the mortgage.

The effect of changes in inflation and real rates with such mortgage contracts, based on the values in Table B, is illustrated in Chart 1. The top panels show how income gearing, the ratio of mortgage payments to income, evolves under the different regimes. This shows that real mortgage payments in the 1980s were not only more volatile, as the result of more volatile nominal and real interest rates, but also that payments were more 'front-end loaded' than in the 1990s.

Unexpected inflation will also affect the real cost of real-world mortgages. Any shock to inflation (or expected inflation) changes the value of the outstanding mortgage liability (negative wealth (w)), expressed as the (negative of) the real *present value* of future payments under the mortgage:

$$w_t = -\sum_{i=t}^{i=T} \frac{M_i}{\prod_{i=t}^{i=i} E_t\left[(1 + \pi_i + r_i)\right]} \qquad \textbf{(4)}$$

For ARMs, the value of the contract will change by approximately $d_t(\pi_t - E_{t-1}\pi_t)$. If inflation is persistent, inflation shocks will have a bigger effect on FRMs. The interest rate under an FRM is fixed at the time the loan is originated, so both parties to the contract are committed on the basis of the expectations of inflation and real interest rates held at the time. But repeated or persistent shocks may take actual inflation a long way from this. Higher future inflation reduces the household's mortgage liability under this contract (wealth is less negative): the expectation of inflation embodied in the nominal mortgage payment will be insufficient to meet the real erosion in the value of debt due to inflation (π), so borrowers will benefit unexpectedly at the cost of lenders. The opposite is true for negative inflation shocks.

The effect of inflation shocks on FRMs is more complex if households have the option to

Chart 1: Behaviour of real-world ARM contract under different monetary regimes

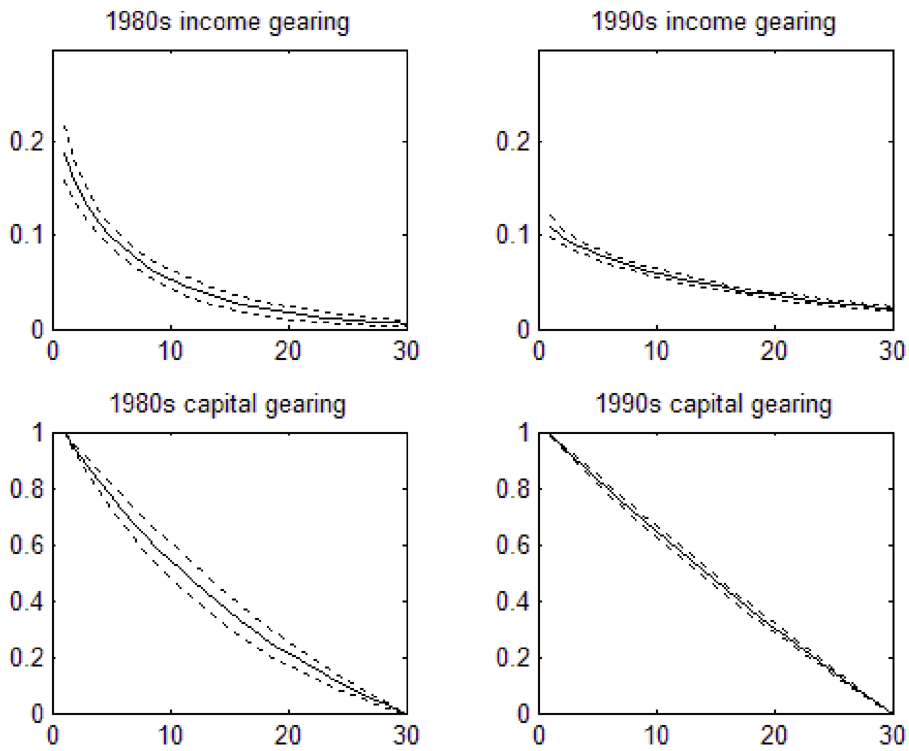

Inflation and real interest rate processes as in Table B. Dashed lines indicate ± 1 standard deviation around the mean. Income gearing defined as mortgage payments over income, where income has the time profile given in Appendix A. Capital gearing defined as the real value of debt outstanding over the real value of the house, where the value of the house is constant and initially set to equal the value of the debt outstanding. This exercise assumes 30-year mortgage contracts.

refinance the mortgage. The household's incentive to refinance is asymmetric: households will never (rationally) refinance if the value of their mortgage commitments is lower under the existing FRM contract than otherwise (ie inflation or real interest rates are higher than anticipated), but there is an incentive to refinance if the cost of the mortgage would be less by switching to a new contract (ie inflation or real interest rates are lower than was expected). As a result, FRM contracts with a refinancing option carry a premium to compensate lenders for providing this insurance against bad outcomes. FRMs also often impose penalties for refinancing. In addition, households may face other costs of refinancing and as a result do not refinance even when it might apparently be profitable (see Deng and Quigley (2003)).

There are a number of secondary features of typical mortgage contracts that also differ from those embodied in simple models. For example, contracts have tended to restrict the ability of households to control the rate of debt repayment. Negative amortisation occurs where mortgage payments are insufficient to cover interest on the loan and hence the principal outstanding increases. In the past, this has typically not been allowed under mortgage contracts, even where the borrower has had sufficient collateral and such borrowing would have been attractive to them. Equally, overpayment of a mortgage (repaying the debt faster than scheduled) has in some cases been difficult or costly. Many of these restrictions have eased with the introduction of 'flexible mortgages' in the UK market and the incorporation of greater flexibility in other contracts.

In practice, households may be able to offset these features of real-world mortgages to some degree through their behaviour. For example, they may save less through other channels than they might otherwise do in order to meet the high initial payments due to 'front-end loading' under high inflation. Equally, households may hold higher precautionary balances of savings to hedge against the risk of higher mortgage payments. However, the market imperfections that give rise to real-world mortgages are likely to be part of a broader class of market failures that imply that households are unlikely to be able to offset the features of these mortgage markets completely through their behaviour or in other markets. Furthermore, the constraints implied by mortgage contracts can be viewed as broadly equivalent to other aspects of behaviour in the mortgage market that depart from standard models. In particular, the behavioural finance literature emphasises how households may choose to ignore wealth maximising opportunities by following the path of payments implied by their mortgage contracts rather than making optimal repayments: households, even given the choice, prefer to allow the schedule set out in the contract to determine their behaviour. Whether due to explicit or behavioural costs, households may end up following the payments schedule specified in the contract rather than seeking to optimise their debt repayment. This may affect their consumption behaviour.

Lenders also impose conditions on the amount borrowed, such as a maximum loan to income and loan to value ratios. These are imposed to ensure that households are able to repay their loans on the terms agreed, and to ensure that there is enough collateral to protect the bank in event of default. This is also a response to asymmetric information on the part of banks about potential borrowers (Stiglitz and Weiss (1981)). These credit constraints on households may not be directly affected by the inflation regime, but the macroeconomic environment may influence how banks set these conditions.

3.3 Risk under 'real-world' mortgages

Mortgage debtors are sensitive to three broad types of shock: to income, inflation and real interest rates. It is the structure of real-world mortgages that determines the relationship between the payoffs under a mortgage and shocks to inflation and real interest rates. It is often assumed that ARMs are riskier than FRMs because the nominal value of payments depends on the current nominal short-term interest rate. However, the overall riskiness of each type of contract depends on the interaction of a range of factors. Consequently, neither type of mortgage is generally or unambiguously less risky.

These shocks and real-world mortgages contribute to two types of risk that have negative effects on the welfare of risk-averse households. First, there is 'cash flow/liquidity' risk. This is the risk that households may need to cut consumption to meet their financial commitments in a given period unless they have hedged appropriately. This risk only arises where households are credit constrained, as otherwise they would borrow rather than engage in a costly reduction of their consumption. Second, there is a 'wealth risk'. This is the risk that the value of households' lifetime resources changes, requiring a change in the level of current and future consumption. Under standard assumptions, this risk does not depend on the availability of credit as it represents a change in the lifetime resources available to a household. These two risks are not entirely independent and both may materialise at the same time - higher interest payments may coincide with lower income, which implies that households face a greater burden on their cash flow as well as having lower wealth.

