Modelling the Demand for Housing over the Lifecycle

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

February 2010

Abstract

This paper models individual demand for housing over the life-cycle, and shows the implications of this behaviour for aggregate demand. Individuals delay purchasing their first home when incomes are low or uncertain. This delay is exaccerbated by downpayment constraints. Higher house prices lead households to downsize, rather than to stop being home-owners. In aggregate, positive house price shocks lead to consumption booms among the old and a fall in aggregate demand for housing, whereas positive income shocks lead to consumption booms among the young and a rise in aggregate demand for housing.

Keywords: housing, uncertainty, credit constraints
JEL codes: D91, R21, G11

1 Introduction

The extent of home ownership varies substantially over the life-cycle and across cohorts. This variation might in part reflect changes in needs and the process of asset accumulation over the life cycle but could also be attributed to differences in credit market conditions, to price fluctuations particularly when young, and to differences in realised incomes. Moreover, the demand for housing also exhibits important fluctuations over the business cycle. Changes to current and future expected income, to interest rates and to house prices are all likely to affect both the propensity individuals have to own rather than to rent and also the size (and value) of the house they would like to live in. The aim of this paper is to model the demand for home ownership over the life-cycle to understand these different factors. There are two aspects to understanding demand for housing: on the one hand, we want to understand the main determinants of individual housing demand. On the other, we also want to characterize the properties of aggregate housing demand, that is, how a particular model of individual housing demand is aggregated in the demand for housing in the macroeconomy and how different shocks to income and price levels are translated into aggregate demand for housing. We pay particular attention to the housing and mortgage market institutional constraints faced by consumers, who are assumed to live in an economy that, in these respects, resembles the UK.

There are several reasons to be interested in the exercises we present. The characterization of the individual demand for housing and how it is affected by different environmental factors is important because housing decisions have a significant impact on individual utility, directly through being a large consumption item as well as indirectly through being the largest asset for most households. Our model helps us to understand how these decisions are affected by policy changes, such as a tax on buying or selling homes (such as the UK Stamp Duty Land Tax) or subsidies to first-time buyers. The characterization of aggregate demand highlights the overall impact of these factors. Furthermore, this is an important first step in understanding the determination of house prices that, in a country like the UK, exhibit a considerable amount of variation. In what follows, we establish how aggregate demand (for different types of houses) moves with different types of idiosyncratic and aggregate shocks.

We construct a model which incorporates a number of realistic features: households choose throughout their lives whether or not to own a home, and choose between houses of different size. Housing services give utility interacting with the consumption of nondurable consumption. Housing and consumption choices are made in the face of uncertainty about earnings and about house prices and in the presence of various capital market imperfections. The earnings and house price processes are calibrated from the data and are taken as given by individuals. Similarly, the features of the capital market are taken as given and are meant to mimic some of the institutional features present in the UK mortgage market: consumers are able to borrow only a fraction of the value of the house, and only able to borrow up to a multiple of their earnings and they are subject to fixed transaction costs. The restrictions on the size of the mortgage relative to earnings and house values are enforced only at the moment of purchasing the house or when consumers re-finance their mortgages. Given that house prices and earnings do fluctuate, these conditions can be violated for a consumer with an existing mortgage. Indeed, a consumer can have negative equity if house prices decline sufficiently. There is no explicit mechanism in the model through which the consumer can insure herself against fluctuations in house prices. However, because we force the consumer to always pay interest and, eventually, her debts, she will never borrow more than she can repay with probability one.

Given the parametrized model, the exogenous processes for earnings and house prices and the rest of the stochastic environment, we solve the model numerically and use it to simulate the demand of many individuals facing a certain set of shocks. We can then aggregate individual behaviour and derive aggregate demand. Our main results follow from two types of exercise. First, we characterize properties of individual demand and, in particular, how housing choices change in response to changes in the features of the mortgage market and of the stochastic processes. Demand among young individuals is particularly sensitive to lifetime income: higher income leads to purchasing housing earlier in the life-cycle, and to an increase in demand for ownership of houses at the expense of ownership of flats, whereas greater uncertainty leads to delays in purchase and to downsizing. Earnings related borrowing constraints are important in determining the level of ownership, while downpayment constraints determine the timing. The price of home ownership affects the ownership of houses in the opposite way to the ownership of flats: as prices (of flats *and* houses) rise, demand for flats rises as individuals downsize rather than moving out of home ownership altogether.

Our second exercise is to take a particular stochastic environment and demographic structure, and show how aggregate demand and aggregate consumption reacts to certain shocks. The realisation of a high house price shock reduces aggregate demand for housing, as the young decide not to buy and the old decide to sell. This leads to consumption rising substantially for the old due to reinforcing wealth and substitution effects, and rising somewhat for the young who substitute out of housing, but without a wealth effect. The realisation of a positive income shock boosts aggregate demand for housing, but the composition differs from the house price shock: the young respond more to the income shock than the old, and consumption rises more for the young than the old. These results suggest that a consumption boom among the young rather than the old indicates a positive aggregate income shock rather a positive aggregate house price shock.

Our analysis is related to a growing set of recent papers that have built a house-type asset into a life-cycle consumption saving framework. One example is Li and Yao (2007), who consider the behavioural and welfare consequences of house price shocks. Our remit is rather wider than theirs, and the nature of our housing asset and our modeling of mortgage borrowing, distinguish our paper from theirs. Distinguishing features of our model are: the modeling of mortgage related borrowing constraints that are only checked when the household buys or has to renegotiate the mortgage with the bank to increase its value; and, the modeling of housing as an asset that takes a discrete number of possible sizes so that there is a housing "ladder" but the dwelling is not continuously adjustable (even at cost). The contribution of Nichols (2005) used a similar two-size structure for the housing asset, but did not have the detail in modeling borrowing constraints that we have. Rios-Rull and Sanchez-Marcos (2008, In Progress) have a similar structure to ours but embed it in an equilibrium setting and concentrate on macroeconomic outcomes, rather than life-cycle decisions and welfare.

The rest of the paper is organized as follows: Section 2 presents and calibrates a life-cycle model of housing choice. Section 3 analyses how lifecycle decisions and welfare are affected by the parameters of the housing market (including credit constraints and fixed costs as well as the house price) and of the income process. Section 4 shows how decisions are affected by shocks that occur during the lifetime. Section 5 concludes.

2 A Life-cycle Model of Consumption and Housing Choices

We start from a relatively standard model of life-cycle consumption in a dynamic stochastic environment. We add to this model several features that capture the complexity of the consumer decision environment with regard to housing and debt choices. We do not consider uncertain life times and bequest motives. (see the discussion in Attanasio, Leicester, and Wakefield (2008)).

2.1 Model Structure

A household lives for T periods. In every period $t \leq T$, the household maximizes lifetime utility by choosing what affects the instantaneous felicity function: consumption, c_t, and housing services. The latter depend on what type of housing the individual lives in and whether she owns or rent the residence. The consumer will decide whether to own a "flat", or to own a "house", or to own neither, with $h_t \in \{0, 1, 2\}$ (where 0 is non-ownership). In addition to housing, the consumer has access to an asset that pays an interest rate $r_{t+1} = R_{t+1} - 1$ between periods t and $t+1$. We write the household value function in period t in recursive form as:

$$V_t \left(A_t, h_{t-1}, p_t, w_t\right) = \max_{\{c_t, h_t\}} u\left(c_t, h_t\right) + \beta E_t V_{t+1}\left(A_{t+1}, h_t, p_{t+1}, w_{t+1}\right) \tag{1}$$

subject to

$$A_{t+1} = R_{t+1} \begin{cases} A_t + w_t - c_t - \kappa p_t(1+F)I(h_t = 1) - p_t(1+F)I(h_t = 2) \\ \quad \text{if } h_{t-1} = 0 \\ \\ A_t + w_t - c_t + \kappa p_t(1-F)I(h_t \neq 1) - p_t(1+F)I(h_t = 2) \\ \quad \text{if } h_{t-1} = 1 \\ \\ A_t + w_t - c_t - \kappa p_t(1+F)I(h_t = 1) + p_t(1-F)I(h_t \neq 2) \\ \quad \text{if } h_{t-1} = 2 \end{cases} \tag{2}$$

where A_t is the start of period asset stock; p_t is the price of housing which is realised at the start of period t; F is a proportional fixed cost which is assumed to be the same for both buying and selling a house or flat; w_t is household earnings in period t. Equation (2) is a standard intertemporal budget constraint, augmented by terms reflecting the house price and transaction costs that must be borne when trading housing. For ease of exposition, in what follows, we distinguish between beginning of period assets A_t and end of period assets s_t.

