

Modelling the Demand for Housing over the Lifecycle

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

March 2011

Abstract

We model individual demand for housing over the life-cycle, and show the aggregate implications of this behaviour. Individuals delay purchasing their first home when incomes are low or uncertain. Higher house prices lead households to downsize, rather than to stop being owners. Fixed costs (property transactions taxes) have important impacts on welfare (a wealth effect) and house purchase decisions (substitution effect). In aggregate, positive house price shocks lead to consumption booms among the old but *falls* in consumption for the young, and reduced housing demand; positive income shocks lead to consumption booms among the young and increased housing demand.

JEL Codes: D91, R21, G11

Keywords: housing, uncertainty, credit constraints

1 Introduction

Houses are complicated objects. They provide housing services and utility and, at the same time, they constitute an important store of value. They are lumpy and their transactions are subject to substantial costs. At the same time they constitute collateral that can be used to obtain credit. As a consequence, modeling the demand for housing is not a trivial exercise. Empirically, the extent of home ownership varies substantially over the life-cycle and across cohorts (Bottazzi et al. (2011)). 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 the importance of these different factors. There are two aspects of the demand for housing we focus on: 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 into 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 could potentially be affected by policy changes, such as a tax on buying or selling homes (such as, the UK Stamp Duty) or subsidies to first-time buyers. The characterization of aggregate demand highlights the overall impact of these factors. This is an important first step in understanding the determination

of house prices, especially since, in a country like the UK with very stable supply of housing, it is likely that demand factors are a major driver of the considerable variation observed in house prices. 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 sizes, which we will call ‘flats’ and ‘houses’. If consumers buy or sell a house, they are subject to (proportional) transaction costs. 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. The features of the capital market are also 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. A novel and realistic assumption we make is that 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, consumers can have negative equity if house prices decline sufficiently. While this complication adds considerably to the numerical burden of the exercise, it constitutes an important element of realism. There is no explicit mechanism in the model through which consumers can insure themselves against fluctuations in house prices. However, because we force consumers to always pay interest and, eventually, all debts, they will never borrow more than they 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 use the model to simulate the demand of many individuals facing a given realization of earnings and house prices 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 par-

ticularly sensitive to lifetime income: higher income leads to the purchase of 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. Borrowing constraints, particularly the earnings related constraint, influence the timing of ownership and especially the choice between a house or a flat. Strikingly, an increase in the price of housing (and of rents) causes the home ownership rate to rise. This reflects a fall in the demand for houses but a rise in the demand for flats as individuals downsize rather than moving out of home ownership and into renting (with its increased price).

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 non-durable consumption rising substantially for the old due to reinforcing wealth and substitution effects, but falling somewhat for the young, who are faced with an increased need to save up to eventually meet a downpayment constraint and are also affected adversely by the comovement of rent with the house price. The realisation of a positive income shock boosts aggregate demand for housing, but the cross sectional pattern differs from the case of the house price shock, particularly with regard to consumption responses: 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 than 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. The nature of the housing asset in our model and our modeling of mortgage borrowing distinguish our paper from theirs. Distinguishing features of our model are: the modeling of mortgage related borrowing constraints, so that these are only checked when the household buys or has to renegotiate the mortgage 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). Nichols (2005) used a similar two-size structure for the housing asset, but did not have the same detail in modeling borrowing constraints. 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 (on the latter, see the discussion in ?).

2.1 Model Structure

A consumer lives for T periods. In every period $t \leq T$, the consumer 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 rents the residence. The consumer decides whether to rent, or to own a “flat”, or to own a “house”, with $h_t \in \{0, 1, 2\}$ (where 0 is renting). 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:

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

subject to

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

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; q_t is the price of renting a home; 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.

In our model, the available housing options differ in two dimensions: they provide different flows of utility and have different costs. 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. Rents are set as a function of house prices, as we describe below. Obviously, the other differences between ownership and renting are that the former gives access to borrowing and constitutes a store of value.