Real-world mortgages can cause 'cash flow/liquidity' risk in a number of ways. The primary mechanism is identified as 'Income Risk' (Campbell and Cocco (2003)): the risk of unexpected changes in the affordability of mortgage payments, either due to a rise in the payments or a shortfall in income or cash . This is similar to 'committed expenditure risk' (Fratantoni (2001)). A transitory increase to required mortgage payments will be costly to households in terms of utility, if they do not have the resources or access to credit, to enable them to smooth their consumption over the shock. They may even be forced to default on the mortgage. Typically, young households will be highly leveraged both with respect to income and wealth. With little collateral to offer, they are likely to be credit constrained and hence sensitive to shocks during the first few years of the mortgage. Cash flow/liquidity risk is increases with the volatility of nominal and real interest rates, and of labour income. This risk is also greater if interest rates and labour income are negatively correlated as households are more likely to experience high debt payments and low income simultaneously. More borrowing relative to income or higher real interest rates also increase this risk, as debt payments represent a higher proportion of income and so households are more exposed to a given negative shock. In addition, the size of this risk increases with the rate of inflation: 'front-end loading' increases the burden of debt in the early years of a mortgage when households are more likely to find themselves credit constrained and vulnerable to this risk. 'Cash flow/liquidity' risk is likely to be greater for ARMs because payments vary with short-term interest rates.[9]

'Wealth risk' also depends on a number of factors. As defined in Campbell and Cocco (2003), the risk is that unexpected inflation causes the real cost of the mortgage to change, as shown earlier (equation **(4)**). This affects all mortgage holders for the immediate period, to the extent that *ex-ante* expected inflation embodied in nominal interest rates differs from *ex-post* inflation. However, if inflation is persistent, this risk is greater under FRMs as they are committed for the whole period of the contract to inflation expectations from the time when the contract was initiated. ARM contracts do not have this feature as mortgage payments and expected inflation are set period by period. To the extent that FRMs can be refinanced, the cost of these mortgages is insulated against the 'bad' outcome (from the debtor's point of view) of an expected fall in inflation. However, risk-averse households care about the uncertainty of their future wealth and consumption, not only the possibility of bad outcomes, and hence this risk remains even under FRMs. A second source of 'wealth risk' in the sense discussed here under both contracts is defined by Campbell and Cocco (2003) as 'real interest rate risk', the risk of unexpected changes in real interest rates. As equation **(4)** shows, this has essentially the same effect on the real present discounted cost of the mortgage contract as shocks to inflation expectations. FRMs insure households against increases in real rates and, to the extent that refinancing is optimal, allows households to benefit from reductions in real interest rates. Hence, FRMs have some real interest rate risk but this will tend to be lower than for ARMs. A third source of 'wealth risk' is revisions

(9) This result is not completely general because FRMs are likely to have, on average, higher real interest rates as a result of the premium to insure households against fluctuations in interest rates. Consequently, households may have higher payments under FRMs and hence be more likely to face affordability problems if they experience a negative temporary income shock (Lehnert and Statman (2003)).

to expected permanent income. This is a general feature of any model with uninsurable permanent income risk, but the commitments under real-world mortgages modify households' portfolios and exposure to other risks. This may in principle have the additional effect of amplifying shocks to permanent income.

Thus, if inflation and real interest rates are uncertain, real-world mortgages can introduce a number of risks relative to the standard textbook real debt accumulation equation. Furthermore, it cannot be said that an ARM is always more or less risky than an FRMs: their relative riskiness will depend on both individual circumstances, such as income uncertainty and how indebted the household is (see Miles (2004)), and the macroeconomic environment. Outcomes under both contracts depend on inflation. The risks a mortgage introduces are a big factor in the overall risk of home purchase and ownership.

4 Model

This section sets out a model, closely following Campbell and Cocco (2003), of consumption at the household level under different types of mortgage contract. The model is rich in some dimensions as it assumes a relatively realistic form for the mortgage contract, but the realism is heavily limited in other respects due to the computational intensity of solving the model. This means that the forms of the contract, the number of periods considered and the role of other variables (such as housing debt) are heavily restricted or curtailed. We briefly discuss setting this model in a partial equilibrium overlapping generations framework, which is somewhat restrictive for this problem, to obtain the properties of aggregate consumption.

4.1 Household optimisation problem

We model each household for $T = 30$ years. This corresponds to a household headed by a 25-year old taking out a mortgage and planning to repay it by age 55. Households derive utility from consumption of housing (hc) and non-housing consumption (c). These are assumed to be additively separable in the utility function. We assume that households own a house of exogenous size ($\overline{H}$). This implies that we can consider both the household's consumption and portfolio decisions without reference to housing.[10] Utility is time-separable, so that current utility does

(10) This version of the model does not allow a role for housing as collateral, unlike some versions of Campbell and Cocco (2003).

not depend directly on past consumption, and discounted at a constant rate δ so that the future utility is discounted by a factor $\beta^t = \frac{1}{(1+\delta)^t}$. γ is the coefficient of relative risk aversion (the inverse of the intertemporal elasticity of substitution). Households derive utility from terminal wealth W_{T+1}, which can be viewed as the value of utility in the period of life beyond T. Households thus solve the following standard intertemporal optimisation problem:

$$\max E_0 \sum_{t=0}^{T} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} + \beta^{T+1} \frac{W_{T+1}^{1-\gamma}}{1-\gamma}$$

Other than the fixed investment in housing, agents can only invest in a one-period safe asset yielding the risk-free rate of interest. Thus cash-on-hand evolves in real terms according to:

$$w_{i,t+1} = w_{i,t}(1+r_t) - c_{i,t} - m_{i,t} + (1-tax)l_{i,y}$$

where w indicates real non-housing wealth (bonds), c denotes real individual consumption, tax is the rate of income tax and l real labour income. Households are precluded from overpaying their mortgages (making repayments larger than required).[11]

Households default when the payments required under the mortgage are greater than their cash-on-hand ($w_{i,t} < m_{i,t}$). Following default, consumption for future periods is set at a low level $\overline{c}$.[12]

4.2 Macroeconomic processes

The main macroeconomic processes follow the assumptions in Campbell and Cocco (2003) and are standard. Inflation is assumed to follow an autoregressive process of order one (AR(1)):

$$\pi_t = \phi\pi_{t-1} + \mu(1-\phi) + \varepsilon_t \quad \textbf{(5)}$$

where ϵ_t is an i.i.d. shock normally distributed with mean zero and variance of σ_ϵ^2.

(11) There is some evidence that overpayment does occur (see Fu *et al* (2000)). As discussed in Subsection 3.2, the introduction of 'flexible mortgages' may have made overpayment easier but it is likely that some barriers remain.
(12) This condition for default does not in itself guarantee that households will always default when it is optimal to do so, as it does not allow households to enter voluntary liquidation. However, $\overline{c}$ is set sufficiently low to ensure that there is no voluntary default.

The *ex-ante* real interest rate is given by:

$$r_t = \bar{r} + \psi_t \qquad \textbf{(6)}$$

where ψ_t is a normally distributed white-noise shock with mean zero and variance σ_ψ^2.

Nominal interest rates are given by the Fisher equation:

$$i_t = r_t + E_t \pi_t$$

The spot interest rate for k-periods ahead is defined under the assumption that the Expectations Hypothesis holds:

$$i_{t,k} = \left(\frac{1}{n}\right) \sum_{i=t}^{i=k} E_t(i_i)$$

Labour income has three components: a deterministic lifetime component that is a function of age ($f(t - t_0)$), an idiosyncratic transitory income shock (ϖ), and a common permanent income shock (ν) . Hence labour income is given by:

$$l_{i,t} = f(T - t) + \varpi_{i,t} + \nu_t \qquad \textbf{(7)}$$

where $\upsilon_t = \upsilon_{t-1} + \eta_t$ and η_t is a normally distributed i.i.d. random variable with mean zero and variance σ_η^2.

4.3 Contracts

We consider three types of mortgage contract (as Campbell and Cocco (2003)): an ARM, a FRM fixed for T with the option to refinance by paying a penalty ζ, and a price-level indexed (INDEX) contract where the payments are specified in real terms (as proposed by Campbell and Cocco (2003)).

For the ARM and FRM contracts, nominal mortgage payments are defined according to equation **(2)**. These contracts therefore capture the full effect of front-end loading. The interest rate under the ARM contract is the current short-term risk-free interest rate i_t plus a premium θ, reflecting the probability of default and other costs to the lender. The interest rate under the FRM contract is set using $i_{REF,t}$, the long-term interest rate in force for the last period in which the loan was refinanced. FRMs carry a premium over the risk-free long rate of θ'. This premium embodies the same credit risk and other costs as θ and an additional premium to reflect the option value of the FRM. Households refinance if the expected liability without refinancing the FRM is greater than the expected liability under a new FRM set assuming the current level of long-term interest rates minus a refinancing penalty ζ. It should be noted that this structure of FRM contract is closer to the US definition of a 'fixed-rate' mortgage than the UK case, where the interest rate is typically fixed for a relatively short period such as two years. However, the general structure of the fixed contract in the model still gives some insights into the differences between types of contract.

Under the Index contract, repayments are set to be constant in real terms across the life of the mortgage, in a manner analogous to equation **(2)**. The interest rate is the current short-term real interest rate augmented by θ.(13) The Index contract is designed to remove the effects of the nominal specification of standard mortgage contracts, including the front-end loading of repayments. It is not, however, necessarily 'optimal' in any wider sense.

4.4 Aggregation

Aggregate consumption is modelled using a standard partial equilibrium overlapping generations (OLG) framework. The economy consists of different cohorts. Each cohort is made up of a single agent but the cohorts alive in any given period are heterogeneous with respect to age and lifetime experience.

