# Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK.*

**Orazio Attanasio**†
(University College London,
Institute for Fiscal Studies, NBER and CEPR)
**Andrew Leicester**‡
(Institute for Fiscal Studies and University College London)
and
**Matthew Wakefield**§
(University of Bologna and Institute for Fiscal Studies)

November 2010

**Abstract**

Three mechanisms have been suggested to explain the strong correlation between house prices and consumption: *wealth effects* and *collateral effects* involve house prices driving consumption changes, but *common causes* may drive both house price and consumption growth simultaneously. In this paper, we use a realistic structural life-cycle model of consumption and housing decisions to understand how data might distinguish between these mechanisms. The driving forces of the model are earnings and price shocks. While the latter have only an aggregate component, the former include both an idiosyncratic and aggregate components. Their properties are estimated from appropriate data. We calibrate our model, which incorporates realistic features of the UK mortgage market, to fit some life cycle facts. The model is then simulated, feeding into it realized aggregate shocks from the late 1970s to the early 2000s as well as counterfactual scenarios. These exercises

*Financial support from the ESRC (Grant RES-000-22-1913) is gratefully acknowledged. This work builds on several previous joint works with a collection of co-authors. We fully acknowledge the contributions to our thoughts of Laura Blow, Renata Bottazzi, Robert Hamilton, Hamish Low, Lars Nesheim and Guglielmo Weber. Special thanks are due to Hamish for his generosity in sharing his knowledge of constructing and solving dynamic models. We also thank Richard Blundell, Christopher D. Carroll, Marjorie Flavin, Costas Meghir, Francois Ortalo-Magné, Victor Rios-Rull for helpful discussions and an anonymous referee and a co-editor for useful comments. This paper was delivered as the Schumpeter Lecture at the Meetings of the European Economic Association held in Glasgow in August 2010. Early versions were presented at the University of Frankfurt, the Board of Governors of the FED, the NBER Summer Institute 2008 (meeting on Aggregate Implications of Microeconomic Consumption Behaviour), the Royal Economic Society Annual Conference 2008, the Bank of England and the University of Padova, where we received useful feedback. Data from the Family Expenditure Survey/Expenditure and Food Survey are Crown Copyright and used with the permission of the Controller of HMSO and the Queen's Printer for Scotland. We acknowledge the Office for National Statistics and Department for Environment Food and Rural Affairs as the original data creators and the UK Data Archive as the source of the data. Data from the British Household Panel Survey are collected by the Institute for Economic and Social Research, funded by the Economic and Social Research Council and downloaded from the UK Data Archive. These organisations bear no responsibility for the analysis or interpretation of the data in this paper. The usual disclaimer applies.

†o.attanasio (at) ucl.ac.uk

‡andrew_l (at) ifs.org.uk

§matthew.wakefield (at) unibo.it

establish: (a) the model fits the aggregate data well; (b) the model confirms the intuition of earlier studies that house price shocks should have a larger effect on the consumption of older households while earnings shocks should have a larger effect on young households; (c) simulation and aggregation of the model yields somewhat counterfactual cross sectional responses of consumption over observed business cycles: while in the data it is the youngest consumers' consumption that fluctuates the most, in the model this is not true as older households' consumption also moves substantially.

Keywords: Consumption, house prices, life-cycle models
JEL codes: D91 E21

# 1 Introduction

There is a widespread consensus that house prices and consumption move together over the business cycle, at least in developed countries. Evidence from many developed countries, including the US, Ireland, Japan and the UK, has made this correlation apparent. In the UK, the synchronization between the cycles of house prices and consumption over the last 35 years has been particularly strong. Figure 1 shows the remarkable co-movements of house price growth and consumption growth. Between the 1970s and today, one can clearly identify four episodes (two in the 1970s, one in the late 1980s and one in the 2000s) where large increases in house prices were mirrored by 'consumption booms'. In each case, these booms were followed by dramatic falls in house prices and were reflected in declines in consumption. The correlation coefficient between the two time series plotted in Figure 1 is 0.74. It should be stressed that this correlation (and the rescaling of the plots in Figure 1 to fit in the same graph) hides the fact that house price growth is much more volatile in the UK than the consumption growth rate: the standard deviation of the house price series is approximately 9.6%, while that of the consumption series is 3.0%. We will come back to these different orders of magnitude in the conclusions.

The correlation between house prices and consumption makes changes in house prices an important indicator for those wishing to judge inflationary pressures within the economy. Indeed, this indicator is closely watched by the Bank of England's Monetary Policy Committee precisely because, in the words of Nickell (2004), "... the evidence suggests that house price inflation is significantly related to household consumption growth and hence to aggregate demand growth and future consumer price inflation in the economy". And as house price booms and busts become prevalent in emerging economies, such as China, understanding the nature of this correlation becomes even more important.

To understand the implications of house price movements for consumption, and therefore the appropriate policy response to conditions in the housing market, it is necessary to identify the drivers of the link between house prices and consumption. The main issue is to establish whether the link is causal or not and, if so, to identify the channels through which it operates.

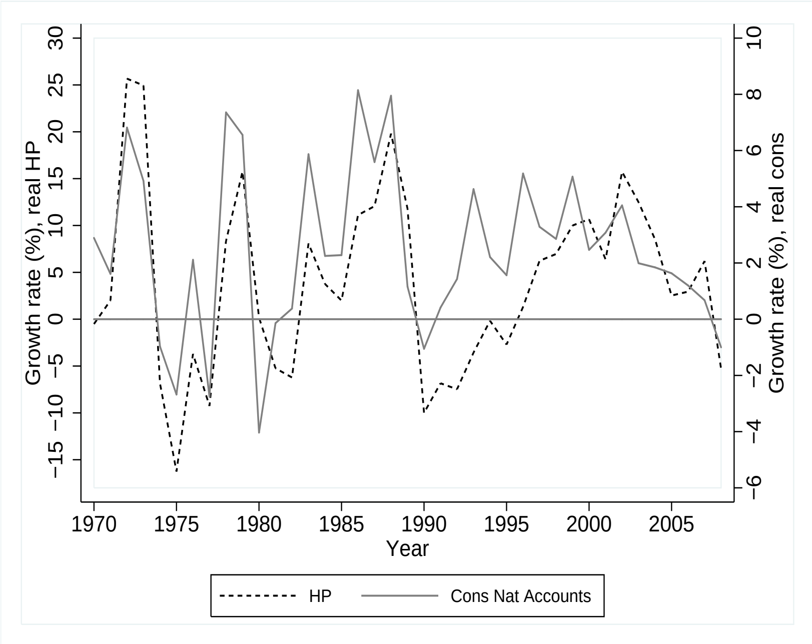


Notes: Consumption series is household final consumption expenditure (national definition) taken from ONS 'Consumer Trends'. House prices are mix-adjusted and taken from DCLG Housing Statistics. Both series show real-terms growth rates, deflating values to 2009 prices using the all-items Retail Prices Index.

Figure 1: Growth in real consumption and house prices

There exists a large literature on the the relationship between house prices and consumption and several hypotheses have been advanced. Some authors have suggested that the relationship might be driven by a wealth effect. According to this hypothesis, as individual consumers experience a capital gain on their real estate - which constitutes an important fraction of household savings - they 'cash it' and increase consumption, especially if the increase in house prices is perceived as permanent. Muellbauer and Murphy (1990, 1997) extensively analyzed the UK consumption boom of the late 1980s and proposed a story along these lines complemented by changes in credit markets. For individual consumers to be able to 'cash in' the capital gains on real estate, they must have access to well-developed financial markets. Muellbauer and Murphy (1990) stress forcefully the importance of the financial market liberalization of the early 1980s, in conjunction with the large increase in house prices, in explaining the consumption boom of the late 1980s in the UK. Recent evidence relating to this issue can be found in Aron et al. (2007). Aron et al. (2010) have proposed a similar narrative to explain the consumption growth of the UK, the US and Japan. International evidence on the role of wealth and (more indirectly) credit markets in explaining consumption growth is Slacalek (2009).

The credit channel can also take more sophisticated and subtle avenues. In markets affected by incompleteness and a variety of imperfections, it might be easier to insure idiosyncratic risk during periods of house prices booms by borrowing against future income using the increased value of housing as collateral. This channel has been analyzed, for the US, by Lustig and Nieuwerburgh (2005) and, more recently, by Piazzesi and Schneider (2010).

King (1990) and Pagano (1990), in commentaries on the Muellbauer and Murphy (1990) paper, questioned the importance of wealth effects in accounting for the link between house prices and consumption on theoretical and empirical grounds. While it is true that houses represent an important component of household wealth, they also provide services. An increase in house prices, therefore, can have many interconnected effects. An increase in the price of housing services would have both a substitution and, possibly, an important income effect. Moreover, while an increase in house prices constitutes a capital gain for home owners or, more precisely, people planning to downgrade their housing stock, the opposite is true for people planning to buy a house or upgrade their housing stock.

An alternative view of the link between house prices and consumption is that the observed correlation simply reflects factors such as changes in expected income growth, tax changes or changes in credit market conditions that may lead to increases in both households' expenditure and house prices. The causality could go from income expectations, to consumption, to house prices.

Different mechanisms might have effects that are observationally equivalent on aggregate data. In other words, a given time series of aggregate consumption and house price data may be consistent with a variety of explanations. Household level data, on the other hand, may provide additional insights if the implications of the various hypotheses considered are different for different households. Attanasio and Weber (1994) (hereafter AW, 1994) made this point and used such micro-data in order to examine the UK consumption boom of the late 1980s, which coincided with a rise in real house prices of slightly more than 40% in the space of four years.

Within a standard life-cycle model, if wealth effects were important they would have the biggest impact on the consumption of older individuals, who are most likely to have equity in any housing assets, and who have a relatively short time horizon over which to enjoy such capital gains. On the other hand, changes in the expected growth rate of future earnings and productivity should have considerably larger effects on young households, as they have a longer horizon over which to enjoy the increased earnings. The evidence from micro data reported in AW (1994) indicated unequivocally that the 1980s consumption boom was driven in large part by strong consumption by the young. AW (1994) interpreted this evidence as indicating that this boom, and the simultaneous boom in house prices, may have owed much to common causes rather than to a causal link running from increased house prices to consumption.[1] AW's exercise was recently extended in a variety of ways by Attanasio et al. (2009, 2005), hereafter ABHL, who, by and large, confirmed the results in AW (1994).[2] ABHL stressed

[1]Financial market liberalization surely played an important role, as AW (1994) discuss. From a conceptual point of view, growth in expected future income can only be translated into consumption via borrowing (or reduced saving). AW (1994) also provide some evidence for the importance of this channel.

[2]We should mention that there is no full consensus about these findings. In a recent paper, Campbell and Cocco (2007) argue that their empirical results suggest wealth effects to be the most likely explanation of the correlation between house prices and consumption. They find that increases in consumption observed during recent house price booms were mainly driven by increases in the consumption of home owners (rather than renters) and older consumers rather than younger ones. Campbell and Cocco (2007) also use a structural model to assess whether endogeneity of the home-ownership decision might bias their empirical results. By contrast, ABHL find the relationship between consumption growth and house prices to be strongest for younger individuals and argue that the common causality channel is most important. Both papers use micro-data from the UK Family Expenditure Survey / Expenditure and Food Survey and yet reach very different conclusions. It is not clear why this is the case, though there are several methodological differences between the papers. A comparative study by Cristini and Sevilla-Sanz (2008) attempted to replicate both studies as closely as possible. The ABHL

the finding that over several consumption booms and busts, the consumption of younger consumers seemed to be the most volatile and the one that mostly explained the overall variation. We reproduce and update some of these findings in Section 2.

To provide intuition about the effect of a change in future income growth on the consumption of young and old consumers, AW (1994) simulated a simple life-cycle model and showed that an upward revision of expectations of future income leads to the consumption of the young (with longer to enjoy the higher income stream) responding more strongly than that of the old. The model used by AW (1994) is, however, extremely simple and, crucially, does not consider houses, either as an asset or as a provider of housing services.