2.1.1 Financial markets

We allow only for collateralized debt: households are only able to have negative financial assets when they are home owners. When they rent ($h_t = 0$) they are subject to the constraint

$$A_t \geq 0. \quad (3)$$

For ease of exposition, we denote end-of-period assets as $s_t = A_{t+1}$. This new piece of notation is convenient for describing borrowing constraints that are imposed only on consumers who own a house and take out a new mortgage or re-mortgage (that is, increase the value of the outstanding mortgage). The net worth of consumers in periods where they do not re-mortgage can violate

these constraints if house prices or incomes decrease sufficiently.

Home owners can borrow. 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:

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

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

$$s_t \geq -\lambda_w w_t \quad (5)$$

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 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 because it implies the presence of ‘kinks’ in the conditional value functions (particularly at points in the state space where the consumer’s choice involves continuing to own their house).²

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

²By 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

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 do not want 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 subsequent falls 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:

$$m_t = r_t s_{t-1} \quad (6)$$

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 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:

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

The parameters θ and μ are housing preference parameters which determine the utility premium that households derive from owning their home. The nonseparable term for the value of ownership is a proportional scaling of the utility from consumption. When $h_t = 0$, the scaling term has value

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.

³Examples 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).

⁴The 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 et al. (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.

1 and the additive term value zero, and thus utility is only derived from non-durable consumption. When $h = 1$, 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 ($\mu < 0$). We assume that owning a house or a flat gives utility which does not arise from renting the same property. This additional utility is motivated primarily by moral hazard considerations, whereby individuals who own their homes capture the benefit of hidden actions that affect the value of their home, whereas renters do not.

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: uncertainty over earnings and uncertainty over house prices.⁵ We assume that house price and earnings shocks are uncorrelated. This reflects our decision to model earnings uncertainty as arising mainly from idiosyncratic shocks, while house prices reflect aggregate shocks. We are effectively assuming that aggregate income risk is negligible relative to idiosyncratic risk and that there is no correlation between realisations of the house price shock and realisations of idiosyncratic income shocks.⁶

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

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

and a_t is the deterministic growth in earnings over the life-cycle and has a hump shape ($a_t =$

⁵In fact, it is not difficult to add i.i.d. noise in the interest rate. However, results were not sensitive to this additional dimension of uncertainty, and so, we removed it from our final analysis. In equation (1) we implicitly assume that the interest rate is not a stochastic state variable.

⁶If the aim of the paper were to capture business cycle movements, it would be necessary to extend the model to allow aggregate income shocks to correlate with aggregate house price shocks, as in Attanasio et al. (Forthcoming 2011); and perhaps for idiosyncratic uncertainty to be correlated with the house price shock (for example, so that in recessions, when house prices fall, idiosyncratic income uncertainty also goes up).

$a_1t + a_2t^2$).

House prices are assumed to evolve as an AR(1) but in this case the deterministic element reflects upwards drift over time:

$$\ln p_t = d_0 + d_1t + \rho_h \ln p_{t-1} + \varepsilon_t \quad \varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \quad (9)$$

The price of a flat is assumed to be a fixed proportion κ of the price of a house and that this is the same for all households, no matter what their age.

If a household chooses not to own their house, they have to rent. The rental price is tied to the price of a flat, and is given by 1% of the flat price, capped at 50% of household income to reflect the state provided support for housing costs for the very poor in the UK.⁷ An extreme alternative assumption to having rental rates tied to house prices would be to have rental rates fixed and the house price process in equation (9) to be estimated as the relative price of owning versus renting. It would be hard to interpret such a process and the assumption that rental rates are independent of house prices, even in the long run, is implausible. However, to show the difference that this assumption would make, we simulated this alternative set-up and we indicate where this qualitatively affected the results.

Our analysis is partial equilibrium in nature. This means that changes in demand for housing do not affect the price of housing. This can be justified in two ways. Our exercise can be viewed as a building block of a general equilibrium model where we describe how individual consumers react to a given house price process. Individual demands should then be aggregated and interacted with supply to generate equilibrium prices. Alternatively, one can justify our framework by the cost of construction of new housing being pinned down by technology and the wider labor market. This is a reasonable assumption when new land is a negligible input into new housing. A further equilibrium effect may be that the relative price of flats versus houses, κ , should vary over time: for instance κ might vary over the business cycle as the relative demand for houses versus flats changes. We abstract from these business cycle effects to focus on life-cycle behaviour in the steady-state.