In our model, there are two differences between a flat and a house: first, owning a house gives, as we shall see below, more utility than owning a flat, and second a house is more expensive than a flat. The price of a flat is a fraction, κ, of the price of a house, and so the fixed cost of buying or selling a flat is a fraction of the fixed cost of buying or selling a house.

2.1.1 Financial markets

We allow only for collateralized debt, such that households are able only 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

$$s_t \geq 0 \tag{3}$$

Home owners can borrow, and when they do so they are subject both to a terminal asset condition, $s_T = 0$, that translates into an implicit limit on borrowing,[1] and to two explicit borrowing constraints. The first of these explicit constraints is a function of the value of the house and the second is a function of household annual earnings. These determine how much a household is able to borrow at the time of purchase or when remortgaging, and translate into the following constraints on saving:

$$s_t \geq -\lambda_h \kappa p_t, \qquad \kappa = \begin{cases} 0 < \kappa < 1 & \text{if } h_t = 1 \\ 1 & \text{if } h_t = 2 \end{cases} \tag{4}$$

where the value $(1 - \lambda_h)$ can be thought of as a down-payment requirement, and:

$$s_t \geq -\lambda_w w_t \tag{5}$$

The explicit constraints on the down-payment and the debt to income ratio only apply when households buy the property or remortgage. That is to say, if at period t the household continues to own the property that they owned at period $t-1$, then as long as they service the interest on any

[1] The specification of marginal utility becoming infinite at 0 consumption means this terminal condition prevents households borrowing more than they can repay with certainty.

outstanding mortgage debt (the next but one paragraph describes this interest repayment), then the debt that they hold will not be limited by the mortgage-related borrowing constraints. This means that these formal borrowing constraints will not force households to shrink their mortgage rapidly, or sell their home, in periods when they are hit by large negative shocks to the house price or income, which would make the formal constraints tighter.

The structure of the constraints just described adds to the computational difficulty of our problem. It means that convexity preserving techniques cannot be used since there are known 'kinks' in the conditional value functions for owning a home, at least at points in the state space where this choice involves continued ownership.[2] This computational difficulty is probably the key explanation of why it has been almost standard in the literature modeling housing and consumption choices to assume that mortgage constraints, (if any, and often represented only by a collateral constraint) must be satisfied in every period.[3] We are not willing to make such an assumption. The ability to borrow more when house prices and income move up and loosen borrowing constraints, without a concern that a subsequent fall will require a large debt repayment, is sure to be of first order importance for young individuals deciding whether or not to buy, and how much to consume, in periods when their incomes fluctuate.[4]

Turning to the cost of servicing the mortgage, the interest due on outstanding debt at the start of period t is defined as:

$$m_t = r_t s_{t-1} \tag{6}$$

There is no fixed mortgage repayment schedule. However, if the household does not repay at least the interest, m_t, on their outstanding debt, they have to remortgage. Remortgaging does not incur

[2]By a kink we mean a point at which the derivative of the value function is not defined. To see why there must be kinks, note that continuing owners who hold some debt will have their assets constrained either by their existing stock of debt, or by the formal borrowing constraints, depending on whether or not they choose to remortgage. Which of these will be binding is a function of the control variable of the dynamic optimization problem, the level of assets (debt). At the point in the asset range where the binding constraint switches between these two, the value function will be kinked. Heuristically, this can be thought of as having constraints on the optimization that switch over within the state space of the problem, at a point at which the Lagrange multipliers on both constraints are strictly non-zero and will not be equal to each other (except by chance). Since one of the constraints ceases to apply without its associated multiplier declining smoothly to zero, this gives a kink in the value function.

[3]Examples include the model of Ortalo-Magné and Rady (2006) in which the assumption is an analytical convenience, and computational contributions such as Li and Yao (2007), Cocco (2005), Yao and Zhang (2005) and Campbell and Hercowitz (2004).

[4]The modeling becomes even more complicated, but the issue perhaps even more pertinent, in a situation in which income is affected labour supply choices as well as random shocks; see Bottazzi, Low, and Wakefield (2007) for a model of this situation which includes mortgage constraints that only apply when buying or increasing the value of the mortgage, as described here.

a cost but, as discussed above, any new mortgage has to satisfy the two formal constraints.

2.1.2 Utility function

The within period utility function is a CRRA function in current consumption, augmented by additive and multiplicative terms to capture the value of home-ownership[5]:

$$u(c_t, h_t) = \frac{c_t^{1-\gamma}}{1-\gamma} \exp(\theta\phi(h)) + \mu\phi(h) \qquad \begin{cases} \phi = 0 & \text{if } h_t = 0 \\ 0 < \phi < 1 & \text{if } h_t = 1 \\ \phi = 1 & \text{if } h_t = 2 \end{cases} \tag{7}$$

The parameters θ and μ are housing preference parameters which determine the utility premium that households derive from owning their home; they are calibrated in our model. The exponential in the multiplicative term for the value of ownership is a convenient way to express that this term represents a proportional scaling of the utility from consumption. When $h_t = 0$, the exponential term has value 1 and the additive term value zero, and thus utility is only derived from non-durable consumption. ϕ determines the relative utility from owning a flat versus a house. The additive term means that we do not impose housing and consumption to be homothetic, and the sign of μ affects whether housing is a luxury ($\mu > 0$) or a necessity ($\mu < 0$).

While the specification of the effects of housing on utility is very simple as we consider effectively only three possible choices, we think that the model captures some essential features of housing services: we have the possibility of non separability between house ownership and non durable consumption, we have the possibility of preference for ownership and we have differences between 'large' and 'small' houses.

2.1.3 Stochastic processes

In the model households face uncertainty in two dimensions: idiosyncratic uncertainty over earnings and aggregate uncertainty over house prices.[6]

[5] A structure where home ownership boosts utility from non-durable consumption is also used by other authors (e.g. Rios-Rull and Sanchez-Marcos (2008, In Progress), and Ejarque and Leth-Petersen (2008)). Our structure, with multplicative and additive terms in home ownership, is flexible in allowing different combinations between whether home ownership is a luxury or a necessity, and whether home ownership and non-durable consumption are complements or substitutes (see also section 2.4).

[6] In fact it is not a difficult extension to include iid noise in the interest rate, but results were not sensitive to this and so we removed this dimension of uncertainty from our final runs; in equation (1) we implicitly assumed that the

Following MaCurdy (1982), the idiosyncratic income process is assumed to follow a random walk:

$$\ln w_t = a_t + v_t \quad \text{ where } \; v_t = v_{t-1} + \xi_t, \quad \xi_t \sim N\left(-\frac{\sigma_\xi^2}{2}, \sigma_\xi^2\right) \text{ and } \rho_w = 1 \tag{8}$$

and a_t is the deterministic growth in earnings over the life-cycle and has a hump shape ($a_t = a_1 t + a_2 t^2$).

The house price is assumed to evolve as an AR(1) but in this case the deterministic element reflects upwards drift over time:

$$\ln p_t = d_0 + d_1 t + \rho_h \ln p_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \tag{9}$$

The price of a flat is assumed to be a proportion κ of the price of a house.

2.2 Model Calibration

We now discuss how we determine the parameters required for the analysis. For some parameters, such as the elasticity of intertemporal substitution of consumption and the discount rate, we use values from elsewhere in the literature. Some parameters, instead, we estimate directly from different data sources. The rest of the parameters we obtain through calibration using the structural model outlined in section 2. These parameter values are reported in table 2.

2.3 External Parameter Values

The ‘externally fixed’ parameter values are reported in table 1. They include the parameters of the income and house prices processes, of the interest rate, the details of the financial market structure and of the utility function. We discuss them in turn.

Borrowing limits. The parameters that determine the fraction of the house price (λ_h) and the multiple of earnings (λ_y) that households can borrow are chosen to match institutional features of the UK mortgage market. Households can borrow up to whichever amount is lower between three times household earnings ($\lambda_y = 3$) and 90% of the house price ($\lambda_h = 0.9$).

House price process. Estimation of the parameters of the house price process is based on

interest rate was not a stochastic state variable.

