

Consumption and Saving Decisions
in the face of
Choices about Housing and Pensions

Matthew John Wakefield

A Dissertation submitted to the Department of Economics
in partial fulfilment of the requirements for the degree of

Doctor of Philosophy (PhD)
University College London (UCL)

London
February 2009

Declaration

I, Matthew John Wakefield, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Matthew Wakefield

Abstract

This thesis presents analyses of households' decisions regarding housing and pension wealth accumulation, in forward-looking models.

The first of three chapters on housing presents a model of housing demand over the life cycle, and examines its sensitivity to prices and borrowing constraints. Demand responses can be unusual: when the price of housing goes up, the demand for starter homes may increase if enough people “downsize” their homes.

The next chapter involves carefully matching a life cycle model of consumption and housing choices, to recent episodes in the U.K., and using this structure to understand why house prices and consumption growth are strongly positively correlated. The model provides a good match to data on home ownership and consumption growth. The analysis gives a firmer theoretical footing to the claim that wealth effects from house prices are *unlikely* to have been the main driver of the correlation with consumption growth.

The third housing chapter presents a model in which the prices of two types of home are endogenous. A perfect foresight set up is used, and transitions between steady states following shocks to income and mortgage markets, are studied. The findings suggest that credit shocks are more promising than income shocks as a potential explanation of large house price fluctuations and housing transactions level that covary positively with prices.

The final substantive chapter uses a difference-in-differences strategy to evaluate the effect on private pension coverage of a recent U.K. reform to private pensions and pension contribution limits. The reform is seen to have had a positive effect on pension coverage for lower earners. This pattern is consistent with forward-looking responses to the financial incentives involved.

Contents

1	Introduction	14
2	Modelling the Demand for Housing over the Lifecycle	19
2.1	Introduction	19
2.2	Life-cycle model	20
2.2.1	Model Structure	21
2.2.2	Model Calibration	25
2.3	Sensitivity to income and house-price parameters	33
2.3.1	Changes related to income	34
2.3.2	Changes relating to house prices	36
2.3.3	Responsiveness to fixed costs	38
2.4	The Response to Shocks	44
2.5	Conclusion	51
A2	Appendix to chapter 2	52
A2.1	Extra tables and pictures	52
A2.2	Computational methods	53
3	The Coincident Cycles of House Prices and Consumption in the U.K.:	
	Do House Prices Drive Consumption Growth?	57
3.1	Introduction	57
3.2	Existing evidence	62

3.3	The model	66
3.3.1	The household maximisation problem	68
3.3.2	The environment: exogenous stochastic processes	72
3.3.3	Calibration and estimation	74
3.3.4	Model fit	78
3.4	Booms and busts in house prices and consumption	81
3.4.1	What is driving consumption growth in the model?	81
3.4.2	How do the shocks affect different age groups?	84
3.5	Conclusions and further work	88
A3	Appendix to chapter 3	90
A3.1	Solution method	90
A3.2	Estimating the idiosyncratic part of the income	93
4	Modelling House Prices with Shocks to Income and Credit	94
4.1	Introduction	94
4.2	The Model	98
4.2.1	Elements of the Model	99
4.2.2	Solving for flat and house prices	106
4.3	Income Shocks and Credit Constraint Shocks	111
4.4	Conclusion and Next Steps	126
5	Tax reform and retirement saving incentives:	
	Evidence from the ‘Stakeholder Pension’ reforms in the U.K.	131
5.1	Introduction	131
5.2	Stakeholder Pensions	133
5.3	Empirical Strategy	138
5.3.1	Data sources and descriptive analysis	138
5.3.2	Modelling the pension take-up decision	142
5.3.3	Implementing the estimator: data considerations	146

5.4	Did the reform increase private pension coverage?	149
5.4.1	Empirical Results	149
5.4.2	Alternative explanations and sensitivity analysis	153
5.5	Conclusions	160
A5	Appendix to chapter 5	162
6	Conclusion: Topics for Further Work	163
	Bibliography	166

List of Figures

2.1	Home Ownership in Baseline Run	31
2.2	Response of ownership to change in fixed costs, high education	40
2.3	Response of ownership to change in fixed costs, low education	41
2.4	Effect of 1% house price shock at 30, 40 or 50, on: a. home ownership and b. consumption; High education.	45
2.5	Effect of 1% house price shock at 30, 40 or 50, on: a. home ownership and b. consumption; Low education.	46
2.6	Effect of 1% income shock at 30, 40 or 50, on: a. home ownership and b. consumption; high education.	49
2.7	Effect of 1% income shock at 30, 40 or 50, on: a. home ownership and b. consumption; low education.	50
A2.1	Effect of HP shock at age 30 on expected HP at older ages	52
3.1	Growth in real consumption and house prices	58
3.2	Consumption function residuals, by age group	64
3.3	Growth in real consumption: Data versus model	79
3.4	Model consumption growth by cohort: 1980s boom	80
3.5	Model consumption growth: different scenarios	82
3.6	Residuals for log consumption: Model Results, Baseline	84
3.7	Residuals for log consumption: Model Results, House Price Shocks Only	85
3.8	Residuals for log consumption: Model Results, Income Shocks Only . . .	86