⁷With this parametrisation, rent accounts for approximately 15% of the income of renters, on average, and this matches the equivalent expenditure share in the data.

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.

2.2.1 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 the Office of the Deputy Prime Minister (ODPM) national and regional house price series for the UK, years 1969-2000.⁸ 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 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. In practice, there are other dimensions of heterogeneity in housing, such as location, that mean κ should vary in these dimensions, but we abstract from this

⁸We use the series reporting average house prices for all dwellings.

Table 1: Parameter Values

Parameter	Value	Source
House Price Process		
ρ_h	0.94	ODPM
σ_ε^2	0.008	ODPM
d	2.32%	ODPM
κ	0.6	BHPS
p_{22}	4.67	BHPS
rentpar	0.01	
Income Process		
ρ_w	1.0	(By assumption)
$\sigma_{\xi HE}^2$	0.035	BHPS
$\sigma_{\xi LE}^2$	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}	
\bar{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>).

and assume only two types of housing are available.⁹

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 et al. (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

⁹It would be interesting to estimate the house and flat prices to be two independent processes not perfectly correlated. This proved computationally too difficult as it introduced an additional state variable.

their less educated counterparts 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 UK Treasury Bill discount rate in years 1968-1997, which gives a rate of 1.8%.

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 (2010)). This corresponds to a curvature parameter $\gamma = 1.43$ for our within period utility function.

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,¹⁰ and we assume that households have zero housing endowments at age 22.

2.3 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 2000,¹¹ by education group.¹² 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, and also the fraction of home-owners who own

¹⁰These data are discussed in Banks et al. (2002).

¹¹Data come from the years 1991-2000, as years prior to 1991 are affected by the large-scale selling off of local authority housing.

¹²We thus pin down parameters against average home ownership over a (short) period of time, and the exercises considered in the next sections consider the sensitivity of ownership to changes in the model environment or shocks to house prices or incomes. Other authors have applied models to disentangle what might have caused observed changes in the aggregate and/or age-group specific home ownership rates in the US - see Chambers et al. (2009) and Fisher and Gervais (Forthcoming).

houses rather than flats. 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 up to 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 the five moments listed above to pin down four parameters: the transactions cost of buying or selling, F , and the parameters specifying the utility benefit of home ownership, μ , θ and φ . Parameter values from the calibration are summarized in table 2. Table 3 presents the calibration moments.

Table 2: Calibrated Parameter Values

Parameter	Value
θ	0.115
μ	0.26
F	0.05
ϕ	0.9

Table 3: Calibration Statistics

Statistic	High Education		Low Education	
	Data	Model	Data	Model
<i>Ownership rate (percentage)</i>				
Age 26 - 35	65.4	58.4	49.5	48.3
Age 36 - 60	81.1	82.4	62.7	71.5
	Data		Model	
% owner w house	58.2		44.2	

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

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 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 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 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. The value of ϕ implies that the extra utility from ownership is only 10% less for a flat than a house.

2.4 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.¹³ 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 education 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 10.4 (7.5), which is a reasonably good match. For the high education group the simulation numbers are 9.0 (6.3). 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 is uncapped, 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.