Houses are complicated objects. They are important stores of value, but also provide housing services that are likely to be non separable from other consumption. Houses are lumpy and subject to transaction and adjustment costs. They can be used as collateral in incomplete credit markets. All these elements imply that changes in the value of houses (current and expected) can affect consumption through a variety of channels, some of which are quite subtle. The complicated nature of the relationship between housing assets and non-durable consumption means that it may be perilous to rely on a stylised model for insights. AW themselves acknowledged this shortcoming.

For this reason, we build and calibrate a realistic and complex life-cycle model of consumption and housing decisions, where individual households, who face uncertain and realistic processes for house prices and earnings, can access asset markets that closely resemble those available to UK households to finance house purchases and to save. With this model, we can derive the implications of shocks to house prices and earnings for the consumption of different groups in the population and for aggregate consumption. We can therefore check in a rigorous fashion whether such a structural model confirms our intuitions about how and why the consumption of different groups might correlate with house prices, intuitions which have formed the basis of past empirical tests.

---

results were found to be robust, whereas the Campbell and Cocco results were sensitive to the specification and the version of the data set used.

Section 3 describes the model of individual behavior that is the main tool of our analysis. We parameterize the stochastic processes faced by our agents to include both idiosyncratic and aggregate components. The former are calibrated using micro data. The latter include *aggregate* shocks to house prices and incomes (which are experienced by everyone in the economy at the same point in time). We estimate the parameters of the time-series processes for house prices and aggregate earnings using data covering the last 35 years in the UK. The estimation using aggregate data also yields estimates of the aggregate shocks for the sample period we analyze, and these are used as inputs into our simulations so that we model the effects of realized shocks to house prices and earnings. We are not aware of other studies that use this combination of aggregate and micro data in the calibration of an individual-level model. And yet it is, in our opinion, important given that we want to understand aggregate fluctuations by aggregating individual consumption.

In Section 4, we choose the remaining free parameters (notably the preference for housing) of our model to reproduce certain facts about home ownership at different ages. Having so obtained the parameters of our model, in Section 5 we simulate individual behavior *given the set of aggregate innovations to house prices and earnings that we estimate from our time series data.* Aggregating up the simulated behavior of heterogenous individuals, we are able to check the extent to which our model is able to reproduce the features of aggregate consumption growth in the UK. Notice that the time series moments of this aggregate data are *not* used to calibrate our individual-level model. Thus this type of comparison to aggregated data constitutes an important test of our model's ability to fit the data, alongside comparisons to micro-data moments.

The next step of our analysis is to simulate behaviour under a set of counterfactual scenarios in which the mechanisms that might drive the link between house prices and consumption are shut down in turn. The construction of these counterfactuals is, by definition, an exercise that cannot be undertaken using historical data and constitutes our main tool to disentangle, and potentially to quantify, the influence of the different mechanisms that might drive the link between house price shocks and consumption growth. As explained in section

5, exploration of cross-sectional patterns of consumption growth for different groups, as well aggregate level modelling, is an important part of this analysis.

We have three main results that we report in Section 5. First, we show that our model is able to fit aggregate consumption growth remarkably well. As the parameters of the model are chosen to fit life-cycle features of the data, this first piece of evidence constitutes an important success of the model. Second, we provide a firmer theoretical grounding for reduced-form empirical analyses of the type conducted by AW and ABHL. We show that in a model with credit constraints and simultaneous housing and consumption choices, a house-price shock that drives consumption changes will lead to the biggest consumption responses from older groups. Third, and on the negative side, when we feed into our model the estimated historical shocks, the cross-sectional pattern of consumption responses the model generates is somewhat different from what we observe in the data: in the simulated data older households' consumption changes as much as that of young consumers, unlike in the data. The main reason for this counterfactual result is the size of the estimated house price shocks, to which older households react.

Our exercise is not without drawbacks. Some simplifications in the way we model the demand for housing and transaction costs are unavoidable. However, the main limitation is the fact that our exercise is *not* a General Equilibrium one. We take the process of house prices as given and derive the aggregate demand for housing (as well as consumption). In reality house prices are determined in equilibrium and would be affected by the level of housing demand (affected, say, by income shocks). While this is obviously very important, and could be the explanation of our third (negative) result, making house price endogenous goes beyond the scope of this paper and is left for future research.

## 2 Some evidence on house prices and the cross sectional distribution of consumption growth

Before discussing the exercise we perform and the model we build, we illustrate the main findings about the relationship between house prices and both the level and cross sectional distribution of consumption in the UK. In doing so, we reproduce and extend some of the main features of the ABHL exercise. We start by specifying a relationship aimed at capturing the life cycle evolution of consumption. In a standard setting, consumption will depend on life-time resources, which we summarize by some group effects (where groups are defined by year-of-birth cohort). In addition, we let consumption change with age and with a number of demographic variables. These variables are supposed to capture the evolution of needs as well as changes in the discount factor (induced, for instance, by changes in mortality probabilities). The baseline specification we adapt from ABHL is therefore:

$$\ln X_t^{gh} = \alpha^g + f(age) + \gamma' \boldsymbol{z}_t^{gh} + u_t^{gh} \tag{1}$$

where $g$ stands for cohort groups and $h$ stands for household; $X$ is non-housing expenditure, $\alpha^g$ is a group-specific intercept, $f(age)$ is a polynomial (quartic) in age, $\boldsymbol{z}_t^{gh}$ are a set of demographics, and $u_t^{gh}$ is an error term.

It should be stressed that versions of the life cycle model that give rise to this specification are extremely simple and unrealistic. The purpose of running this baseline consumption function is to identify which groups deviate most markedly from predicted consumption levels in times of consumption boom or consumption bust, after controlling for age, simple cohort effects and other observable variables. ABHL look at this question by estimating residuals from equation (1), averaging these residuals by year and age-group, and then plotting the resulting average residuals. We have updated this analysis using data from 1978 to 2005/6 (ABHL had data running until 2001/2), and the results are shown in figure 2.[3]

The figure shows that in times of large consumption booms or busts - particularly the late 1980s boom and the mid 1990s slump - it is the consumption of the young that deviates most from expected levels. These are also times of unusually strong house price growth. Overall, consumption growth of the youngest consumers seems to be more volatile: the standard deviation of the consumption

[3]Starting in 1978 implies that we miss the house price boom and bust of the early 1970s. Micro data from the FES are available on a consistent basis since 1974. However, they do not contain indicators about the education of the household head (or other household members) until 1978.

growth series plotted in Figure 2 is 0.042 for the youngest consumers, 0.026 for the middle aged ones and 0.032 for the oldest group. Furthermore, the correlation coefficient between consumption growth and house prices growth equals 0.341 for the youngest group, 0.292 for the middle and 0.243 for the oldest. The picture and these figures are a good indication that it is the consumption of the young that is most strongly associated with house price growth. One way in which we will assess how our model relates to the data will be to recreate figures similar to figure 2 using data generated by our simulated model.


Figure 2: Consumption function residuals, by age group

To look in more detail at the role that house price shocks have in explaining these patterns in residuals, ABHL extend the specification in equation (1) to control explicitly for the influence of (regional) house prices, $hp$ . They consider several specifications that incorporate house prices into equation (1) in different ways. In particular, they first consider the possibility that the level of consumption is directly affected by the level of house prices, as would be the case for a straightforward wealth effect. They then allow *innovations* to house prices to affect the level of consumption. These innovations are identified either as house price changes (as if house prices followed a random walk) or as residuals of a prediction equation that models house prices as a function of earnings and interest rates. In this last case, ABHL allow both the predicted and the 'surprise' on house prices to have an impact on consumption. Finally, for all three specifications, the coefficients on the relevant house prices variables are allowed to be different for different age groups.

The specifications estimated by ABHL are:

$$\ln X_t^{gh} = \alpha^g + f(age) + \gamma' \boldsymbol{z}_t^{gh} + \theta^y \ln hp + \theta^m \ln hp + \theta^o \ln hp + u_t^{gh} \quad (2)$$

$$\begin{aligned} \ln X_t^{gh} &= \alpha^g + f(age) + \gamma' \boldsymbol{z}_t^{gh} + \phi^y \ln hp^e + \phi^m \ln hp^e + \phi^o \ln hp^e + \\ &\quad \psi^y \ln hpu^u + \psi^m \ln hp^u + \psi^o \ln hp^u + u_t^{gh} \end{aligned} \quad (3)$$

$$\ln X_t^{gh} = \alpha^g + f(age) + \gamma' \boldsymbol{z}_t^{gh} + \lambda^y \Delta \ln hp + \lambda^m \Delta \ln hp + \lambda^o \Delta \ln hp + u_t^{gh} \quad (4)$$

where the superscripts $y, m, o$ refer to three age groups: young (aged less than 35), middle-aged (aged between 35 and 60) and old (aged over 60), and the superscripts $e$ and $u$ on log house prices ($\ln hp$) in equation (3) refer to 'expected' and 'unexpected' respectively.[4] We report the coefficients of interest, as estimated by ABHL, in Table 1.

The main message that comes out of these results is that, while house prices seem to be correlated with consumption even after one controls for life cycle factors (including age and cohort effects and family composition changes), there is no evidence that the consumption of older consumers is more sensitive to changes in house prices than that of younger consumers, as one might expect if a substantial wealth effect was operative. The *level* of house prices attracts effectively the same coefficient for the three age groups considered in equation (2). The same is true when we consider the *expected level* of house prices in equation (3). In this equation, however, unexpected shocks to house prices seem to be more important for younger consumers than for older ones. Indeed the effect of unexpected shocks to house prices declines with age. A similar pattern emerges when we consider equation (4), where we consider the *changes* in (log) house prices: the coefficient is largest for young consumers, it is halved for middle aged ones and is not significantly different from zero for older consumers.

[4]The full set of independent variables in the ABHL regressions is: a quintic in age of head of household; cohort dummies; occupation dummies; region and month dummies; controls for family composition and educational attainment dummies.

The empirical evidence, therefore, indicates that, in the recent past, it has been households headed by young individuals (those aged under 35) who have had the biggest deviations (on average) of consumption from expected levels in times of boom and bust. The young group are also the group for which swings in house prices are most able to explain such deviations of consumption from expected levels. This group contains individuals who are unlikely to hold large amounts of housing wealth. Moreover, this is also the group for which (due to the planning horizon) wealth would have the smallest immediate effect on consumption.

These results could be interpreted as suggesting that wealth effects have not been the main driver of the correlation between house prices and consumption. To confirm this intuition, however, we need to check it within the context of a well specified model calibrated to fit micro data. In such a model, changes in house prices can have complex and subtle effects on consumption and the intuition derived from a simple framework can be misleading. In what follows we build a coherent and realistic life cycle model that we calibrate to the UK context. In our model, there are two drivers of consumption behaviour: earnings and house prices. Moreover, the extent to which innovations to these exogenous shocks are reflected (or not) in consumption depends on the nature of financial markets to which individual households have access. We will model the intertemporal budget constraint so to reflect the main features of the UK mortgage markets. Given the demographic composition of the population, these individual responses can be simulated and aggregated to yield aggregate consumption both under estimated historical shocks and under counterfactual scenarios.

## 3 The model

As discussed in section 1, our main objective is the use of a realistic structural life-cycle model of consumption, savings and housing choices, to disentangle and distinguish the various mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. As such, our work is building on several recent contributions that have examined the relationship

Table 1: ABHL house price terms by age groups: predicted price specification

| | Estimated Coefficients | | |
|---|---|---|---|
| **House Price Terms** | *Young* | *Mid-age* | *Old* |
| *log house price: equation (2)* | | | |
| $ln(hp)$ | 0.161* | 0.163* | 0.165* |
| | (0.009) | (0.009) | (0.009) |
| *expected and unexpected log house price: equation (3)* | | | |
| $\ln(hp)^e$ | 0.291* | 0.292* | 0.294* |
| | (0.013) | 0.013) | 0.013) |
| $\ln(hp)^u$ | 0.188* | 0.088* | -0.012 |
| | (0.022) | (0.015) | (0.020) |
| *changes in log house price: equation (4)* | | | |
| $\Delta \ln(hp)^e$ | 0.209* | 0.127* | 0.042 |
| | (0.029) | (0.020) | (0.026) |

Notes: Number in parentheses are standard errors. * means significant at 5%. These results are taken from Attanasio et al. (2009).

between housing (or durables) and consumption in structural frameworks, including Campbell and Cocco (2007), whose exercise is similar to some aspects of ours, and Cocco (2005) who studies portfolio choices in the presence of housing. Many such studies have found that the nature of housing as consumption good, asset, source of collateral, and potential intergenerational heirloom, makes the relationship complex. Fernandez-Villaverde and Krueger (2005), for example, argue persuasively that a housing asset may have an important role in explaining the observed "hump shape" of profiles of consumption and durable consumption over the lifecycle. Flavin and Nakagawa (2004) discuss the fact that when the amount of housing can only be adjusted at cost, this may explain why both housing and non-durable consumption tend to be smooth over time apart from at infrequent periods of large adjustment. Similarly, the analysis of Li and Yao (2007) shows how a housing asset can result in consumption behaviour (including for housing consumption) that is very insensitive to income in some ranges, but very sensitive in others.

