



IFS

WHY DO HOME OWNERS WORK LONGER HOURS?

Renata Bottazzi
Hamish Low
Matthew Wakefield



THE INSTITUTE FOR FISCAL STUDIES
WP10/07

Why do home owners work longer hours?*

Renata Bottazzi[†] Hamish Low[‡] Matthew Wakefield[§]

July 11, 2007

Abstract

This paper uses a structural model to address the question of why home-owners with large mortgage debt work longer hours than those without such debt. We consider whether this is due to lower net wealth or to capital market imperfections, including mortgage constraints that depend on current earnings and, therefore, labour supply choices. We show that the need to meet current mortgage commitments can generate the observed correlation, and this impact of current commitments arises from the institutional borrowing constraints. We also show that labour supply as a function of household debt is highly non-linear: those with greater debt are more likely to face binding borrowing constraints and their labour supply is more variable.

JEL Codes: J22, D91.

Keywords: Housing, labour supply, life-cycle models, borrowing constraints.

*This paper arose out of joint work with Orazio Attanasio and Lars Nesheim and we are very grateful to both for many helpful discussions. For helpful comments, we also thank Chris Carroll, Ken Judd and seminar participants at the Finance and Consumption Workshop at the European University Institute; at CAM, University of Copenhagen; at the COST meeting at CREST, Paris, at the Department of Economics, University of Reading, at the Netspar Workshop at the University of Utrecht, and at the Department of Economics, University of Bologna. The British Household Panel Survey (BHPS) data used in this paper were made available through the UK Data Archive (<http://www.data-archive.ac.uk/>). Neither the original collectors of the data nor the Archive bear any responsibility for the analyses or interpretations of the data that are presented here. The paper is part of the program of research of the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS. Financial support from the ESRC is gratefully acknowledged (Grant no. RES-000-23-0283).

[†]University of Bologna and Institute for Fiscal Studies.

[‡]University of Cambridge and Institute for Fiscal Studies.

[§]University College London and Institute for Fiscal Studies.

1 Introduction

The aim of this paper is to understand why home-owners with large mortgage debt work longer hours than those without such debt. Interpreting this observed correlation is problematic because labour supply and debt holdings are jointly determined, and these decisions are made in an environment with imperfect capital markets. We use a structural model of the joint labour supply / housing choice to disentangle three broad reasons which can explain this observed correlation: income effects, capital market constraints and unobserved heterogeneity. Further, we use our model to explore the effect of changing capital market restrictions on labour supply and debt holdings.

For many households, housing wealth comprises a large fraction of total household wealth, but it is also often offset by substantial mortgage debt. For example, in the British Household Panel Survey (BHPS) for 2000, among the 57% of families who owned their homes, housing assets represented approximately 83% of all non-pension assets but more than a quarter of this gross housing wealth was offset by outstanding mortgage debt.¹ Among homeowners aged less than 45, housing assets represented 89% of non-pension assets but the value of outstanding mortgage debt was equal to more than half of gross housing wealth. The interaction between these mortgage commitments and labour supply behaviour has been studied by Fortin (1995), Del Boca and Lusardi (2003), and Bottazzi (2007) for home-owners in Canada, Italy, and the UK respectively. They each show that women in households with greater mortgage commitments are more likely to participate in the labour force and/or work longer hours.

This correlation can be explained in a number of ways. Conditioning on the house price, greater mortgage commitments suggests a lower level of wealth and this will itself lead to a wealth effect of higher labour supply. Further, mortgage repayments can be seen as committed expenditure (see Chetty and Szeidl, 2007): in the presence of fixed costs of buying and selling a home, households will adjust in other dimensions than home-ownership in order to meet debt repayments. Additionally, the ability to borrow is constrained by household income. For example, in the UK, mortgage debt is usually constrained to be less than a fixed multiple of household income, and so households may work longer hours in order to overcome this borrowing constraint. A second borrowing constraint is the downpayment requirement on home-buyers, which may induce households to increase labour supply and save their extra earnings to meet such a requirement. The

¹BHPS numbers cited in this paragraph are the result of authors' calculations based on data that were set up for and manipulated by Banks, Smith and Wakefield (2002).

correlation between debt and labour supply may also be driven by unobserved heterogeneity in preferences: some individuals may have a high (unobserved) utility from housing, and a low utility cost of work, or the marginal disutility of work may be decreasing in home-ownership.

This brief discussion serves to highlight the complex interactions between the housing, debt and labour supply choices. To address these, both Fortin (1995) and Bottazzi (2007) test for the endogeneity of mortgage payments to the labour supply decision, the latter using panel data to also control for unobserved heterogeneity, but neither finds this endogeneity to be important. Similarly, Del Boca and Lusardi (2003) find only a small effect of the propensity to participate in the labour market on the likelihood of having a mortgage in a simultaneous equation model. In this paper, we adopt an alternative strategy of modelling the joint labour supply-housing decision in a structural model to examine the nature of the interaction between these different decisions.

We develop a life-cycle model where households decide whether or not to buy a house, how much to borrow or save and how much labour to supply in each period. Households make these decisions facing uncertainty about wages and about house prices. In modeling these decisions, we pay particular attention to the financial instruments that households have access to. We model the details of the debt instruments available to the households in our model after the type of instruments available in the UK in the 1990s. We calibrate our model to match the life-cycle profile of home ownership in UK data.

Campbell and Hercowitz (2004) also study a dynamic model with housing, consumption, and labour supply. In an infinite horizon framework, they show that introducing a downpayment constraint and allowing for accelerated repayment leads to labour supply being correlated with debt. In our framework, we focus instead on the life-cycle characteristics of home-ownership and labour supply, and we model income-related borrowing constraints. We believe these income related constraints to be important because labour supply choices affect income and so have a direct impact on borrowing possibilities. We also explicitly introduce house price uncertainty.

Our conclusions on the central question of why we observe the correlation between labour supply and debt emphasise three points: first, greater current mortgage commitments lead to greater labour supply, even when controlling for wealth; second, the solution of our life-cycle model shows that the labour supply of households with high debt is very sensitive to changes in the amount of debt and changes in family income, but that this sensitivity lessens if we look at households less in debt and so is not apparent if we average across

all households; third, tightening borrowing constraints leads to substantial effects on the timing of home-ownership over the lifetime, delaying purchases, but such tightening leads only to small effects on labour supply on average. These points suggest that for most households, wealth effects and current commitments are driving the correlation between labour supply and debt, but that for some highly indebted households, capital constraints imply there will be a stronger correlation between debt and labour supply.

The rest of the paper is structured as follows. Section 2 reports life-cycle patterns of home ownership and labour supply in the data. Section 3 presents the structural model, discusses the calibration strategy and shows the calibrated solution. Section 4 addresses the question of how labour supply and home ownership interact and the question of how changes to the capital market affect home ownership and labour supply. Section 5 concludes.

2 Empirical background

In this section we present empirical lifecycle profiles of home ownership and labour supply, and of the relationship between them, for households in the U.K. These profiles both put our discussion in context, and also underlie the calibration of our model in the next section.

Home ownership Home ownership rates are relatively high in the UK, with around 4 in 5 couples aged 26 - 60 owning their home. To show how this varies over the life cycle, Figure 1 reports home-ownership rates by age and cohort for couples. Whereas in the remainder of the paper we use panel data from the British Household Panel Survey (BHPS) which has been available since 1991, in this section we want to display longer time series and so draw home ownership and labour supply profiles from repeated cross-sections of the Family Expenditure Survey (FES).² To create a pseudo panel from these data, groups are defined by date-of-birth and education,³ where, as throughout this paper, education groups are defined on the basis of academic achievement of the household head, and “high education” means that the household head remained in education beyond the compulsory period.⁴

²The FES is an annual cross section of around 7,000 households, who record a two-week diary of their spending and information about recent purchases of durables and/or expensive items. The survey also provides a detailed information on demographic factors including homeownership, education, family structure and income, and we use data for the years 1978 - 2002.

³Since the FES is a series of cross-sections rather than a true panel, we construct home-ownership rates by averaging across individuals from a particular cohort, separately for each year of data. This generates a pseudo-panel showing how a particular cohort behaves over time (Browning et al., 1985). Our analysis focuses on couples and so our results should be interpreted recognising that we are conditioning on marital status which is not constant over the life-cycle.

⁴The age up to which schooling was compulsory was 14 until 1946, 15 until 1972, and 16 afterwards.