Table 1: Parameter Values

Parameter	Value	Source
	House Price Process	
ρ_h	0.94	ODPM
σ^2_ε	0.008	ODPM
d	2.32%	ODPM
κ	0.6	BHPS
p_{22}	4.67	BHPS
	Income Process	
ρ_w	1.0	(By assumption)
$\sigma^2_{\xi HE}$	0.035	BHPS
$\sigma^2_{\xi LE}$	0.044	BHPS
(a_{1HE}, a_{2HE})	(0.042, - 0.00082)	BHPS
(a_{1LE}, a_{2LE})	(0.022, - 0.00037)	BHPS
w_{22HE}	1.0	BHPS
w_{22LE}	0.8	BHPS
	Preference Parameters	
γ	1.43	(Attanasio and Weber, 1995)
ϕ	0.6	BHPS
	Other parameters	
λ_y	3.0	
λ_h	0.9	
β	1.02^{-1}	
$\overline{r}$	0.018	B.o.E

Notes: BHPS indicates that source is survey data from the British Household Panel Study; ODPM indicates data that were provided by the then Office for the Deputy Prime Minister, these data are now available from Communities and Local Government (http://www.communities.gov.uk/index.asp?id=1165366).

the Office of the Deputy Prime Minister (ODPM) national and regional house price series for the UK, years 1969-2000.[7] We estimate an AR(1) process, with linear trend (equation (9)), for the logarithm of real house prices, where the conversion from nominal house prices was made using the Retail Price Index (RPI, all items). The result of the estimation is a persistence parameter (ρ_h) of 0.94 and a standard deviation of the shock (σ_ε) equal to 0.089. A unit root test on the persistence parameter does not reject the null hypothesis ($\rho_h = 1$). When running simulations to estimate the elasticity of aggregate demands we treat house price shocks as aggregate.

The ratio of the price of a flat to the price of a house, κ, is set by dividing all houses and flats in the data into two categories by the number of rooms. The ratio κ is therefore the ratio of the

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

average price of a home with less than 5 rooms (including kitchens and bathrooms) to the price of a home with more than five rooms. While it would have been interesting to allow the house and flat prices to be two independent processes not perfectly correlated, this would have increased the number of state variables of the dynamic programming problem solved by the consumer.

Income process. We estimate the parameters of the income process (σ_ξ and a_1 and a_2) using data from the British Household Panel Survey (BHPS) for the years 1991-2002. Since the decision making units in our model are best thought of as families, the process that we estimate is for household (non-investment) income. To obtain an estimate of the variance of the permanent shock in the income process (σ_ξ^2), we follow the estimation procedure proposed in Blundell, Pistaferri, and Preston (2004).

We estimate the parameters for high and low education groups separately. The results in table 1 show firstly that high education individuals can expect a more hump-shaped income profile than their less educated couterparts during their working lives (both a_1 and a_2 have a bigger magnitude for the high education group), and secondly that the high education group have a lower variance in permanent shocks to their income.

We model retirement as being a period of 15 years in which households' income is given by a replacement rate of 70 percent of their last annual income. Income is not subject to risk during retirement. Having retirement income allows households to continue owning their house when they stop working, and therefore home ownership in our calibration is still close to the levels observed in the data around age 60.

Interest rate process. For interest rates we use the average 90 day Treasury Bill discount rate in years 1968-1997, which gives a rate of 1.5%.

Utility function. The preference parameter γ in the utility function is set to match the consumption elasticity of intertemporal substitution of 0.7 in the data (see Attanasio and Weber (1995) and the survey by Attanasio and Wakefield (2008 and forthcoming)). This corresponds to a curvature parameter $\gamma = 1.43$ for our within period utility function. The parameter ϕ indicating the relative utility value of a flat to a house is set at 0.6.

Initial Wealth. We set the initial distribution of financial assets for the two education groups to match data on 22-26 year olds in the 2000 wave of the BHPS,[8] and we assume that households

[8] These data are discussed in Banks, Smith, and Wakefield (2002).

have zero housing endowments at age 22.

2.4 Calibrated parameters

Given the parameters above, we set the remaining parameters so that the model reproduces some features of life-cycle home-ownership profiles for household heads aged 26-60 between 1991 and 2000,[9] by education group. Our approach is to choose the parameters to minimise the sum of squared deviations between moments calculated in the data and corresponding simulated moments. The moments we use are the average home-ownership rates for households in low and high education groups, for those aged 26-35 and those aged 36-60. The statistics from our model are measured across 40 different simulation runs (i.e. 40 different realized sequences of the aggregate house price process), each of which simulates the behaviour of 1000 individuals. We set the calibrated parameters to be common across the two education groups. Our calibration and comparative static exercises focus on home-ownership behaviour upto age 60 because we do not model bequests, and so households run down all assets by the end of life, leading to an overestimate of the amount of selling of homes towards the end of life.

We use this calibration process to pin down the transactions cost of buying or selling, F, and the parameters specifying the utility benefit of home ownership, μ and θ. Parameter values from the calibration are summarized in table 2. Table 3 presents the calibration statistics, showing home-ownership rates predicted by the model match those observed in the data.

Table 2: Calibrated Parameter Values

Parameter	**Value**
θ	0.026
μ	0.11
F	0.05

When assessing the plausibility of our (proportional) fixed cost parameter, it is necessary to bear in mind that there is a fixed cost of buying and of selling, so an agent who trades up or down while continuing to own will pay a transaction cost equal to five percent of the sale price *plus* 5

[9]Data come from the years 1991-2000, as years prior to 1991 are affected by the large-scale selling off of local authority housing.

Table 3: Calibration Statistics

Statistic	High Education Data	High Education Model	Low Education Data	Low Education Model
Ownership rate (percentage)				
Age 26 - 35	65.4	68.2	49.5	49.2
Age 36 - 60	81.1	76.5	62.7	62.8

Notes: The data figures for home-ownership rates are based on the years 1991-2000 of the FES.

percent of the purchase price. Our fixed cost parameter of 5% seems plausible given the costs of employing estate agents, lawyers, surveyors, removal companies, and other specialists, when moving house house in the UK. In addition, residential property transactions incur stamp duty, a transactions tax which has rates varying between zero and 4% of the price of the property (the rate increases with the house price) and which is formally paid by the house buyer. Since the fixed cost is affected by tax policy, it is interesting to know how it affects behaviour and welfare, and we shall consider these issues in the next section.

Unlike with the fixed cost, we do not have strong priors about the plausible level for the calibrated utility parameters. The figures in Table 3 show that the parameters we chose fit extremely well the statistics for the low educated individuals. However, for the individuals with higher education the profile in the model is 'flatter' over the life cycle than what we see in the data. The fit for the high education group could have been improved by considering slightly higher values of μ and θ. However, for these larger values the parameter the fit of the model for the low education group would worsen considerably. The point chosen minimized the overall criterion function we chose. We did not consider different taste parameters for the two education groups.

The housing utility parameters determine whether home ownership is a luxury or a necessity, and whether home ownership and consumption are substitutes or complements in utility. As mentioned above, μ being positive means that home ownership is a luxury in utility: given that the risk aversion parameter γ is greater than one, the utility from consumption is negative and decreasing in absolute value as consumption increases; the constant, positive additive term is thus a bigger proportional shift in utility at higher levels of consumption, and this is the sense in which home ownership is a luxury. The positive value of θ implies that home ownership and consumption are complements in utility, in the sense that the cross-partial derivative of utility with respect to

c and h is positive.

2.5 Baseline life-cycle profiles

Figure 1 shows profiles of flat and house ownership for both education groups in our baseline simulation. These plotted profiles are not completely smooth because in the calibration we have averaged over a relatively small number (40) of realizations of sequences of house prices. This is done because the house price is thought of as an aggregate variable (and the parameters of the house price process are calibrated from aggregate data). Since this price is aggregate, data for a given set of years contain information on only a relatively small number (equal to the number of cohorts observed) of realizations of the house price at each age.

To assess our calibration, it would be useful to compare the predictions of our model concerning the pattern of non-housing wealth holdings over the lifetime, to patterns observed in data. Unfortunately it is hard to obtain such numbers from the data reliably, since private pension wealth is rarely well measured in survey data. One exception to this is the English Longitudinal Study of Ageing, which started in 2002 and measures detailed information on different elements of wealth portfolios for households that include individuals aged 50 or older.[10] Data from this survey in 2002/03 indicate that the mean (median) family wealth held in financial assets and private pensions was approximately 9.2 (4.7) times median income for low income individuals aged 51 -60, and for high education individuals the figures were 16.8 (11.0). The nearest equivalent measures for the low education group in our baseline simulation are 9.5 (7.2), which is a reasonably good match. For the high education group the simulation numbers are 8.2 (5.8). The worse match for the high education group may be partly due to the fact that the pension replacement rate in our model is 70% for all individuals, and this is higher than the replacement rate that higher income individuals can expect from state pensions in the UK. The absence of inheritances (apart from some initial wealth) and bequests from the simulations may also be a simplification that is less realistic for the high education group.