### 3.1 Non-technical description of the model

Given the nature of the exercise we perform in this paper, we could not use a standard off-the-shelf life-cycle model, as we wanted to incorporate the features that are potentially important to explain the correlation between consumption and house prices. In particular, we wanted a realistic description of housing so that it provides utility, it has an element of lumpiness, and can be used as a store of value. We also wanted to incorporate adjustment costs and, more importantly, a realistic description of the UK mortgage market. Finally, we wanted households to face both idiosyncratic and aggregate shocks. In this section, we describe the model in detail. Before presenting the formal specification of the model, we provide a verbal description of its main features and components.

Households in the model are heterogeneous ex-ante and ex-post in various dimensions. Households differ in terms of educational attainment of the household head (which determines differences in earning profiles and, possibly, in preferences) and because demographic composition varies with age (and this affects preferences). Both education status and demographic composition of the households are considered as exogenous.

The household decision maker in the model aims, given available resources, to maximize expected (discounted) lifetime utility choosing, in each period, the level of non-durable consumption and housing. Utility is additively separable over time and within each period is derived non-separably from consumption and from home ownership. Households also derive utility from leaving bequests, which helps us capture behaviour regarding the (non) sale of housing assets to fund expenditure late in the life-cycle.

Housing services are acquired either by renting or owning, with the latter being more attractive. Conditional on owning, the household can choose between a flat and a house. It is important to distinguish between owning and renting for several reasons. Young consumers, especially if they have low income, might be unable to afford to own. Older consumers have the possibility to cash in their housing wealth while still using housing services. It is important to distinguish between flats and houses so as to consider a 'housing ladder' and to have the possibility of increasing then decreasing (albeit in a coarse fashion) housing wealth over the life cycle.

House purchases and sales are subject to a proportional fixed cost, which constitutes an important friction in our model. In the simulations presented in this paper, the fixed cost of selling a home, and the utility from owning one's home, are both age dependent and increase steadily after age 55. This age-dependence was introduced in order to allow us to capture incomplete downsizing of housing assets during the last part of the life-cycle. It is difficult to capture this element of behaviour using only the bequest motive. Increasing the bequest motive alone to induce some households not to sell before the end of the life-cycle would result in counter-factually high ownership earlier in the lifetime. Capturing incomplete downsizing is not critical to our arguments, but helps to reinforce some of our conclusions regarding the relationship between house-price movements and consumption growth: we will be careful to describe where it is important.

Conditional on owning, the household might receive a 'moving shock' (maybe because of a work-related relocation decision) that forces it to sell the house or flat at the beginning of the period and incur the associated adjustment cost we describe below. This shock obviously makes owning relatively less attractive because it imposes some constraints on mobility.[5]

Households maximize expected utility subject to various credit market imperfections that limit the extent to which resources can be shifted over the life-cycle. The model allows for mortgage debt and, building on the models presented in Bottazzi et al. (2007) and in Attanasio et al. (2007), we are careful to model mortgage-related borrowing constraints as realistically as possible since these are likely to have first-order effects on the behavior we are modeling. As is typical in the UK, mortgage size is restricted through a down-payment constraint and a multiple-of-income constraint. These constraints apply only in periods

[5]Li and Yao (2007) have a similar moving shock, which helps to capture moves based on factors other than optimal choices regarding house size, in their model.

when households take out a new mortgage to buy a home, or when they remortgage their current property (increasing the value of mortgage debt held against the property in which they already live).

The simulations we present aggregate up the behavior of many individuals. The process generating income for each individual is made of two components, one idiosyncratic and one aggregate. The idiosyncratic component is itself made of a deterministic (hump-shaped) expected profile, and a stochastic element with shocks that have persistent (but not permanent) effects. However, it is the aggregate element of the income process, and its relationship to the stochastic process generating house prices (also aggregate), that is crucial to determine aggregate fluctuations in consumption growth. The presence of an aggregate component is our way of capturing correlation in income across individuals. A strong assumption we make is that individuals are able to distinguish between aggregate and idiosyncratic shocks.

An issue we considered in the design of our model is the definition of 'aggregate'. In particular, one could think of aggregate shocks as being common to all households in the UK or to a specific region. Moreover, for earnings, one could think of education or age specific shocks. We experimented with different definitions. For the results we report, house prices are common across the UK, while earnings shocks are age-group specific (although generated by a process with identical properties across age groups).

Since we want to capture how aggregate shocks feed in to aggregate consumption, we are careful to model aggregate income and house price uncertainty in a way that reflects the data. We are particularly careful to model correlation between the shocks in the two processes since house prices and incomes both tend to go up (down) when the economy is performing strongly (poorly), and this must be reflected in agents' expectations. To our knowledge, building these aggregate shocks and their empirical correlation into a model of the kind we construct has not been done previously.

As we mention below, the data indicate that the aggregate processes for both the aggregate element of income and for house prices are subject to permanent shocks. Moreover, the shocks to the two variables seem strongly positively correlated. Individuals in the model are assumed to be rational in the sense that they are fully aware of the process generating house prices, and of the separate elements of the income generating process.

We now describe the technical details of the model and its calibration to UK data.

## 3.2 Model details

The model is effectively made of two blocks: the household maximization problem, which includes the specification of preferences and the budget constraint, and the environmental factors that, ultimately, drive consumption: earnings and house prices. We describe the two blocks in turn. Further details on the model, its solution and its calibration, are given in the Web Appendix.

### 3.2.1 The household problem: preferences and budget constraints.

In our exercise we will simulate the behavior of many households who live for a finite and known number of years and maximize expected utility when faced with uncertainty about an exogenous earning process and house prices. As we discuss in detail below, households receive aggregate (house price and earnings) and idiosyncratic (earnings) shocks and choose consumption and housing. Heterogeneity in the level of education implies the calibration of different earnings processes for households with different levels of education. Markets are incomplete and are modeled to mimic some feature of the UK mortgage market.

A household lives $T = 59$ periods (ages 22-80). In every period $t \leq T$, the household maximizes expected utility by choosing consumption $c_t \in \mathbb{R}_+$, whether to live in rented accommodation (there is one type of rental property) or to own a flat or a house, with $h_t \in \{0, 1, 2\}$, and next period assets $A_{t+1}$. The household value function in period $t$ is given by:[6]

[6]The notation for this description of the model has only one index, $t$, that captures both time and age (for an individual born at a given date). In the simulations that produce the results reported from section 4, a key feature is that observed aggregate shocks are input for the date at which they hit each cohort.

$$
\begin{aligned}
V_t\left(A_t, h_{t-1}, P_t, Y_t, Z_t\right) \quad &= \max_{\{c_t, h_t\}} \{u_t\left(c_t, h_t\right) + \\
&+ \beta[(1-\eta) E V_{t+1}\left(A_{t+1}, h_t, P_{t+1}, Y_{t+1}, Z_{t+1}\right) + \\
&+ \eta E V_{t+1}\left(A_{t+1}, 0, P_{t+1}, Y_{t+1}, Z_{t+1}\right)]\}
\end{aligned}
$$

subject to

$$
A_{t+1} = \begin{cases}
\begin{array}{r}
R_{t+1}[A_t + W_t - c_t - \text{rent}_t I(0) - \\
-\kappa P_t(1+F_t) I(1) - P_t(1+F_t) I(2)] + \\
+D^m_{t+1}\{\kappa P_{t+1}(1-F_t) I(1) + P_{t+1}(1-F_t) I(2)\} \\
\text{if } h_{t-1} = 0 \\
\\
R_{t+1}[A_t + W_t - c_t - \text{rent}_t I(0) + \\
+\kappa P_t(1-F_t) I(\neg 1) - P_t(1+F_t) I(2)] + \\
+D^m_{t+1}\{\kappa P_{t+1}(1-F_t) I(1) + P_{t+1}(1-F_t) I(2)\} \\
\text{if } h_{t-1} = 1 \\
\\
R_{t+1}[A_t + W_t - c_t - \text{rent}_t I(0) - \\
-\kappa P_t(1+F_t) I(1) + P_t(1-F_t) I(\neg 2)] + \\
+D^m_{t+1}\{\kappa P_{t+1}(1-F_t) I(1) + P_{t+1}(1-F_t) I(2)\} \\
\text{if } h_{t-1} = 2
\end{array}
\end{cases} \tag{5}
$$

$$
W_t = D_t Y_t Z_t \tag{6}
$$

where $A_t$ is the start of period asset stock and $R_{t+1} = 1 + r_{t+1}$ where $r_{t+1}$ is the (real) interest rate on the liquid asset; $P_t$ is the price of housing which is realized at the start of period $t$; $F_t$ is the cost of selling or buying a house, which is proportional to the price and may change with age reflecting changes in adjustment costs; $W_t$ is household income in period $t$; $D_t$ and $Y_t$ are both idiosyncratic components of income of which the former is deterministic and the latter stochastic; $Z_t$ is the aggregate component of income; $D^m_t$ is a dummy variable recording an exogenous shock, realized at the very beginning of the period, that forces a family to move home with probability $\eta$; $I(x)$ is an indicator function that takes value one if the housing choice ($h_t$) takes value $x$ and $I(\neg x)$ takes value one whenever $h_t$ is not equal to $x$; $\text{rent}_t$ is the rental price paid by those who do not own their home at period $t$; and $\kappa$ is the ratio of flat to house prices.

It is useful to interpret the intertemporal budget constraint (equation 5) term by term. Financial assets at the end of period $t$, gross of interest, become assets held at the beginning of $t+1$. The end of period $t$ assets are assets held at the beginning of that period (after any incomes due to the exogenous moving shock), net of all incomes and expenditures during the period. That is, they are (taking each term on the RHS of (each version of) equation (5) in turn): beginning of period financial assets; plus labour income; less consumption; less rent if paid; less the cost of buying, or plus the revenue from selling, a flat, if applicable; less the cost of buying (or plus the revenue from selling) a house, if such a transaction is chosen. The final terms in the expression for beginning of period assets are revenues from a house or flat sale at the very beginning of period $t+1$ if a family that owns a flat or a house at the beginning of the period is forced through the exogenous shock to sell it. Revenues received due to such a forced sale occur after interest payments (and price realisations), and so this wealth does not accrue interest.

Notice that the rental price and the price of flats, which appear in the budget constraint, do not appear as additional state variables upon which the consumer's value function depends. This is because flats are assumed to be worth a constant fraction $\kappa$ of the value of houses and the level of rent is given by the lesser of a fraction ($\iota$) of the expected house price in the relevant period,[7] and half of total income:

$$rent_t = \min\{\iota \overline{P_t}, 0.5W_t\} \tag{7}$$

We may think of the capping of the rent at half of total income as a stylized representation of a housing benefit system, and it means that renting can never be unaffordable (implying zero resources for nondurable consumption), even for

[7] In fact the process for the house price is non-stationary (it is a random walk), so the "expected price" in a given year is not well defined. The price that we use is that given by supposing that prices were at trend level in 1971 (when the house-price cycle was mid way between a bust and a boom), and then applying the expected rate of growth in prices. We experimented with setting the rent equal to a proportion of the actual house-price, and results were very similar.

those with very low incomes.