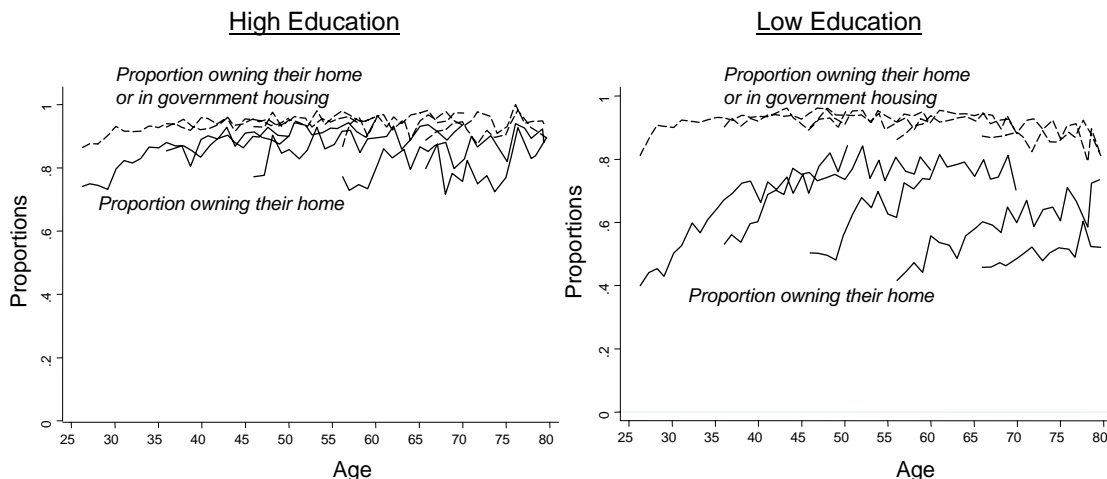


Figure 1: Home-Ownership over the Life-Cycle

Figure 1 shows increases in home-ownership with age (the solid lines). This increase is particularly marked for the low educated, and each of the low education cohorts shows an increase in home ownership during the early years of our data (the 1980s).⁵ This increase over time is not evident when we plot the combined proportion of households owning or in government housing⁶ (the dashed lines), indicating a transfer over time from government housing to owner occupation. In fact, in the 1980s many tenants who rented government housing were given the opportunity to buy their homes, often at below market prices (the so-called “right-to-buy” policy). Thus the increase in home ownership rates across the low education group during the 1980s was largely a time rather than an age effect. For this reason, we calibrate our model to data on home ownership at different ages during the 1990s, which we take to be representative of life cycle profiles. As reported in the calibration section 3.2, these data show home ownership increasing during early working age to a peak, during middle-age, of around 75% and 90%, respectively for the low and the high education groups.

Female labour supply Using the same pseudo-panel underlying Figure 1, Figure 2 shows how average hours of work of women in couples vary over the life cycle, where non-participants have hours set to zero. The figures show the well-known decline in labour supply around the age of fertility across cohorts.⁷ Since these

⁵This pattern is even clearer if we plot home ownership against year, which we omit for the sake of brevity.

⁶Socially assisted rental housing in the UK is provided by local authorities and by housing association. We refer to them jointly as “government housing.”

⁷Further analysis of the data (not shown) suggests that this fall in hours of work is largely due to a fall in the participation rate.

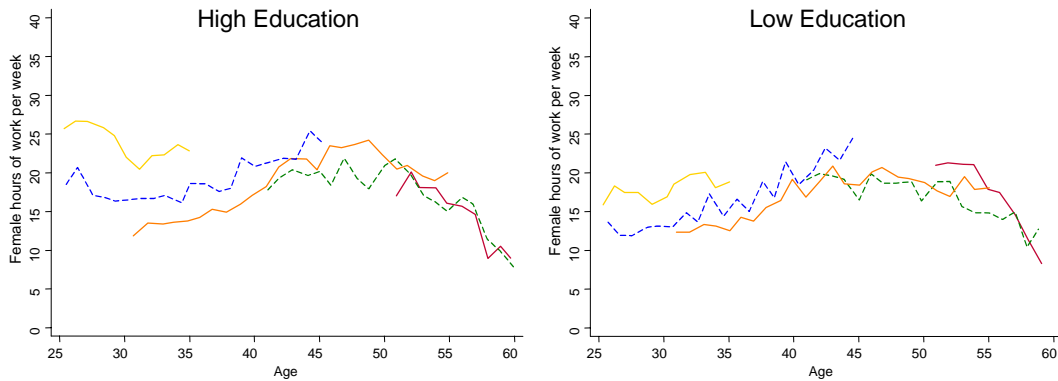


Figure 2: Female Hours of Work Over the Life-Cycle by Cohort

figures report hours of work averaged over all women in couples, they mask the considerable heterogeneity in labour supply among women of the same age that is related not only to family circumstances but also to housing decisions and mortgage commitments.

Labour supply and home ownership Evidence of how this heterogeneity in female labour supply in the UK is explained by mortgage commitments comes from Bottazzi (2007), and is summarised by Figure 3. The picture is based on BHPS data⁸ and shows that women in households that own their homes, and particularly those with greater mortgage commitments, supply more labour.⁹ However, as discussed in the introduction, we cannot treat the extent of debt or home-ownership as exogenous. Instead, we resort to calibration of a structural model of housing and labour supply which enables us to model explicitly the labour supply decision when households are making a decision about housing.

3 Life-Cycle Model of Housing and Labour Supply

In order to model the interaction between housing choices, mortgage debt and labour supply, we start from a standard model of lifecycle consumption and labour supply in a stochastic dynamic environment. We add to this model several features that capture the complexity of the consumer decision environment with regard to housing and debt choices.

With regard to housing, we assume that households make a discrete choice to own a house or not. This

⁸The BHPS is currently available for 1991 – 2004 and its original (representative) sample tracks around 10,000 adults in 5,000 households. The survey collects information on a wide range of topics, including detailed questions on income, employment, household composition, education and, importantly, housing choices.

⁹The picture is based on a sample of women whose partners work full time and a household is said to have a high mortgage commitment if its mortgage repayments are in the top two thirds of the sample distribution.

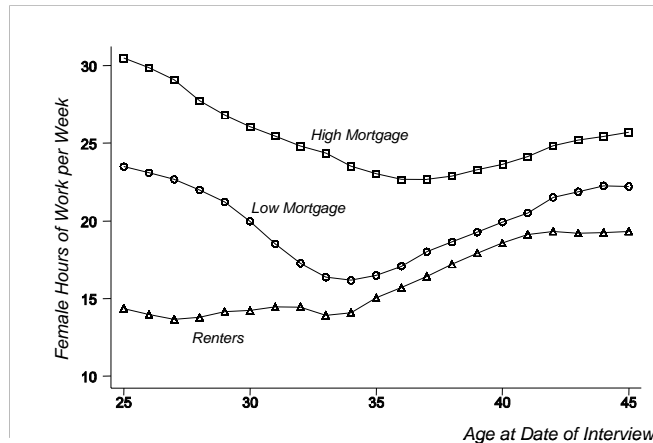


Figure 3: Female Hours of Work by Housing and Mortgage Status (BHPS)

captures the lumpiness inherent in housing market choices. In addition, we assume there are transaction costs associated with changing housing market status.

Since our focus is particularly on mortgage debt, we assume in the model that households without collateral cannot borrow.¹⁰ On the other hand, we model the mortgage market as accurately as we can, particularly in relation to the constraints on borrowing which exist at the time that the mortgage is taken out (whether this is when a house is bought or when a house is remortgaged). The explicit constraints on mortgage borrowing we impose are a downpayment constraint, which specifies that households can borrow only a fraction of the value of the house, and an income-related constraint, such that households can borrow only up to a multiple of household income. On the repayment side, we impose that, in order to avoid selling or remortgaging the house, households must repay at least the interest on their mortgage debt. Given these constraints, households remortgage for two reasons: when it is optimal for them or when they do not meet the minimum mortgage repayment.