[10]For more details on this, see Banks, Emmerson, Oldfield, and Tetlow (2005). We are grateful to Gemma Tetlow for help dealing with the ELSA data cited here.

Figure 1: Simulated Home Ownership in the Baseline

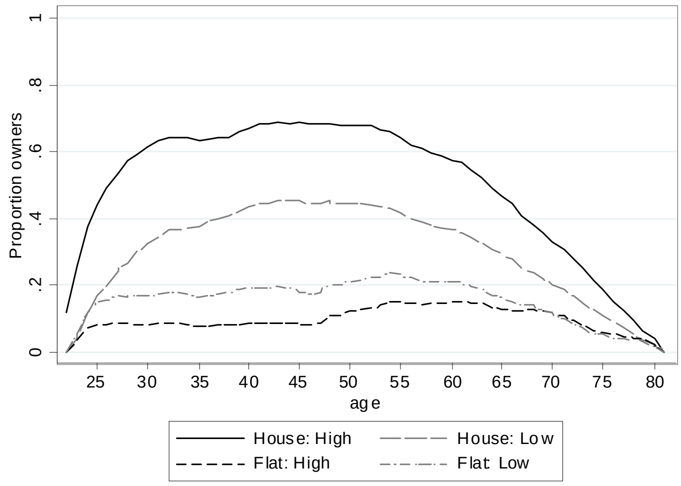

3 What Determines Individual Housing Demand?

In this section, we use our calibrated model to show how sensitive the demand for housing is to the level of income and house prices, to capital market imperfections and to other frictions in the economy, such as the fixed transactions cost. We focus in this section on changes to parameters that affect individuals across their entire lifetime and which are fully anticipated. This is in contrast to the effect of particular realisations of stochastic variables across the life-cycle, which we focus on in the next section. The analysis in this section can be split into three parts: first, we characterize the effects due to differences in the level and variability of income and in the borrowing constraint relative to income; second, we show the importance of factors related to house prices: the level and variability of prices as well as the borrowing constraint relative to the value of the house. Finally, we show the importance of the fixed cost.

Throughout this section we will be looking at changes induced by changes in the various parameters we study on housing and consumption demand. There are two different aspects to housing demand that we focus on: demand for ownership of a house of any size, and demand for housing of a particular size, which in our simple framework is characterized by the choice of flats versus houses.

To look at housing and consumption demand responses, we utilize elasticity measures. Since consumption is a continuous variable, we can calculate the elasticity of consumption to changes in the parameters of our model. However, since home-ownership is discrete, we cannot compute its derivative. Instead we calculate changes in ownership rates for a given cohort of ex-ante identical individuals. These individuals will be ex-post different because of the shocks they will experience. We will simulate the behaviour of a large number of such individuals and compute ownership rates under our baseline parameters and in an alternative scenario which incorporates the specific change we want to consider. The changes we consider, therefore, are not shocks but changes in environment. For instance, to compute the responsiveness of home-ownership to a 1% change in income, we simulate the life-cycle of ownership for the 40,000 ex-ante identical individuals in our baseline case and again in a different economy where income has increased by 1%. We then compare home-ownership rates in the two economies and measure their proportional change relative to the change in income.[11] This elasticity measure captures the effect of an anticipated increase in lifetime income. To aid interpretation of the elasticities for the housing good, in table 5 in the appendix, we report home ownership rates in our baseline run, since these enable the calculation of quantity effects from the elasticity measures.

3.1 The importance of income

Tables 4, 5 and 6 show the extent that home ownership and consumption respond to a change in average lifetime income (ϵ_Y), to a change in the variance of earnings (ϵ_{vY}) and to a change in the borrowing constraint that depends on income (ϵ_{λ_Y}). In these as in following tables, the first three columns refer to high education individuals, while the last three to low education individuals. The tables report, for each of the three elasticities and each of the two education groups, the averages for three age groups: 26 to 35, 36 to 50 and 51 to 60.

The top panel of Table 4 reports the responsiveness of home-ownership, while the second panel looks at changes in the 'quantity' of housing. We define quantity by counting a flat as 0.6 of a house, as is defined in the price process. The main message from the results in this table is that home ownership of the young is the most sensitive to anticipated income changes and to changes in the constraints relating to income. Higher income or a loosening of the income constraint leads to greater home ownership among the young but makes less difference to the older groups. This

[11] Apart from being scaled to take account of the change in income levels in the economy, the income and house price shocks experienced by each of the 40,000 simulated individuals are held fixed across experiments.

Table 4: Elasticities for ownership: Factors relating to income

	High Education			**Low Education**		
	ϵ_Y	ϵ_{vY}	$\epsilon_{\lambda Y}$	ϵ_Y	ϵ_{vY}	$\epsilon_{\lambda Y}$
			Ownership elasticities			
Age 26 - 35	0.95	-0.06	0.34	0.98	-0.26	0.24
Age 36 - 50	0.36	-0.11	0.03	0.51	-0.10	0.02
Age 51 - 60	0.43	-0.10	-0.00	0.48	-0.10	0.02
Lifetime	0.67	-0.08	0.14	0.69	-0.13	0.12
			Quantity elasticities			
Age 26 - 35	1.12	-0.05	0.42	1.39	-0.27	0.39
Age 36 - 50	0.49	-0.12	0.08	0.72	-0.14	0.09
Age 51 - 60	0.54	-0.12	0.03	0.68	-0.12	0.05
Lifetime	0.80	-0.09	0.20	0.92	-0.14	0.19

message remains the same if we consider the quantity of housing bought rather than just the effect on ownership, although the quantity elasticities tend to be somewhat larger than the ownership elasticities.

Comparing quantity and ownership elasticities is one way to think about the importance of including different sizes of property in the model.[12] Another is to consider separate elasticities of demand for flats and for houses, as in table 5: the income elasticity of demand for flats is negative while for houses it is positive, suggesting that flats are, in our model, an 'inferior good'.[13] Similarly, the elasticity of demand for flats with respect to the income related borrowing constraint is negative, while for houses is positive: a relaxed borrowing makes housing more affordable. As housing becomes more affordable, more people upgrade from a flat to a house than from not owning

[12] In a version of our model with only a single property type (with value (utility and price) equal to the average of the flat and house values in the current model, or with value equal to the value of the house), the qualitative patterns of elasticities (both for changes related to income and changes related to house prices) were similar to those reported for the two size model. Such a 'one-size' model would not allow scope to show that quantity elasticities are generally amplified (rather than dampened) relative to ownership elasticities, nor to consider separate flat and house markets.

[13] This feature is reliant on the fact that a change in the price of housing affects flats and houses equally.

to owning a flat, and thus the demand for flats falls.

Table 5: Elasticities for flats and houses: Factors relating to income

	High Education			**Low Education**		
	η_Y	η_{vY}	$\eta_{\lambda Y}$	η_Y	η_{vY}	$\eta_{\lambda Y}$
	Elasticities for Ownership of a Flat					
Age 26 - 35	-2.43	-0.30	-1.27	-1.65	-0.18	-0.74
Age 36 - 50	-2.27	0.17	-1.02	-1.09	0.20	-0.51
Age 51 - 60	-0.96	0.21	-0.38	-0.77	0.00	-0.19
Lifetime	-1.29	-0.04	-0.64	-0.84	-0.07	-0.33
	Elasticities for Ownership of a House					
Age 26 - 35	1.41	-0.28	0.56	2.30	-0.30	0.73
Age 36 - 50	0.71	-0.14	0.17	1.18	-0.22	0.24
Age 51 - 60	0.74	-0.17	0.09	1.14	-0.15	0.13
Lifetime	1.04	-0.09	0.29	1.44	-0.16	0.34

Table 6 looks at the effects of the income related changes on non-durable consumption. We see that a change in income each period has an effect on consumption that is close to one-for-one in each age range. An increase in the variance of income shocks tends to reduce consumption, particularly for the youngest group. More generally, this reinforces the picture of the young as being most willing to adjust in all margins.

3.2 The importance of house prices

Tables 7, 8 and table 9 show the extent that home ownership and consumption respond to an increase in the level of house prices at all ages, to a change in the variance of shocks to house prices, to a change in the downpayment constraint and to a change in the fixed cost.