We only allow for collateralized debt, i.e. households are only able to have negative financial assets when they are home owners, so that when they do not own a house ($h_t = 0$) they are subject to the constraint

$$A_t \geq 0. \tag{8}$$

Home owners can borrow, and when they do so they are subject both to a terminal asset condition that translates into an implicit borrowing constraint, and to two explicit borrowing constraints. In particular, we impose the terminal condition[8]

$$A_{T+1} + \kappa P_T(1 - F_{T+1})I(1) + \kappa P_T(1 - F_{T+1})I(2) \geq 0. \tag{9}$$

This constraint ensures non-negative total wealth at the end of life (the indicator functions are for ownership in period $T$). The specification of marginal utility becoming infinite at zero consumption means this terminal condition prevents households borrowing more than they can repay with certainty.

In addition to this implicit borrowing constraint, we model two explicit constraints. The first is a function of the value of the house and the second is a function of household annual income. They determine how much a household is able to borrow at the time of purchase or when remortgaging, and translate into the following constraints in the period after the new mortgage is agreed:

$$A_{t+1} \geq -\lambda_h \varkappa P_t(1 + r), \quad \varkappa = \begin{cases} \kappa & \text{if } h_t = 1 \\ 1 & \text{if } h_t = 2 \end{cases} \tag{10}$$

The value $(1 - \lambda_h)$ can be thought of as a downpayment requirement.

$$A_{t+1} \geq -\lambda_w W_t(1 + r) \tag{11}$$

The explicit constraints on the downpayment and the debt-to-income ratio only apply when households buy the property or remortgage. Formulating the constraints in this way complicates the model solution. In particular, it makes

[8]In the simulations the fixed cost of selling increases with age after age 55, but the fixed cost $F_{T+1}$ if the home is sold by the late householder's estate (immediately) after the death of the homeowner, remains at the level for those aged below 55.

the model more complicated to solve numerically than, for example, having only a downpayment constraint that must be satisfied every period (as, for example, in Campbell and Cocco (2007), Li and Yao (2007), Cocco (2005), Campbell and Hercowitz (2004)). However it seems important to us to capture the institutional features of the UK mortgage market, since these are likely to affect how house price shocks feed into consumption. A direct implication of applying the constraints only in periods of buying or remortgaging is that when a household continues owning without remortgaging, they can keep their existing debt even if the value of the debt would exceed what would be permitted if the house were bought in that period. Although there is no mortgage repayment schedule, the household does have to pay off mortgage interest each year in which it does not remortgage.

To complete the description of the maximization problem, we must specify the within period utility function and the function describing how utility is derived from bequests. Households get utility from consumption and from home ownership. The within-period utility function is CRRA in equivalized non-durable consumption, augmented by a term reflecting the value of home ownership:

$$u(c_t, h_t) = \exp(\theta_t \phi h_t) \frac{(c_t/eq_t)^{1-\gamma}}{1-\gamma} \qquad \begin{cases} \theta_t,\ \phi \in \mathbb{R}\backslash\{0\} & \text{if } h_t = 0 \\ \theta_t \in \mathbb{R},\quad 0 < \phi < 1 & \text{if } h_t = 1 \\ \theta_t \in \mathbb{R},\quad \phi = 1 & \text{if } h_t = 2 \end{cases} \qquad (12)$$

The parameter $\theta_t$ is a housing preference parameter which determines the utility that households obtain from owning a house rather renting it. Notice that we allow it to be time varying. When $\gamma$ is greater than one, home ownership raises utility if $\theta$ is negative. In that case (our case), housing and non-durable consumption are substitutes in utility in the sense that a higher level of $h_t$ reduces the marginal utility of a given level of non-durable consumption. The parameter $\phi$ determines the relative utility from owning a flat versus a house.[9] The age-specific value $eq_t$ is an equivalence scale used to convert consumption into per adult equivalent terms. We use a simple equivalence scale so that a couple has

[9]Similar (though not age dependent) forms for preferences regarding housing are used by Ejarque and Leth-Petersen (2008) and Rios-Rull and Sanchez-Marcos (2008).

consumption needs of 1 and each child adds 0.35. Family size is assumed to be the same for all households, and reflects the average number of children of couples at different ages in UK data.[10]

The utility from leaving a bequest is described by a second iso-elastic function:

$$b(\omega) = \tau * \frac{(\omega/\tau)^{1-\gamma}}{1-\gamma} \tag{13}$$

Where $\omega$ is the value of wealth (both financial and housing wealth, net of the fixed cost of selling a property) left over at the end of life $T$, after all shocks to resources (income and the house price) have been realized, and all consumption decisions have been made.[11] The parameter $\tau$ is calibrated. A bequest motive, although not central to our analysis, is crucial to match certain features of the life cycle profile of home ownership, particularly at later ages.

Even though the rental price is not a state variable, the number of state variables in the problem (with four continuous states - assets, idiosyncratic and aggregate income and house prices - plus current home ownership and time) means it is computationally demanding to solve.

### 3.2.2 The environment: exogenous stochastic processes

Households face three dimensions of uncertainty: shocks to house prices, which are aggregate (i.e. common across all properties in the economy); aggregate shocks to income; and idiosyncratic shocks to income. In the present version of the model, the interest rate on liquid assets and debt is fixed.[12]

If we take the income generating process first, this may be thought of as being composed of three parts:

$$\ln W_t = d_t + y_t + z_t \tag{14}$$

[10]Data come from the 2002 British Household Panel Survey. The average is smoothed using a quartic in parent's age and all children are assumed to leave home by the time parents are 60. This set up and polynomial are the same as in Attanasio and Wakefield (2010) for their results concerning family.

[11]For more motivation of the modelling of bequests, see, for example, De Nardi (2004).

[12]It is a relatively simple extension to add i.i.d shocks to the interest rate. Experiments indicated that such a change has no qualitative impact on our results. Considering persistent interest rates is challenging because it adds a state variable to the problem.

where lower case has been used for logs, and $d_t$ is a deterministic part to the income generating process, $y_t$ is a persistent idiosyncratic stochastic element and $z_t$ is the aggregate stochastic component. The household is assumed to observe these components separately, which explains why $Y_t$ and $Z_t$ enter the value function as two separate state variables.

The deterministic component $d_t$ is hump-shaped over the working lifetime and is captured using a polynomial in age. The coefficients of this polynomial are calibrated for the two education groups we consider separately. In practice, to fit the observed data, and particularly differences between education groups around the age of fertility, we need a cubic specification (see Appendix B).

The idiosyncratic stochastic component $y_t$ is modelled as an AR(1) process[13]:

$$y_t = \rho_y y_{t-1} + \xi_t, \quad \xi_t \sim N\left(0, \sigma_\xi^2\right) \tag{15}$$

The stochastic aggregate component to income is modelled jointly with the stochastic house price price using a first-order vector auto-regression with correlated innovations. When we used data to estimate this process, we could not reject the hypothesis that lagged income does not affect current house prices and similarly that lagged house prices do not explain current income. However, we find a significant and sizeable correlation between house price and earnings innovations. As for the effects of own lags, we could not reject that these coefficients were equal to unity. We therefore impose this value.

If we let HP stand for house prices and again use lower case letters to represent logs, this can be written as:[14]

[13]In implementing the model numerically, we use finite state Markov chains that mimic underlying continuous-valued autoregressive or vector autoregressive processes. This is done as described in Tauchen (1986), since the analysis of Floden (2007) suggests that this technique is relatively robust to high persistence in the process modelled. Very recent analysis by Kopecky and Suen (2010), drawn to our attention since we completed our modelling work, indicates that the method proposed in Rouwenhorst (1995) may have further improved our approximation. Fuller details of the numerical solution and simulation of our model are provided in the online appendix to the paper.

[14]Given the unit persistence in this equation, the constant terms $\alpha_0^z$ and $\alpha_0^h$ capture drift over time even though there is no time trend in equation 16 - see Davidson and Mackinnon (2004), pp.606f.

$$\begin{bmatrix} z_t \\ hp_t \end{bmatrix} = \begin{bmatrix} \alpha_0^z \\ \alpha_0^h \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} z_{t-1} \\ hp_{t-1} \end{bmatrix} + \begin{bmatrix} u_t^z \\ u_t^h \end{bmatrix} \tag{16}$$

The value of modelling these aggregate processes jointly as a VAR, rather than as separate autoregressions, is that the joint distribution of the error terms will capture correlation between aggregate shocks to house prices and incomes. Capturing this correlation is valuable since it will affect the degree to which individuals will choose to modify their asset accumulation as a means of self insurance against shocks. The joint distribution of the shocks is assumed to be normal:

$$\underline{u}_t \sim \mathbb{N}(0, \underline{\Omega}) \tag{17}$$

where

$$\underline{\Omega} = \begin{bmatrix} \sigma_z^2 & \pi\sigma_z\sigma_h \\ \pi\sigma_h\sigma_z & \sigma_h^2 \end{bmatrix}$$

with $\pi$ measuring the correlation between the shocks.

# 4 Calibration, estimation and simulation

The parameters of our model can be divided into three categories: those we take from other studies, those we estimate outside the model and those we calibrate to fit some moments of the micro data we consider. We estimate the parameters of the exogenous stochastic processes faced by the agents in our model using time-series data. Our calibration exercise involves matching moments for life cycle levels of home ownership status in the UK for those aged 26-70 in 1990-2006. Since we simulate cohorts born between the 1910s and 1970s, we match the ownership levels for the cohorts in the relevant age range in the 1990s and early 2000s.

## 4.1 Parameters fixed or estimated outside of the model

For inputs into the calibrated model, we need to use data on earnings, the house price process and the interest rate on liquid assets. Values for parameters fixed or estimated outside the model are summarized in Table 2.

Table 2: Estimated / Fixed Parameters

| Parameter | Value | | Source |
|---|---|---|---|
| | Utility Parameters | | |
| $\gamma$ | 1.43 | | (Attanasio and Weber, 1995) |
| | High Edu | Low edu | |
| $\beta$ | $1.04^{-1}$ | $1.045^{-1}$ | |
| | Aggregate House Price and Income Process | | |
| $\alpha_0^z$ | 1.66% | | FES |
| $\sigma_z$ | 0.033 | | FES |
| $\alpha_0^h$ | 3.58% | | DCLG |
| $\sigma_h$ | 0.091 | | DCLG |
| $\tau$ | 0.645 | | FES / DCLG |
| $\kappa$ | 0.6 | | BHPS |
| $\iota$ | 0.006 | | FES |
| $\eta$ | 0.01 | | |
| | Idiosyncratic Income Process | | |
| Deterministic component: cubic in age | | | BHPS |
| | High Edu | Low edu | |
| $\rho_y$ | 0.76 | 0.77 | BHPS |
| $\sigma_\xi$ | 0.39 | 0.41 | BHPS |
| $\frac{MedianP_{22}}{MedianY_{22}}$ | 3.3 | 4.4 | BHPS |
| | Credit market Institutions | | |
| $\lambda_y$ | 3.0 | | |
| $\lambda_h$ | 0.9 | | |
| $\bar{r}$ | 0.03 | | BoE. |

Notes: FES indicates that the source is Family Expenditure Survey / Expenditure and Food Survey data, BHPS means British Household Panel Survey, DCLG indicates the use of data published by the Department for Communities and Local Government and BoE is Bank of England.

### 4.1.1 Utility function

The preference parameter $\gamma$ in the utility function is set to match the consumption elasticity of intertemporal substitution of 0.7 found in data (see Attanasio and Weber (1995)). This corresponds to a curvature $\gamma = 1.43$ for our within-period utility function. We make households impatient in the sense that we set the discount rate to be slightly higher than expected return on any asset (including the housing asset). We also found that it was hard to match differences between low and high education groups if the discount rates were equal, and so (in line with Attanasio et al. (1999)) we set this parameter slightly higher for the low educated. The discount rates for high and low education are, respectively, 4% and 4.5%.