3.9	Growth in real house prices and "aggregate income"	92
4.1	Housing transactions and house price inflation in the U.K	96
4.2	Ownership by Age: Steady states before and after income shock	113
4.3	Ownership by Age: Steady states before and after credit market shock	113
4.4	Log flat price following income and credit market shocks	114
4.5	Log house price following income and credit market shocks	115
4.6	Flat transactions, relative to after SS level, following income and credit market shocks	116
4.7	House transactions, relative to after SS level, following income and credit market shocks	117
5.1	Maximum annual gross contribution limit, by annual gross relevant earn- ings and age, personal pension tax regime from 2001 to 2006	136
5.2	Percentage of individuals receiving limit increase by age and household type	147

List of Tables

2.1	Parameter Values	26
2.2	Calibrated Parameter Values	29
2.3	Calibration Statistics	30
2.4	Elasticities for ownership: Factors relating to income	34
2.5	Elasticities for flats and houses: Factors relating to income	35
2.6	Consumption elasticities: Factors relating to income	36
2.7	Elasticities for ownership: Factors relating to house prices	37
2.8	Elasticities for flats and houses: Factors relating to house prices	38
2.9	Consumption elasticities: Factors relating to house prices	39
2.10	Effects on consumption, homeownership, and welfare, due to changing fixed costs	42
A2.1	Ownership rates	52
3.1	ABHL house price terms by age groups: predicted price specification	65
3.2	Estimated / Fixed Parameters	75
3.3	Calibrated Parameters	78
3.4	Calibration Statistics	79
3.5	Estimated parameters of the idiosyncratic income process	93
4.1	Parameter Values	104
4.2	Parameters before and after shocks	111
4.3	House and flat prices before and after shocks	112

4.4	Ownership Rates Around Income Shock	118
4.5	Financial Asset Distribution: Non-owners aged 30 and under, around income shock	119
4.6	Financial Asset Distribution: Owners aged over 60, around income shock	120
4.7	Ownership Rates Around Credit Market Shock	122
4.8	Financial Asset Distribution: Non-owners aged 30 and under, around credit market shock	122
4.9	Financial Asset Distribution: Owners aged over 60, around credit shock	123
5.1	Pre-2001 tax reliefs for defined contribution pension plans: Maximum contributions as a % of earnings by age	135
5.2	Pension coverage by type of pension and earnings band: 1999–2000 to 2002–03	140
5.3	Results: Impact of reform of contribution limits on take-up of private pensions for selected groups, using Blundell et al. (2004) procedure . . .	151
A5.1	Sample sizes for analysis in Table 5.3	162

Acknowledgements

Funding Arrangements

Much of the work in this thesis has formed part of the Economic and Social Research Council (ESRC) Centre for the Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (IFS), and I am grateful to the centre for research funding. Chapters 2 and 3 also benefited from additional ESRC specific research grants (respectively RES-000-23-0283 and RES-000-22-1913), and again the support is gratefully acknowledged.

Data Provision

Publicly available data from the British Household Panel Survey, the Family Expenditure Survey / Expenditure and Food Survey, and the Family Resources Survey, were made available through the ESRC data Archive at the University of Essex. Neither the ESRC nor the data archive bear any responsibility for the analyses presented here.

Conjoint Work

Some of the chapters of this thesis are based on joint work:

1. Some of the work underlying chapter 2 was conducted as part of a project in which I collaborated with Orazio Attanasio, Renata Bottazzi, Hamish Low and Lars Nesheim.
2. The work for chapter 3 was conducted jointly with Orazio Attanasio and Andrew Leicester.
3. The work for chapter 5 was conducted jointly with Richard Disney and Carl Emmerson. The version presented in the thesis is a modified version of Disney et al. (Forthcoming (2009)).

I thank these co-authors and collaborators for allowing me to draw in these works in the thesis. The remaining chapters (1, 4 and 6) are entirely my own, as of course is the responsibility for any errors.

Personal Thank Yous

My greatest personal, intellectual debt is to my supervisors, Richard Blundell and Orazio Attanasio. Their input and ideas have been vital and without them, this thesis, in its present form, would not have been possible. They have also provided unfailing support, as well as the occasional bit of constructive criticism, or of prompting and prodding. I could not have wished for better supervisors.

Alongside Richard and Orazio, I have been lucky enough to work with many other extremely able colleagues at the IFS, a number of whom have become firm friends. Here I mention only those who have particularly influenced my thinking on the topics discussed in the thesis.

Hamish Low deserves special thanks for sharing his knowledge of numerical techniques (much of the notation of my Fortran code still bears his mark) as well the results of his precise economic reasoning. I hope that some of his clarity rubbed off.

Carl Emmerson has been an excellent guide to the policy issues and datasets that I have confronted during my time at IFS, but it is his striking ability to incisively cut to the core of an economic or policy argument, that has most benefitted my thinking. Richard Disney also took me under his wing when I arrived at IFS, and has tried to make sure that I didn't lose sight of the issues among the details. I hope that he doesn't despair too much about the direction my research has taken.