We believe that two features of the constraints that we impose are particularly important and, to our knowledge, represent an advance relative to other papers in the literature. First, the income-related constraint is endogenous, i.e. it depends on the current choice of labour supply, rather than on an exogenous income process (an example of this latter setup is Campbell and Hercowitz, 2004). This seems to us particularly relevant to our question since the income-related constraint can be relaxed by increasing labour supply. Second, in line with practice in the housing market, we impose the mortgage constraints only when

¹⁰We could in principle relax this borrowing constraint, and introduce a difference between collateralised and non-collateralised debt through interest rate differences.

the mortgage is taken out (ie when the house is bought or remortgaged), rather than whenever the household owns a house regardless of remortgaging. Imposing the constraint regardless of remortgaging is computationally convenient and typical in the literature (see, for example, Li and Yao, forthcoming, Ortalo-Magne and Rady, 2006, Cocco, 2005, and Campbell and Hercowitz, 2004). These features of the model make the set of constraints dependent on housing and on labour market behaviour. This makes the problem non-convex and complicates the analysis considerably. We explain our solution method in the appendix.

3.1 Model specification

We now specify the model in detail. A unitary household lives for T periods. In every period $t \leq T$, the household maximises utility by choosing consumption, c_t , and housing $h_t \in \{0, 1\}$. In addition, prior to retirement at age $R < T$, the household chooses female labour supply $l_t \in [0, 1]$. After retirement at age R , we assume that $l_t = 0$ and that income is not subject to risk.

For ease of exposition, we distinguish between beginning-of-period assets A_t and debt held at the end-of-period D_t . Mortgage market constraints place restrictions on D_t . The household value function in period $t < R$ is given by

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

where, in addition to the variables already defined, the variables (p_t, w_t, x_t) are the price of housing, the female wage, and male earnings, and the values of these stochastic outcomes are realised at the beginning of the period. $u(\cdot)$ is the within-period utility function, β is the discount factor and E_t denotes expectations at time t over the three stochastic outcomes. The time series evolution of the stochastic processes and the form of the utility function are described at the end of this section.

The value function is subject to a per period budget constraint:

$$A_{t+1} = (1 + r_{t+1}) \left\{ \begin{array}{l} A_t + x_t + w_t l_t + p_t h_{t-1} \\ -c_t - p_t h_t - F p_t |h_t - h_{t-1}| \end{array} \right\} \quad (1)$$

where the variable r_{t+1} is the interest rate on the liquid asset or on debt, which is assumed to be constant,¹¹ and F is the transaction cost of selling or buying a house, which is a proportion of the house price. In addition, optimisation is subject to constraints on borrowing, which can be described by considering three separate scenarios:

¹¹A constant interest rate implies that the current minimum mortgage repayment to avoid remortgaging is pre-determined.

1. If the household owns their house at the end of period t (*ie.* $h_t = 1$) and has either bought their house in period t or remortgaged in period t , debt is constrained by

$$D_t \leq \min [\lambda_y (w_t l_t + x_t), \lambda_h p_t h_t]. \quad (2)$$

2. If the household owns their house at the end of period t (*ie.* $h_t = 1$) and is not remortgaging, debt is constrained by

$$D_t \leq D_{t-1}. \quad (3)$$

3. If the household does not own their house at the end of period t , debt is constrained by

$$D_t \leq 0 \quad (4)$$

Finally, we impose the terminal condition $A_{T+1} \geq 0$.

Equation (1) is the per period budget constraint: total resources equal beginning-of-period assets A_t plus male and female earnings, x_t and $w_t l_t$, plus housing wealth $p_t h_{t-1}$. Total expenditure equals consumption c_t , plus the resources allocated to home ownership, $p_t h_t$, plus transaction costs $F p_t$ whenever home ownership status changes. If $h_{t-1} = h_t = 1$, all the terms in the budget constraint that involve housing cancel out.

The second set of constraints (equations 2, 3 and 4) are the liquidity constraints. If $h_t = 1$, then if the consumer wants to borrow more than the previous period (either because they are taking out a new mortgage or because they have an existing mortgage but want to re-mortgage and borrow more), then constraint (2) imposes that the amount they can borrow is limited to (the smaller of) λ_h times the value of the house and λ_y times current income. The value $(1 - \lambda_h)$ can be thought of as a downpayment requirement. If the household wishes to continue owning and does not repay at least the interest on their outstanding debt, as required by constraint (3), they have to remortgage and satisfy the two constraints in equation (2).

The specification of marginal utility becoming infinite at zero consumption means that the terminal condition prevents households borrowing more than they can repay with certainty. The terminal condition therefore translates into an implicit additional borrowing constraint.

The within period utility is specified as

$$u(c, h, l) = e^{\theta h_t} \frac{\left[c^{-\rho} + \eta (1-l)^{-\rho} \right]^{-\frac{1-\gamma}{\rho}}}{1-\gamma} + \mu h_t \quad (5)$$

This specification contains a Constant Elasticity of Substitution (CES) function between consumption and leisure embedded in a Constant Relative Risk Aversion function. This is supplemented by multiplicative and additive terms that capture the value of home ownership. As we discuss in Section 3.2.2, having these two terms allows us to consider housing as a luxury or a necessity and to allow consumption to be a complement or substitute for consumption and leisure, depending on the outcome of the calibration. If there were no additive term (ie if $\mu = 0$), then $\frac{\partial u}{\partial h} > 0$ would require $\theta < 0$ when $\gamma > 1$ and this would restrict consumption and housing to be substitutes ($\frac{\partial^2 u}{\partial c \partial h} < 0$). We use the more general specification to avoid this apriori restriction.

We assume that the house price is an AR(1) process, with aggregate shocks. That is

$$\ln p_t = d_t + \zeta_t \quad \text{where} \quad \zeta_t = \rho_h \zeta_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \quad (6)$$

where d_t is the deterministic trend, assumed to be linear.

We assume that the female wage process is a random walk¹² with idiosyncratic shocks:

$$\ln w_t = a_t + v_t \quad \text{where} \quad v_t = v_{t-1} + \xi_t, \quad \xi_t \sim N\left(-\frac{\sigma_\xi^2}{2}, \sigma_\xi^2\right) \quad (7)$$

where a_t is the deterministic growth in wages, assumed to be quadratic: $a_t = a_1 t + a_2 t^2$. Similarly, the process for male wages follows a random walk with idiosyncratic shocks:

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

where a_t^h is the deterministic growth in wages, assumed to be quadratic: $a_t = a_1 t + a_2 t^2$. x_t is defined as $\exp(w_t^h)$ times normal (full-time) hours.¹³

3.2 Parameter Selection

We use our model to calibrate the parameters governing the preference for home ownership and the transaction cost. The remaining parameters of our model are external: estimated outside the model either by us or in the existing literature, or chosen to match institutional features. External parameters are summarised in Table 1 and discussed below. Calibrated parameters are discussed in section 3.2.2.

¹²This assumption of a random walk for shocks to female and male wages follows MaCurdy (1982).

¹³We treat male earnings as exogenous because nearly all male workers in the sample underlying our calibration work full time. With exogenous earnings, the assumption that male labour supply does not enter the utility function is equivalent to introducing it into the utility function through an additive term.

Table 1: Estimated/Fixed Parameter Values

Preference Parameters			
γ	1.8	Attanasio and Weber (1995)	
η	0.25	and Pistaferri (2003)	
ρ	-0.2		
β	1.02^{-1}	fixed	
Credit Market Parameters			
λ_y	3.0	fixed	
λ_h	0.9	fixed	
\bar{r}	0.018	B.o.E	
Wage Processes			
	High Education	Low Education	
σ_ξ	0.124	0.133	BHPS
σ_χ	0.132	0.130	BHPS
a_1^f	0.055	0.016	BHPS
a_2^f	-0.0011	-0.00017	BHPS
a_1^m	0.054	0.034	BHPS
a_2^m	-0.0011	-0.00063	BHPS
$\frac{\text{Median } p_t}{\text{Median } y_t}$	2.8	3.5	BHPS
House Price Process			
ρ_h	0.94	ODPM	
σ_ε	0.089	ODPM	
d	2.32%	ODPM	

3.2.1 External Parameter Values

Preference Parameters The preference parameters ρ , η and γ in the utility function are set to match estimated elasticities in the data and the average fraction of time available spent working. The consumption elasticity of intertemporal substitution is set at 0.7 (from Attanasio and Weber, 1995), the hours of work elasticity of intertemporal substitution is set at 0.3 (from Pistaferri, 2003) and the average fraction of time spent working is set at 0.7. Using our CES utility function (5), the consumption elasticity of intertemporal substitution is given by:

$$\varepsilon_c = \frac{u_c}{u_{cc}c} = \frac{1}{(1 - \gamma + \rho) \left[\frac{c^{-\rho}}{c^{-\rho} + \eta z^{-\rho}} \right] - (1 + \rho)} \quad (9)$$

and the hours of work intertemporal elasticity by:

$$\varepsilon_{hrs} = -\frac{u_z}{u_{zz}} \frac{1}{1 - z} = \left[\frac{1}{(1 - \gamma + \rho) \left[\frac{\eta z^{-\rho}}{c^{-\rho} + \eta z^{-\rho}} \right] - (1 + \rho)} \right] \frac{z}{1 - z} \quad (10)$$

To pin down the preference parameters, we assume assets are zero and this allows us to write leisure and consumption as functions of ρ, η and the wage rate, w . We can then use the average fraction of hours to pin down η as a function of ρ and w . Finally, we substitute this expression for η and the functions for z and c into the elasticities (9) and (10) and solve for γ and ρ so that the elasticities at average hours match the elasticities in the data. The resulting numbers correspond to $\gamma = 1.8$, $\rho = -0.2$ and $\eta = 0.25$. The preference parameters over housing are set through calibration.