As with income, home ownership among the young is the most sensitive when these parameters relating to the housing market are changed: ownership among the young falls as the house price increases, as uncertainty about the house price increases, as the downpayment requirement becomes

Table 6: Consumption elasticities: Factors relating to income

	High Education			Low Education		
	ϵ_Y	ϵ_{vY}	$\epsilon_{\lambda Y}$	ϵ_Y	ϵ_{vY}	$\epsilon_{\lambda Y}$
Age 26 - 35	1.02	-0.25	0.01	1.02	-0.27	0.01
Age 36 - 50	1.00	-0.09	0.00	1.01	-0.11	0.01
Age 51 - 60	0.98	-0.01	0.00	0.99	-0.05	0.01
Lifetime	0.98	-0.05	0.00	0.98	-0.08	0.00

Table 7: Elasticities for ownership: Factors relating to house prices

	High Education				Low Education			
	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}
				Ownership elasticities				
Age 26 - 35	-0.74	-0.18	0.15	-0.23	-0.68	-0.17	0.02	-0.31
Age 36 - 50	-0.24	-0.09	0.01	-0.07	-0.39	-0.08	0.01	-0.10
Age 51 - 60	-0.34	-0.03	0.01	-0.10	-0.42	-0.00	0.01	-0.12
Lifetime	-0.53	-0.06	0.10	-0.19	-0.55	-0.04	0.02	-0.21
				Quantity elasticities				
Age 26 - 35	-0.92	-0.15	0.17	-0.22	-1.06	-0.11	0.03	-0.36
Age 36 - 50	-0.37	-0.06	0.02	-0.09	-0.58	-0.03	0.01	-0.15
Age 51 - 60	-0.44	-0.01	0.01	-0.10	-0.59	0.03	0.01	-0.14
Lifetime	-0.67	-0.05	0.11	-0.17	-0.76	-0.00	0.02	-0.21

greater and as fixed costs go up. Another pattern that is similar to that from the changes related to income is that - with the partial exception of the responsiveness to the variance in income shocks - quantity elasticities again tend to be slightly larger than ownership elasticities.

We can look in more detail at how changes in housing demand are allocated across the two types of housing, by considering elasticities for houses and flats separately (see table 8). In some cases, the elasticities for flats and houses have opposite signs. This is most evident for the elasticities

Table 8: Elasticities for flats and houses: Factors relating to house prices

	High Education				**Low Education**			
	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}
			Elasticities for Ownership of a Flat					
Age 26 - 35	2.79	-0.68	-0.15	-0.34	1.81	-0.56	-0.02	-0.04
Age 36 - 50	2.39	-0.72	-0.17	0.16	1.05	-0.51	0.01	0.25
Age 51 - 60	1.03	-0.33	0.02	-0.11	0.60	-0.19	0.01	-0.01
Lifetime	1.52	-0.28	-0.02	-0.52	0.84	-0.29	0.01	-0.17
			Elasticities for Ownership of a House					
Age 26 - 35	-1.23	-0.11	0.19	-0.21	-1.92	0.03	0.04	-0.45
Age 36 - 50	-0.59	-0.01	0.03	-0.11	-0.99	0.09	0.00	-0.25
Age 51 - 60	-0.64	0.04	0.01	-0.10	-0.96	0.10	0.02	-0.18
Lifetime	-0.92	-0.02	0.13	-0.12	-1.23	0.08	0.02	-0.22

when the price of housing goes up: in response to this change, the demand for houses declines but the demand for flats increases. This property is due to the fact that when the price of housing increases more people shift from house ownership to flat ownership, than shift from flat ownership out of the housing market.

Table 9: Consumption elasticities: Factors relating to house prices

	High Education				**Low Education**			
	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}
Age 26 - 35	-0.02	0.01	0.01	-0.02	-0.03	0.01	0.00	-0.02
Age 36 - 50	-0.00	0.01	0.00	-0.01	-0.01	0.01	0.00	-0.02
Age 51 - 60	0.01	0.01	0.00	-0.02	0.00	0.02	0.00	-0.02
Lifetime	0.02	0.01	0.00	-0.01	0.02	0.02	-0.00	-0.02

Table 9 reports the consumption response to changes in house price factors. All effects are economically very small.

3.3 Responsiveness to fixed costs

The effect of a change in fixed costs is particularly important, partly because it can be affected directly by tax policy.[14] We consider in this section the effect of halving this cost from 5% to 2.5%. We look at how fixed costs affect patterns of home ownership over the life-cycle, and at the impact of the level of fixed costs on expected lifetime welfare.

We know from tables 7 and 8 that increases in fixed costs lead to a reduction in home ownership and to a delay in buying, particularly among the young. In Figure 2 we show the effect on ownership of flats and houses of reducing the fixed cost from 5% to 2.5%. A lower fixed cost means households are more likely to buy flats before buying houses, and more likely to trade down to a flat after owning a house. A lower transactions cost has a number of effects: first, a substitution effect leading the demand for housing to rise, and consumption to fall; second, a positive wealth effect since, for a given number of home purchases at given house prices, more income is available for consumption.

In this subsection as well as looking at how fixed costs affect patterns of home ownership, we explore which of the substitution and wealth effects is most important in shaping behaviour.

Figure 2: Response of Ownership to Change in Fixed Costs

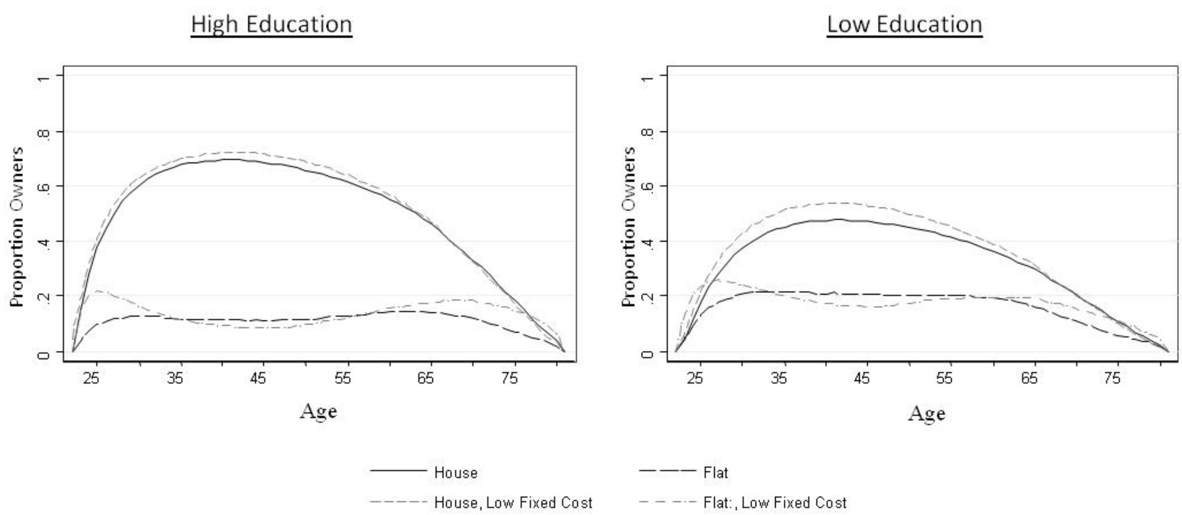

We separate out the wealth and substitution effects by first running the model with a low fixed cost that is known in advance (column 'c.' in table 10). The impact of this low fixed cost relative

[14]In the U.K., by stamp duty land tax.

to the baseline scenario (column 'a.') combines both substitution and wealth effects. We then run the model with households expecting to pay the baseline fixed cost, but when transactions occur the household pays only the low fixed cost (column 'b.'). There is therefore no substitution effect associated with the lower fixed cost.[15] The results of this exercise are reported in table 10, which reports statistics for the level of consumption and home-ownership, and for the average number of different housing market transactions during the lifetime, in the different scenarios that are simulated.