Throughout my time at IFS James Banks has been a source of inspiration, enthusiasm and ideas. Thanks must also go to Andrew Dilnot, who first nudged me towards seriously contemplating enrolling at UCL.

Finally, Renata. I owe her thanks for helping me to learn many of the skills I have used in this thesis, and for being an excellent sounding board for ideas, and in recent days even for some proof reading and copy editing. But her constant support runs to far more than the sum of things I could list, and it is for this that I am truly grateful. This thesis is for Renata, and for Emma: they have always had all my love, now they deserve more of my time.

Chapter 1

Introduction

The work presented in this thesis separately analyses the decisions that households make about whether to own the house that they live in, and whether to save in a private pension. The behavioural model underpinning the analyses is the “lifecycle model” of Modigliani and Brumberg (1954), which has become a workhorse of modern macroeconomics and public finance. This model shares with the Permanent Income Hypothesis of Friedman (1957) the idea that individuals allocate resources over time (and therefore determine current consumption and saving - or, equivalently, present and future consumption) taking into account the resources available over a long period of time and the possibility of moving them over time through saving or borrowing.¹ The model provides a structure within which we can think about what would be the optimal choices for households to make regarding their savings and asset choices.

Analysis of the lifecycle framework quickly becomes analytically complex. In particular, with a realistic specification for uncertainty regarding the arrival of resources, for the markets within which intertemporal trades are possible, and for the way in which consumption is converted into ‘utility’, the model of optimal choices does not have any known closed form solution. At least since the contributions of Deaton (1991)

¹A discussion of the model, and a survey of the literature describing our understanding of its implications, is available in Attanasio and Wakefield (2008 and forthcoming).

and Zeldes (1989), numerical methods have been used to analyse models of this kind.

A recent literature has begun the task of adding a housing good/asset to such a computational framework, and the analyses of chapters 2 - 4 represent a contribution to this literature.² While these chapters are written to be self-contained and free-standing, it is worth saying a few words about their content, how they relate to each other and how the models used are modified for the particular research questions analysed.

The aim of chapter 2 is to model carefully the demand for home ownership over the life-cycle and to show how sensitive demand is to price fluctuations and to changes in capital market imperfections. A related recent paper regarding the price fluctuations topic is Li and Yao (2007), which considers the behavioural and welfare consequences of house-price shocks.³ As explained in the chapter, as well as its broader remit, the analysis presented here is distinguished from theirs by modelling choices regarding the nature of the housing asset and a focus on capturing mortgage constraints.

In addition to being interesting in themselves, the set up and results of chapter 2 are useful as background to the subsequent two chapters. The model shares many features with the set ups used in the following chapters, including that the housing good is modelled as taking one of two sizes (a ‘flat’ or a ‘house’), and a realistic form for the mortgage borrowing constraints. The chapter looks at ‘elasticities’ of housing and consumption demands to various parameters in the model and to shocks in housing prices and incomes. An interesting finding is that when the price of housing (i.e. for both flats and houses) goes up permanently, the demand for flats can increase as the number of households that ‘downsize’ from houses exceeds the number who exit the

²The most related papers are mentioned in this introduction, and fuller references are provided in each chapter.

³Other papers that build a housing good/asset into a lifecycle framework have focussed on different topics. Fernandez-Villaverde and Krueger (2005) and Flavin and Nakagawa (2004), for example, address how patterns of non-durable consumption over time are affected by the existence of a durable or housing asset, while Yao and Zhang (2005) and Cocco (2005) both consider how financial asset portfolios interact with decisions over housing wealth, and Campbell and Hercowitz (2004) and Bottazzi et al. (2007) consider the relationship between housing and labour-supply decisions.

flat market. The results of this chapter regarding how house prices and income affect consumption choices are useful background to the analysis of chapter 3, while the results on how housing demand responds to incomes, credit conditions and house price fluctuations, is relevant to the model of housing prices presented in chapter 4.

Chapter 3 involves applying a model similar to that outlined in chapter 2, to the analysis of the extent to which movements in house prices might be driving the changes in aggregate consumption growth with which they are strongly correlated. As such the chapter is an addition to a hotly contested literature on whether there are strong wealth effects from house prices to consumption (see Muellbauer and Murphy (1990) and Campbell and Cocco (2007)), or whether it is more likely that common causes are at work (King (1990), Attanasio and Weber (1994) and Attanasio et al. (2009)).

To address this issue, the model is adapted to include aggregate income and house price shocks, and to capture (both in the modelling of expectations, and in the simulation of the model) correlation between these shocks. Simulations of the model are carefully matched to recent episodes in the U.K. economy by inputting observed aggregate shocks when simulating the model for individuals in different cohorts. These simulations can then be aggregated up to assess the extent to which the fluctuations in the input aggregate processes can explain fluctuations in aggregate consumption. To my knowledge, at least in the consumption literature, it is an innovative contribution to aggregate up simulations of individual behaviour in this way, to check whether the model of individual choices is consistent with aggregate evidence. The model provides a good match to U.K. data on home-ownership and aggregate consumption growth. The second contribution of this paper lies in the use of counterfactual simulations to help us understand whether house price movements are likely to have driven a substantial part of fluctuations in aggregate consumption. The findings support the argument that since the observed correlation between house-price growth and consumption is particularly strong for the young, it is unlikely that wealth effects have been the key factor driving this correlation.