#### 4.1.2 Aggregate processes: house prices and aggregate earnings

As described by equations (16) and (17), the aggregate house price and income processes are specified in logs as a vector autoregression of order one (a VAR(1)), with drift. To estimate this process we use data on the national mix-adjusted house price series for the UK[15], which are the same data underlying figure 1 and are constructed using series freely available from the Department for Communities and Local Government (DCLG). Income data is taken from the UK Family Expenditure Survey / Expenditure and Food Survey (FES). Since our model is set up in terms of real prices, we deflate both series to prices in the latest year using the all-items Retail Prices Index, and we use data for the years 1969-2005. As anticipated in section 3.2.2, estimation of this process does not reject unit persistence, and so we impose that aggregate shocks are permanent as in equation (16). The other parameters returned from this estimation are trend growth rates for the house prices (3.58%) and income (1.66%), and standard deviations of the shocks to the processes of 0.091 and 0.033 respectively for the house price and income.[16] The correlation of the shocks for the two processes is approximately 0.65, and our model assumes that people know this correlation and build it in to their expectations.

The level of rents in the model is set as a fraction of the house price. Its level of 0.6% (or 1% of the flat price) is chosen to match the statistic from our FES data that rents account for approximately 15% of the income of renters, on average.

[15] We use the series for average house prices for all dwellings.

[16] These are standard deviations of residuals and need not match the standard deviations in growth rates that were mentioned in the first paragraph of this article.

#### 4.1.3 Idiosyncratic earning process

On top of the aggregate income process, there are two further elements to the household income generating process (see equation (14)). These are a deterministic process, which captures the hump-shaped profile of household incomes over the working life and a 50% replacement rate in retirement, plus a persistent stochastic component which is described by a first order autoregression (equation (15)). The deterministic component plus the process generating shocks to the stochastic component are both education group specific, and the realized shocks to the stochastic component are idiosyncratic at the household level.

Parameters for both processes are estimated simultaneously. Since this involves estimation of a dynamic process at the household level, we require panel data and so use the British Household Panel Survey (BHPS) for the years 1991-2005. The estimation suggests that the deterministic component can be approximated as a cubic that shows a hump-shape over the working life which is slightly more pronounced (in particular with a steeper slope at the beginning of the working life) for the college educated group than for those with only compulsory level education. The process generating stochastic innovations is quite similar across groups with a persistence parameter of slightly more than 0.75 and a variance of the shock of around 0.16. Further details are available in the Web Appendix.

#### 4.1.4 Credit market institutions

The parameters that determine the fraction of the house price ($\lambda_h$) and the multiple of earnings ($\lambda_y$) that households can borrow are chosen to match the UK institutional features. At the time of taking out a new mortgage (i.e. of buying or remortgaging) households can borrow whichever amount is lower between three times household earnings ($\lambda_y = 3$) and 90% of the house price ($\lambda_h = 0.9$). The interest rate on the liquid asset / debt is taken to be fixed and is set to match the average of 3% for the real interest rate on 90 day Treasury bills for the period since 1990 (i.e. the period for our calibration).

## 4.2 Calibrated parameters

We select the preference parameters for housing, $\theta_t$ and $\phi$ plus the parameter $\tau$ that determines the bequest motive, and the fixed cost of housing, $F_t$, by matching average life-cycle home-ownership rates between ages 26 and 35, 36 and 55, and 56-70, for our two education groups, as well as the overall proportion of house (rather than flat) owners. We assume that these parameters are common across the two education groups.

Table 3: Calibrated Parameters

| **Parameter** | **Value** |
|---|---|
| $\theta$ | -0.04 |
| $\phi$ | 0.25 |
| $\tau$ | 4 |
| $F_b = F_s$ | 0.04 |

Table 4: Calibration Statistics

| | **High Education** | | **Low Education** | |
|---|---|---|---|---|
| **Statistic** | **Data** | **Model** | **Data** | **Model** |
| *Ownership rate* | | | | |
| Age 26 - 35 | 0.559 | 0.579 | 0.473 | 0.471 |
| Age 36 - 55 | 0.787 | 0.814 | 0.623 | 0.681 |
| Age 56 - 70 | 0.821 | 0.830 | 0.632 | 0.607 |
| Owners w House | 0.582 | | | 0.601 |
| **Sum Abs. Devs** | | | | **0.161** |

Notes: The data figures for home-ownership rates are based on the years 1990-2006 of the FES.

As mentioned in sections 3.1 and 3.2.1, in the simulations we present here the parameter $\theta_t$, and the fixed cost of selling $F_t$, are both kept constant up to age 55 and increase steadily thereafter. In table 3 we report the values that apply until age 55. Our parameter values imply that owning a house raises utility by approximately 4% and owning a flat raises utility by approximately 1%, and the fixed cost of buying and selling is 4% of the property price. The parameter $\tau$ is set at a value of 4. From age 56 the fixed cost of selling, and the parameter $\theta_t$ both increase (in magnitude) by a constant amount each year to reach respective values of 40% and -0.4, by the final year of life (i.e. both parameters show a tenfold increase).

Table 4 compares home-ownership rates predicted by the model to those observed in the data for different education and age groups[17]; we do a good job of matching the moments of interest, although the home-ownership profile for the low-educated is slightly too hump-shaped.

Given all the fixed parameters of the model (notably the discount rate), the parameter $\theta_t$ is the dominant factor for home-ownership in middle-age (36-55). The parameter $\phi$ then largely determines the proportion of house-owners and, together with the fixed cost parameter, also influences the rate at which those aged 35 and under get into the housing market.[18] The bequest parameter $\tau$ has some influence on the rate at which families leave the housing market after age 55. When we add the consideration of increasing utility from ownership and fixed costs of selling after age 55, these largely offset each other so that relatively modest changes to the other parameters are needed to maintain a good calibration. The exception is that as the amount of downsizing late in life falls, this tends to be associated with increased holding of houses rather than flats.[19] Thus the simulations we present have a rather higher relative utility from flat ownership than is the case when $\theta_t$ and the fixed cost of selling are not age-dependent.

### 4.3 Simulations

With our model, we simulate the behaviour of several cohorts designed to resemble cohorts of UK consumers. The model is calibrated to reproduce some facts about the home-ownership experience of these cohorts. In our simulations, we aggregate the behaviour of many individuals that comprise the cohorts we consider and compare the behaviour of aggregate consumption growth generated by our model with actual consumption data. In the model, alongside idiosyn-

[17]Data come from the years 1990-2006, as years prior to 1990 are affected by the large-scale selling off of local authority housing.

[18]This is a simplification: fixed costs and the relative utility from flat ownership also influence ownership levels after age 35, but the effects are second order relative to the influence of $\theta_t$.

[19]The longer a property is held, the more valuable is the bigger utility boost provided by holding a house.

cratic shocks, individual households face some aggregate stochastic processes - the process for house prices, and a component of the income generating process - which we estimate from aggregate data. In addition to the parameters of the processes generating shocks to these aggregate variables, we also construct estimates of the path of aggregate shocks that have been experienced in the UK economy since the beginning of the 1970s. We use this path as an input into our simulations, including those for the calibration described above.

More precisely, we take the data on actual growth rates in house prices and aggregate earnings in the economy each year, and input these into our simulations at the correct age for the particular cohort that we are simulating. We repeat this process for each 5-year cohort born between 1915-19 and 1975-79, and create simulated data for a large number of different households (i.e. realizations of the idiosyncratic shocks) in each cohort given the house prices and the aggregate income that they actually faced. These simulated data are the basis for our comparisons to data on the UK economy, including to household survey data which incorporate the same set of cohorts.

It is important to stress that our simulations take as an input the actual path of prices and incomes in the UK economy. Since the relationships between choices, prices and income in our model are not linear, simulating for a specific path of prices and aggregate income as we have done gives different results from simulating behaviour on average (across different levels of the house price and aggregate income) for particular shocks. Our simulations based on actual shocks to house prices and incomes also form the "baseline" against which we can compare counterfactual scenarios in which the mechanisms that might drive the link between house prices and consumption are shut down in turn. These scenarios and how they compare to our calibrated baseline are the subject of the next section.

# 5 Booms and busts in house prices and consumption

In this section, which contains the main results of our analysis, we use the calibrated life-cycle model we have built and perform three types of exercise. First, we check whether the model 'fits the aggregate data'. Second, we quantify, given the observed innovations over the period we study, the relative importance of house price shocks and earning shocks in generating fluctuations in consumption growth (in the simulations). Third, we use the model to check and quantify the intuition that young consumers would be more sensitive to earning innovations, while older consumers should be more sensitive to house price shocks. We also consider how patterns of consumption growth vary with age when house price or income shocks are more important.

### 5.1 Model fit

In calibrating our model, we did not match any time series moment of consumption or its covariance with aggregate variables. Instead, the parameters of the utility function were chosen either from pre-existing studies or by matching facts such as the shape of the profile of home ownership across different ages. As a consequence, the comparison of the properties of aggregate consumption as generated by our model to the actual evolution of consumption is a meaningful and powerful test of the ability of the model to fit the data.

The way we construct our simulated data is the following. Having estimated the aggregate time-series processes for house prices and aggregate earnings, we estimate the actual shocks to these two time series over the relevant period and feed them into our simulations. The consumers in our model also experience idiosyncratic shocks, which are assumed to be uncorrelated to both house prices and aggregate earning shocks. We simulate a large number of idiosyncratic earning shocks paths and, therefore, simulate a large number consumers, to match the distribution of ages in the UK. Having done that, we can then aggregate these simulations and construct aggregate consumption.

The implementation of this strategy requires making certain assumptions and dealing with a number of technical issues that are discussed in the Web Appendix. Here we only want to stress that to the best of our knowledge, at least in the consumption literature, this type of exercise (to aggregate up simulations of individual behavior based on observed aggregate shocks, in order to check whether the model of individual choices is consistent with aggregate evidence) has not been performed. In the present context, the exercise we perform is particularly useful, since the question we wish to address concerns the evolution of aggregate consumption.

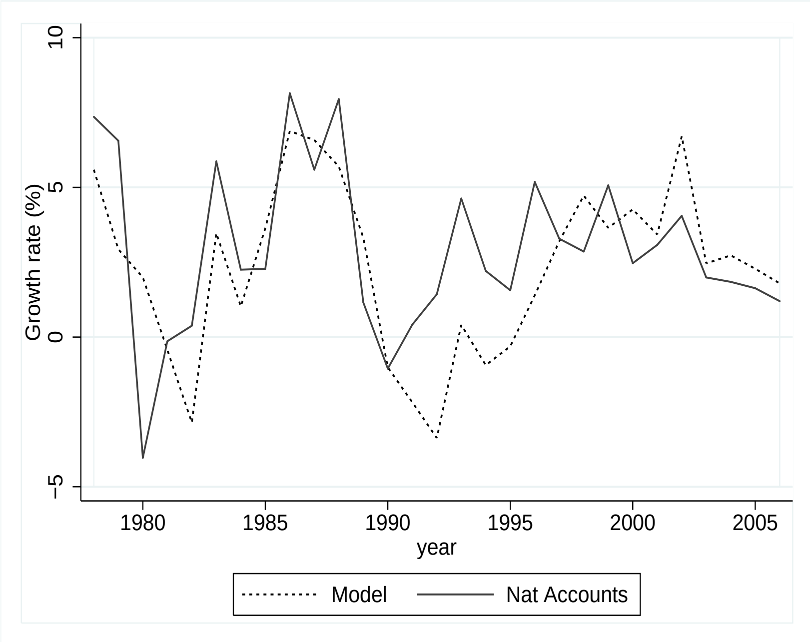


Figure 3: Growth in real consumption: Data versus model

Figure 3 reports the results of this exercise and plots, for the period from 1975 to 2006, the time series for aggregate consumption growth in the UK (the same series in Figure 1) and aggregate consumption growth as generated by the simulation of our model with the estimated aggregate earnings and house prices shocks. The model fit is remarkable: the correlation coefficient between the two series plotted is 0.62.[20] Moreover, the model does a good job at matching the timing and variability of the consumption booms and busts. This is particularly true for the boom of the late 1980s. Given that the parameters of the model were not chosen to match the time series properties of aggregate consumption growth, this exercise can effectively be interpreted as an out of sample prediction of the model and the figure shows that the predictive performance of our model

[20] A key innovation in our model is the addition of the housing asset so as a benchmark we can compare the model fit to the case with only income shocks: in that case the correlation is 0.45.

is extremely good.