The final three rows in each panel of table 10 report measures of the welfare change induced by the reduction in fixed costs. The welfare change is reported in terms of the amount of resources needed to compensate the agents, on average, for not being in the low fixed cost scenario. The compensation is measured by a transfer of assets at the beginning of life (third from last row in each panel) and by a "consumption equivalent" which is the proportional increase in consumption at every age that is needed to make agents indifferent (ex ante) between the two scenarios being compared. Since the utility function with housing terms, and the credit constraints, means that the model does not have a solution that is homothetic throughout, the consumption equivalent can not be calculated analytically. Instead it is found through an iterative procedure that relies on expected discounted lifetime utility as derived in the numerical solution to the model for comparisons between the scenarios reported in columns a. and c. For comparisons involving the scenario for column b. in which the low fixed cost is a "surprise", there is no numerical solution for expected utility (rather the solution for the relevant policy functions is that with the higher fixed costs). For these cases the iterative procedure for finding the consumption equivalent relies on simulated expected lifetime utility. Fuller details of the computational procedure for finding the welfare measures are reported in the appendix.

The table shows that the halving of the fixed cost leads to an increase in consumption, and also to an increase in the expected level of home ownership (or equivalently, in the expected number of periods of ownership), across the lifecycle. The results also show the extent to which the cut in fixed cost leads to an increase in the number of housing market transactions during the lifetime: with the lower fixed cost the likelihood of moving directly from flat to house ownership (i.e. of

[15] The part of the wealth effect which would induce an ex-ante change in behaviour is also excluded from this experiment and so only the ex-post effect of higher wealth is calculated.

Table 10: Effects on consumption, homeownership, and welfare, due to changing fixed costs

High Education	**a. Calibrated Run**	**b. Low FC as "surprise"**	**c. Low FC**
E(Consumption)	1.421	1.432	1.430
E(ownership rate)	0.59	0.60	0.64
E(flat ownership)	0.11	0.11	0.13
E(house ownership)	0.48	0.49	0.50
E(# Transacts in life)	2.55	2.57	3.35
% trade up at least once	8.5	8.6	21.1
% downsize at least once	13.5	13.6	42.2
Equivalent variations, wrt c.:			
Initial asset measure	0.246	0.043	0
Consumption equiv (%)	1.106	0.195	0
Equivalent variations, wrt b.:			
Consumption equiv (%)	0.909	0	N/A

Low Education	**a. Calibrated Run**	**b. Low FC as "surprise"**	**c. Low FC**
E(Consumption)	0.988	0.998	0.997
E(ownership rate)	0.45	0.45	0.50
E(flat ownership)	0.15	0.15	0.16
E(house ownership)	0.30	0.30	0.33
E(# Transacts in life)	2.39	2.42	3.21
% trade up at least once	8.6	8.8	17.4
% downsize at least once	9.4	9.4	31.6
Equivalent variations, wrt c.:			
Initial asset measure	0.219	0.035	0
Consumption equiv (%)	1.263	0.196	0
Equivalent variations, wrt b.:			
Consumption equiv (%)	1.045	0	N/A

Notes: 'Trade up' is a direct movement from a flat to a house without an intervening period of non-ownership, and 'downsizing' is the converse. The initial asset measure is expressed as a percentage of expected initial income.

'trading up') at some age is more than doubled compared to in the high fixed cost case, while the likelihood of the converse change in ownership ('downsizing) approximately trebles. Also evident from the table is that consumption behaviour is almost the same in the scenarios of columns 'b.' and 'c.', while housing market behaviour is almost constant between scenarios 'a.' and 'b.'. That is, the change in the consumption behaviour happens even when the change in the fixed cost is not expected, but the change in behaviour in the housing market only occurs when decisions are made on the basis of knowledge of the lower transactions cost.

The welfare measures in table 10 show that while a consumption increase of slightly more than one percent at each age is required to compensate individuals for facing a 5% fixed cost instead of a 2.5% fixed cost, only a 0.2% increase in consumption is required to provide compensation for being surprised by the fixed cost, rather than being able to plan on the basis of it. Overall, then, we see that a cut in the fixed cost leads to an increase in welfare primarily through a wealth effect that allows increased lifetime consumption once the overall cost of each housing market transaction is reduced. That is not to say that the substitution effect is unimportant. Being able to plan on the basis of the lower fixed cost leads to an important reshaping of individuals' housing market choice functions, and the increased flexibility of trading in housing markets also increases welfare appreciably.

4 The Aggregate Demand for Housing

In the previous section, we saw how the level of housing demand and consumption differ with the level of house prices and the level of income over the entire lifecycle. In this section, we contrast those results with the effect of a particular realisation of the house price or income that strikes during the lifecycle. That is to say, we calculate $\eta_{H,t}$, $(\eta_{Y,t},)$ the effect on home ownership at t of realised house prices (income) being 1% higher than in the baseline.[16] We consider the effect on ownership and consumption of house price and income shocks occurring at age 30, 40 or 50.

Figure 3 shows the effect of a realised high house price on home ownership rates of the high and low educated. Figure 4 shows the effect on consumption.[17] The results show larger percentage impacts for the low education group, as the fact that the house price is a bigger proportion of lifetime wealth in the low education group (on average) outweighs the fact that fewer low education individuals own.

Though the size (in percentage terms) of effects is different across the two education groups, patterns of results are rather similar. The effect of the shock on home-ownership rates is an initial

[16] Since income is stochastic, the experiment is that the realisation of income is 1% higher than a draw from the distribution would imply.

[17] Each of the lines plotted in figures 3 and 4 shows the percentage difference in the home ownership rate or average consumption in two simulation runs between which the only exogenous change is to the house price shock at the given age such that the level of the house price at that age is increased by one percent. The simulation runs are based on 5000 sequences of house price shocks, and each of the 5000 drawn prices at the relevant age are increased by exactly one percent.

Figure 3: Effect of 1% house price rise at 30, 40 or 50, on Home Ownership

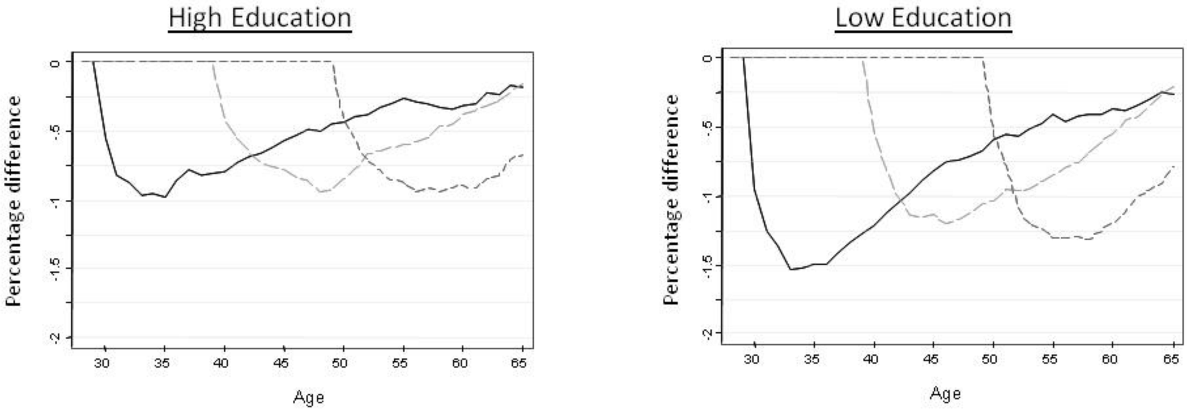

Figure 4: Effect of 1% house price rise at 30, 40 or 50, on Consumption

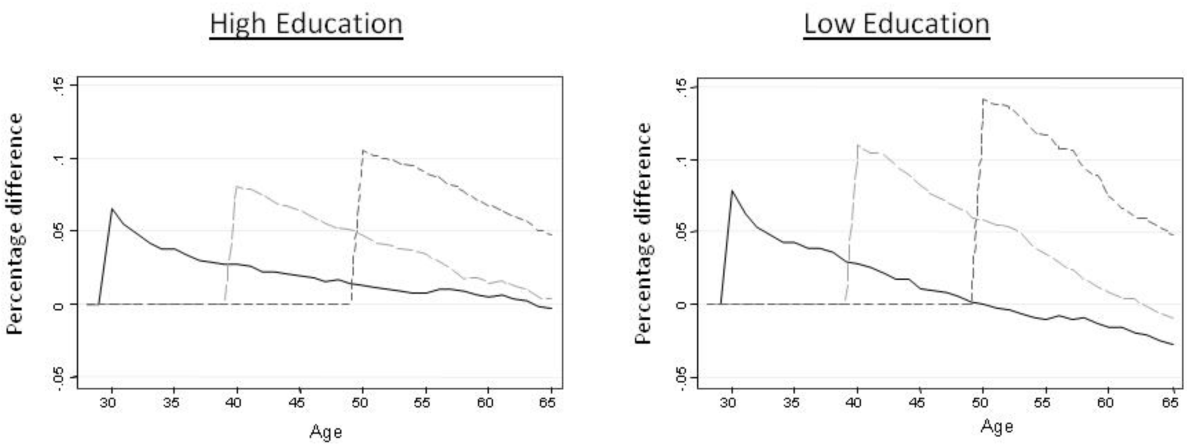

decline, with the largest reduction in home-ownership (of around 1% for the high education group, and 1.2% to 1.5% for the low education group) experienced around five years after the shock. The peak decline does not occur in the year of the shock because an initial contraction in the number of buyers has a lasting effect on the stock of home owners, and this is compounded by the persistence in the house price process meaning that the shock continues to discourage new purchases for several years after it is initially felt. In the longer term the home-ownership rate recovers towards where it would have been without the shock, and this reflects the fact that the house price displays persistence, but not permanence in shocks (figure 7 in the Appendix shows the rate at which the effect of a shock dissipates from the expected level of the house price). It also reflects that over time agents can modify their saving behaviour to mitigate the shock.