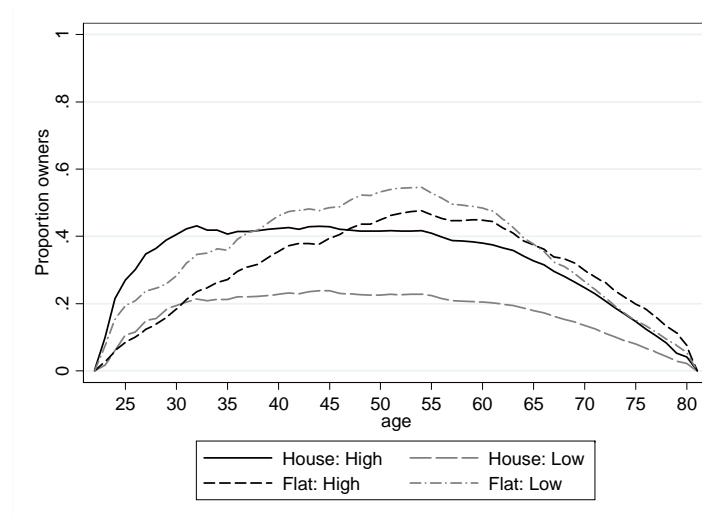


Figure 1: Simulated Home Ownership in the Baseline

3 What Determines Individual Housing Demand?

Throughout this section, we will be looking at the impact of changes in the various parameters of our model on housing and consumption demand. The changes we consider, therefore, are not

¹³For more details on this, see Banks et al. (2005). We are grateful to Gemma Tetlow for help dealing with the ELSA data cited here.

realisations of shocks (which are the focus of the next section) but changes in environment that affect individuals across their entire lifetime and which are fully anticipated. The analysis in this section is divided 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.

There are two different aspects to housing demand that we focus on: demand for ownership of a home 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. However, since home-ownership is discrete, we cannot compute its derivative. Instead we calculate changes in ownership rates for a large number of individuals under our baseline parameters and in an alternative scenario which incorporates the specific change we want to consider. We then calculate “elasticities” as proportional changes in ownership rates relative to the proportional change in the parameter of interest. 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.¹⁴ This elasticity measure captures the effect of an anticipated increase in lifetime income.

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 ($\epsilon_{\lambda Y}$), for the high and low educated individuals. The tables report results for each of the three elasticities, for each of the two education groups and for three age groups: 26 to 35, 36 to 50 and 51 to 60. The tables also reports home ownership rates (labeled “Rate”) in our baseline run to enable the calculation of quantity effects from the

¹⁴Apart 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.

elasticity measures.

Table 4: Elasticities for ownership: Factors relating to income

	High Education				Low Education			
	ϵ_Y	ϵ_{vY}	$\epsilon_{\lambda Y}$	Rate	ϵ_Y	ϵ_{vY}	$\epsilon_{\lambda Y}$	Rate
<i>Ownership elasticities</i>								
Age 26 - 35	0.86	-0.14	0.08	0.58	0.42	-0.03	0.07	0.48
Age 36 - 50	0.24	-0.08	0.02	0.80	0.33	-0.10	0.03	0.70
Age 51 - 60	0.18	-0.11	0.01	0.86	0.39	-0.19	0.02	0.74
Lifetime	0.44	-0.11	0.05	0.62	0.47	-0.13	0.07	0.51
<i>Quantity elasticities</i>								
Age 26 - 35	1.21	-0.14	0.25	0.51	0.91	-0.08	0.17	0.36
Age 36 - 50	0.51	-0.07	0.10	0.65	0.65	-0.09	0.07	0.51
Age 51 - 60	0.44	-0.11	0.06	0.68	0.67	-0.15	0.04	0.53
Lifetime	0.72	-0.10	0.14	0.50	0.80	-0.11	0.11	0.38

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, to income uncertainty and to changes in the constraints relating to income. Higher income, lower uncertainty or a loosening of the income constraint leads to greater home ownership among the young but makes less difference to the older groups. This message remains the same if we consider the quantity of housing bought rather than just the effect on ownership, and indeed 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.¹⁵ Another is to consider separate elasticities

¹⁵In a version of our model with only a single property type, with value (utility and price) equal to those of the flat in the current model, 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. However, such a “one-size” model does