Aggregate consumption AC_t is defined as the sum of consumption over all N individuals at each period t:

(13) Assuming the same premium for the ARM and Index contracts is a simplifying assumption, reflecting the fact that there is no optionality implied by either contract as there is with the FRM contract. In general, these premia would depend on the costs to the lenders of providing the different contracts, depending partly on the structure of their liabilities, and would therefore not necessarily be the same.

$$AC_t = \sum_{i=1}^{i=N} c_{i,t}$$

Hence, aggregate consumption represents the response to common real interest rate, inflation and permanent income shocks but also captures the heterogeneity of response across different cohorts. Ideally, we would want to allow for all combinations of permanent/transitory and idiosyncratic/common shocks. This would require two state variables, one for each of the permanent states of the idiosyncratic and common shocks. To economise on one state variable, we follow Campbell and Cocco (2003) and identify permanent shocks as common and transitory shocks as idiosyncratic.[14]

The volatility of aggregate consumption growth is measured as the average standard deviation of the change in the log of aggregate consumption. This is the average across $nsim$ simulations of the standard deviation of consumption growth over t periods. This is calculated as (where ac indicates the log of AC):

$$\sigma_{\overline{AC}} = \frac{1}{nsim} \sum_{i=1}^{i=nsim} \sqrt{\frac{\left(\sum_{t=1}^{t=T} (ac_t - ac_{t-1})^2\right)}{t-1}} \qquad \textbf{(8)}$$

This gives the unconditional uncertainty of aggregate consumption.[15]

4.5 Model solution

The household-level problem reduces to three control variables: consumption ($c_{i,t}$), the decision about default ($def_{i,t}$), and (with FRMs) the decision about refinancing ($ref_{i,t}$). The state vector $X_{i,t}$ is given by:

$$X_{i,t} = (t, i_t, w_{i,t}, P_t, i_{REFt,k}, \upsilon_{i,t}, def_{i,1,t})$$

(14) This is perhaps not the most economically plausible interpretation but offers the best method for reconciling the relatively low variance of aggregate income growth and the relatively high variance of individual income growth, while not having excessively high permanent income uncertainty. The cost of this approach is that permanent income uncertainty is constrained to match aggregate income uncertainty and is consequently likely to be unrealistically low.
(15) It would also be valid over a large number of simulations to use the level of ac as aggregate consumption is stationary. Although permanent income is common and has a unit root, incomes for each cohort at the beginning of the life cycle are the same so that over a period of time there will be no trend in aggregate income or consumption.

where P_t is the price level and other variables are as defined above.

There is no known analytical solution to the household's optimisation problem. The programs solve the individual household's optimisation problem by backwards recursion from the terminal period, searching over discretised grids and interpolating the continuation value function using a spline.(16) The model has 7 state variables and 3 control variables. Calculating a solution is therefore extremely computationally intensive as a large number of gridpoints must be considered and the interpolation is very time consuming. For these computational reasons, the programs model the life of a mortgage as 15 two-year periods rather than 30 one-year periods, scaling variables appropriately, to reduce the number of points to be considered. In addition, the number of state variables is necessarily limited and it would be difficult to work with a model with additional state variables (such as housing) as the model would take much longer to solve. Given the derived decision rules, household-level consumption is then simulated using the model.

Aggregate consumption is obtained by simulating the OLG model of the economy for a set of macroeconomic and idiosyncratic shocks with one homogenous cohort in each different phase of the life cycle at each period in time. We simulate this economy $nsim$ times. Each simulation lasts t periods. Each simulation is initialised randomly.(17) We set $nsim$ sufficiently high to give convergence of the estimates of aggregate consumption uncertainty and other measures reported in the paper. Aggregate consumption variability is calculated according to equation **(8)**.

5 Calibration

This section discusses the calibration of the model to the United Kingdom. The unobservable parameters of the utility function (δ, γ) are set to the standard values assumed in the literature (see Appendix A for a full list of parameters). A drawback of this method is that we cannot model the high-frequency variability in payments on the ARMs common in the United Kingdom, where the interest rate adjusts almost immediately in response to changes in the policy rate.(18) All parameters are calibrated using annual data and scaled appropriately in line with the model

(16) The model is solved with code adapted from that used in the NBER working paper version of Campbell and Cocco (2003), kindly provided by the authors.
(17) As there is no systematic relationship between the macroeconomic shocks experienced during the initialisation and the simulation, this procedure will not result in any bias.
(18) The majority of mortgages in the UK are variable rates mortgages of this form (Standard Variable Rate, Discounted and Tracker mortgages). Furthermore, a high proportion of 'fixed-rate' mortgages are relatively short fixes. Hence, the 'ARM' in this model corresponds most closely to a UK fixed-rate mortgage.

assumptions.

5.1 *Interest rate and inflation processes*

Inflation is measured as the annual change in the retail prices index excluding mortgage interest payments (RPIX). The other parameters of the inflation process are estimated by OLS using equation **(5)**. Hence, σ_{ϵ}^{2} is the conditional variance of inflation.

Households base their decisions on *ex-ante* real interest rates. We measure the nominal interest rate as the UK policy rate. As we cannot observe inflation expectations directly, we use the model of inflation above to calculate expected inflation and subtract this from the nominal interest rate to obtain by the Fisher equation a measure of the *ex-ante* real interest rate. So we are effectively assuming that agents know the 'true' inflation process for each regime.[(19)] It is important to use *ex-ante* rather than *ex-post* real rates in order to avoid spurious correlation between inflation surprises and real rates. The real rate parameters are then derived using equation **(6)**.

In simulations where we allow for the cross-correlation of variables,[(20)] we use the estimated contemporaneous correlation of annual magnitudes using quarterly data (see Table C).

5.2 *Spreads and refinancing costs*

The ARM interest rate is set at a fixed spread (θ) of 1 percentage point over the risk-free interest rate. We abstract from any countercyclicality of this spread, which would tend to mitigate the effect of increases in the policy rate, as this variation in the spread is relatively small compared with the overall level of mortgage interest rates.[(21)]

It is more difficult to calibrate the difference between the ARM and the FRM spreads. This cannot be estimated empirically as there are very few long-term mortgages in the United Kingdom for either period we consider. Furthermore, even relatively short-term fixed rates have only been

---

(19) Hence, the mean and conditional standard deviation of inflation are the same as those given in Appendix A. Although highly stylised, this model fits actual inflation reasonably closely.

(20) Inflation, real interest rates and labour income growth. Labour income in these cases is given as aggregate real post-tax labour income.

(21) For example, evidence of fairly complete pass-through from policy to mortgage rates can be found in Mizen and Hofman (2002) (which also includes existing mortgages with fixed rates, where there is no pass-through).

available since the late 1980s so we do not have observations for the whole of the 1980s. As discussed in the Miles Review (2004), mortgage pricing in the UK mortgage market has been affected by cross-subsidisation and a lack of transparency. It is consequently far from clear which mortgage rates should be taken as 'representative', partly because the choice of mortgage depends on how well households are able to benefit from the cross-subsidy.

It is possible to derive the fair value of the option embedded in FRMs and calibrate the appropriate spread analytically using this and a profitability condition.(22) However, this is far from straightforward.(23) In addition, studies show that actual incidence of remortgaging depends on a number of household-specific factors and is hence an empirical issue (Lacour-Little (1999)). This suggests that the analytically implied spread may not be appropriate.

We assume that the spread of the FRM over the risk-free rate (θ') is 1.3 percentage points. This is higher than the spread on ARMs, to compensate the lender for the value to the borrower of the optionality embedded in the FRM contract. This is based on the average difference for the period 1998-2003 between the average spread on a discounted variable-rate mortgage over the base rate and the average spread on a three-year fixed rate mortgage over a swap rate of the same maturity. We hold this spread and the refinancing penalty constant over the two periods. One would expect mortgages to have been refinanced more frequently in the (more volatile) 1980s, raising the FRM spread in that period. So our calibration will tend to favour FRMs in the 1980s, relative to the 1990s.(24)

We set the cost of refinancing an FRM (ζ) to 1% of the original loan. This represents both a charge imposed by the lender to limit the value of the option to refinance, and other costs of refinancing such as administrative costs and time. Setting this to 1% of the loan yields a cost that is likely to be typical of what households face. The refinancing penalty varies with loan size, reflecting the relationship between loan size and the value of exercising the option.(25) There is clearly a relationship between θ' and ζ, as the amount households are willing to pay for the option

(22) This applies equally to the default spread on both ARMs and FRMs.
(23) For example, there is a complex interaction between the option to default and the refinancing option.
(24) Campbell and Cocco (2003) use a spread of 1.0 percentage point over the base rate for ARMs and 1.6 percentage points for FRMs. They also find difficulties in calibrating this given the specific features of the US mortgage and the 60 basis points difference in spread between the contracts is based on their judgement, rather than their lower empirical estimates.
(25) The assumption that the cost is fixed over the life of the mortgage makes the cost of exercising the option relatively higher later in the loan, when the outstanding balance is lower, but it is necessary to fix the size of the premium to limit the computational size of the problem.

to refinance will depend on how costly it is to exercise the option. Simulations suggest that the results are relatively insensitive to the size of the refinancing penalty over a plausible range of values for ζ.

5.3 Labour income

We calibrate the labour income process using existing studies for the United Kingdom. The labour income process in the model (see equation **(7)**) has a deterministic age-related income profile and transitory and permanent shocks. The profile of labour income is based on estimates for real age-specific labour income derived in Miles (1997) and reported in Miles (1999).(26) We use these estimates to generate the profile for the model and fix initial income.(27) The labour income profile is given in Appendix A. Estimating the other parameters is more complex, both due to the ambiguity of the empirical evidence for this purpose and the need to choose values that can be plausibly interpreted from both household and macroeconomic perspectives. The appropriate calibration is important as the overall year-on-year variability of income, as well as the degree to which shocks are transitory or permanent, has an impact on costs of different types of contract and the consumption behaviour they generate. It is also important that common and idiosyncratic shocks are set and identified appropriately, as only the latter will tend to cancel out on average across households.