The third chapter on housing - chapter 4 - involves constructing a model in which the prices of flats and houses are endogenous, determined by the decisions of non-owners concerning whether to buy, and the decisions of owners about whether to supply their property to the market. This model does not include aggregate fluctuations, and so steady states involve constant house and flat prices. However, we do use the framework to consider dynamics in the housing market, by modelling the transition path for the economy following unforeseen shocks to income and mortgage market conditions. This analysis is a contribution to the literature on whether housing market fluctuations can be explained by the interaction between aggregate shocks and the institutions of the housing market (see Ortalo-Magné and Rady (2006), who suggest one mechanism that might generate the right properties for volatility in prices and transactions, and also Rios-Rull and Sanchez-Marcos (In Progress)). The findings suggest that credit shocks might be more promising than income shocks as a potential explanation of large house price fluctuations and transaction levels that are pro-cyclical. This chapter is a preliminary application of the model with endogenous house and flat prices, and the final part of the chapter considers how work with this model might be extended.

The final substantive chapter of the thesis is concerned with pension saving rather than housing choices, and involves a different set of techniques as policy evaluation methods are used to analyse the effects of a particular reform to the U.K. pension system. In April 2001, the U.K. government introduced Stakeholder Pensions - a new private pension arrangement. The reform also changed the structure of tax-relieved pension contribution ceilings, increasing their generosity for lower-earners. The chapter examines the impact of these changes on private pension coverage using individual level data. The analysis therefore provides U.K. based evidence on the important policy question of whether adjustments to financial incentives can induce extra retirement saving, a question which has been extensively addressed in a north-American literature (the paper that is most related to chapter 5 in terms of the type

of reform analysed, is Milligan (2003)).⁴ The assessment presented in chapter 5 uses a difference-in-differences strategy with an estimator that is modified to allow for dichotomous outcomes (the method follows Blundell et al. (2004a)). Contrary to the conventional wisdom that the Stakeholder Pension reforms had little or no impact on saving behaviour, our results indicate that the change to the contribution ceilings affected private pension coverage rates among lower-earners, especially among women. This is in line with the nature of the reform to incentives implicit in the change in tax relief. While we do not formally test these findings against the predictions of the lifecycle model, this pattern is at least consistent with the idea that savers make forward-looking decisions about pensions provision.

The final chapter of the thesis is a conclusion. Rather than repeating the conclusions of each of the substantive chapters (which have been summarised in this introduction), the conclusion focuses on some developments and further work that would build on the analyses presented in the thesis.

⁴Other recent related papers include Benjamin (2003) and Chernozhukov and Hansen (2004), and the literature builds on a stream of research summarised in Journal of Economic Perspectives (1996).

Chapter 2

Modelling the Demand for Housing over the Lifecycle

2.1 Introduction

The aim of this chapter is to model carefully the demand for home ownership over the life-cycle and to show how sensitive demand is to price fluctuations and to changes in capital market imperfections. As explained in chapter 1, the exercise is of interest in its own right, but it also provides useful background to the material in the following two chapters.

We present a model where households choose throughout their lives whether or not to own a home, and choose between houses of different size. These choices are made in the face of uncertainty about earnings and about house prices and in the presence of various capital market imperfections. Capital market imperfections are widespread in the housing market: individuals 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; there is very little insurance provided against house price risk, and the lack of insurance against earnings risk makes housing more risky.

Our analysis is related to a growing set of recent papers that have built a house

type asset into a lifecycle consumption saving framework.¹ Distinguishing features of our model are: careful modelling of mortgage-related borrowing constraints so that these are only checked when the household buys or has to renegotiate the mortgage with the bank to increase its value (see subsection 2.2.1 for a fuller discussion of this point); and, the modelling of housing as an asset that takes a discrete number of possible sizes so that there is a housing “ladder” but the dwelling is not continuously adjustable (even at cost). The contribution of Nichols (2005) used a similar two-size structure for the housing asset, but did not have the detail in modelling borrowing constraints that we have. Rios-Rull and Sanchez-Marcos (2008, In Progress) have a similar structure to ours but embed it in an equilibrium setting and concentrate on macroeconomic outcomes, rather than lifecycle decisions and welfare.

Section 2.2 presents the life-cycle model in detail, and, where relevant, relates it more closely to existing literature. Section 2.3 discusses 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 2.4 instead discusses how decisions are affected by shocks that occur during the lifetime. A final section concludes.²

2.2 Life-cycle model

We start from a standard model of lifecycle 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.

¹One example is Li and Yao (2007), which addresses issues related to those that we consider as they consider the behavioural and welfare consequences of house price shocks. Our remit is rather wider than this, and the nature of our housing asset and our realistic modelling of mortgage borrowing, distinguish our model from theirs.

²There is also an appendix with some supplementary material and computational details

2.2.1 Model Structure

A household lives for T periods. In every period $t \leq T$, the household maximizes lifetime utility by choosing consumption, c_t , and whether to own a “flat”, or to own a “house”, or to own neither, with $h_t \in \{0, 1, 2\}$ (where 0 is non-ownership).