## 5.2 What is driving consumption growth in the model?

Having built a model that has the ability to fit aggregate dynamics, we can use it to assess the relative importance of its two driving forces for the dynamics of consumption growth. This can be useful to understand why house price and consumption growth are strongly correlated in the data as well as the simulations. In particular, we can perform counterfactual experiments in which aggregate shocks to house prices and to incomes are switched off in turn and compare the resulting aggregate consumption to the baseline simulations plotted in Figure 3. This method of decomposing the drivers of fluctuations is very much related to that outlined by Chari et al. (2007) in a business cycles context.[21]It is important to stress that by switching off a given shock in the counterfactual simulations, we mean feeding into the model a specific realization where the shock is assumed to take its expected value of zero. Households, however, behave as if the shock was operational and, in forming expectations about the future, take into account the possibility that all shocks can be operative.

In Figure 4 we plot the results of four different simulations. The first series, plotted as a dashed line, is the the baseline simulation already plotted in Figure 3. We saw that this simulation is very similar to the actual data on consumption growth in the UK over this period. In the second simulation, we set house-price shocks to zero and let income shocks be the only driver of aggregate consumption fluctuations. The resulting aggregate consumption growth is plotted as a dotted

[21]Those authors emphasize that in computing the decomposition, the analyst wants to switch off shocks to variables in turn, but not to eliminate any effect of the level of other aggregate processes on expectations of shocks to the process whose influence is to be measured: one wants to isolate the effect of shocks each period, given rational expectations based on the observed state of the economy. In many cases this involves switching off the shocks, but having expectations based on the (multivariate) aggregate state of the economy that was actually observed. In our case with unit persistence (and zero cross terms) in the process generating aggregate states, the level of the house price or of aggregate income does not affect expectations about how the other aggregate state variable will move, and so our simulations do not need to keep track of an aggregate state measuring “where the economy would have been”. Thanks to Roman Sustek (Bank of England) for pointing out the parallel with this paper.


Figure 4: Model consumption growth: different scenarios

line. The series tracks the baseline simulation relatively closely in early years of the period modeled, with bigger deviations between the two series from the end of the 1980s. The standard deviation of consumption growth in this series is 0.015 (or 1.5 percentage points), compared to 0.028 in the baseline simulation. The correlation between this series and the one from the baseline simulation is 0.79.

In the third set of simulations, we set to zero the aggregate income shocks and let house prices be the only aggregate driver of consumption. Aggregate consumption for this simulation is plotted in the picture as a pale grey continuous line. This series seems to be more nearly related to the baseline simulation in the period after the late 1980s, compared to in the early part of the sample. For this simulation the variability of aggregate consumption is reduced to 0.019 and the correlation coefficient of this series with the baseline simulation is now 0.84.

Finally, in the last simulation, plotted as a thin continuous line we set both house price and aggregate income shocks to zero. In this case, consumption grows annually at the constant rate of approximately 1.7%.

Comparing across simulations, we notice that in the late 1970s and early

1980s it is clear that income shocks are driving much of the movement in consumption growth: the growth rates when only house price shocks hit are typically closer to the constant level whilst growth rates when only income shocks hit are closer to those in the baseline (i.e. full shocks) economy. However, by the latter 1980s, the effect of changes in house prices on consumption growth becomes more substantial. By the house price boom of the late 1980s, within our model, a larger fraction of total consumption growth is driven by house prices than by income shocks; similarly the house price bust of the early 1990s drives almost all of the negative consumption growth rates observed. Since then, house price shocks appear to have been the predominant factor driving aggregate consumption growth in the model. The relatively modest influence of income shocks in the later years modelled is due to the fact that aggregate income shocks, set to match the aggregate historical data, have become rather weak since the late 1990s, compared to the volatile years of the 1970s and early 1980s.

One way to quantify "how much" of consumption growth, *within our model*, is being driven by house price or income shocks, is to calculate the proportion of the difference between consumption growth in the baseline and no shock scenarios which is accounted for by the deviation from baseline when the relevant shock is turned off. Doing this over the whole period shown in figure 4, we find that house price shocks explain about 50% of the deviation, while income shocks explain about 36%.[22] This difference is stronger after 1990 (for which period the proportions are 54% and 30% respectively), but much smallerl (47% and 43%) for the period 1978 to 1990. To check whether this feature of the model (and the importance of house price shocks in generating consumption fluctuations) reflects something that is in the data or not, we look at the cross sectional implications of our model and how they compare to what we observe in the data.

[22]Though the income and house price shocks are the only aggregate shocks in the model, these two factors might also have a joint effect on consumption, and so there is no reason why these to numbers should sum to 100%.

## 5.3 How do the shocks affect different age groups?

The analysis above looked at aggregate trends in consumption growth under different versions of the model where house price and income shocks were present or absent. We now turn to a more disaggregate analysis looking at the results for different age groups. This analysis will allow us to examine how the response to shocks varies for households at different points in the life cycle in our model and compare it to the evidence in the data.

To do this, we use the kind of regression analysis described in Section 2 and that was used to construct Figure 2, and generate similar figures based on model data. In particular, we plot residuals from a consumption function, averaged at each point in time, for each of three age groups. The estimated consumption function controls for variation in demographic factors and skill levels.[23]

The bottom panel of Figure 5 shows such averaged residuals for simulated data from the baseline run of the model that featured both aggregate-income and house-price shocks. The top panel shows the corresponding picture obtained from the data. This plot is equivalent to that in figure 2, except that we are now smoothing the data by computing a three-year moving average.

The main difference between the model and actual data is that the latter shows the young group experiencing the biggest swings in deviations from expected consumption. In the model data, instead, we see that the deviations for different age groups in general move together. During the boom of the late 1980s the middle-age group seems to be moving faster, while during the boom of the mid 2000s it is actually the oldest whose consumption grows faster. In the case of the bust of the early 1990s, on the other hand, the oldest group present the mildest reduction in consumption. Moreover, the fluctuations generated by the model are considerably larger than those observed in the data (though this is partly due to the smoothing of the data). Overall, however, the figure seems to indicate a failure of the model to generate some important features of the data: namely that the consumption of the youngest cohorts seems to be much more

[23]In the model we do not have all the demographic variables that can be measured in data, and the the regressors are a subset of those mentioned in footnote 4. Specifically, we include a polynomial in age, plus measures of education and cohort.

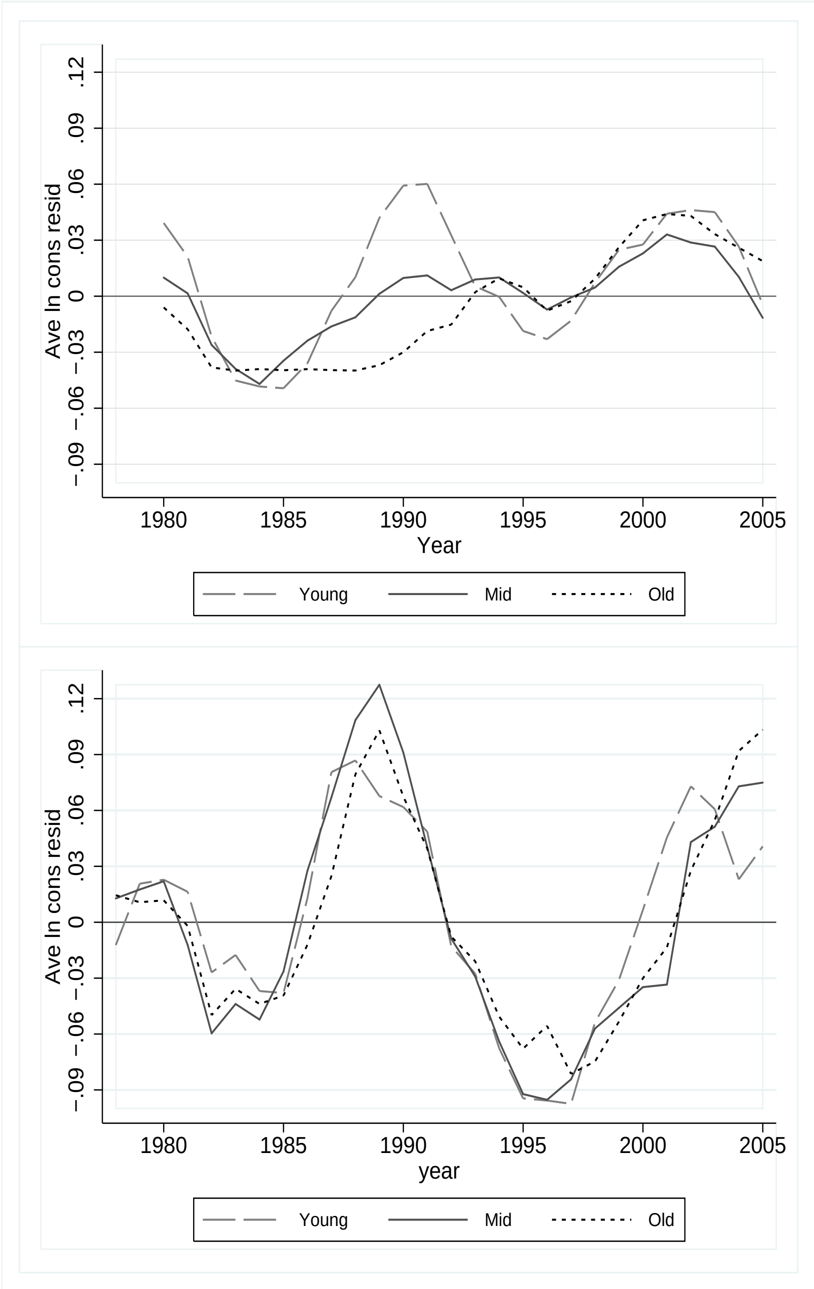


Figure 5: Residuals for log consumption: Data (top) and Baseline Model Results

variable than that of the oldest cohorts over recent episodes; and the magnitude of the fluctuations.

Before discussing this evidence further, it is worth verifying how the behavior of the various age groups differs in the same counterfactual exercises we considered in Section 5.2, namely when we shut off in turn income shocks and house

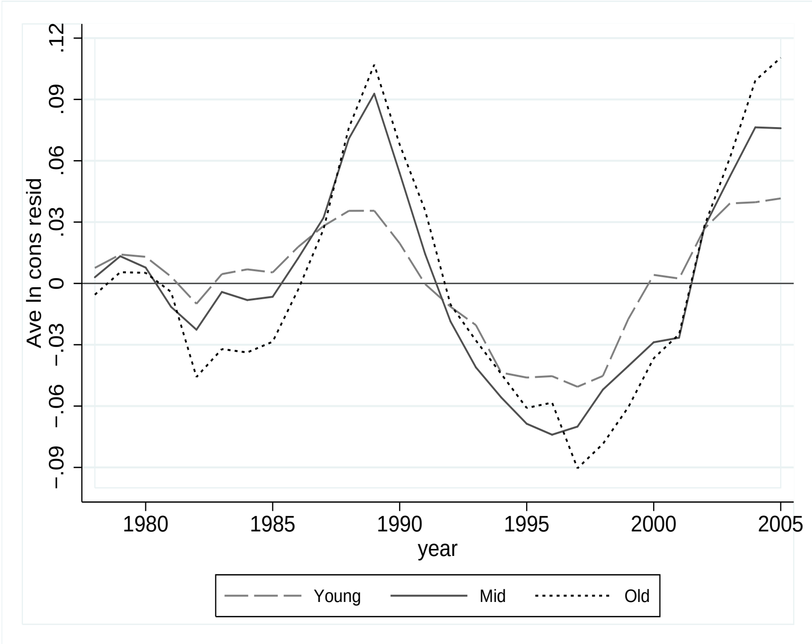


Figure 6: Residuals for log consumption: Model Results, House Price Shocks Only

price shocks. The results of these exercises are reported in Figures 6 and 7.

Figure 6 shows the averaged residuals for simulated data based on an economy in which house price movements are the only aggregate shocks. The results clearly show that, in the model economy, it is the consumption of the older group that responds most strongly to house price shocks. Individuals in this group tend to have high equity in housing, and have a short horizon over which to spread any new wealth. In times of large shocks to house prices, the average consumption residual in the older group can be as much as two to three times larger than the same statistic for the young.