Turning to the effect on consumption, we see that in the model the positive house price shock increases average consumption regardless of the age at which the shock is experienced. The increase in consumption is seen to be strongest for the eldest group who are most likely to have large equity in their home, and who have the shortest horizon (on average) over which to spread extra wealth. This pattern across age-groups confirms the intuition behind the empirical analyses of Attanasio and Weber (1994) and Attanasio, Blow, Hamilton, and Leicester (2009, 2005) investigating why house price shocks have been strongly correlated with consumption growth in the U.K. over the past 30 years.

The effect on consumption tapers off over time. It is worth noting that the positive effect when the house price shock occurs becomes a negative difference some time in the 50s for those experiencing a shock aged 30, particularly in the low education group. Many of this group would be credit constrained at the time of the shock and so the opportunity to borrow more when the house price increases can be exploited by young individuals taking out mortgages and bringing forward some of their lifetime consumption in order to flatten out the lifetime consumption profile and so move towards the non-constrained optimum.

Figures 5 and 6 show the effect on average on home ownership and on consumption of a 1% upwards shock to income at age 30, 40 and 50.[18] As with the house price shocks just discussed,

[18] Each of the lines plotted in figures 5 and 6 shows the percentage difference in the home ownership rate or average consumption in two simulation runs between which the only exogenous change is to the income shock at the given age, such that the level income at that age is increased by exactly one percent for every individual. The simulation runs are based on 5000 individuals each with their own realisation of the (idiosyncratic) income process.

the patterns of results for these income shocks are similar for the two education groups.

Figure 5: Effect of 1% income shock at 30, 40 or 50, on home ownership

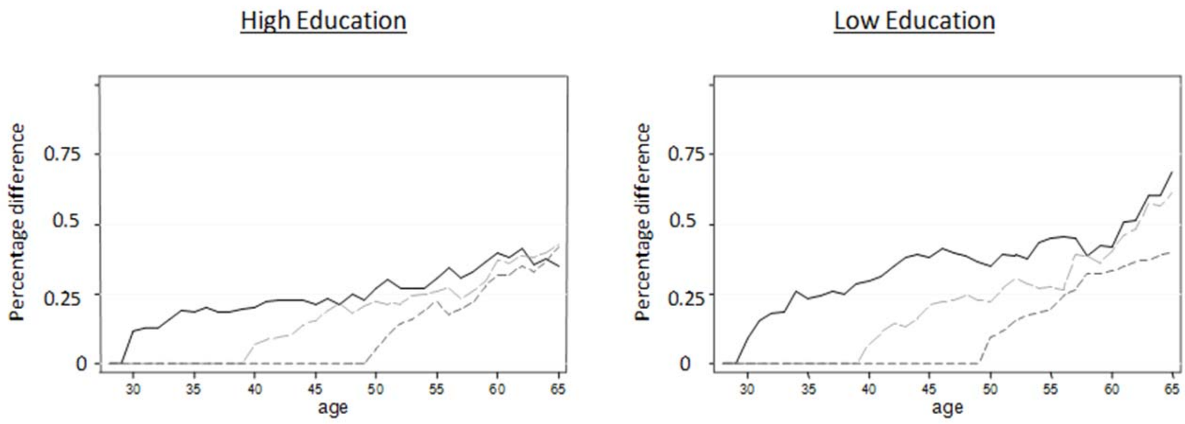

Figure 6: Effect of 1% income shock at 30, 40 or 50, on consumption

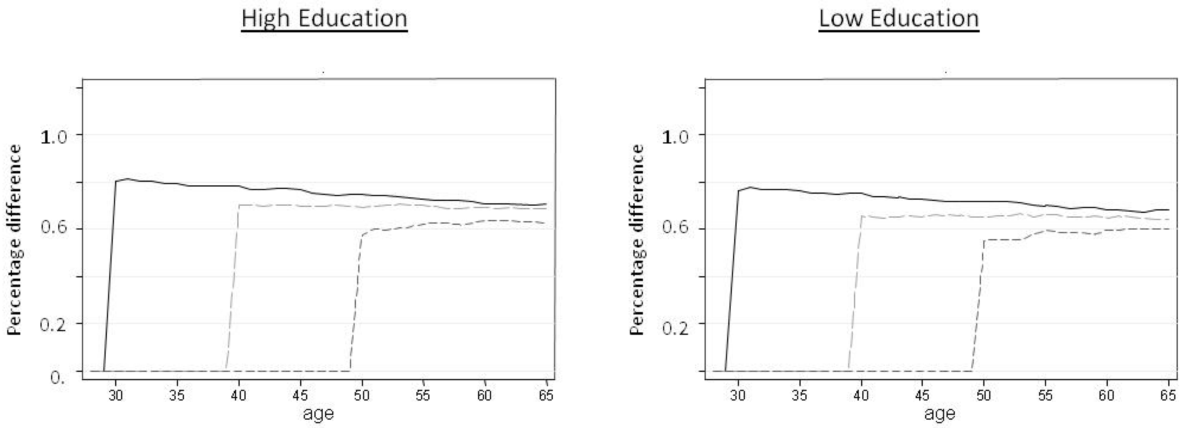

Figure 5 and 6 show that income shocks at the given ages lead to increases in both home-ownership rates, and non-durable consumption. The increase in consumption is immediate and (since income shocks are permanent[19]) persistent. The proportionate increase in consumption is somewhat less than the proportionate increase in income. One reason for this is that some of the increase in lifetime wealth is spent on housing. A second reason is that income only accrues while individuals are working. Older individuals enjoy the income shock for a smaller proportion of their

[19]Income is a random walk.

remaining lifetime, and experience it when they have already built up some positive (financial plus housing) wealth. This wealth means that these individuals will consume more than the sum of their remaining income during the remaining part of their life. Thus the one-percent increase in (expected) remaining lifetime income can fund an increase in remaining lifetime spending of somewhat less than one percent. Since this effect gets more marked as age and wealth increase, we see smaller proportionate increases in per period consumption, the later in life that the income shock occurs.

Whereas the increase in consumption after the income shock is rapid, the increase in the level of home ownership is more gradual. The random walk nature of the income process means that the income shock at a given age affects the income level throughout the rest of the lifetime. Thus, while for some individuals the initial shock is enough to induce them to become owners or to upsize straight away, for others some years of increased resources are required to allow a change in housing market choices. This explains why the ownership rate tends to drift up, rather than immediately jumping to a given higher level, after the shock to incomes.

5 Conclusions

We presented a life-cycle model of households choosing how much to consume or save, and whether to own a flat, a house or no housing, in each period of life. The model was constructed to capture realistically mortgage related borrowing constraints. Simulations of the model show how sensitive individual demand for housing is to the parameters of house price process, and the income process, and also to tightness of mortgage related borrowing constraints. The level of the house price and income have marked effects on behaviour: when incomes go up (all else unchanged) or the price of housing (i.e. houses and flats) goes down, this leads to an increase in demand for houses, but to a *fall* in the demand for flats. The level of fixed costs in the housing market can have a noticeable effect on consumption and welfare (through a wealth effect), and on housing market behaviour (a substitution effect). In terms of aggregate behaviour, particular realisations of shocks increasing house prices or decreasing income will have differing effects on behaviour depending on the age at which they occur: house price shocks lead to a consumption boom for the old, income shocks lead to a consumption boom for the young.