The household value function in period t is given by:

$$V_t(A_t, h_{t-1}, p_t, w_t) = \max_{\{c_t, h_t\}} u(c_t, h_t) + \beta E_t V_{t+1}(A_{t+1}, h_t, p_{t+1}, w_{t+1}) \quad (2.1)$$

subject to

$$A_{t+1} = R_{t+1} \begin{cases} A_t + w_t - c_t - \kappa p_t(1 + F)I(h_t = 1) - p_t(1 + F)I(h_t = 2) \\ \quad \text{if } h_{t-1} = 0 \\ A_t + w_t - c_t + \kappa p_t(1 - F)I(h_t \neq 1) - p_t(1 + F)I(h_t = 2) \\ \quad \text{if } h_{t-1} = 1 \\ A_t + w_t - c_t - \kappa p_t(1 + F)I(h_t = 1) + p_t(1 - F)I(h_t \neq 2) \\ \quad \text{if } h_{t-1} = 2 \end{cases} \quad (2.2)$$

where A_t is the start of period asset stock and $R_{t+1} = 1 + r_{t+1}$ and r_{t+1} is the interest rate on the liquid asset; p_t is the price of housing which is realised at the start of period t ; F is the proportional fixed cost which is assumed to be the same for both buying and selling a house or flat; w_t is household earnings in period t . Equation (2.2) is a standard intertemporal budget constraint, augmented by terms reflecting the house price and transaction costs that must be borne when trading housing. For ease of exposition, we distinguish between beginning of period assets A_t and end of period assets s_t .

In our model, there are two differences between a flat and a house: first, owning a house gives more utility than owning a flat, and second, a house is more expensive than a flat. However, there is no explicit modelling of the size of flats or houses and

so the increased utility associated with housing in our model can be thought of as a reduced form for a preference for space driven by, for example, demographics. 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.

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

$$s_t \geq 0. \quad (2.3)$$

Home owners can borrow, and when they do so they are subject both to a terminal asset condition, $s_T = 0$, that translates into an implicit limit on borrowing,³ 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:

$$s_t \geq -\lambda_h \kappa p_t, \quad \kappa = \begin{cases} 0 < \kappa < 1 & \text{if } h_t = 1 \\ 1 & \text{if } h_t = 2 \end{cases} \quad (2.4)$$

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

$$s_t \geq -\lambda_w w_t \quad (2.5)$$

The explicit constraints on the downpayment 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

³The specification of marginal utility becoming infinite at 0 consumption means this terminal condition prevents households borrowing more than they can repay with certainty.

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

The structure of the constraints just described adds to the computational difficulty of our problem. It means that convexity preserving techniques cannot be used since there are known ‘kinks’ in the conditional value functions for owning a home, at least at points in the state space where this choice involves continued ownership.⁴ This computational difficulty is probably the key explanation of why it has been almost standard in the literature modelling 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.⁵ We were not willing to make such an assumption. The ability to borrow more when house prices and income move up and loosen borrowing constraints, without a concern that a subsequent 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.⁶

⁴By 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 optimisation problem, the level of assets (debt). At the point in the asset range where the most binding constraint switches between these two, the value function will be kinked. Heuristically, this can be thought of as having constraints on the optimisation that switch over within the state space of the problem, at a point at which the Lagrange multipliers on both constraints are strictly non-zero and will not be equal to each other (except by chance). Since one of the constraints ceases to apply without its associated multiplier declining smoothly to zero, this gives a kink in the value function.

⁵Examples 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).

⁶The modelling becomes even more complicated, but the issue perhaps even more pertinent, in a situation in which income is affected labour supply choices as well as random shocks; see Bottazzi, Low, and Wakefield (2007) for a model of this situation which includes mortgage constraints that only apply when buying or increasing the value of the mortgage, as described here.

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

$$m_t = r_t s_{t-1} \quad (2.6)$$

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.

(a) 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:

$$u(c_t, h_t) = \frac{c_t^{1-\gamma}}{1-\gamma} \exp(\theta\phi(h)) + \mu\phi(h) \quad \begin{cases} \phi = 0 & \text{if } h_t = 0 \\ 0 < \phi < 1 & \text{if } h_t = 1 \\ \phi = 1 & \text{if } h_t = 2 \end{cases} \quad (2.7)$$

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

(b) Stochastic processes

In the model households face uncertainty in two dimensions: idiosyncratic uncertainty over earnings and aggregate uncertainty over house prices.⁷

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

$$\ln w_t = a_t + v_t \quad \text{where } v_t = v_{t-1} + \xi_t, \quad \xi_t \sim N\left(-\frac{\sigma_\xi^2}{2}, \sigma_\xi^2\right) \text{ and } \rho_w = 1 \quad (2.8)$$

and a_t is the deterministic growth in earnings over the life-cycle and has a hump shape ($a_t = a_1t + a_2t^2$).