We find that the annual standard deviation of real aggregate post-tax labour income growth for the United Kingdom is about 1.6% for the 1980s and 2% for the 1990s. Dickens (2000) estimates the annual standard deviation of real hourly earnings as around 16% (based on full-time male employees aged 22-59), using data from the New Earnings Survey (NES) from 1975 to 1995. Using this methodology, the serial correlation of these income shocks appears to be fairly low (in the range from 0.05 to 0.30). Although this paper finds an increase in income uncertainty over the sample, this occurred largely in the 1979 to 1982, which mostly precedes the 1980s period we consider. This paper also finds that permanent income uncertainty declines with age. These results are similar to those implicit in Jarvis and Jenkins (1995) in their estimated transition matrix for UK net income between 1991 and 1992 using data from the British Household Panel Survey (BHPS). This matrix can be used to derive an average (unweighted) standard deviation of 11%

(26) Compared with the US labour income profile used in Campbell and Cocco (2003), the UK profile is similar albeit somewhat flatter until age 40 but then declining less rapidly in later years than in the United States.
(27) This scaling is essentially arbitrary but set so that initial income is around £33,000 and the initial size of the loan is £100,000 for a LTIR ratio of 3.

across all income groups, although there is a relatively high proportion of movements between the extremes of the income distribution. Dutta, Sefton and Weale (2002) construct a model to explain the features of the changes in annual income between 1991 and 1992, which is reported in Jarvis and Jenkins (1995), and generate a higher estimate of the variance of income (around 25%). We assume that the labour income process is the same across both periods. There is some evidence from the New Earnings Survey that labour income uncertainty for males has increased in the 1990s (Nickell, Jones and Quintini (2000)) but we do not consider this possibility further.

On the basis of these previous results in the literature and the variability of aggregate income growth, we assume a standard deviation of permanent income shocks (σ_η) of 2% and standard deviation for transitory income shocks (σ_ω) of 11%.(28) This value of the uncertainty of permanent income is probably somewhat low, implying a probability of only 5% that permanent income in five years' time will be more than 10% either side of its expected level. This is likely to be considerably lower than the uncertainty of idiosyncratic permanent income facing households, particularly the young.

5.4 Loan to income ratios

We examine the effect of varying the loan to income ratio (LTIR) between 2.25, 3 and 3.75. This covers a plausible range of LTIRs for new borrowers, as suggested by evidence from the Council of Mortgage Lenders' Survey of Mortgage Lending database (see Chart 3).

5.5 Taxation

The rate of income tax is set to 0.3 for the 1990s and 0.45 for the 1980s. This is intended to represent both the average and marginal tax rate of middle-income households in the United Kingdom. For the 1980s, nominal mortgage interest payments are assumed to be deductible from tax at this rate, as was the case under the MIRAS (Mortgage Interest Relief At Source) scheme. MIRAS was progressively abolished during the 1990s and hence we do not allow for it in the 1990s regime. The difference in the tax rates during the two periods is set so that the tax take in both periods is approximately equal. The higher income tax rate in the 1980s regime compared

(28) Campbell and Cocco (2003) set $\sigma_\eta = 2\%$ (the standard deviation of the growth rate of average income in the US Panel Study of Income Dynamics data set). Transitory income uncertainty is set at $\sigma_\omega = 14.1\%$ and they also consider $\sigma_\omega = 25\%$.

with the 1990s reflects the high value to mortgage debtors of MIRAS in a regime of relatively high inflation (such as the 1980s): MIRAS applied to nominal interest rates so tax relief was essentially granted on a large part of the real amortisation of the loan under high inflation.

6 Results

This section considers the results of simulations of the model for a number of different cases (i) at household level and (ii) in terms of aggregate consumption. We consider the effect of changes in the key macroeconomic variables between the regimes, the relationship between the different variables in the new economic environment, and the effect of households having increased their borrowing since the 1980s.

6.1 Household-level consumption under different regimes

This subsection considers the implications of the change in monetary regime on household-level consumption.

Household welfare provides one way of summarising information about household-level consumption. Welfare is measured by calculating certainty-equivalent consumption, the constant stream of consumption that yields the same utility as actual consumption under each regime. For the purposes of comparison, we index the level of consumption under each contract to consumption in the 1990s under the indexed contract with LTIR=3. This gives a simple measure of the welfare costs for each type of contract/different house sizes. It is important to note that this measures welfare from non-housing consumption *only*. This qualification is important in two senses. First, the overall proportional difference in welfare across contracts is less than the difference in welfare derived from non-housing consumption, given that utility from housing is fixed for each LTIR. Second, the LTIR and housing consumption are exogenous in this framework, whereas households may in reality substitute between the housing and non-housing consumption as the shadow cost of housing consumption varies. As a result, this measure and comparisons based on it tend to understate welfare when the shadow cost of housing is relatively high compared to when the cost is low.

The calibration results suggest that the change in regime has had a substantial effect on household circumstances and behaviour (see Table D).[29] Welfare is higher in the 1990s under all contracts, often by a substantial margin. The gain in welfare is most marked for the ARM contract. In terms of the differences between the three types of contract, the Index contract would have been best in both regimes. In the 1980s, there would have been large advantage to FRM contracts over ARMs. This is because the cash flow risk induced by the volatility of nominal interest payments under ARMs would have been very costly. FRMs would have also been more costly in the 1980s than the 1990s, primarily due to greater wealth risk from less stable inflation and real interest rates as well as stronger front-end loading, but this would have had a lower quantitative impact on the welfare than the risks and costs of the ARM contract. In contrast to the 1980s, the two contracts have relatively similar welfare costs in the 1990s, although the ARM is marginally preferred.[30] In practice, this implies that the optimal choice of mortgage contract between these two would depend on individual household preferences and circumstances not captured by this model, and that a substantial proportion of households might find FRMs attractive. Interestingly, using Index contracts would give households a benefit worth around 8% of certainty-equivalent consumption even in the relatively stable regime of the 1990s.

Table D: Welfare under different mortgage contracts and regimes

	1980s	1990s
ARM	35.1	92.4
FRM	45.7	92.1
INDEX	94.1	100.0

Certainty-equivalent consumption, indexed to welfare under index contract in 1990 (=100) for LTIR = 3.

Table E shows welfare comparisons taking into account the estimated change in the correlation of the macroeconomic variables in our simulation (as discussed earlier in Table C). In most cases, allowing for these correlations reduces welfare compared to the case where the variables are uncorrelated. It also changes the ranking in terms of welfare under the different contracts.

The results for the 1980s regime include the effects of mortgage-interest tax deductibility and a tax rate of 0.45. In fact, MIRAS played an important role in terms of welfare in mitigating the effects of the high inflation 1980s regime. Table F shows a number of different tax scenarios.

(29) It is important to emphasise that these results are sensitive to the choice of parameter values, particularly in making comparisons with the three types of contract.
(30) Results from a similar model were considered in the Miles Review (Chapter 2, Miles (2004)) to assess whether households would find ARMs or FRMs more attractive.

Table E: Welfare under different mortgage contracts allowing for changes in the relationship between macroeconomic variables

	1980s	1990s
ARM	46.5	65.5
FRM	30.3	65.8
INDEX	22.3	60.5

Certainty-equivalent consumption, indexed to welfare under index contract in 1990 (=100) for LTIR = 3.

The 1980s is the baseline case discussed above. Welfare is somewhat lower in the 'no MIRAS' case (0.45 tax rate). The '1980s, 0.3 tax rate' case imposes MIRAS and the (lower) 1990s tax rate.[(31)] The key point is that, for a given level of revenue, a higher tax rate and MIRAS delivers higher welfare than a lower tax rate: MIRAS is an efficient tax relief system from the mortgage holder's perspective because it gives relief at the time when liquidity is potentially most highly valued, when mortgage payments are high.[(32)]

Table F: Welfare under different mortgage contracts in the 1980s with different tax regimes

	1980s	1980s, no MIRAS	1980s, 0.3 tax rate
ARM	35.1	35.8	67.9
FRM	45.7	41.6	80.5
INDEX	94.1	74.6	115.1

Certainty-equivalent consumption, indexed to welfare under index contract in 1990 (=100) for LTIR = 3.

The changes in welfare described above summarise a number of effects on household-level consumption induced by the change in regime. First, the economic environment has become more certain so agents are less exposed to both 'cash flow/liquidity' and 'wealth' risks. Second, the change in environment has altered the degree to which agents are credit constrained. The lower average level of inflation has reduced 'front-end loading' of mortgage payments, reducing the real value of mortgage commitments in the early years of mortgage. However, the higher real interest rate in the 1980s gave households a strong incentive to postpone consumption, and the constraint introduced by greater front-end loading of mortgage payments might not have bitten so hard. Third, lower real interest rates in the 1990s made housing less costly so, for a given size of mortgage, more non-housing consumption would have been feasible. This effect would hold even under perfect foresight and complete markets.

(31) The higher welfare in the second column for ARMs is due to the approximation used in setting the revenue neutral tax rate (0.45) for the 1980s.

(32) This should not be taken to imply anything about the desirability of MIRAS as a part of the tax system.

The interaction between these effects complicates the assessment of how much each of these factors contributes to the changes in welfare reported above. The change in certainty-equivalent consumption is much larger than the change in actual consumption. This implies that some of the change in welfare is attributable to either greater uncertainty of consumption, lower real interest rates or a smoother profile of lifetime consumption. The role of these different factors can be assessed using Chart 2. This shows the mean and ± 1 standard deviations of household-level consumption for an ARM contract in the 1980s and 1990s across $nsim$ simulations of each of the regimes. There are a number of striking aspects of the two consumption profiles. The profile of consumption in the 1990s is considerably flatter than in the 1980s. The average annualised growth rate of consumption is 2.4% for the 1990s, compared to almost 5.4% in the 1980s. This is likely to be the result of lower real interest rates, reduced front-end loading of debt repayments and a less risky macroeconomic environment. These effects reduce the need for buffer stock saving in the early years of mortgage, allowing households initially to save less and consume more. The chart also shows that, even with these higher levels of buffer stock savings, the unconditional *ex-ante* uncertainty of consumption is higher for the 1980s regime than in the 1990s.