A Appendix

A1 Extra tables and pictures

Appendix Table A1: Ownership rates

	Quantity owned	Proportion owners
High education		
Age 26 - 35	0.65	0.68
Age 36 - 50	0.73	0.76
Age 51 - 60	0.72	0.77
Low education		
Age 26 - 35	0.43	0.49
Age 36 - 50	0.55	0.63
Age 51 - 60	0.54	0.63

Table A1 shows ownership rates by age group in the calibration run.

Figure 7: Effect of HP shock at age 30 on expected HP at older ages

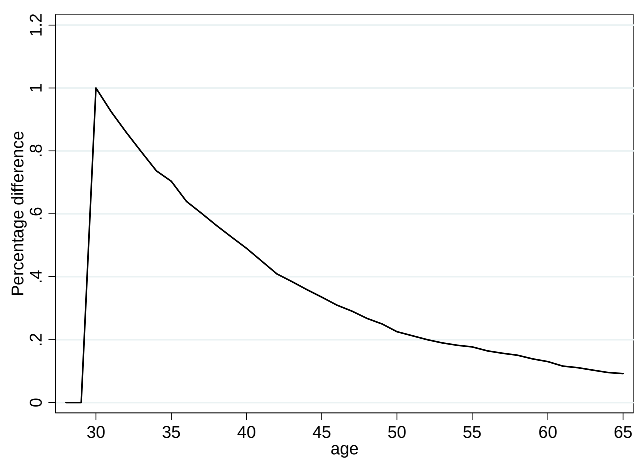

Figure 7 shows the effect on the expected house price at ages 31 to 60 of a 1% house price shock at age 30.

A2 Computational methods

The setup of our model, with a discrete choice concerning home-ownership, coupled with fixed costs and the particular form of the borrowing constraints, means that the functions of the household's optimisation problem are not '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.[20] Instead we rely on 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.[21]

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, and behaviour in this period is constrained by the necessity that assets at the end of life 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 and the house price. This involves modelling these processes using finite state Markov chains that mimic the underlying continuous-valued univariate 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 has been shown to be more robust to very high persistence in the modelled processes (Floden, 2007).

[20] 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.

[21] The development of these techniques was a key part of the research for the project RES-000-23-0283, funded by the U.K. Economic and Social Research Council (ESRC), (Attanasio, Low, and Nesheim, 2005).

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 grids, to make sure that non-convexities in the value function are not overlooked in the maximisation process. Income and the house price are each represented by a grid of fifteen nodes. Monte-Carlo experiments showed that these grid sizes were sufficient to capture the modelled processes to a high degree of accuracy. With this set up, the model solution and simulation takes around 3 hours on a desktop PC.

Method for deriving consumption equivalent measures

When constructing welfare measures in subsection 3.3 we adopt a standard approach of expressing these as a "consumption equivalent". This measures the proportionate increase in consumption every period that is required to compensate the individual in terms of ex-ante expected utility for some welfare reducing factor, in this case higher fixed costs in the housing market. In a lifecycle consumption/savings model with constant relative risk aversion utility and without the complications of a discrete housing choice and its associated borrowing constraints, such a measure can be found analytically. It is derived from the "target expected utility" $\overline{U_0}$ and the expected utility in the less preferred state of the world before compensation, $\underline{U_0}$, as:

$$\tau = (\overline{U_0}/\underline{U_0})^{(1/[1-\gamma])} \tag{10}$$

where γ is the coefficient of relative risk aversion.

In our case with the housing choice and mortgage borrowing constraints, the (expected) value function does not inherit homogeneity properties from the utility function, and so equation (10) cannot be applied and instead the value must be found numerically. As mentioned in the main text, for the case of comparing two fully optimised solutions, this iterative procedure can compare numerical solutions for ex ante expected utility, but in the case where one of our comparison cases is low fixed costs arriving as a 'surprise' the values of expected utility must come from averaging across many simulations of the lifecycle. Whichever way we arrive at the expected utility levels, the iterative procedure we adopt to find the consumption equivalent is as follows:

1. Find the baseline expected utility levels $\underline{U_0}$, for the case in which compensation is needed, and

$\overline{U_0}$, the target utility. Use these to calculate an initial guess for the consumption equivalent (τ) by applying equation (10).

2. Re-solve the compensated case, with consumption in every period multiplied up by the factor τ.

3. Find (from the solution or from simulation as appropriate) the expected lifetime utility in the compensated case.

4. Compare this to the target expected utility. If it is sufficiently close, stop.

5. If the level of expected utility is not sufficiently close to the target, then readjust τ. To readjust, calculate the amount by which expected utility exceeds (or falls short of) the target, as a proportion of the change in expected utility due to the latest change in τ, and use this proportion to scale τ to its new value. For example, if the change in expected utility has been two percent too large to hit the target, then reduce τ by two percent.

6. Return to step '2.', and keep repeating the procedure until step '4.' is satisfied.

This procedure turned out to be reliable and quite rapid. For the cases we considered the initial guess for τ based on equation (10) was a good guide to the final value of the compensating equivalent (the initial adjustment being less than ten percent inaccurate compared to the correct adjustment) and iteration to the solution generally took no more than half a dozen steps.

Method for implementing later shocks

When implementing one percent shocks in income or the house price at ages 30, 40 and 50 (see section 4), we had to set up the discrete approximation to the relevant price in such a way that an increase of exactly one percent could be implemented. This required altering the grids underlying our Markov chain approximations, and in particular a move away from the evenly spaced grids that we adopted for our baseline solution and simulation. In order to maintain the accuracy of the approximation, we increased the number of points in the relevant grid.

References

Attanasio, O. P., L. Blow, R. Hamilton, and A. 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, O. P., A. Leicester, and M. Wakefield (2008): "The Coincident Cycles of House Prices and Consumption Growth in the UK: Do House Prices Drive Consumption Growth?," Paper prepared for the NBER Summer Institute (Consumption).

Attanasio, O. P., H. W. Low, and L. Nesheim (2005): "Modelling Life-Cycle Labour Supply and Asset Holdings with Fixed Costs," ESRC specific grant.

Attanasio, O. P., and M. Wakefield (2008 and forthcoming): "The Effects on Consumption and Saving of Taxing Asset Returns," paper prepared for the report of a commission on reforming the tax system for the 21st century, chaired by sir james mirrlees, Institute for Fiscal Studies, (http://www.ifs.org.uk/mirrleesreview/reports/consumption.pdf).

Attanasio, O. P., and G. 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.

Banks, J., C. Emmerson, Z. Oldfield, and G. Tetlow (2005): "Prepared for retirement? The adequacy and distribution of retirement resources in Britain," IFS Report 67, Institute for Fiscal Studies.

Banks, J., Z. Smith, and M. Wakefield (2002): "The distribution of financial wealth in the UK: evidence from 2000 BHPS data," IFS Working Paper 02/21, The Institute for Fiscal Studies.

Blundell, R., L. Pistaferri, and I. Preston (2004): "Consumption inequality and partial insurance," IFS Working Paper 04/28, The Institute for Fiscal Studies.

Bottazzi, R., H. W. Low, and M. Wakefield (2007): "Why Do Home Owners Work Longer Hours?," IFS Working Paper WP10/07, IFS, (http://www.ifs.org.uk/wps/wp1007.pdf).

Campbell, J. R., and Z. Hercowitz (2004): "The Dynamics of Work and Debt," NBER Working Paper 10201, NBER.

Cocco, J. (2005): "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*.

Ejarque, J., and S. 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.

Floden, M. (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).

Li, W., and R. Yao (2007): "The life-cycle effects of house price shocks," *Journal of Money, Credit and Banking*.

MaCurdy, T. (1982): "The use of time series processes to model the error structure of earnings in a longitudinal data analysis," *Journal of Econometrics*, pp. 82–114.

Nichols, J. (2005): "Mortgage Contracts and The Definition of Demand For Housing Wealth," Ph.D. thesis, University of Maryland.

Ortalo-Magné, F., and S. Rady (2006): "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints," *Review of Economic Studies*, 73(2), 459–485.

Rios-Rull, J.-V., and V. Sanchez-Marcos (2008): "An Aggregate Economy with Different Size Houses," *Journal of the European Economic Association, Papers and Proceedings*, 6(2-3), 705–714.

——— (In Progress): "House Price Movements," Presentation.

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20(177-81).

Tauchen, G., and R. Hussey (1991): "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59, 371–396.

Yao, R., and H. H. Zhang (2005): "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints," *Review of Financial Studies*, 18, 197–239.