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

	High Education				Low Education			
	η_Y	η_{vY}	$\eta_{\lambda Y}$	Rate	η_Y	η_{vY}	$\eta_{\lambda Y}$	Rate
<i>Elasticities for Ownership of a Flat</i>								
Age 26 - 35	-1.45	-0.13	-1.02	0.19	-1.07	0.12	-0.21	0.30
Age 36 - 50	-0.93	-0.09	-0.35	0.38	-0.53	-0.14	-0.09	0.47
Age 51 - 60	-0.77	-0.10	-0.20	0.46	-0.32	-0.27	-0.03	0.52
Lifetime	-0.73	-0.16	-0.31	0.30	-0.43	-0.19	-0.05	0.35
<i>Elasticities for Ownership of a House</i>								
Age 26 - 35	2.00	-0.14	0.62	0.39	2.82	-0.28	0.52	0.19
Age 36 - 50	1.29	-0.07	0.34	0.42	2.12	-0.04	0.27	0.23
Age 51 - 60	1.28	-0.11	0.24	0.40	2.07	0.03	0.14	0.22
Lifetime	1.52	-0.07	0.39	0.32	2.31	-0.01	0.31	0.17

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	1.07	-0.25	0.01	1.08	-0.26	0.00
Age 36 - 50	1.01	-0.08	0.00	1.03	-0.11	0.00
Age 51 - 60	1.00	0.00	0.00	1.00	-0.04	0.01
Lifetime	1.00	-0.04	0.00	1.01	-0.07	0.01

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’. 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 constraint makes buying houses more feasible. As housing becomes more accessible, more people upgrade from a flat to a house than from not owning to owning a flat, and thus the demand for flats falls.

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 9 show the extent that home ownership and consumption respond to an increase in the level of house prices, 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.

Table 7 shows that ownership rates rise, especially among the young, when house prices rise. On the other hand, the quantity of housing bought falls. Quantity falling while ownership rises reflects a reduction in the ownership of houses but an increase in the ownership of flats.¹⁶ This pattern, which makes flats appear to be a Giffen good, is evident from table 8.¹⁷ The substitution out of houses is a natural response to the higher house price. It is less obvious why the increase in the price of housing (of houses *and* flats) does not also translate into lower overall ownership rates. There are two reasons. First, the increase in the house price implies increased rents and so discourages renting. If instead, the rental price was independent of the house price and was constant, homeownership rates would be much more likely to fall along with the quantity of housing

not allow us to show that quantity elasticities are generally amplified (rather than dampened) relative to ownership elasticities, nor to consider separate flat and house markets.

¹⁶In the version of our model with only a single property type, there cannot be this difference in sign between quantity and ownership elasticities. In that version of the model, the ownership rate rose for the young and fell for the old when the house price rose. Over the whole life time, the ownership elasticity was slightly negative. Full details are available on request.

¹⁷This feature is reliant on the fact that a change in the price of housing has the same proportional effect on the prices of flats and houses.

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

	High Education					Low Education				
	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}	Rate	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}	Rate
<i>Ownership elasticities</i>										
Age 26 - 35	0.47	-0.27	0.04	-0.13	0.58	0.67	-0.21	-0.01	-0.18	0.48
Age 36 - 50	0.14	-0.12	0.00	-0.01	0.80	0.04	-0.11	0.00	-0.07	0.70
Age 51 - 60	-0.01	-0.06	0.00	-0.04	0.86	-0.19	-0.05	-0.00	-0.06	0.74
Lifetime	0.04	-0.10	0.02	-0.07	0.62	-0.01	-0.08	-0.02	-0.09	0.51
<i>Quantity elasticities</i>										
Age 26 - 35	-0.23	-0.04	0.01	-0.36	0.51	0.00	0.06	-0.01	-0.43	0.36
Age 36 - 50	-0.37	0.06	-0.02	-0.23	0.65	-0.40	0.11	-0.00	-0.29	0.51
Age 51 - 60	-0.37	0.06	-0.02	-0.21	0.68	-0.50	0.09	-0.01	-0.24	0.53
Lifetime	-0.40	0.04	-0.00	-0.24	0.50	-0.40	0.09	-0.02	-0.29	0.38

bought.¹⁸ Second, a higher house price makes house ownership financially less feasible for some. Some individuals respond to this by buying flats when young rather than remaining renters and saving up for the downpayment to buy a house as their first property owned.

Tables 7 and 8 also show that, as with the responses to parameters relating to 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 uncertainty about the house price increases, as the downpayment requirement becomes greater and as fixed costs go up.