The discount factor β is also external to the model, and we assume discounting at a rate of 2% per year.

Credit Market Parameters 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.¹⁴ We allow households to borrow whichever amount is lower between three times household earnings ($\lambda_y = 3$) and 90% of the house price ($\lambda_h = 0.9$). In the final part of section 4 we will analyse the sensitivity of the model to different borrowing limits.

For interest rates we use the average 90 day Treasury Bill discount rate in years 1968-1997¹⁵, which gives a real rate of 1.8%.

¹⁴The Financial Services Authority Guide to Mortgages (2004) reports that “typically, the maximum mortgage a lender offers is three times the main earner’s income plus one times any second earner’s income, or two-and-a-half times your joint income. Some lenders offer more, some less.” Moreover, that “It is possible to borrow up to 100% of the property’s value. But a loan of more than 75% of the property value often costs extra.”

¹⁵We stop in 1997 since in that year the interest rate setting regime was changed when the Bank of England became independent with a remit to set interest rates to achieve a target inflation rate.

Wage processes We estimate the female wage and male earnings processes using data from the BHPS for the years 1991-2002 (see Appendix A.3 for details). We estimate separately the parameters for high and low education groups. The results in table 1 show that high education individuals can expect a more hump-shaped income profile than their less educated counterparts during their working lives (both a_1 and a_2 have a bigger magnitude for the high education group). The variance of permanent shocks is similar across education groups for males and for females with low education, but lower for females with high education.

We assume households work from age 22 to 65 and are retired from age 66 to age 81, when they die. In retirement, households receive a pension which is equal to 50% of male earnings in the final year plus 50% of female earnings in the final year assuming the woman is working for H hours where H is the average number of hours worked by all women over the life-cycle.

The initial level for (expected) income relative to the (expected) house price is set to match data from the BHPS. From the data, we calculate the education specific ratios of the median house price to median household earnings for heads of household aged 22-26. We match these to the equivalent ratios for initial income and the initial house price in the model.

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 1969-2000.¹⁶ We estimate an AR(1) process, with linear trend (equation 6), 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 0.94,¹⁷ a standard deviation of the shock (σ_ε) equal to 0.089, and a trend growth rate (d) of a little over 2% per year.

Initial Wealth We set the distribution of initial financial assets for the two education groups to match data on 22-26 year olds in the 2000 wave of the BHPS, and we assume that households have zero housing endowments at age 22.

¹⁶We use the series reporting average house prices for all dwellings.

¹⁷A unit root test on the persistence parameter does not reject the null hypothesis ($\rho_h = 1$).

3.2.2 Calibrated parameters

Given the parameters above, we set the remaining parameters to fit the model to data on life-cycle home-ownership profiles for household heads aged 26-60 between 1991 and 2000,¹⁸ by education group. Since we model income in retirement, we can quite accurately match home ownership until near the age of retirement. On the other hand, the model implies that households run down all assets by the end of life (age 81). This means that we will not obtain a good match for home-ownership among the retired since in the simulations home-ownership declines to zero in a way not observed in the data.¹⁹ Our calibration and comparative static exercises focus on home-ownership behaviour up to age 60. We do not calibrate the labour supply interactions with housing shown in figure 3, rather we use our calibrated model to explore these interactions in the next section.

The life-cycle ownership profiles we will be matching are calculated as an average of the ownership decisions of different cohorts who face different house price realisations. To match these average profiles using our model, we simulate 20 different realised sequences of the aggregate price process, each of which is faced by 1,000 individuals. We construct our simulated home ownership profiles by averaging across the resulting 20,000 individuals (see Figure 4). At the end of this section we will discuss the sensitivity of home ownership to the set of house price sequences, and in particular the effect on life-cycle home ownership of initial house prices being especially high or low.

Our calibration approach is to choose the transactions cost of buying or selling, F , and the parameters specifying the utility benefit of home ownership, μ and θ , to minimise the sum of absolute deviations of simulated moments from corresponding data moments. The moments we use are the average home ownership rates for households in low and high education groups, for those aged 26-35, 36-45 and 46-60. We set the calibrated parameters to be common across the two education groups and since we have 3 parameters to match 6 moments we cannot achieve a perfect fit. Parameter values from the calibration are summarised in Table 2. Table 3 presents the calibration statistics and Figure 4 corresponding life-cycle profiles, showing the extent that home-ownership rates predicted by the model match those observed in the data.

Our fixed cost parameter of 6% 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

¹⁸Data come from the years 1991-2000, as years prior to 1991 are affected by the large-scale selling off of local authority housing.

¹⁹To match home ownership in retirement would require a specification of the bequest motive, of imperfections in the annuity market and possibly of the market for households annuitising their home while remaining home-owners.

Table 2: Calibrated Parameter Values

Parameter	Value
θ	0.028
μ	0.022
F	0.060

Table 3: Calibration Statistics

Statistic	High Education		Low Education	
	Data	Model	Data	Model
Ownership Rate (%) Age 26 – 35	80.9	80.6	64.5	64.5
Ownership Rate (%) Age 36 – 45	87.7	89.0	73.8	74.0
Ownership Rate (%) Age 46 – 60	91.8	87.5	77.4	75.2
Sum of absolute deviations	6.7			

Home ownership rates are measured across couples in the British Household Panel Survey (BHPS), and data come from the years 1991-2002.

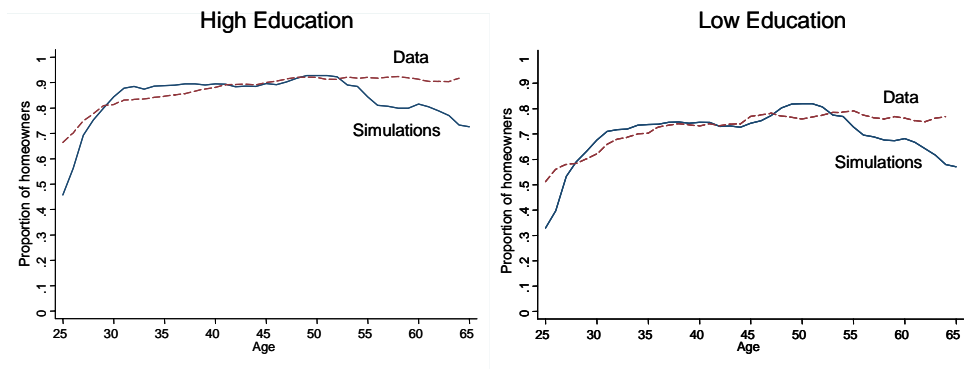


Figure 4: Simulated and Actual Home-Ownership Rates

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.

The fact that the parameters specifying the utility benefit of home ownership, μ and θ , are positive implies that home ownership is a luxury good: at low levels of consumption the multiplicative term (which reduces utility) is relatively important, but as consumption increases the additive term (through which home ownership increases utility) becomes increasingly dominant. Given $\gamma > 1$, the sign of θ specifies whether consumption and housing are complements or substitutes. In our calibration, $\theta > 0$ and they are complements: at a given level of consumption, the marginal utility of consumption is higher when individuals own their own homes.