To assess the marginal contribution of lower inflation,(33) we assess what welfare would have been in the 1980s if the level and volatility of real interest rates had been at 1990s levels (see Table G). This shows that welfare would have been lower in the 1980s than in the 1990s, even if the real interest rate process had been the same. Put another way, the results suggest that lower and more stable inflation accounts for around half of the gains of the move to a low inflation environment, with the rest accounted for by lower real interest rates.(34)

Table G: Welfare under different mortgage contracts in the 1980s with lower real interest rates

	1980s	1980s with 1990s real interest rate process	1990s
ARM	35.1	68.0	92.4
FRM	45.7	70.5	92.1
INDEX	94.1	103.7	100.0

Certainty-equivalent consumption, indexed to welfare under index contract in 1990 (=100) for LTIR = 3.

One other important difference between the 1980s and 1990s is that households have become

(33) This also includes the effects of the change in taxation.

(34) The difference welfare between the 1990s scenario under the index contract and the 1980s scenario with the 1990s real interest rate process is due to a shortcoming of the approximation used in implementing the different tax regimes.

Chart 2: Profile of household-level consumption under ARM

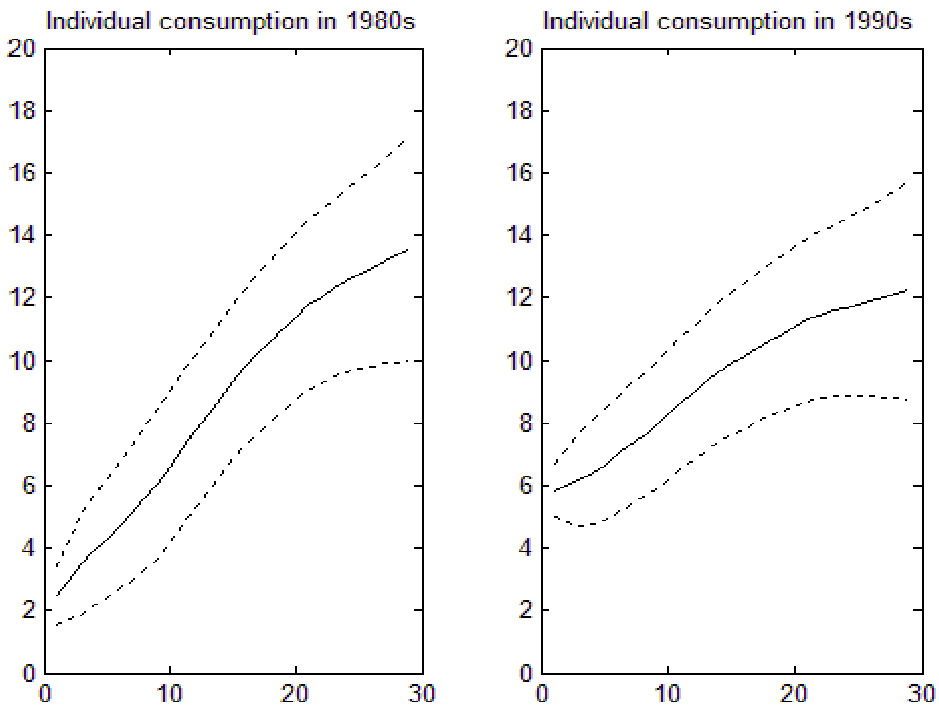

Dashed lines indicate ± 1 the annualised percentage standard deviation from the mean across simulations.

more indebted. The average LTIR for the 1990s regime is 2.3 compared to 2.1 for the 1980s. In addition, there has been a sharp increase in LTIRs towards the end of the 1990s sample with the mean LTIR reaching 2.8 during 2003. The whole distribution of new mortgages relative to borrowers' income has also shifted up with many more households having historically high LTIRs (Chart 3).

This increase in borrowing may itself reflect the change in monetary policy regime. The reason for this is implicit in Tables D and E: the utility costs in terms of non-housing consumption of a given loan have changed. Financing ownership of a house of a given value has become cheaper, in utility terms. The loss of certainty-equivalent non-housing consumption can be interpreted as an indicator of the expected cost of financing a house of a given value in each regime. This measure of cost captures both measurable costs, such as the change in the mean level of real interest rates, and also the shadow cost of financing a house in terms of the postponement of and uncertainty about consumption.

Chart 3: Size of new mortgages

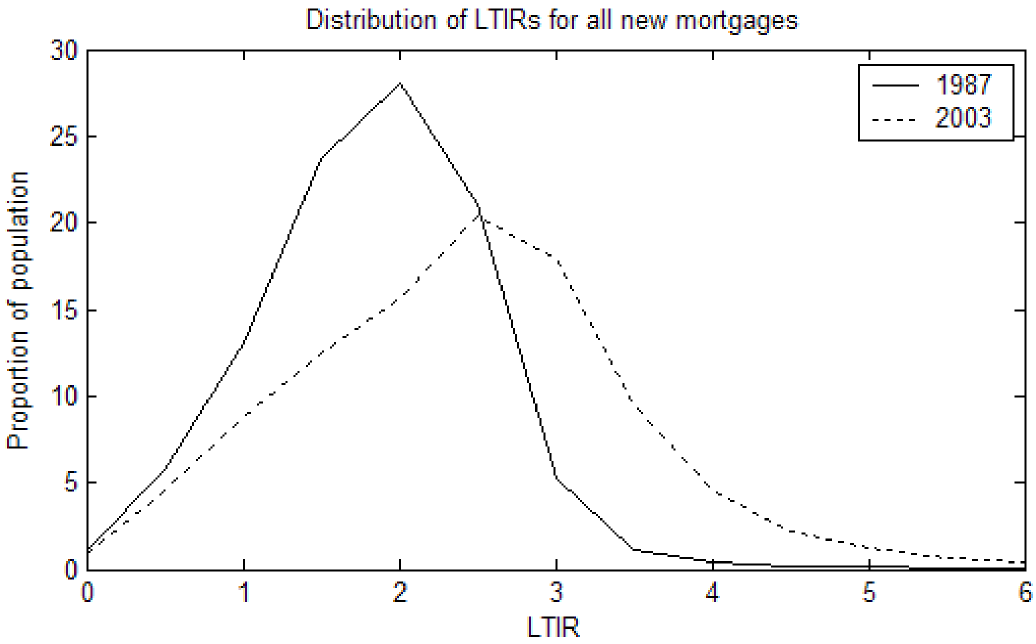

Sources: Bank of England and Council of Mortgage Lenders.

We initially consider the increase in LTIRs as exogenous and examine the impact of a higher average LTIR on consumption volatility between the two regimes. We find that the size of house (loan) makes a material difference to the cost of different types of mortgage contract. Table H shows that certainty-equivalent non-housing consumption declines rapidly in the size of loans. The size of loan size may also determine the choice of contract: FRMs are more attractive to households with larger debts relative to income, as the higher debt increases risk in general and makes finding a safer mortgage contract relatively important. To the extent that different households have different sizes of loan, this may partly explain why some households but not others would chose FRMs. The level of welfare under each contract is higher with smaller loans. As a result, the increase in the average LTIR between the 1990s and the 1980s would have somewhat mitigated the effects on non-housing consumption of moving to a more stable regime. By comparing columns (1) and (2), we see an increase in the LTIR that broadly corresponds to the rise between the 1980s regime and 2003.[(35)] This suggests that the combined effect of this increase in average LTIR to 2003 and the change in macroeconomic conditions is that

(35) The end of the 1990s sample.

certainty-equivalent consumption with current levels of indebtedness would still be somewhat higher in the 2003. Hence, households are currently sacrificing somewhat less non-housing derived utility to finance the ownership of houses as in the 1980s, although they are now borrowing larger amounts.

Table H: Welfare for different loan sizes

	1980s			1990s		
	2.25	3	3.75	2.25	3	3.75
ARM	76.7	35.1	22.3	114.7	92.4	65.1
FRM	80.4	45.7	23.6	113.8	92.1	66.0
INDEX	107.1	94.1	79.5	118.8	100.0	80.0

Certainty-equivalent consumption, indexed to welfare under index contract in 1990 (=100) for LTIR = 3.

The welfare statistics reported above can be interpreted as the opportunity cost in terms of non-housing welfare foregone of a given level of housing. This relationship can be inverted to show the amount of housing (relative to income) consistent with a given level of non-housing welfare. Chart 4 shows that the welfare costs of an LTIR of 2.25 in the 1980s is about the same as an LTIR of around 3.3 in the 1990s. This is a limiting case because it assumes that households would use all of the welfare benefits of the change in regime to spend more on housing, where in practice it is plausible that household preferences are such that they would substitute some non-housing consumption for housing consumption.[36] To consider substitution between housing and non-housing consumption is beyond the scope of this paper. However, this may provide at least an indication that the change in inflation regime may have had a substantial effect on housing demand.[37]

Overall, the shift in monetary policy regime appears likely to have had a substantial effect on household-level consumption. In particular, the consumption path is considerably flatter and less variable in the 1990s regime for a given LTIR. This change has generated substantial welfare benefits. These effects are even larger when one takes account of the change in the correlation between different macroeconomic variables. Around half of the gains can be attributed to the changes in the inflation process. Furthermore, the increase in debt levels observed in the 1990s

(36) Other caveats to this analysis are also important. First, housing is not treated in a realistic way in this model. For example, households are not allowed to change from their initial choice of housing, whatever shocks they experience. Second, in general equilibrium and with inelastic supply of housing, the more favourable macroeconomic conditions might increase the price of a house with given properties and thereby mitigate the effects on housing demand.