Figure 7 shows the averaged residuals based on a simulated economy in which income changes are the only aggregate shocks. In this case the average residuals tend to be smaller than when house price shocks operated, and the pattern is also reversed as the younger group have the largest deviations of consumption from expected levels. The contrast of this pattern to the case with house price shocks, is important for the analyst who wants to understand why aggregate consumption and house prices tend to be positively correlated.

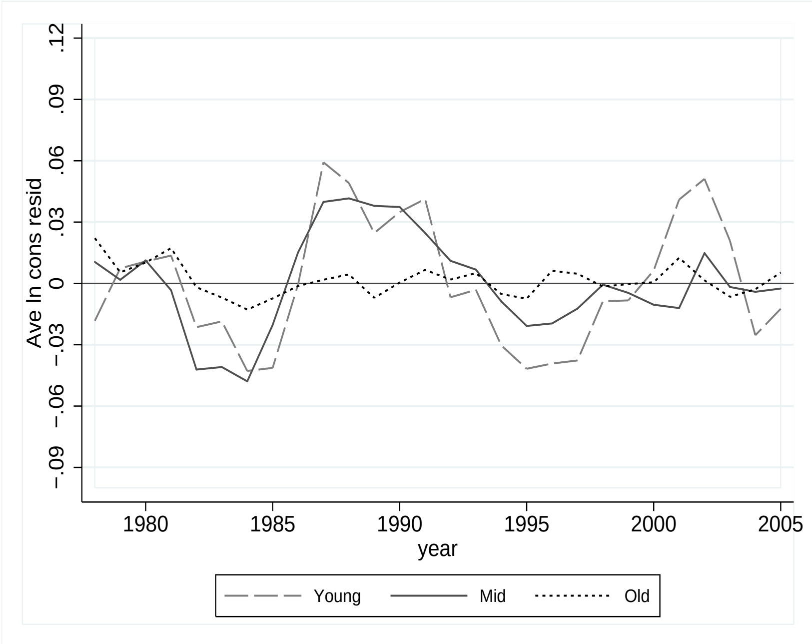


Figure 7: Residuals for log consumption: Model Results, Income Shocks Only

In summary, the analysis of consumption residuals by age group shows that different shocks translate into consumption responses that differ across groups. Younger individuals, within the model, tend to respond more to aggregate income shocks whilst older groups respond more to house price shocks. In a stylized life-cycle model, it is possible to see an intuition for why this occurs. A permanent income shock effectively acts as a shock to lifetime wealth, and younger cohorts have a longer horizon over which to enjoy a positive shock (or suffer a negative shock) and so adjust their consumption by more. Older cohorts respond more to house price shocks as they are more likely to have positive equity in their homes and so benefit from a wealth channel that allows them to adjust consumption. They would also respond more to this one-off wealth shock simply because they expect fewer periods of life over which they can spread extra consumption.

However, without the kind of analysis that we have conducted, one could not be sure that these intuitions would carry over to a more complicated model with credit constraints and simultaneous housing and non-durable consumption decisions. In particular, the possibility that house-price shocks might affect consumption through altering collateral, and so borrowing possibilities, would seem more pertinent for the young than the old. This implies that relative to the case of just a wealth effect, collateral effects might complicate the pattern across cohorts of changes in consumption following a house price shock. Our analysis shows that even with this factor in play, it is the old whose consumption is most responsive to the house-price shocks.

In this sense our analysis justifies the intuition used by Attanasio and Weber (1994) and Attanasio et al. (2009). However, the model we have have used induces counterfactual dynamics in the cross-section of consumption. In particular, we have seen that feeding aggregate earnings and price shocks into our model induces consumption responses from the younger consumers that are not much larger (if not smaller) than those from the other age groups. This pattern is not what we observe over several cycles of house price fluctuations in the actual data, where the behaviour of the youngest consumers seems to play a very important role in the aggregate.

The main reason for this pattern is the fact that the aggregate price shocks that we estimate from the data are relatively large. Although our older consumers are endowed with a special taste for home ownership, they still experience large capital gains and cash them in more than we see in the data. The fact that older consumers do not respond much to large changes in house prices in the data is reflected both in the lack of large consumption responses and in the limited evidence of downsizing of housing during the last part of the life-cycle. This type of phenomena calls for a slightly different model from the one we have considered.

To summarize, the results of the exercises we have conducted with our model indicate that: (1) simulation and aggregation of the model fits the aggregate data on consumption growth well even though none of the model parameters were chosen with reference to moments of aggregate consumption; (2) counterfactual simulations confirm the intuition that older individuals are much more sensitive to shocks to house prices, while younger ones react more to innovations in aggregate earnings; (3) in the simulations, house price shocks are, if anything, more important than income shocks in generating fluctuations in the rate of growth of aggregate consumption; (4) partly as a consequence of this and partly given the size of the house price shocks observed over the period we are studying, the model yields a somewhat counterfactual cross-sectional response of consumption - while in the data it is the youngest consumers' consumption that fluctuates the most, in the model this is not true as older households' consumption also moves substantially.

## 6 Conclusions and further work

Our analysis provides a better understanding of the mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. Through our systematic use of a realistic life-cycle model, we have investigated the generality of intuitions about how and why the consumption of different groups might correlate with house prices. We have examined whether these intuitions, which come from some simple life-cycle models, hold true in more complex and realistic frameworks. In this respect, the consideration of a mortgage market that resembles that available to UK consumers and where collateral effects can be sizeable, is particularly important. Our analysis has involved some methodological innovations in the way we have constructed our model and particularly in the way we have applied it to simulate the behavior of a series of cohorts, designed to resemble cohorts of the UK economy.

The analysis we conducted has in fact confirmed the intuitions in question. In particular, we have seen that if a house price shock is driving changes in aggregate consumption growth, then these changes will be most evident in the consumption paths of older groups. In our simulations the average deviations of consumption from expected levels *caused by* a house price shock were almost twice as strong for those at the end of their careers as they were for those at the start of their careers. To get the opposite pattern of stronger deviations for younger groups required some other type of aggregate shock. In the model, and plausibly for many episodes in the data that have coincided with periods of house price booms or busts, this could be a shock to incomes and expected permanent incomes throughout the economy. Thus we have given a more solid footing to the view that micro-data can be useful for disentangling the possible mechanisms underlying the correlation between house prices and consumption growth, and to the interpretation of Attanasio and Weber (1994) and Attanasio et al. (2009) that a stronger correlation for young groups is evidence against the hypothesis that wealth effects from house price changes have been the main mechanism driving the correlation.

The exercise we have performed also yields some new and interesting results: in particular, we show that our model, when fed the observed aggregate shocks (to house prices and earnings) and then simulated for a large number of individuals whose composition by age and education mimics the composition of the UK population, is able to generate patterns of aggregate consumption growth that closely resemble the dynamics of aggregate UK data. This result is remarkable as the model was not estimated or calibrated to fit such data. On the negative side, the model generates cross-sectional patterns of consumption growth across different age groups that are not too similar to those observed in the data. In particular, the model seems unable, in the version we have used, to generate relatively muted consumption responses of older consumers in the face of sizable swings in house prices. One consequence of this fact is that consumption is extremely sensitive, in the model, to changes in house prices. In other words, house prices in our model are 'too important', especially for older consumers, to be consistent with observed behavior.

We see the exercise we have performed, with its virtues and its limitations, as a first step of a research agenda that aims to explain the behavior of house prices as an equilibrium process. We have, in effect, modeled the demand side of the housing market. For such a purpose, we have made a set of assumptions about the process generating house prices which have facilitated an empirical specification that could be approximated in our simulations. In doing this we have learnt that the precise shape and parameterization of this process is an extremely important driver of the evolution of the joint decisions of housing demand and consumption expenditure, both at the household and the economy-wide level.

It is possible that building a setup with endogenous prices would help us to understand why consumption, and particularly the consumption of the old, seems to be less responsive to house-price shocks than is predicted by the current model.

But more than that, it is a major research challenge to assess whether the volatile house price process we observe, and other features of aggregate consumption and housing-market behavior, can be obtained from a unified framework in which house prices are determined endogenously.

## References

Aron, Janine, John Muellbauer, and Anthony Murphy (2007) 'Housing wealth, credit conditions and consumption.' mimeo, Oxford University

Aron, Janine, John V. Duca, John Muellbauer, Keiko Murata, and Anthony Murphy (2010) 'Credit, housing collateral and consumption: Evidence from the uk, japan and the us.' Discussion Paper, University of Oxford Department of Economics 487, University of Oxford, May

Attanasio, Orazio P., and Guglielmo Weber (1994) 'The uk consumption boom of the late 1980's: aggregate implications of microeconomic evidence.' *The Economic Journal* 104(427), 1269–1302

_ (1995) 'Is consumption growth consistent with intertemporal optimization? evidence for the consumer expenditure survey.' *Journal of Political Economy* 103, 1121–1157

Attanasio, Orazio P., and Matthew Wakefield (2010) 'The effects on consumption and saving of taxing asset returns.' In *Dimensions of Tax Design: The Mirrlees Review,* ed. Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba (Oxford University Press) chapter 7, pp. 675–736

Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber (1999) 'Humps and bumps in lifetime consumption.' *Journal of Business and Economic Statistics* 17(1), 22–35

Attanasio, Orazio P., Laura Blow, Robert Hamilton, and Andrew Leicester (2005) 'Booms and busts: Consumption, house prices and expectations.' IFS Working Paper WP05/24, Institute for Fiscal Studies, (http://www.ifs.org.uk/wps/wp2405.pdf)

_ (2009) 'Booms and busts: Consumption, house prices and expectations.' *Economica* 76(301), 20–50

Attanasio, Orazio P., Renata Bottazzi, Hamish W. Low, Lars Nesheim, and

Matthew Wakefield (2007) 'Explaining life-cycle patterns of home ownership.' mimeo, IFS

Bottazzi, Renata, Hamish W. Low, and Matthew Wakefield (2007) 'Why do home owners work longer hours?' IFS Working Paper WP10/07, IFS, (http://www.ifs.org.uk/wps/wp1007.pdf)

Campbell, Jeffrey R., and Zvi. Hercowitz (2004) 'The dynamics of work and debt.' NBER Working Paper 10201, NBER

Campbell, John Y., and João F. Cocco (2007) 'How do house prices affect consumption? evidence from micro data.' *Journal of Monetary Economics*

Chari, Varadarajan V., Patrick. J. Kehoe, and Ellen. R. McGrattan (2007) 'Business cycle accounting.' *Econometrica* 75(3), 781–836

Cocco, João F. (2005) 'Portfolio choice in the presence of housing.' *Review of Financial Studies*

Cristini, Annalisa, and Almudena. Sevilla-Sanz (2008) 'Do house prices affect consumption and why? a replication and comparison exercise.' mimeo, University of Oxford, (http://www.eea-esem.com/files/papers/EEA-ESEM/2008/2656/comparison_ccabhl_140208.pdf)

Davidson, Russell., and James G. Mackinnon (2004) *Econometric Theory and Methods* (Oxford: Oxford University Press)

De Nardi, Mariacristina. (2004) 'Wealth inequality and intergenerational links.' *Review of Economic Studies* 71(3), 743–68

Ejarque, João, and Søren Leth-Petersen (2008) 'Consumption and savings of first time house owners: How do they deal with adverse income shocks?' CAM Working Papers 2008-08, University of Copenhagen. Department of Economics. Centre for Applied Microeconometrics, November

Fernandez-Villaverde, Jesus, and Dirk Krueger (2005) 'Consumption and saving over the life cycle: How important are consumer durables.' Working Paper, University of Pennsylvania, (http://www.econ.upenn.edu/~dkrueger/research/durables12192005sec.pdf)

Flavin, Marjorie, and Shinobu Nakagawa (2004) 'A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence.' Working Paper 10458, NBER

Floden, Martin (2007) 'A note on the accuracy of markov-chain approximations to highly persistent ar(1)-processes.' Working Paper in Economics and Finance 656, SSE/EFI, (http://swopec.hhs.se/hastef/papers/hastef0656.pdf)

King, Mervyn A. (1990) 'Discussion (of "is the u.k. balance of payments sustainable?").' *Economic Policy* 5(11), 383–387

Kopecky, Karen, and Richard Suen (2010) 'Finite state markov-chain approximations to highly persistent processes.' *Review of Economic Dynamics* 13(3), 701–714