Table 9 reports the consumption response to changes in the parameters of the house price process. All effects are economically very small. The largest (in absolute value) effects are observed for young consumers with low education in response to a shock to house prices. Even for this group, however, the elasticity of consumption is only -0.05.

¹⁸To explore this effect, we solve, calibrate and simulate a model where there is no explicit cost of renting and so the house price process can be interpreted as the stochastic mark up of the house price over the rental price. In this model, ownership rates and the quantity owned falls as house prices rise. The details of these results are available on request.

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}	Rate	ϵ_{HP}	ϵ_{vHP}	$\epsilon_{\lambda HP}$	ϵ_{FC}	Rate
<i>Elasticities for Ownership of a Flat</i>										
Age 26 - 35	5.07	-1.81	0.24	1.37	0.19	2.70	-1.06	0.00	0.58	0.30
Age 36 - 50	2.34	-0.87	0.09	0.92	0.38	1.23	-0.70	0.00	0.53	0.47
Age 51 - 60	1.31	-0.49	0.07	0.57	0.46	0.59	-0.40	0.00	0.41	0.52
Lifetime	1.88	-0.69	0.12	0.63	0.30	1.06	-0.56	-0.02	0.44	0.35
<i>Elasticities for Ownership of a House</i>										
Age 26 - 35	-1.81	0.49	-0.06	-0.87	0.39	-2.60	1.14	-0.02	-1.40	0.19
Age 36 - 50	-1.82	0.57	-0.07	-0.85	0.42	-2.43	1.12	-0.01	-1.30	0.23
Age 51 - 60	-1.53	0.43	-0.08	-0.74	0.40	-2.04	0.81	-0.02	-1.16	0.22
Lifetime	-1.67	0.44	-0.07	-0.72	0.32	-2.22	0.90	-0.01	-1.18	0.17

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	-0.03	0.00	0.00	-0.02	-0.05	0.00	-0.00	-0.02
Age 36 - 50	-0.02	0.01	0.00	-0.01	-0.04	0.01	-0.00	-0.02
Age 51 - 60	-0.01	0.01	0.00	-0.02	-0.03	0.02	-0.00	-0.02
Lifetime	0.01	0.01	0.00	-0.02	-0.01	0.02	-0.01	-0.02

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.¹⁹ 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%. With a lower fixed cost there is a clear move towards house ownership. Table 10 shows that this arises because a lower fixed cost makes households 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.

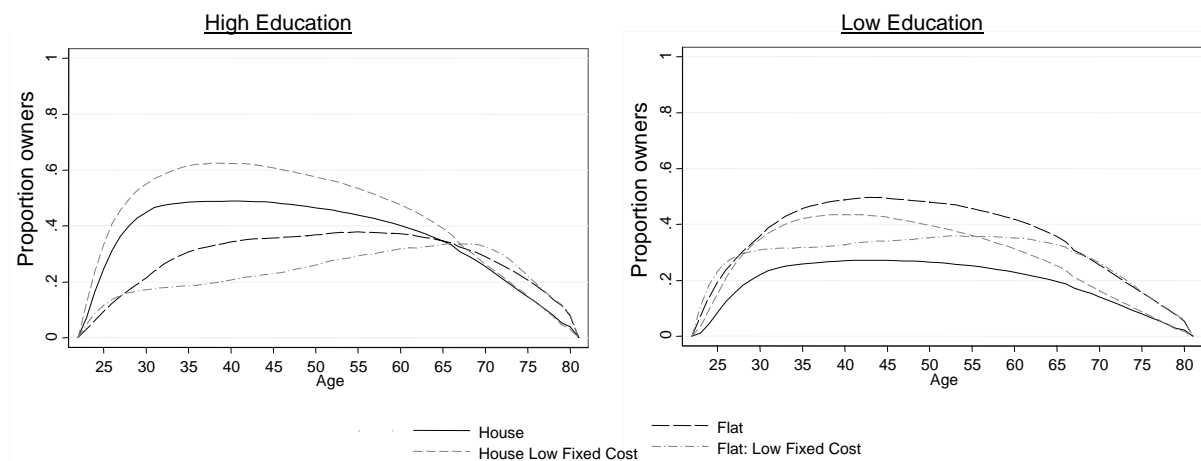


Figure 2: Response of Ownership to Change in Fixed Costs

We separate out the wealth and substitution effects by first running the model with a low fixed