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

$$\ln p_t = d_0 + d_1t + \rho_h \ln p_{t-1} + \varepsilon_t \quad \varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \quad (2.9)$$

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

2.2.2 Model Calibration

We now estimate the parameters required for analysis. We impose values for some parameters, such as the elasticity of intertemporal substitution of consumption and the discount rate, using values from elsewhere in the literature. Some parameters we estimate directly from the data. These ‘externally fixed’ parameter values are reported in table 2.1. The rest of the parameters we obtain through calibration using the structural model outlined in section 2.2. These parameter values are reported in table 2.2.

⁷In fact it is not a difficult extension to include iid noise in the interest rate, but results were not sensitive to this and so we removed this dimension of uncertainty from our final runs; in equation (2.1) we implicitly assumed that the interest rate was not a stochastic state variable.

Table 2.1: 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
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>).

(a) External Parameter Values

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 2.9), 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 ($\rho_h = 1$). 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.

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. (2004b). We separately estimate the parameters for high and low education groups. The results in table 2.1 show firstly 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), 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

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

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. However, since we do not model bequests, households run down all assets by the end of life (age 81), leading to an overestimate of the amount of selling of homes towards the end of life. Our calibration and comparative static exercises focus on home-ownership behaviour up to age 60.

Interest rate process

For interest rates we use the average 90 day Treasury Bill discount rate in years 1968-1997,⁹ which gives a rate of 1.5%.

Utility function

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

Initial Wealth

We set the initial distribution of financial assets for the two education groups to match data on 22-26 year olds in the 2000 wave of the BHPS,¹⁰ and we assume that households have zero housing endowments at age 22.

⁹We 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.

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

(b) 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 2000,¹¹ by education group. Our approach is to choose the parameters to minimise the sum of squared deviations between moments calculated in the data and corresponding simulated moments. The moments we use are the average home-ownership rates for households in low and high education groups, for those aged 26-35 and those aged 36-60. The statistics from our model are measured across 40 different simulation runs (i.e. 40 different realised 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.

We use this calibration process to pin down the transactions cost of buying or selling, F , and the parameters specifying the utility benefit of home ownership, μ and θ . Parameter values from the calibration are summarised in table 2.2. Table 2.3 presents the calibration statistics, showing how home-ownership rates predicted by the model match those observed in the data.

Table 2.2: Calibrated Parameter Values

Parameter	Value
θ	0.026
μ	0.11
F	0.05

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

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

Table 2.3: Calibration Statistics

Statistic	High Education		Low Education	
	Data	Model	Data	Model
<i>Ownership rate (percentage)</i>				
Age 26 - 35	65.4	68.2	49.5	49.2
Age 36 - 60	81.1	76.5	62.7	62.8
Sum of squared deviations	28.74			

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

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 calibration runs showed that a range of different combinations of μ and θ (with higher μ associated with higher θ) would result in a very good fit for home-ownership for the low education group. Finding the values that worked best across the two education groups involved finding a (μ, θ) pair that produced a reasonably good fit for the high education group while also being close to a point of good fit for the low education group.

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 that cross-partial derivative of utility with respect to c and h is positive.

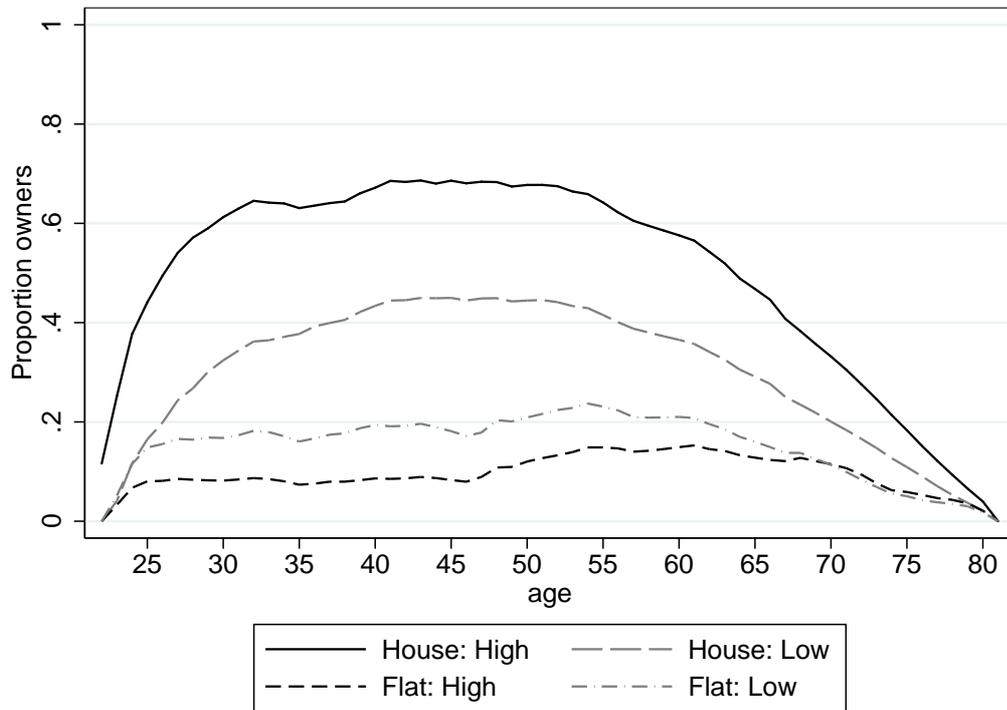


Figure 2.1: Home Ownership in Baseline Run

Figure (2.1) 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 realisations 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 realisations of the house price at each age. In the next chapter we will be even more careful about the aggregate nature of the house-price process, inputting observed re-

alisations of the house price in each year into our simulations for the calibration and thus capturing that different cohorts experience the same shocks at different ages. As well as figure 2.1 charting ownership by age, in appendix table A2.1 we also report ownership rates in certain age bands. This table supplements the figure and the reported ownership rates are useful when interpreting the elasticities described in later sections.