As mentioned above, the simulated profiles of home ownership used in the calibration are obtained by averaging across many individuals facing 20 different sequences of house price realisations. We interpret this as being similar to averaging across different cohorts in the data, since different cohorts face different house prices when they first enter the housing market, and different subsequent house price shocks. To show the extent to which this averaging matters, in Table 4 we compare the calibration statistics for home ownership to those arising from different initial conditions for the house price. The “Low P_0 ” and “High P_0 ” columns refer respectively to profiles generated by groups of individuals that began their adult (i.e. modelled) lives at the time of low and high house prices. The table shows that “lucky” individuals that enter the labour market when house prices are low, tend to purchase earlier. Further, although this earlier purchase does give rise to greater average home-ownership over the life-time, by about age 45, the difference across lucky and unlucky groups has mostly disappeared. In our model, the earlier purchase by those whose faced a low initial price occurs for two reasons. Firstly, given our stochastic assumptions on the house price, a low initial house price implies a high expected rate of return on housing; and secondly, a low initial house price means that more households have savings and income sufficient to overcome the downpayment and income related borrowing constraints. By middle-age, these effects are less important.

4 Implications for Labour Supply

We now use our model to understand why homeowners work longer hours. We break our analysis into four stages. First, we show that our model reproduces the correlation between mortgage debt and labour supply

Table 4: Simulated Homeownership by Initial House Price

Statistic	High Education			Low Education		
	Low p_0	Baseline	High p_0	Low p_0	Baseline	High p_0
Ownership Rate (%) Age 26 – 35	89.9	80.6	71.3	78.9	64.5	52.7
Ownership Rate (%) Age 36 – 45	91.7	89.0	85.2	81.0	74.0	68.4
Ownership Rate (%) Age 46 – 60	88.0	87.5	86.0	78.1	75.2	73.9

that is observed in the data. Second, we show that this correlation reflects differences in the ratio of debt to husband’s earnings rather than differences in the ratio of debt to the house price, suggesting that current income and liabilities rather than total wealth are driving female labour supply. Third, we draw out the heterogeneity in behaviour by looking explicitly at labour supply functions. We show that labour supply is highly nonlinear in the level of debt, with labour supply being very responsive to debt only at high levels of debt. Finally, we show that changes in capital market imperfections affect the timing of house purchases, but have only a limited effect on average labour supply.

The Correlation between Labour Supply and Mortgage Debt In Figure 5 we divide our simulated data by the extent of mortgage obligations. We define the obligation ratio for each household as the ratio of mortgage repayments to other household income (ie excluding female earnings). For each age, we split the simulated households according to this obligation ratio: a high ratio is defined as being a ratio above the 33rd percentile. We characterise households as either high or low obligation at each age. This means that households may move from one group to another across their life-time. The reason we do not hold group composition constant is because there is substantial variation in debt repayments over the life-cycle in our simulations and so the per-period definition shows the extent that current obligations matter.

Figure 5 shows that among both the young and the old, greater mortgage repayments are associated with higher labour supply. This pattern mimics the pattern in the data shown in Figure 3. Further, compared to non-homeowners, average labour supply is higher for home owners: for the high education group, simulated average labour supply as a fraction of available hours for non-homeowners is 0.72, compared to an average

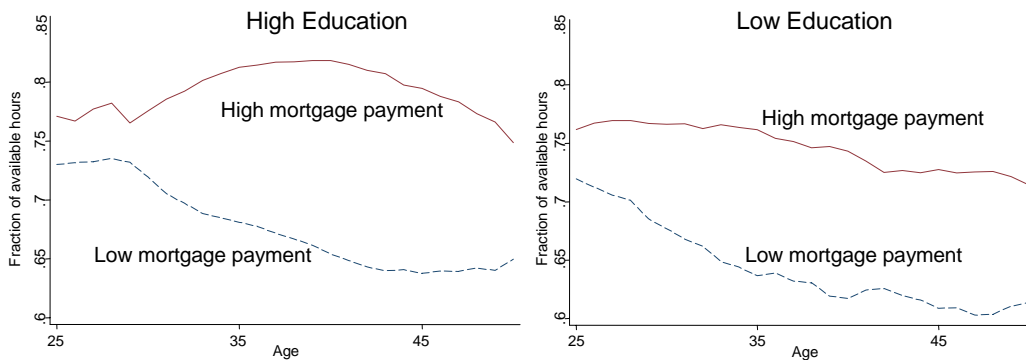


Figure 5: Simulated Labour Supply by Extent of Mortgage Repayment

of 0.74 for home-owners. For the low education group, average labour supply for non-homeowners is 0.67 compared to 0.69 for home owners.

Although Figure 5 reproduces the pattern in the data shown in Figure 3, it also shares the problem that it does not disentangle the endogeneity of the mortgage repayment to labour supply choices. In particular, there are two sources of endogeneity: first, the high obligation ratio and high female labour supply may both be due to low husband’s earnings, which directly increase the obligation ratio and give the wife a financial incentive to work more; second, higher levels of debt may be taken out in period $t - 1$ in anticipation of working longer hours in period t .

Debt-to-house price vs debt-to-husband earnings To understand the correlation, we now show the extent to which differences in labour supply are driven by differences in household wealth or by differences in household income. In Tables 5 and 6, we show, from our simulations, how labour supply varies with the ratio of debt to the house price and with the ratio of debt to husband’s earnings, respectively. Both tables show that the greater the indebtedness, the greater labour supply. However, labour supply varies very little with the debt-to-house price ratio, but varies substantially with the ratio of debt-to-husband’s earnings. This suggests that the higher labour supply of those with greater debt is being driven not by a low level of net wealth per se, but rather by the household’s difficulty in servicing its debt obligations: it is the combination of substantial liabilities and low current non-female income that seems to induce higher female labour supply.

Table 5: Labour Supply and Debt-to-House Price

Low education		High education	
Debt to House Price	Mean of Female Labour Supply	Debt to House Price	Mean of Female Labour Supply
$x \leq 0.45$	0.73	$x \leq 0.39$	0.76
$0.45 < x \leq 0.53$	0.72	$0.39 < x \leq 0.58$	0.77
$0.53 < x \leq 0.65$	0.73	$0.57 < x \leq 0.72$	0.76
$0.65 < x \leq 0.76$	0.73	$0.71 < x \leq 0.83$	0.75
$0.76 < x \leq 0.85$	0.74	$0.82 < x \leq 0.90$	0.75
$x > 0.85$	0.75	$x > 0.90$	0.77

The sample is restricted to age 26-35. Debt refers to home owners and is measured at the beginning of period. Intervals reported in rows 1-6 are defined on the basis of, respectively, the 10th, 25th, 50th, 75th and 90th percentile of the distribution of the debt-to-house price ratio.

Table 6: Labour Supply and Debt-to-Husband's Earnings

Low education		High education	
Debt to Husband's Earnings	Mean of Female Labour Supply	Debt to Husband's Earnings	Mean of Female Labour Supply
$x \leq 2.28$	0.63	$x \leq 1.36$	0.72
$2.28 < x \leq 3.04$	0.67	$1.36 < x \leq 2.26$	0.72
$3.04 < x \leq 4.03$	0.71	$2.23 < x \leq 3.25$	0.74
$4.03 < x \leq 5.23$	0.76	$3.21 < x \leq 4.39$	0.77
$5.23 < x \leq 6.57$	0.80	$4.33 < x \leq 5.60$	0.79
$x > 6.57$	0.83	$x > 5.60$	0.82

The sample is restricted to age 26-35. Debt refers to home owners and is measured at the beginning of period. Intervals reported in rows 1-6 are defined on the basis of, respectively, the 10th, 25th, 50th, 75th and 90th percentile of the distribution of the debt-to-husband's earnings ratio.

Nonlinear Labour supply functions The conclusion about the importance of current income and liabilities is reinforced by looking at the labour supply functions. The labour supply functions also highlight the degree of nonlinearity in the responsiveness of labour supply.²⁰ In Figure 6 we show how optimal hours of work in period t depend on mortgage debt held at the end of the previous period and how this dependence varies with the wife’s wage and with husband earnings. The figure shows the solution for households at age 30 who own a house at the start of period t and continue to own through to period $t + 1$. The striking point about these policy functions is the sudden increase in gradient that occurs as debt passes a certain level. At debt beyond this cut-off, women increase their labour supply sharply in order to make required mortgage repayments or to satisfy the constraints on remortgaging. Furthermore, the darker parts of the labour supply functions in the figure, which indicate that the household chooses to own in period t , show that households do not always sell their homes in order to avoid these steep sections of the labour supply function. For example, there is a small range of debt over which a wife with low wages is prepared to increase her labour supply by between 25 and 40 percentage points rather than choosing to sell their home (see left panel in Figure 6).