¹⁹In the U.K., by stamp duty land tax.

cost that is known in advance (column (c) in table 10). The impact of this low fixed cost relative 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 they pay only the low fixed cost (column (b)). There is therefore no substitution effect associated with the lower fixed cost.²⁰ 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”, the solution for the relevant policy functions is that with the higher fixed costs and 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

²⁰The aim of our procedure is to isolate an ex-post wealth effect. An alternative way to separate the wealth and substitution effects would be to rebate the extra income through an anticipated lump sum payment. The difficulty with this is that in a dynamic model with liquidity constraints, such a lump-sum payment will affect individuals’ attitudes towards risk and their borrowing: the anticipation of such a payment acts to reduce uncertainty and changes the ability to borrow in a way that would not happen were the income to come through reduced transactions costs. There is also a redistributive issue: since the allocation of the lump-sum is decided ex-ante, it cannot depend on choices about housing, and so some individuals will gain from the lump-sum, particularly if they are individuals who would never have bought a house. In this way, the lump-sum acts as if it is providing insurance against the bad realizations which mean home ownership is not possible.

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

	(a.)	(b.)	(c.)
High Education	Calibrated Run	Low FC as “surprise”	Low FC
E(Consumption)	1.416	1.428	1.428
E(ownership rate)	0.63	0.63	0.65
E(flat ownership)	0.28	0.28	0.23
E(house ownership)	0.35	0.35	0.42
E(# Transacts in life)	2.72	2.74	3.35
% trade up at least once	13.4	13.5	20.9
% downsize at least once	26.3	26.3	47.2
Equivalent variation, baseline (c):			
Initial asset measure	0.230	0.040	0
Consumption equiv (%)	1.047	0.143	0
Equivalent variation, baseline (b):			
Consumption equiv (%)	0.909	0	N/A
Low Education	(a.) Calibrated Run	(b.) Low FC as “surprise”	(c.) Low FC
E(Consumption)	0.980	0.989	0.991
E(ownership rate)	0.54	0.54	0.56
E(flat ownership)	0.35	0.35	0.28
E(house ownership)	0.19	0.19	0.28
E(# Transacts in life)	2.43	2.44	3.13
% trade up at least once	13.0	13.1	23.3
% downsize at least once	16.0	15.8	35.5
Equivalent variation, baseline (c):			
Initial asset measure	0.157	0.027	0
Consumption equiv (%)	1.241	0.220	0
Equivalent variation, baseline (b):			
Consumption equiv (%)	0.825	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.

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 ‘trading up’) at some age is almost doubled compared to in the high fixed cost case, while the likelihood of the converse change in ownership (‘downsizing’) approximately doubles. 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 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.²¹ 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

²¹Since income is stochastic, the experiment is that the realisation of income is 1% higher than a draw from the distribution would imply.

and low educated. Figure 4 shows the effect on consumption.²² 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.




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

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 decline, with the largest reduction in home-ownership (of around 1% for the high education group,

²²Each of the lines plotted in figures (3) and (4) shows the percentage difference in, respectively, 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.

and 1.2% to 1.5% for the low education group) experienced around five years after the shock. With the exception of the shock hitting the less educated at age 30, 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 (appendix figure 7 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 if the shock hits at age 40 or 50, but decreases consumption on average when the shock hits at age 30. 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. The decline for those who experience the shock when young reflects a need to save more for a downpayment, and the fact that the cost of renting increases with the house price. This pattern across age-groups confirms the intuition behind the empirical analyses of Attanasio and Weber (1994) and Attanasio et al. (2009, 2005) investigating why house price shocks have been strongly correlated with consumption growth in the U.K. over the past 30 years, an intuition which is investigated more thoroughly in Attanasio et al. (Forthcoming 2011).