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 get such numbers from the data, 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.¹² Data from this survey in 2002/03 indicate that the mean (median) family wealth held in financial assets and private pensions was approximately 9.2 (4.7) times median income for low income individuals aged 51 -60, and for high education individuals the figures were 16.8 (11.0). The nearest equivalent measures for the low education group in our baseline simulation are 9.5 (7.2), which is a reasonably good match. For the high education group the simulation numbers are 8.2 (5.8).

The less good match for the high education group may be partly due to the fact that the pension replacement rate in our model is 70% for all individuals, and this is higher than the replacement rate that higher income individuals can expect from state pensions in the UK. The absence of inheritances (apart from some initial wealth) and bequests from the simulations may also be a simplification that is less realistic for the high education group.

¹²For more details on this, see Banks et al. (2005). Thanks to Gemma Tetlow for help dealing with the ELSA data cited here.

2.3 Sensitivity to income and house-price parameters

In this section, we use our calibrated model to show how housing demand depends on the level of income, the house price and on capital market imperfections and on frictions in the economy such as the fixed transactions cost. We focus on changes to parameters that affect individuals across their entire lifetime, and first consider parameters that relate to the income process, before addressing parameters of the housing market. In the first two subsections we consider how the parameters in question affect housing and consumption demand. There are two different aspects to housing demand that we consider: one is the demand for ownership of any size house; the second is the demand for flats versus houses. A final subsection looks in more detail at the importance of fixed costs, focussing on the welfare consequences of changes in this parameter.

To look at housing and consumption demand responses, we utilise elasticity measures. Since consumption is a continuous variable, we can calculate the responsiveness of consumption to changes in the parameters of our model using a standard elasticity measure. Since home-ownership is discrete, we calculate the responsiveness of home-ownership through “quasi-elasticities”, measuring the proportional change in the homeownership rate to proportional changes in a given parameter. For example, a (lifetime) income elasticity (ϵ_Y) measures the proportionate change in homeownership in our simulations, to a given proportionate change in incomes within our simulated economy. We calculate this by comparing behaviour in the baseline economy to that in an economy with all incomes increased by one percent.¹³ We calculate elasticities within particular age groups, as well as for lifecycle home-ownership. This elasticity measure captures the effect of an anticipated increase in lifetime income. To aid interpretation of the elasticities for the housing good, in table A2.1 in the appendix, we report home ownership rates in our baseline run, since these enable the calculation of quantity effects from the elasticity measures.

¹³Apart 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.

2.3.1 Changes related to income

Tables 2.4, 2.5 and 2.6 show the extent that home ownership and consumption respond to a change in lifetime income, to a change in the variance of earnings and to a change in the borrowing constraint that depends on income.

Table 2.4: Elasticities for ownership: Factors relating to income

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

The main message from the results in table 2.4 is that home ownership of the young is the most sensitive to income parameters. Higher income or a loosening of the income constraint leads to greater home ownership among the young but makes less difference to the older groups. This message remains the same if we consider the quantity of housing bought rather than just the effect on ownership,¹⁴ although the quantity elasticities tend to be somewhat larger than the ownership elasticities.¹⁵

¹⁴We define quantity by counting a flat as 0.6 of a house, as is defined in the price process.

¹⁵Comparing ownership and quantity elasticities is one way to assess the importance of including

Table 2.5: Elasticities for flats and houses: Factors relating to income

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

Comparing quantity and ownership elasticities is one way to think about the importance of including different sizes of property in the model. Another is to consider separate elasticities for flats and houses (table 2.5). The pattern here is revealing: for both the change in the income level, and in the income-related borrowing constraint, the elasticity of demand for flats is negative while for houses it is positive. These elasticities indicate that the flat is an ‘inferior good’,¹⁶ and that demand for flats falls as the borrowing constraint makes housing more affordable. These two properties of the

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

¹⁶This feature is reliant on the fact that a change in the price of housing affects flats and houses equally.

demand for flats arise because, as the change in question makes housing more affordable, more people upgrade from a flat to a house than from not owning to owning a flat, and thus the demand for flats falls. These properties of the demand for flats would be important in a model in which the price of housing is endogenous (see chapter 4). The reason why this is true is discussed in the next section.

Table 2.6: Consumption elasticities: Factors relating to income

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

Table 2.6 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.

2.3.2 Changes relating to house prices

Tables 2.7, 2.8 and table 2.9 show the extent that home ownership and consumption respond to an increase in the level of house prices at all ages, to a change in the variance of shocks to house prices, to a change in the downpayment constraint and to a change in the fixed cost. The consumption responses are rather small and we shall not comment on them further.