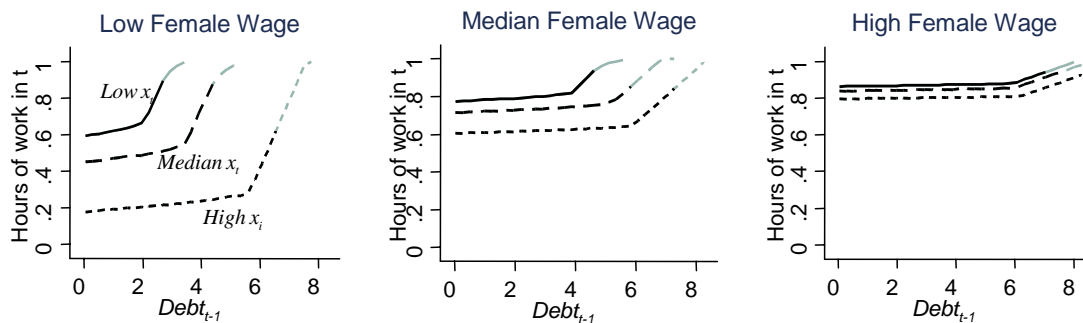
Figure 6 also shows that female labour supply increases with the woman’s wage and decreases with their husband’s earnings. These effects interact with the effect of the borrowing constraints on labour supply: the change in gradient of the labour supply function is most pronounced at low values of the woman’s wage.

In Figure 7 we show how optimal hours of work in period t depend on asset holdings at the start of period t , for households who do not own their home at the start of the period but who buy in period t . Labour supply varies significantly with assets only at low levels of asset holdings where greater labour supply leading to greater income will ease the house purchase. On the other hand, whereas the steep segments of the labour supply functions were relevant for those who already owned their homes, the dark segments of the labour supply functions for those who are thinking about buying indicate that households rarely choose to purchase a house when assets are in a range where the labour supply function is steeper.

While we have highlighted segments of the policy functions in which labour supply varies dramatically with assets, for large ranges of asset holdings, average simulated labour supply does not vary much with debt. This nonlinear nature of labour supply means that sharp effects of debt on behaviour may be underestimated in looking solely at average behaviour and this underlines the value of analysing labour supply functions.

²⁰We focus in this subsection on the low educated, but very similar results are available for the high educated.

Figure 6: Labor Supply for Home-Owners by Debt (Low Education)



The graphs show labour supply functions for those who own their home in $t - 1$ and continue to own in t . The darker part of each line indicates the range of debt where ownership in t is optimal (compared to not owning in t). The scale of debt on the x-axis is as a ratio of debt to mean income at age 22.

Figure 7: Labor Supply for New Home-Owners by Initial Assets (Low Education)



The graphs show labour supply functions for those who did not own their home in $t - 1$ but who buy in t . The darker part of each line indicates the range of assets where ownership in t is optimal (compared to not buying in t). The scale of assets on the x-axis is as a ratio to mean income at age 22.

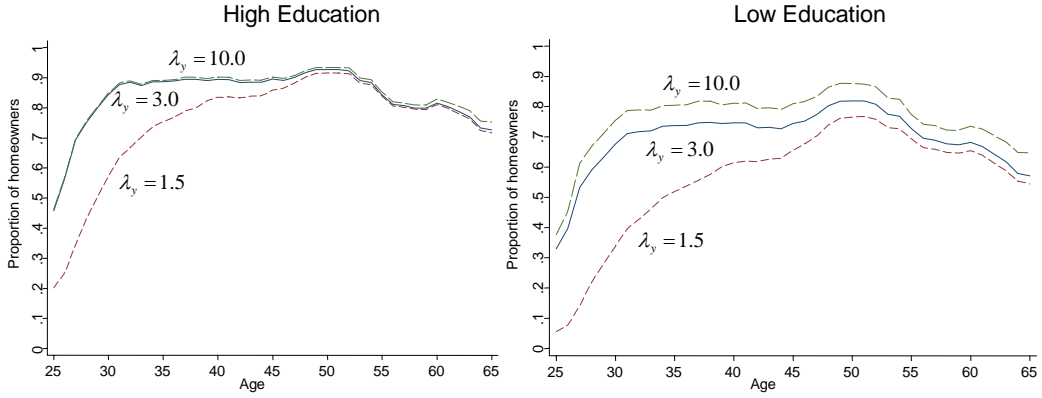


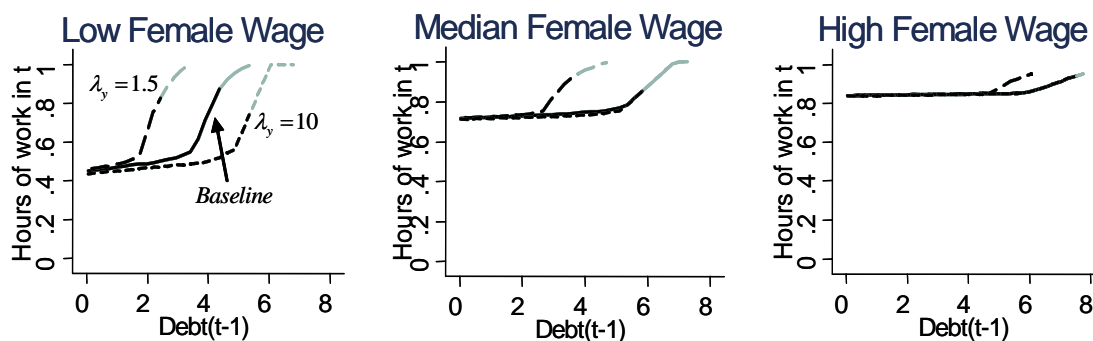
Figure 8: Simulated Home-Ownership Rates Varying Income Related Borrowing Constraint, λ_y

Capital market imperfections While the effects of indebtedness discussed above indicate the extent that labour supply is affected by debt, the interaction comprises both the effect of capital market imperfections and the effect of the debt per se. In this subsection, we try to disentangle these effects further by showing how the effects on labour supply of indebtedness change with the severity of the liquidity constraints faced.

First, we consider the effect of varying the income related borrowing constraint. The effect on home-ownership of varying λ_y is shown in Figure 8. Relaxing this constraint leads to households buying their homes earlier in the life-cycle. By middle age, there is less difference in the level of ownership, although some effect remains throughout the life cycle for the low educated. For the sake of space, we do not show the corresponding labour supply profiles as the differences in these profiles are rather small. The fact that we observe delayed house purchase rather than increased labour supply on average, is consistent with the evidence from figure 7 that households will rarely purchase a home when this requires a substantial increase in labour supply.

This small effect on average labour supply could be masking substantial effects on individual behaviour since we know from the labour supply functions presented in Figures 7 and 6 that labour supply is nonlinear in debt in a way that is heterogenous between individuals with different incomes, and between home owners and non-homeowners. To examine this, we show how the labour supply functions vary with the severity of the borrowing constraints. Figure 9 reports labour supply functions for households who continue to own between t and $t + 1$, varying the severity of the income related constraint. A more severe earnings related borrowing constraint (lower λ_y) results in labour supply being very responsive to debt at a lower range of

Figure 9: Labor Supply for Home-Owners by Debt varying λ_y



The graphs show labour supply functions for those who own their home in $t - 1$ and continue to own in t . The darker part of each line indicates the range of debt where ownership in t is optimal (compared to not owning in t). The scale of debt on the x-axis is as a ratio of debt to mean income. Husband earnings are set at the median value

debt holdings. Since the earnings related constraint is the only factor explaining differences in the three labour supply functions plotted on each panel of the figure, the shift of these functions as the constraint changes clearly indicates that this constraint has a role in shaping the steep segments of the functions. The labour supply functions for $\lambda_y = 3$ and $\lambda_y = 10$ overlap in the middle and right panels of the figure, indicating that the earnings related constraint has become loose relative to the other constraints in the model and so no longer influences behaviour. Similar non-linear labour-supply functions, with a steep segment that shifts as the borrowing constraint is altered, exist for households who become home owners in period t (not shown). However, as in figure 7, it is again the case that non-homeowners rarely choose to buy where the labour supply function is steeper.