The effects on consumption taper 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 40, particularly in the low education group. Some in 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 remortgaging or taking out mortgages and bringing forward some 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 home ownership and on consumption of a 1% upwards shock

to income at age 30, 40 and 50.²³ As with the house price shocks just discussed, 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 persistent since income follows a random walk and so shocks are permanent. 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

²³Each of the lines plotted in figures (5) and (6) shows the percentage difference in, respectively, the home ownership rate or average consumption in two simulation runs between which the only exogenous change is to the persistent income shock at the given age, such that the level of 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.

income only accrues while individuals are working. Older individuals enjoy the income shock for a smaller proportion of their 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 have presented a life-cycle model of households choosing how much to consume or save, and whether to rent or own a flat or a house, in each period of life. We built our model to capture realistically the main choices faced by consumers when faced with earnings and house price uncertainty and by mortgage related borrowing constraints. We calibrated the important parameters of the model to reproduce some life cycle facts of home ownership in the UK.

Simulations of the model show how sensitive individual demand for housing is to the parameters of the house price process, and the income process, and also to the tightness of mortgage related borrowing constraints. The level of house prices and earnings have marked effects on behaviour. When incomes go up the demand for housing shows the expected increase. More surprisingly, when the price of housing (i.e. of houses and flats) goes up, this leads to a fall in demand for houses, but to *an increase* in the demand for flats. This somewhat unusual behaviour is explained by the interaction of demand for rented and owned housing services. As we mentioned in the

introduction, we see our exercise as a first step towards the construction of a general equilibrium model that, considering simultaneously demand and supply of housing, could generate house prices endogenously. In this respect, the unusual responses of aggregate housing demand to house prices is interesting because it could imply extremely volatile price levels, even when the ‘fundamentals’ of the problem do not move much.

The level of fixed costs in the housing market has 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 Effect of House Prices shock on subsequent expected House Prices




Figure 7: Effect of HP shock at 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.²⁴ 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.

As is standard for these dynamic problems, the solution for consumption and home-ownership is found recursively from the last period of life, T . 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 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 (see (Floden, 2007)). A 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 the approximation.

For the runs presented in this draft of the paper, we have used 100 nodes in each 'conditional'

²⁴The 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.

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 1 hour on a 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:

$$\tau = (\overline{U}_0/\underline{U}_0)^{1/[1-\gamma]} \quad (10)$$

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 doubled 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, 20–50.
- ATTANASIO, O. P., A. LEICESTER, AND M. WAKEFIELD (Forthcoming 2011): “Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK,” *Journal of the European Economic Association*, -, -.
- ATTANASIO, O. P. AND M. WAKEFIELD (2010): “The Effects on Consumption and Saving of Taxing Asset Returns,” in *Dimensions of Tax Design: The Mirrlees Review*, ed. by S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba, Oxford University Press, chap. 7, 675–736.
- 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, 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., T. F. CROSSLEY, AND M. WAKEFIELD (2011): “House Prices and Home Ownership: a Cohort Analysis,” *Paper prepared for the Royal Economic Society annual conference, 2011*.
- 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.
- CHAMBERS, M., C. GARRIGA, AND D. E. SCHLAGENHAUF (2009): “Accounting for Changes in the Homeownership Rate,” *International Economic Review*, 50, 677–726.
- COCCO, J. (2005): “Portfolio Choice in the Presence of Housing,” *Review of Financial Studies*.
- FISHER, J. D. M. AND M. GERVAIS (Forthcoming): “Why has home ownership fallen among the young?” *International Economic Review*, -, -.
- 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>).
- KOPECKY, K. AND R. SUEN (2010): “Finite State Markov-chain Approximations to Highly Persistent Processes,” *Review of Economic Dynamics*, 13, 701–714.
- 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*, 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, 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, 705–714.

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

ROUWENHORST, K. G. (1995): “Utility analysis and the consumption function: an interpretation of cross-section data,” in *Frontiers in Business Cycle Research*, ed. by T. F. Cooley, Princeton University Press, 294–330.

TAUCHEN, G. (1986): “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20.

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.