(37) This assumes that the LTIR constraint is binding and, more implausibly, that households can either costlessly vary the loan to value (LTV) ratio of their loan or the deposit required.

Chart 4: Possible effect on housing demand

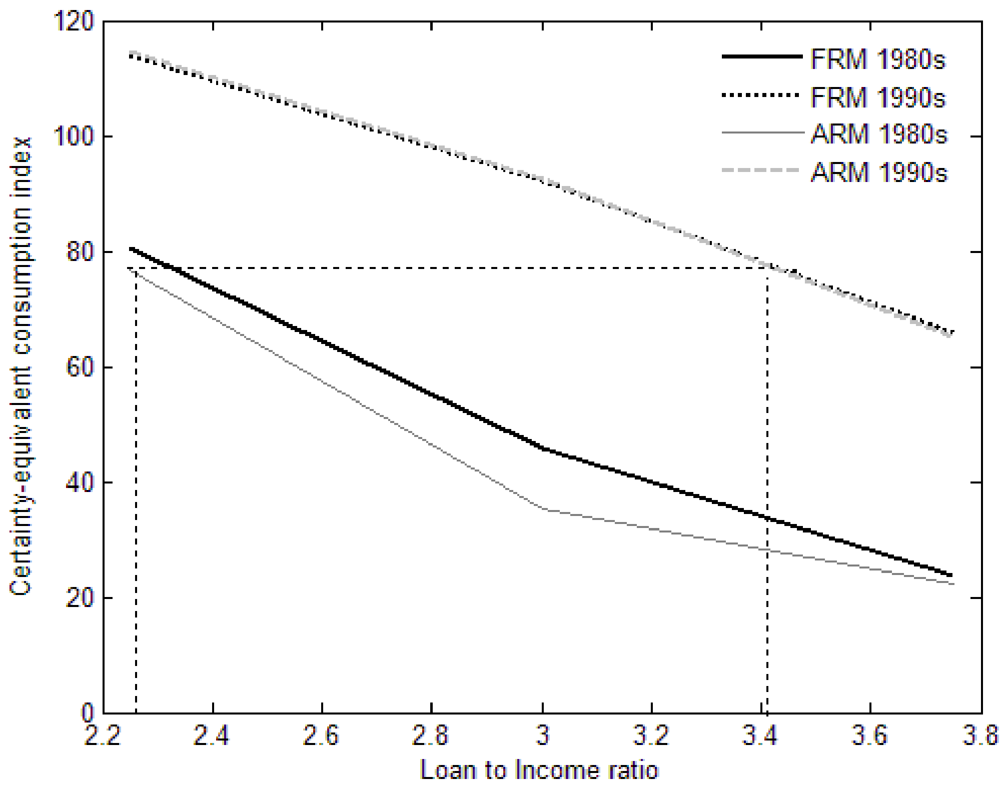

could potentially be rationalised as a response to the fall in the utility cost of mortgage debt. The simulations suggest that with the level of LTIRs reached by the late 1990s, households would have similar utility from non-housing consumption to the 1980s: higher indebtedness has partially offset the effects of the more stable regime.

6.2 Aggregate consumption volatility under different regimes

We now examine the behaviour of aggregate consumption under the two regimes. Translating knowledge of household-level consumption into the behaviour of aggregate consumption is not a trivial problem. The properties of individual consumption do not necessarily hold in the aggregate if the economy is made up of agents with heterogeneous preferences or circumstances. The analysis of aggregate consumption is further complicated by the non-linearity of individual consumption functions, changes in the correlation of shocks across individuals, and the relative strength of responses to aggregate and idiosyncratic shocks.

Our measure of volatility is the standard deviation of the difference of the log of aggregate consumption.[(38)] Aggregate consumption in this model is highly persistent and hence this is an appropriate measure because the (log) difference of consumption growth is stationary.[(39)] This also justifies converting the two-year growth rates from the model into annual terms in the standard way.[(40)]

Table I shows the volatility of aggregate consumption growth for each regime under each contract.[(41)] The volatility generated by the model under this parameterisation is somewhat higher than is observed in the data (see Table A). This may in part reflect the parameterisation of the labour income process,[(42)] and modelling only a subset of households. Table I shows that the annualised percentage standard deviation of consumption growth in the 1980s is 0.047 under ARMs, compared with 0.055 for FRMs.[(43)] Consumption volatility in the 1990s is higher under both contracts, although it remains lower for ARMs than FRMs.

Table I: Volatility of aggregate consumption

	1980s	1990s
ARM	0.047	0.053
FRM	0.055	0.058
INDEX	0.049	0.056

Annualised percentage unconditional standard deviations, LTIR=3.

Table J presents analogous results in the case where the correlation of macroeconomic variables is allowed to vary. This provides a measure of the sensitivity of our results to different specifications

(38) It is important to note that this is not equal to the sum of log growth rates.
(39) Note that aggregate consumption growth is equal to zero in this model as there is no aggregate labour income growth (each new generation begins with the same initial income).
(40) In the aggregate case, this change is exact (due to the stationarity of the series) while it is only an approximation for household-level consumption growth.
(41) This is the volatility over t years averaged over $nsim$ simulations so presents the average volatilty of consumption growth over a large sample of sequences of draws of the shocks.
(42) In calibrating the model, we matched the uncertainty of common shocks to the observed variance of shocks to aggregate labour income. However, the idiosyncractic labour income shocks only cancel on average so aggregate labour income in this calibration is somewhat higher than observed. Furthermore, the assumed homogeneity of each cohort tends to exacerbate this. There is some evidence that allowing for several members of each cohort, each facing individual transitory income shocks, would approximately halve aggregate consumption uncertainty but the qualitative results would not change. An alternative explanation for this divergence is that the population for observed aggregate consumption includes non-mortgage holders and hence may be less volatile than for those with mortgages considered in this model.
(43) Following Campbell and Cocco (2003), we annualise two-year values by division by the square root of 2. This is only an approximation because consumption in this case is not statistically independent between time periods. However, the degree of autocorrelation of consumption growth is not constant across time periods, being the outcome of the interaction of many factors. For this reason, we use the standard approximation.

of the macroeconomic processes, in addition to showing the specific effect of the identified change in the relationship between these variables. There are a number of interesting differences between these results and those in Table I.(44) For the 1980s, volatility is higher under the ARM and FRM contracts relative to the baseline case, narrowing the difference between the two types of contract. For the 1990s, volatility is higher under all contracts than in the baseline case and is more similar to the variability in the 1980s (although it remains marginally higher).

Table J: Volatility of aggregate consumption incorporating changes in cross-correlations

	1980s	1990s
ARM	0.067	0.074
FRM	0.071	0.073
INDEX	0.047	0.072

Annualised percentage unconditional standard deviations, LITR=3.

The finding that aggregate consumption volatility in the 1990s regime is slightly higher than in the 1980s (Tables I and J) seems to run counter to intuition and what we have observed. Intuition would suggest that a more stable economic environment, where individuals enjoy higher welfare for a given level of income, would also be characterised by more stable aggregate consumption growth through real-world mortgages. In contrast, we find a less stable outcome. It is not, however, unusual to find such apparently counterintuitive results for this type of problem in the literature (see Feigenbaum (2004)).

The underlying economic situation at each point in time our simulated economy is complex, consisting of many individuals with different characteristics (eg age, wealth), facing some common shocks (eg inflation) and some idiosyncratic shocks (transitory income uncertainty). The unconditional volatility of aggregate consumption will depend not only on the one-step-ahead variance(45) of each individual's consumption, but also the degree to which consumption is correlated across households.

One approach to understanding this is to use a simple accounting framework to decompose

(44) This emphasises the sensitivity of some of the results to the parameter values and functional forms of the main macroeconomic processes.
(45) Note the important distinction here between uncertainty in the next period and over the life cycle. Chart 2 shows the unconditional *ex-ante* uncertainty that households face about their future consumption at the beginning of their lives: it cannot be interpreted as directly showing the contribution to aggregate consumption volatility, which is related to the marginal uncertainty of consumption at each point in time. In short, individual-level consumption becomes more uncertain the further away it is in time from the present but it remains the same for aggregate consumption at any horizon.

aggregate consumption volatility into the variability of individual consumption growth and the covariance of consumption growth across cohorts (see Appendix B). Table K shows this decomposition applied to the ARM and FRM contracts in the 1980s and 1990s regimes. This shows greater synchrony in consumption across cohorts in the 1990s than the 1980s. There is also higher synchrony for FRMs than ARMs. This helps to explain both why aggregate consumption volatility is higher in the 1990s regime than the 1980s, and why the contracts are ordered as they are in terms of aggregate consumption volatility. Although the uncertainty in the economic environment is less correlated in the 1990s, the responses across cohorts are more closely synchronised.

Table K: Decompositions of aggregate consumption

		Covariance
1980s	ARM	0.07
	FRM	0.16
1990s	ARM	0.11
	FRM	0.18

Average covariance (excluding own variance) of changes in consumption across cohorts.

Aggregate consumption *functions* are also complex. Underlying individual consumption functions are not linear, either with respect to time or to shocks, complicating aggregation across cohorts. For example, one important source of non-linearity is credit constraints, whether hard or soft. Theory suggests that credit-constrained households may have relatively volatile consumption growth, and respond strongly to cash-on-hand, while unconstrained households have weak responses (no response to purely transitory shocks, proportional response to changes in permanent income). Furthermore, the weight of each cohort/generation in aggregate consumption will affect aggregate consumption growth and its volatility. There are material differences between the two regimes in the shares of aggregate consumption of different cohorts/generations: the higher real interest rate in the 1980s regime encouraged households to postpone consumption until later in their lives, as did the incentives to build up a precautionary balance of buffer-stock savings to counteract transitory shocks to nominal interest rates. As a result, the share of aggregate consumption in the 1980s regime was more weighted to older households than in the 1990s so, for a given set of household-level responses, the aggregate marginal propensity to consume may have been lower: older households are typically less credit constrained and hence have a lower marginal propensity to consume.