In Figure 10 we show simulated home ownership rates, varying the downpayment requirement, $(1 - \lambda_h)$. Reducing downpayment requirements (increasing λ_h) leads to households buying their homes earlier in the life-cycle, but again, by middle-age, there is little difference in home-ownership across different values of the downpayment requirement. Similarly to changes in λ_y , changes in λ_h do not result in substantial changes in the profile of average labour supply (not shown).

Figure 11 reports labour supply functions for different values of the downpayment constraint. As with the earnings-related constraint, an increase in the downpayment constraint makes labour supply sensitive to debt at lower ranges of debt holdings. Again, a shift of the functions in response to a change in the constraint indicates when the downpayment requirement is important in generating the steep sections of the labour

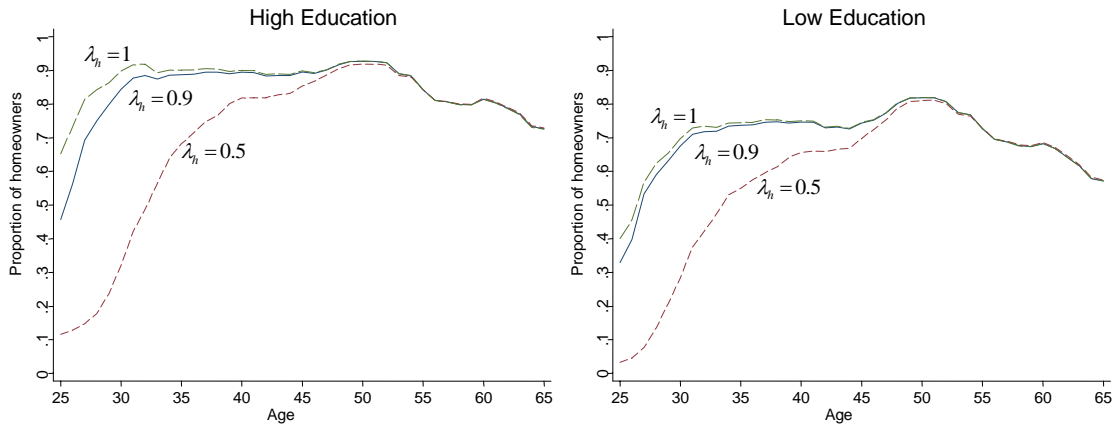


Figure 10: Simulated Home-Ownership Rates Varying Downpayment Constraint, λ_h

supply functions. However, the differences between the different policy functions is somewhat less marked than when varying the earnings constraint, especially for the downpayment requirements of 10% and zero for which the labour supply functions are overlapping or almost overlapping. This reiterates the pattern from the home-ownership simulations in Figures 8 and 10 which show that, for parameters considered and for the low educated, the earnings-related constraint has a bigger effect on behaviour.

Figure 11: Labor Supply for Home-Owners by Debt varying λ_h



The graphs show labour supply functions for those who own their home in $t - 1$ and continue to own in t . The darker part of each line indicates the range of debt where ownership in t is optimal (compared to not owning in t). The scale of debt on the x-axis is as a ratio of debt to mean income. Husband earnings are set at the median value

5 Conclusion

This paper used a structural life-cycle model to address the question of why homeowners with greater debt work longer hours. Two main conclusions arise from our model: first, this correlation is driven more by

current liabilities than by a lower level of net wealth. This impact of current liabilities arises because it is less costly to adjust labour supply than to adjust housing stock in order to meet mortgage repayments. This effect might be less important if capital markets allowed greater borrowing as this would allow borrowing to adjust following shocks rather than labour supply. We show, however, that relaxing borrowing constraints can have a substantial effect on the timing of house purchases, but has relatively small effects on the level of home-ownership from middle-age onwards, and has only small effects on average labour supply. The second conclusion of our model is that for highly indebted households or households with low wages, labour supply is very sensitive to changes in debt, whereas for average households, there is much less sensitivity.

References

- [1] Attanasio, O. and G. Weber (1995) "Is consumption growth consistent with intertemporal optimisation? Evidence from the consumer expenditure survey." *Journal of Political Economy*, 103: 1121-1157.
- [2] 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, <http://www.ifs.org.uk/wps/wp0221.pdf>.
- [3] Blundell, R., M. Browning and C. Meghir (1994) "Consumer Demand and the Life-Cycle Allocation of Household Expenditure" *Review of Economic Studies*, 61: 57-80.
- [4] Blundell, R., L. Pistaferri. and I. Preston, (2004) "Consumption Inequality and Partial Insurance", *IFS Working Paper WP04/28*, (<http://www.ifs.org.uk/wps/wp0428.pdf>)
- [5] Browning, M., A. Deaton. and M. Irish (1985) "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle" *Econometrica*, 53(3): 503-544.
- [6] Bottazzi, R. (2007) "Labour market participation and mortgage-related borrowing constraints" *Institute for Fiscal Studies, mimeo*
- [7] Campbell J.R., and Z. Hercowitz (2004) "The Dynamics of Work and Debt", *NBER working paper 10201*.
- [8] Carroll, C. (1997) "Buffer-stock saving and the life cycle/permanent income hypothesis", *Quarterly Journal of Economics*, 112(1): 1-55.
- [9] Chetty, R. and A. Szeidl (2006) "Consumption Commitments and Risk Preferences", *Quarterly Journal of Economics*, May 2007
- [10] Cocco, J. (2005) "Portfolio Choice in the Presence of Housing" *Review of Financial Studies*, 18: 535-567.
- [11] Deaton, A. (1991) "Saving and Liquidity Constraints" *Econometrica*, 59(5): 1221-1248.
- [12] Del Boca, D. and A. Lusardi (2003) "Credit Market Constraints and Labor Market Decisions" *Labor Economics*, 10: 681-703.
- [13] Disney, R., A. Henley and D. Jevons (2003) "House price shocks, negative equity and household consumption in the 1990s" *mimeo*.

- [14] Fernandez-Villaverde, J. and D. Krueger (2005) "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?" *Proceedings of the 2002 North American Summer Meetings of the Econometric Society: Macroeconomic Theory*: 1- 55.
- [15] Financial Services Authority (2004) "Guide to Mortgages" www.fsa.gov.uk
- [16] Fortin, N. (1995) "Allocation Inflexibilities, Female Labour Supply, and Housing Asset Accumulation: Are Women Working to Pay the Mortgage" *Journal of Labor Economics*, 13(3): 524-557.
- [17] Li W. and R. Yao (forthcoming) "The life-cycle effects of house price shocks" *Journal of Money, Credit and Banking*.
- [18] Low, H. W. (2005) "Self-insurance in a life-cycle model of labour supply and saving" *Review of Economic Dynamics*, 8(4): 945-975.
- [19] Low, H.W., C. Meghir. and L. Pistaferri (2006) "Wage risk and employment risk over the life-cycle", *Institute for Fiscal Studies Working Paper W06/27*, <http://www.ifs.org.uk/wps/wp0627.pdf>
- [20] MaCurdy, T.E. (1982) "The use of time series processes to model the error structure of earnings in a longitudinal data analysis" *Journal of Econometrics*, 18: 82-114
- [21] Nordvik, V., (2001) "A Housing Career Perspective on Risk" *Journal of Housing Economics*, 10: 456-471.
- [22] Ortalo-Magne, F. and S. Rady (2002) "Tenure Choice and the riskiness of Non-housing Consumption" *Journal of Housing Economics*, 11: 266-279.
- [23] Ortalo-Magne, 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

A Appendix

A.1 Computational methods: general

The solution for consumption, labour supply and home-ownership is found recursively from the last period of life, T , backwards. Since there is no bequest motive in our model, and since households are subject to the constraint of non-negative assets at the end of life, they spend all that it is available to them in the last period and allocate this optimally between consumption and labour supply, given the optimal choice of home-ownership. Given the optimal choices at $t+1$, $t < T$, they then need to choose home ownership, consumption, leisure and saving that maximise period t 's value function, subject to the borrowing constraints. We solve our model via value function iteration rather than using the Euler equation because of the non-convexity of the value function.²¹

In order to compute the solution, we discretise the state space for the wages of the husband and wife, and for the house price. The wage and house price processes are modeled as finite-state Markov chains that mimic the underlying continuous-valued AR(1) processes, as described in Tauchen (1986). The wage processes are represented by 11 nodes, whereas the house price process is represented by 13 nodes. For assets, we use a grid of 100 nodes but maintain a continuous underlying variable. Points are more dense in the lower range of the asset grid where the curvature of the value function is likely to be changing more rapidly. Given the solution of the household optimisation problem at each point of the grid, we approximate the policy functions and the value functions at points off the asset grid using linear interpolation.