Li, Wenli, and Rui Yao (2007) 'The life-cycle effects of house price shocks.' *Journal of Money, Credit and Banking*

Lustig, Hanno N., and Stijn G. Van Nieuwerburgh (2005) 'Housing collateral, consumption insurance, and risk premia: An empirical perspective.' *Journal of Finance* 60(3), 1167–1219

Muellbauer, John, and Anthony Murphy (1990) 'Is the uk balance of payments sustainable?' *Economic Policy* 5(11), 348–395

_ (1997) 'Booms and busts in the uk housing market.' *The Economic Journal* 107(445), 1701–1727

Pagano, Marco (1990) 'Discussion (of "is the u.k. balance of payments sustainable?").' *Economic Policy* 5(11), 387–390

Piazzesi, Monika, and Martin Schneider (2010) 'Inflation and the price of real assets.' mimeo, Stanford University

Rios-Rull, Jose Victor, and Virginia Sanchez-Marcos (2008) 'An aggregate economy with different size houses.' *Journal of the European Economic Association, Papers and Proceedings* 6(2-3), 705–714

Rouwenhorst, K. Geert (1995) 'Utility analysis and the consumption function: an interpretation of cross-section data.' In *Frontiers in Business Cycle Research,* ed. Thomas F. Cooley (Princeton University Press) pp. 294–330

Slacalek, Jiri (2009) 'What drives personal consumption? the role of housing and financial wealth.' *The B.E. Journal of Macroeconomics*

Tauchen, George (1986) 'Finite state markov-chain approximations to univariate and vector autoregressions.' *Economics Letters*

Tauchen, George, and Robert Hussey (1991) 'Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models.' *Econometrica* 59, 371–396

# Appendix A Solution Method and Simulation Details

## Solution Method

The setup of our model, with a discrete choice concerning home-ownership, coupled with fixed costs and the particular form of the borrowing constraints, mean that the functions of the household's optimisation problem will not be 'well behaved' and we cannot rely on the existence of smooth first-order conditions that could otherwise have been exploited to improve efficiency in solving the problem.[24] Instead we rely on more robust techniques which involve solving using iteration on the value function (rather than the first order condition), and finding different "conditional value functions" (one for each of the current choices of house ownership, flat ownership, and non ownership) which can be compared in order to determine the discrete choice.

As is standard for these dynamic problems, the solution for consumption and home-ownership is found recursively from the last period of life, T, backwards. In the final period of life the value function consists of current utility from home ownership and consumption, plus the utility from leaving a bequest, and behavior in this period is constrained by the necessity that assets at the end of life (after leaving the bequest) be non-negative. Given the optimal choices at $t+1$, $t<T$, the backwards recursion then proceeds to choose home ownership, consumption and saving that maximise period $t$'s value function, subject to the borrowing constraints.

In order to compute the solution, we solve at a finite number of points in the asset dimension. We store optimal decisions and value functions at grid points but in our simulations households' choices are not restricted to coincide with these points. We perform linear interpolation in all the cases in which choices lie between points.

We also use discrete approximations to the specified continuous processes for idiosyncratic income, the house price and aggregate income. This involves modelling these processes using finite state Markov chains that mimic the underlying

[24]The combination of the two borrowing constraints mean that we can show not only that the value function will not be universally concave, but also that the derivative (with respect to assets) will not be defined at all points in the support of this function.

continuous-valued univariate or bivariate processes. This is done as described in Tauchen (1986). We preferred Tauchen's method of equally spaced nodes over the quadrature based method proposed by Tauchen and Hussey (1991), because this is more robust to very high persistence in the modelled processes (see Floden (2007)). Avery recent analysis by (Kopecky and Suen (2010)), drawn to our attention since we completed our modelling work, indicates that the method proposed in Rouwenhorst (1995) may have further improved the approximation.

For the runs presented in this draft of the paper, we have used 105 nodes in each 'conditional' asset grid (we have separate grids underlying each conditional value function, since assets are limited by different borrowing constraints depending on the home-ownership choice). Points are more dense in the lower range of the asset grid, to make sure that non-convexities in the value function are not overlooked in the maximization process. Idiosyncratic income is represented by a grid of eight nodes, while there are 13 nodes in each of the grids for the house price and aggregate income (or, effectively 169 nodes for the joint process). Monte-Carlo experiments showed that these grid sizes were sufficient to capture the modeled processes to a high degree of accuracy. With this set up, the model solution and simulation takes 30-40 hours on a desktop PC.

### Simulation details

Once we have the model solution to hand, it is relatively straightforward to simulate modeled behavior for a large number of individuals. Simulation proceeds forwards through the lifecycle. Initial draws for wealth and prices (the house price and idiosyncratic and aggregate income) determine choices about consumption (and saving) and home ownership in period 1 (at age 22). These choices contribution to wealth and home ownership realizations at the beginning of period two. These realizations, together with the exogenously determined values of the prices in period 2, are the state variables that drive optimal choices in that period. A similar loop determines state variables, and thus choices, in period three, and the process continues until choices throughout the lifetime have been modeled.

As explained in the body of the paper, we simulate a series of cohorts designed to resemble cohorts of the UK population. This means that our simulation procedure is different from that of drawing realizations of the shocks at the individual level in order to investigate behavior "on average". In our case, every member of a particular cohort experiences the same shocks to aggregate income and the house price at each age. Members of subsequent or earlier cohorts experience the same series of realizations of the house-price shocks, but at the relevant ages so that every cohort experiences the same shocks at the same point in time. For the aggregate income shocks, the results we present are based on realizations of the shock that are different for the under 36s, for those aged 36 to 55, and for those aged 56 plus. Only the realizations are assumed to differ by age group: the parameters of process generating the shocks are assumed to be common to all.

To make our simulated cohorts resemble cohorts of the UK population we set realizations of the model's aggregate shocks to match observed outcomes. The realizations that are input in our baseline run are designed to match actual shocks to house prices and aggregate incomes in the UK between the 1970s and 2005. Practically, instead of determining realizations based on the draw of a random number, this involves setting the house price and aggregate (age-group) income to match the growth rates in Figure A.1. The upper panel of this figure shows growth rates in house prices and "aggregate" income for 1970 to 2005, the lower panel shows the excess growth rates (residuals from the VAR described in equation 16) which are the shocks that are input into the baseline run of the model. It should be emphasized that the realizations are not known to households ahead of time: expectations of the aggregate processes are formed on the basis stochastic process with permanent shocks that is described in Section 3.2.2.

Alongside aggregate shocks to the house price and income, households also receive idiosyncratic income shocks, and the exogenous shock that may induce a move of house for home owners, each period. The initial level of financial assets is also a household specific feature so that within education group, households differ according to their initial financial wealth and the idiosyncratic income and moving shocks that they face at each period.

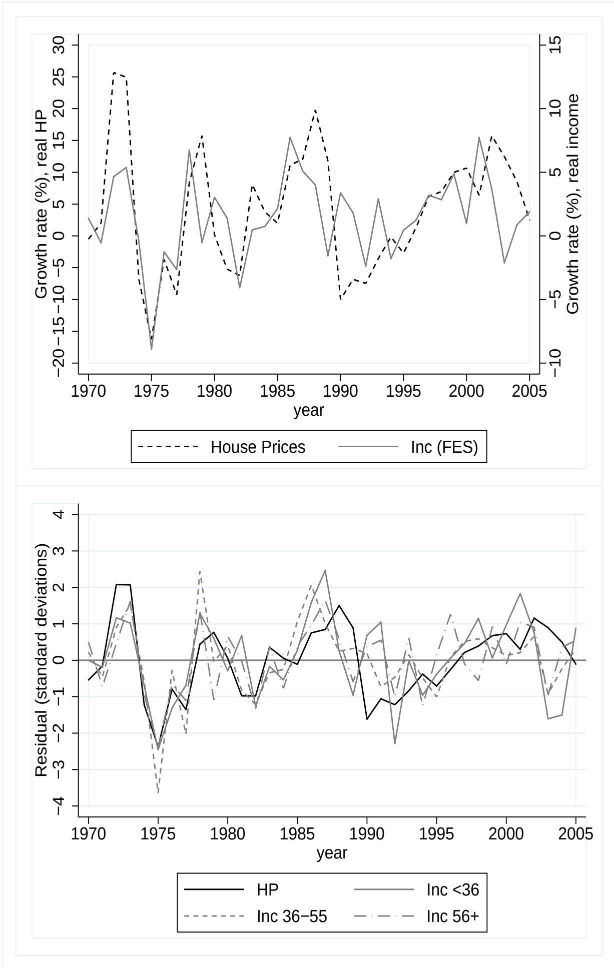


Figure A.1: Growth in real house prices and "aggregate income", and model shocks

As well as determining the stochastic shocks each year, in the simulations we must fix "initial conditions", in particular the initial house price (relative to income) and the initial level of financial wealth. Data considerations meant that these were fixed using data from the British Household Panel Survey (BHPS), which is the dataset used to estimate the parameters of the idiosyncratic income process (see Appendix B).

The house price at the start of adult life is set such that, on average for the years 1991-2002, the expected flat-price to income ratio matches the average house value to income ratio reported by those aged 22-26 in the BHPS for those years. These ratios are approximately 3 for the high education group and 3.5

for those with only compulsory schooling. The nature of the aggregate shocks in our model mean that this ratio is not matched each year: Those who enter the model in 1990, or in the mid-2000s, faced a higher ratio as house-prices had risen more than proportionally with income in those boom years; in the mid-1990s house-price bust, by contrast, the ratio was (is modeled to be) lower. Having fixed the initial house-price to income as described in this paragraph, all cohorts for each education group are simulated based on a single solution of the model with the expected level of income at age 22 normalized to 1 (0.79) for the high-education (low-education) group. That is not to say that our simulated datasets assume no real income growth between cohorts. Instead, when we aggregate the simulated data we gross up income (and consumption) realizations for all but the first (1915-19) cohort according to the expected annual growth-rate of aggregate incomes (1.66%) compounded by the age-difference compared to the first cohort (i.e. by 10 for the 1925-29 cohort).

The 2000 BHPS dataset[25] was our source for the distribution of financial assets for the two education groups at the beginning of adult life. This distribution was matched to that of individuals aged 22-26 in the relevant education group. The matching used eleven points of support for the distribution: households could draw zero wealth (with a probability matching the proportion that report zero assets), or the mean wealth of any decile group in the distribution of those reporting positive assets. Consistent with the credit constraints in the model, since households do not own a home at the start of their modeled life, initial assets must be non-negative.

The profiles of behavior reported in the paper are obtained by simulating 13 5-year-of-birth cohorts (born between 1915-19 and 1977-79). We simulate 2000 individuals, split evenly between two education groups, for each cohort. When we compare our simulated cohorts to FES survey data, we use weights so that the relative size of the different cohorts, and the proportions of high and low educated households within a cohort, match the data. In particular, weighted group sizes match group sizes implied by using FES survey weights to gross up that sample to match the UK population. Thus there are separate weights for groups defined

[25]Wealth data are not reported in every survey.

by cohort and education, and within cohort-education groups there are different weights at each age (or, equivalently, each year). The weights are matched to a smoothed (three-year moving average) of the grossed up population cellsizes for the FES.

# Appendix B Estimating the idiosyncratic part of the income process

The estimation of the idiosyncratic element of the income process required panel data on family incomes. We used data from the British Household Panel Survey 1991-2005, and a measure of family (or, more properly, tax unit) non-investment income.

The estimation proceeded in two steps and was carried out separately for each education group. The first step was to regress income on a polynomial in age, and it turned out that in order to capture noticeable differences across education groups a cubic specification was required, but higher order terms did not enter significantly. Having fitted this regression, the persistent element of the process was then estimated as an AR(1) on the residuals from this regression. Results from this exercise are summarized in table B.1.

Table B.1: Estimated parameters of the idiosyncratic income process

| | **Low Education** | **High Education** |
|---|---|---|
| age | -0.01068 | 0.01724 |
| age squared | 0.00184 | 0.00037 |
| age cubed | -0.00005 | -0.00002 |
| $\rho_y$ | 0.76771 | 0.76365 |
| $\sigma_\xi^2$ | 0.16580 | 0.15471 |