Given both non-linearity in individual consumption functions and the additional considerations introduced by aggregation, we cannot consider the marginal effects of particular determinants of aggregate consumption simply by subtracting one effect from another. But Table L, the analogue of Table G, provides an indication of the effect of the change in inflation by using the 1990s real interest rate process with the 1980s inflation and tax processes. The table shows that consumption volatility is lower in this case than with the actual 1980s interest rate process. This suggests that the 1990s inflation process might more than account for the higher consumption volatility in the 1990s.

Table L: Volatility of aggregate consumption in the 1990s with lower real interest rates

	1980s	1980s with 1990s real interest rate process	1990s
ARM	0.047	0.043	0.053
FRM	0.055	0.051	0.058
INDEX	0.049	0.046	0.056

Annualised percentage unconditional standard deviations, LTIR=3.

The changes in household borrowing, discussed above (Subsection 6.1), have material consequences for the volatility of consumption under different policy regimes. The variability of aggregate consumption increases with the LTIR so the increase in the average LTIR from 2.1 to 2.3 from the 1980s to the 1990s regime would tend to increase the volatility of aggregate consumption.

Table M: Aggregate consumption in different regimes including housing

	1980s			1990s		
	2.25	3	3.75	2.25	3	3.75
ARM	0.041	0.047	0.055	0.046	0.053	0.063
FRM	0.047	0.055	0.064	0.050	0.058	0.069
INDEX	0.043	0.049	0.056	0.049	0.056	0.067

Annualised percentage unconditional standard deviations.

The observed rise in indebtedness during the 1990s (with the average LTIR reaching 2.8 in 2003) suggests that the volatility of mortgage-holders' consumption currently would be somewhat higher than if LTIRs had remained the same as in the 1980s, although the effect is smaller than the change in macroeconomic regime itself (see Table M). Furthermore, the non-linear relationship between debt and consumption uncertainty shown in Table M suggests that comparing the uncertainty at the mean understates the actual effect of the wider distribution of debt in the 1990s regime (Chart 3). In the model, the greater number of households with large initial debts would contribute disproportionately to the uncertainty of aggregate consumption.

The result that aggregate consumption volatility, counterintuitively, is often found to be higher in the more stable macroeconomic environment of the 1990s is not completely general, although it remains an interesting possibility. Not only does it depend on the specification of the model but it also depends on the parameterisation. In particular, we consider the sensitivity of these results to the parameterisation of the labour income process given that, as discussed in Section 5.3, there are important limitations to the way this is modelled and calibrated. We consider two cases where aggregate consumption volatility is higher in the 1980s than the 1990s regime based on two alternative labour income processes but retaining all other features of the model and parameterisation as above.

First, we consider a case where there are no common permanent income shocks. In the results reported above, the common permanent income shocks are identical across regimes. In a linear model, this would suggest that we could ignore this source of variation when comparing across time,[46] but this reasoning does not apply in a non-linear framework. To investigate this effect, we consider a calibration where permanent income is certain. Table N shows that excluding this source of uncertainty has both a quantitative and qualitative effect on the volatility of aggregate consumption. As expected, consumption growth volatility is lower when this source of uncertainty is removed. More importantly, the result in the baseline case is reversed for this parameterisation: aggregate consumption growth volatility under ARMs and FRMs is actually higher in the 1980s and the 1990s.

Table N: Volatility of aggregate consumption with certain permanent income

	1980s	1990s
ARM	0.039	0.035
FRM	0.038	0.033
INDEX	0.028	0.030

Annualised percentage unconditional standard deviations, LTIR=3.

(46) This is because the different factors would be separable. In this non-linear framework, there may be some interaction between the effect of changes in policy regime and the way permanent income shocks are specified. Hence, the effect on consumption volatility of permanent income uncertainty is not necessarily the same across the two period and so cannot necessarily be ignored. Of course, the results from ignoring permanent income uncertainty would differ from those with this source of uncertainty even in a linear model.

Second, we consider results for the case where there is no transitory income uncertainty as a test of the sensitivity of the results to the parameterisation of the labour income process. Table O shows that removing transitory income uncertainty also changes both the quantitative and qualitative implications of the model given this parameterisation. As with certain permanent income (see Table N), eliminating the volatility of transitory income implies that aggregate consumption volatility in the 1990s would have been lower, not higher, than in the 1980s.

Table O: Volatility of aggregate consumption with certain transitory income

	1980s	1990s
ARM	0.059	0.047
FRM	0.042	0.033
INDEX	0.043	0.034

Annualised percentage unconditional standard deviations, LTIR=3.

6.3 Explaining changes in aggregate consumption volatility

Empirically, consumption was more stable in the 1990s period than previously. But our model predicts the opposite (see Table A), despite the more stable and favourable economic environment and the large welfare benefits this generated for mortgage debtors. In addition, allowing for the increased level of borrowing in the 1990s regime (consistent with the lower non-housing welfare costs of mortgages) suggests that aggregate consumption should have been more volatile still. One reason for this result could be the restrictive framework of the model used and particular parameterisation chosen, particularly for the labour income process. This is also a partial equilibrium framework so there are important effects this model does not account for at all, such as how policymakers might respond in their interest-rate setting behaviour to the changing properties of consumption. In addition, there are some important features of the world such as endogenous housing demand that model is unable to incorporate due to computational constraints. However, there are also other explanations which can rationalise the discrepancy between these results and observed behaviour.

First, it is possible that the fall in the volatility of consumption is due primarily to *non*-mortgage holders (with consumption of mortgage holders, modelled here, remaining unchanged or having fallen by less).[47] This could be tested in further work.

Second, the assumption that households cannot borrow other than for house purchase has become increasingly unrealistic given continued financial liberalisation (Fernandez-Corugedo and Price (2002)). As a result, it is plausible to assume that fewer households are credit constrained in the 1990s regime than the 1980s in the sense that banks are more willing to lend other things being equal. This may have offset the effects identified above and hence explain the reduction in the observed volatility of aggregate consumption.

Third, the characterisation in the model of the key macroeconomic processes is extremely stylised. This applies particularly to labour income. In practice, the change of regime is also likely to be have been more subtle than our model describes. Interest rates are not set exogenously but in response to economic conditions. As a result, one possible explanation of lower aggregate consumption volatility is that policymakers in the 1990s were more responsive than in the past to consumption, lowering rates to ease negative shocks to consumption rather than using interest rates to rein in consumption.

Fourth, the model describes the steady-state properties of aggregate consumption volatility under two regimes, but does not model the transition from one regime to another. It is possible that households have reacted slowly to the change in regime, only gradually adjusting their behaviour to become less 'cautious' in response to the more stable economic environment. As a result, observed consumption volatility in the 1990s might have been lower than the 1990s steady state described by the model. An important aspect of the transition is how indebtedness evolves. We know that households have only gradually adjusted the size of the mortgages they take out initially. The increase in the *flow* of new borrowing implies that the *stock* of debt will increase but only gradually, as new larger loans replace smaller older loans. In addition, debt is repaid more slowly in real terms in 1990s than the 1980s regime due to weaker 'front-end loading' (as shown in Chart 1), so this will tend to cause the debt stock to rise during the transition. All of these factors imply that the debt stock in the 1990s, which may take several decades to adjust, would

(47) In principle, it would also be possible to take into account the effect of changing proportions of ARM and FRM mortgages in overall mortgage debt. It would, however, be unsatisfactory to do this in the context of this model given that the FRM contract in this case is much longer than those common in the United Kingdom.

have been below its new long-run equilibrium (see Hamilton (2003)). To the extent that lower debt is associated with greater stability of consumption, the average consumption volatility observed in the 1990s shown in Table A may understate the true volatility of the new regime.

7 Conclusions

This paper considers the relationship between changes in the macroeconomic environment and household consumption behaviour, at both household and aggregate levels. We find that there have been important changes in the time-series properties of key macroeconomic variables between the 1980s and 1990s regimes, consistent with the change in the UK monetary policy regime (among other explanations).

We show that properties of 'real-world' mortgages, rather than simple real debt accumulation equations, imply that the timing and variability of mortgage payments depend on the inflation and real interest rate processes. In particular, the shift from the 1980s to the 1990s regime suggests that debt repayments would have been more 'front-end loaded' in the earlier regime, and that real mortgage payments would have been more variable.

We use a model closely related to Campbell and Cocco (2003) to explore the effect of the change in the economic environment on household consumption and portfolio decisions. We consider household-level consumption under adjustable-rate mortgages (ARMs) and long-term fixed-rate mortgages (FRMs). Comparing the two regimes, we find that individual households would have experienced more certain consumption in the 1990s than in the 1980s. In addition, the path of life-cycle consumption would have been considerably steeper for the 1980s: higher uncertainty and more severe front-end loading would have induced greater buffer stock saving. This greater uncertainty would have depressed consumption early in the life cycle. We find that there are substantial welfare gains from the switch from the 1980s to the 1990s regime when measured in terms of certainty-equivalent consumption. We also find that allowing for the observed relationships between variables in the two periods has some impact on these results, particularly in increasing the consumption uncertainty in the 1980s under ARM contracts relative to FRMs.