A.2 Computational methods: endogenous liquidity constraint

We would like to solve for consumption and labour supply using two-stage budgeting, solving first for total within period spending (on goods and leisure), and then for consumption and labour supply given that level of total within period spending and wages. However, in our model the income-related constraint is endogenous with respect to labour supply. This means that the choice of labour supply can affect how much the household can borrow and that we cannot treat total spending as given when choosing labour supply. This situation arises only for those who choose to own a home in (at the end of) the current period t : for those who sell or choose not to buy borrowing is bounded by an amount (0) that does not vary with labour supply, and so two-stage budgeting is feasible. In this section we describe the computational approach used

²¹Non-convexities in the value function arise from having transaction costs associated to buying and selling the house and also from the interactions of different credit constraints.

for finding consumption and labour supply for home-owners.

Within the group of home-owners it can be conceptually useful to distinguish between those who buy their property in the current period, and those who are continuing owners. The former group must take out a mortgage in the current period if they want to borrow, whereas the latter group has the choice between carrying over their existing debt (provided that they repay at least the new interest on this debt - see equation (3) of main text) or taking out a new mortgage. In this appendix we refer to a case in which there is an option of carrying over existing debt. The computational procedure for a new buyer²² is the special case of the setup described in which the debt that can be carried over is zero.

The procedure for finding consumption and labour supply for home-owners proceeds in a series of steps:

- 1. Solve the model as if two-stage budgeting were possible. Set limits on total expenditure at this step such that end-of period debt can be no bigger than the greater allowed by: beginning-of-period debt minus current interest paid on debt (see (3) of the main text); or the debt associated with whichever constraint is binding between the downpayment constraint and the earnings-related constraint when the latter is a function of *maximum possible* female labour supply ($l_t = 1$).²³ We will refer to this solution as the ‘partially constrained solution’.
- 2. Check whether the debt level implied by the ‘partially constrained solution’ exceeds the level allowed by carrying over existing debt (ie D_{t-1}).²⁴ If it does not the ‘partially constrained solution’ is feasible without remortgaging and is the constrained solution. If this level of debt is exceeded, proceed to step 3.
- 3. Check whether the debt level implied by the ‘partially constrained solution’ exceeds the level allowed by the earnings-related borrowing constraint when this debt limit is calculated based on the level of labour supply associated with the ‘partially constrained solution’. If this constraint is not violated, then due to the way in which step 1 was set up, we know that the ‘partially constrained solution’ does not violate the borrowing constraints associated with remortgaging. Since remortgaging is costless, this implies that the ‘partially constrained solution’ is the constrained solution. If this constraint is violated, proceed to step 4.

²²Or an existing owner who had paid off their mortgage by the start of period t .

²³In fact we also take account of the implicit constraint due to the no bankruptcy condition here, but this is rarely the binding constraint until late in the household’s lifetime.

²⁴This step is effectively redundant for new buyers unless they are wealthy enough to buy a property without borrowing.

4. Search for a new (higher) constrained level of labour supply, exploiting the information that the earnings-related constraint is binding. At each iteration of this search, we have to check whether the higher level of l_t associated with the constraint still results in the earnings-related constraint being binding, or whether this in fact becomes less tight than the downpayment constraint.²⁵ End-of-period assets are given by the binding constraint and consumption is found residually through the budget constraint. For example, if the earnings related constraint binds and the household is buying a house, then substituting from (2) into (1) (equations are in main text) and rearranging yields:

$$c_t = A_t + (1 + \lambda_y)(w_t l_t + x_t) - p_t h_t - F p_t$$

Having found labour supply, consumption and saving, we can obtain the value associated with remortgaging with constrained labour supply. We also compute the value associated with not remortgaging and instead repaying (at least) the interest on existing debt (if this permits a feasible outcome). The household's behaviour is determined by choosing the maximum of the values for the constrained solution with remortgaging, and the solution with interest repayment.

A.3 Estimation of wage variance and deterministic trend

We estimate the variance of the wage processes for the UK by drawing on the methodology developed in Blundell, Pistaferri and Preston (2004) for estimating the variance of the permanent and transitory shocks to income. Blundell et al. (2004) adopt the following income process for household i :

$$\log Y_{i,a,t} = Z'_{i,a,t} + H_{i,a,t} + v_{i,a,t}$$

where a and t index age and time respectively, Y is real income, Z is a set of observable characteristics, H and v are, respectively, a permanent and a transitory income component. They then assume that the permanent component $H_{i,a,t}$ follows a martingale process: $H_{i,a,t} = H_{i,a,t-1} + \varsigma_{i,a,t}$ and that the transitory component $v_{i,a,t}$ follows an MA(q) process: $v_{i,a,t} = \sum_{j=0}^q \theta_j \varepsilon_{i,a-j,t-j}$, ($\theta_0 \equiv 1$). They can therefore obtain an expression for “unexplained” income growth as follows: $\Delta y_{i,a,t} = \varsigma_{i,a,t} + \Delta v_{i,a,t}$, where $y_{i,a,t} = \log Y_{i,a,t} - Z'_{i,a,t}$ (real income net of predictable individual components), and derive covariance restrictions in panel data, as follows:

$$\text{cov}(\Delta y_{a,t}, \Delta y_{a+s,t+s}) = \begin{cases} \text{var}(\varsigma_{a,t}) + \text{var}(\Delta v_{a,t}) & \text{for } s = 0 \\ \text{cov}(\Delta v_{a,t}, \Delta v_{a+s,t+s}) & \text{for } s \neq 0 \end{cases}$$

²⁵Or indeed, than the no bankruptcy constraint.

where $\text{var}(\cdot)$ and $\text{cov}(\cdot)$ denote cross-sectional variances and covariances. In the particular case of serially uncorrelated transitory shocks, $\text{var}(\zeta_{a,t})$ becomes only a function of “unexplained” income growth:

$$\text{var}(\zeta_{a,t}) = \text{cov}(\Delta y_{a,t}, \Delta y_{a-1,t-1} + \Delta y_{a,t} + \Delta y_{a+1,t+1}) \quad (11)$$

Recalling that we only allow for permanent shocks in our wage process (equation (7)), and assuming that in our data the transitory shocks are uncorrelated, we base the estimation of the variance of permanent shocks on equation (11). In particular, we need to obtain real wages net of predictable individual components and use their lags to construct the covariance between a one-period lag and a three-period lag, as follows:

$$\text{var}(\xi_{a,t}) = \text{cov}(\Delta \tilde{w}_{a,t}, \Delta \tilde{w}_{a-1,t-1} + \Delta \tilde{w}_{a,t} + \Delta \tilde{w}_{a+1,t+1})$$

where \tilde{w} is the logarithm of real wages net of predictable individual components.

For males, this is obtained as the residual from a regression, for each education group, of individual log-real hourly wages on a quadratic term in age, 5-year cohort and region dummies, a dummy for being married, household size, and number of children in the household. The resulting wage variance (and standard errors in parentheses²⁶) is 0.0170 (0.00226) for the group with low education and 0.0174 (0.00177) for the group with high education.

For females, we deal with selection into the labour market by estimating a Heckman model. \tilde{w} is therefore the residual from a regression of log-real hourly wages on a quadratic term in age, 5-year cohort dummies and dummies for region, where selection is based on the following "excluded variables": other household income (i.e. annual household income excluding female labour income), number of children, household size, and dummy for being married. The resulting wage variance (and standard errors in parentheses, bootstrapped as for males) is 0.0178 (0.00221) for the group with low education and 0.0153(0.00273) for the group with high education.

The coefficients for the deterministic component of both the male and the female wage processes are given by the quadratic terms in age in the respective regressions. Estimated coefficients (and robust standard errors in parentheses) for, respectively, the linear and the quadratic term, are as follows: 0.034 (0.0036) and -0.00063 (0.00009) for males with low education, 0.054 (0.0043) and -0.0011 (0.0001) for males with high education, 0.016 (0.0027) and -0.00017 (0.00006) for females with low education, 0.055 (0.0033) and -0.00114 (0.00009) for females with high education.

²⁶Standard errors computed by bootstrapping 500 samples of size n, where n is the number of individual clusters in our data.