We model aggregate consumption from household-level consumption using a partial equilibrium overlapping generations framework, where macroeconomic shocks are common across different generations but there are idiosyncratic (cohort-specific) transitory shocks to labour income. We find, perhaps surprisingly, that the variance of aggregate consumption growth is predicted to be higher in the 1990s regime than the 1980s. Consumption is more variable with FRMs under both regimes under this model specification and parameterisation, although the more intuitive case (consumption volatility lower in the 1990s steady state as was actually observed) holds under two alternative labour income specifications. We consider the effect on consumption of the considerable increase in indebtedness since the 1980s, as higher levels of debt tend to increase the volatility of consumption. The increase in initial loan to income ratios (LTIRs) between the 1980s and the 1990s regime would have tended to exacerbate the effect of the change in regime on consumption at both household and aggregate level.

In conclusion, the model suggests that household-level consumption would be smoother over the life cycle and less variable in the 1990s than the 1980s regime. But aggregate consumption would not necessarily have been less volatile in the 1990s regime. The difference between this result and observed changes in aggregate consumption volatility may either be due to shortcomings of the model (such as disregarding the behaviour of non-mortgage holders) and the parameterisation we employed, or because the observed consumption volatility during the 1990s has been below its new steady state as households have benefited simultaneously from a more stable macroeconomic environment and the lower levels of indebtedness inherited from the past.

Appendix A: Calibrated parameter values

Table A.1: Model parameters

	Common	1980s	1990s
γ	3	-	-
δ	0.02	-	-
μ	-	0.056	0.025
σ_{ϵ}	-	0.02	0.008
ϕ	-	0.89	0.75
$\bar{r}$	-	0.061	0.034
σ_{ψ}	-	0.026	0.013
σ_{η}	0.02	-	-
σ_{ω}	0.11		
D	LTIR = {2.25, 3, 3.75}		
ζ	$\frac{D}{T}$		
θ	100 bps	-	-
θ'	160 bps	-	-
tax	0.3	Deductibility of mortgage interest	- -

All parameters expressed as annual values (appropriately scaled up two-year values are used for the simulations).

Table A.2: Labour income profile

Period	Labour income
1	33.3
2	35.5
3	37.1
4	38.6
5	40.0
6	41.3
7	42.3
8	43.2
9	43.9
10	44.4
11	44.7
12	44.8
13	44.6
14	44.3
15	43.7
16	42.9

Figures indicate average income over each two-year period in £000s, where period 1 is set to 33.3 (as discussed above).

Appendix B: Decomposition of aggregate consumption volatility

This shows a method for the decomposition of the volatility of aggregate consumption. Denote the consumption in simulation i of cohort j at time time by c_{ijt}. As discussed previously, aggregate consumption volatility σ_C is given by aggregating over cohorts, calculating the log difference for each simulation, taking the standard deviation over time, and averaging across simulations to reduce sampling error. That is:

$$\begin{aligned}
\sigma_c &= I^{-1}\sum_i \surd\left(T^{-1}\sum_t \left[D\log\left(\sum_j Dc_{ijt}\right)\right]^2\right) \\
&\backsimeq I^{-1}\sum_i \surd\left(T^{-1}\sum_t \left[\frac{\sum_j Dc_{ijt}}{\sum_j Dc_{ijt-1}}\right]^2\right) \\
&= I^{-1}\sum_i \surd\left(T^{-1}\sum_t \left[\Sigma_j D\widetilde{c}_{ijt}\right]^2\right)
\end{aligned}$$

where the last step introduces $D\widetilde{c}_{ijt}$ as the change in individual consumption normalised by the previous period's aggregate consumption. With a little algebra we can express σ_C in terms of the covariance of consumption growth across cohorts:

$$\begin{aligned}
\sigma_c &= I^{-1}\sum_i \surd\left(T^{-1}\sum_t \left[\Sigma_j D\widetilde{c}_{ijt}\right]^2\right) \\
&= I^{-1}\sum_i \surd\left(T^{-1}\sum_t\sum_j\sum_k D\widetilde{c}_{ijt}\widetilde{c}_{ijkt}\right) \\
&= I^{-1}\sum_i \surd\left(T^{-1}\left[\sum_j Var_t\left(D\widetilde{c}_{ijt}\right) + 2\sum_{j\neq k} Cov\left(D\widetilde{c}_{ijt},\widetilde{c}_{ijkt}\right)\right]\right)
\end{aligned}$$

This expression shows that aggregate consumption volatility depends on both the volatility of each cohort's consumption growth and how it covaries with other cohorts' consumption growth.

References

Alm, J and Follain, J R (1984), ‘Alternative mortgage instruments, the tilt problem and consumer welfare’, *Journal of Financial and Quantitative Analysis*, Vol. 19, pages 113–26.

Benati, L (2004), ‘Evolving post-world war II U.K. economic performance’, *Journal of Money, Credit and Banking*, Vol. 36, No. 4, August, pages 691–718.

Brueckner, J K (1997), ‘Consumption and investment motives and the portfolio choices of homeowners’, *Journal of Real Estate Finance and Economics*, Vol. 15, No. 2, pages 159–80.

Campbell, J Y and Cocco, J F (2003), ‘Household risk management and optimal mortgage choice’, *Quarterly Journal of Economics*, Vol. 119, No. 4, November, pages 1,449–94.

Carroll, C D (2001), ‘A theory of the consumption function, with and without liquidity constraints’, *Journal of Economic Perspectives*, Vol. 15, No. 3, pages 23–45.

Cocco, J F (2001), ‘Hedging house price risk with incomplete markets’, *London Business School - Institute of Finance and Accounting Working Paper No. 317*.

Deng, Y and Quigley, J M (2003), ‘Woodhead behaviour and the pricing of residential mortgages’, USC Lusk Center for Real Estate, *Working Paper No. 2003-1005*.

Dickens, R (2000), ‘The evolutions of individual male earnings in Great Britain: 1975-95’, *Economic Journal*, Vol. 110, January, pages 27–49.

Dutta, J, Sefton, J A and Weale, M R (2002), ‘Income distribution and income dynamics in the United Kingdom’, *NIESR Discussion Paper No. 116*.

Feigenbaum, J (2004), ‘The contribution of borrowing constraints and uncertainty to aggregate saving’, *unpublished mimeo*, www.pitt.edu/ jfeigen/compaggsaving.pdf.

Fernandez-Corugedo, E and Price, S (2002), ‘Financial liberalisation and consumers’ expenditure: ‘FLIB’ re-examined’, *Bank of England Working Paper no. 157*.

Flavin, M D and Yamashita, T (2002), ‘Owner-occupied housing and the composition of the household portfolio’, *American Economic Review*, Vol. 92, No. 1, March, pages 345–62.

Fratantoni, M C (2001), ‘Homeownership, committed expenditure risk, and the stockholding puzzle’, *Oxford Economic Papers*, Vol. 53, pages 241–59.

Fu, Q, Lacour-Little, M and Vandell, K D (2000), ‘Commercial mortgage prepayments under heterogeneous prepayment penalty structures’, *Real Estate Research Institute Working Paper No. 82*.

Hamilton, R (2003), ‘Trends in households’ aggregate secured debt’, *Bank of England Quarterly Bulletin*, Vol. 43, No. 3, Autumn, pages 271–80.

Jarvis, S and Jenkins, S (1995), 'Do the poor stay poor?', *ESRC Research Centre on Micro-Social Change Occasional Paper No. 95-2*.

Kau, J B and Keenan, D S (1995), 'An overview of option-theoretic pricing of mortgages', *Journal of Housing Research*, Vol. 6, No. 2, November, pages 217–44.

King, M A (2002), 'The inflation target ten years on', *Bank of England Quarterly Bulletin*, Winter, pages 459–74.

Lacour-Little, M (1999), 'Another look at the role of borrower characteristics in predicting mortgage prepayments', *Journal of Housing Research*, Vol. 10, No. 1, pages 45–60.

Lehnert, A and Statman, M (2003), 'Mortgage choice', *unpublished mimeo*, http://andreas.marginalq.com/Research/ARM/mortchoice.pdf.

Miles, D M (1997), 'A household level study of the determinants of incomes and consumption', *Economic Journal*, Vol. 107(440), January, pages 1–25.

Miles, D M (1999), 'Modelling the impact of demographic change upon the economy', *Economic Journal*, Vol. 109(452), January, pages 1–37.

Miles, D M (2004), 'The UK mortgage market: taking a longer view', *HM Treasury*, March.

Mizen, P H and Hofman, B (2002), 'Base rate pass-through: evidence from banks' and building societies' retail rates', *Bank of England Working Paper no. 170*.

Nickell, S, Jones, T and Quintini, G (2000), 'A picture of job insecurity facing British men', *CEP Discussion Paper*, March.

Rotemberg, J R and Woodford, M (1999), 'Interest rate rules in an estimated sticky price model', in Taylor, J B (ed), *Monetary Policy Rules*, University of Chicago Press.

Sanders, A B and Slawson, C V (2001), 'Pricing shared appreciation mortgages', *unpublished mimeo*.

Shilling, J D, Dhillon, U S and Sirmans, C F (1987), 'Choosing between fixed and adjustable rate mortgages', *Journal of Money Credit and Banking*, Vol. 19, pages 260–67.

Stiglitz, J E and Weiss, A (1981), 'Credit rationing in markets with imperfect information', *American Economic Review*, Vol. 71, No. 3, June, pages 393–410.

Taewon, K (1987), 'A contingent claims analysis of price-level adjusted mortgages', *AREUEA Journal*, Vol. 15, No. 3.