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

Table 2.7: Elasticities for ownership: Factors relating to house prices

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

the downpayment requirement becomes greater and as fixed costs go up. The response to a tightening of the credit constraint is again greater for the low educated. Another pattern that is similar to that from the changes related to income is that - with the partial exception of the responsiveness to the variance in income shocks - quantity elasticities again tend to be slightly larger than ownership elasticities.

As when considering changes in factors relating to income, we can look in more detail at how housing demand changes are allocated across the two types of housing, by considering elasticities for houses and flats separately (see table 2.8). Again we see that in some cases the elasticities for flats and houses have opposite signs. This is most evident for the elasticities when the price of housing goes up: in response to this change the demand for houses declines but the demand for flats increases. Given the parameterisation, the flat in our model is a *Giffen good*.¹⁷ This property is due to the

¹⁷For this conclusion it is important that there is a single price for housing, so that the prices of

Table 2.8: Elasticities for flats and houses: Factors relating to house prices

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

fact that when the price of housing increases more people shift from house ownership to flat ownership, than shift from flat ownership out of the housing market. In line with the comments of the previous subsection, this property of the demand for flats would be important in a model in which the price of housing is endogenous (see chapter 4). The ‘Giffen good’ property emphasizes how in an equilibrium setting quite different price movements will be required to choke off (stimulate) the changes in demand for different types of property that arise from aggregate economic shocks.

2.3.3 Responsiveness to fixed costs

In the previous two subsections we have looked at how home ownership and consumption are affected by the parameters of our model. In this subsection we extend the analysis for one particular parameter, the level of the fixed cost, and think about halv-
flats and houses go up at the same time.

Table 2.9: Consumption elasticities: Factors relating to house prices

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

ing this cost from 5% to 2.5%. The fixed cost is the one institutional parameter that was calibrated in our moment matching exercise, and it is also a feature of the housing market that is directly affected through by tax policy.¹⁸ 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 2.7 and 2.8 that increases in fixed costs lead to a reduction in home ownership and to a delay in buying. Further, Figures 2.2 and 2.3 suggest that higher fixed costs mean households are less likely to buy flats before buying houses, and less likely to trade down to a flat after owning a house. Clearly increasing the fixed costs increases transactions costs, causing a substitution effect (demand for housing falls, and consumption increases) and a wealth effect (for a given number of home purchases at given house prices, less 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 and the welfare consequences of a change in fixed costs.

We separate out the wealth and substitution effects by first running the model with a low fixed cost that is known in advance (column ‘c.’ in table 2.10). The impact of this low fixed cost relative to the baseline scenario (column ‘a.’) combines both

¹⁸In the U.K., by stamp duty land tax.

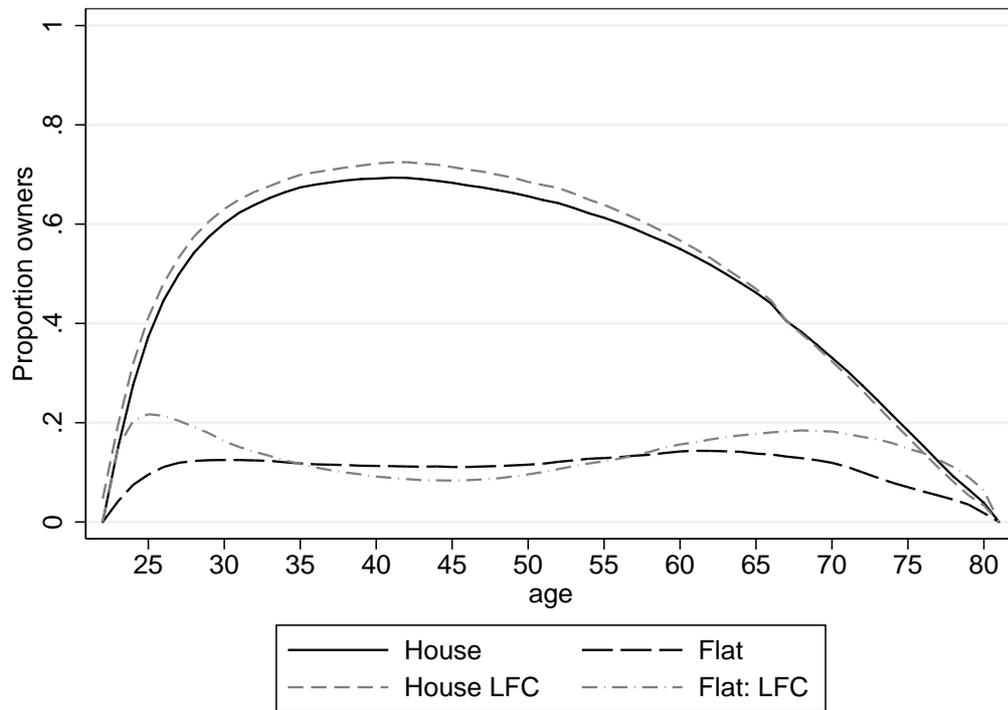


Figure 2.2: Response of ownership to change in fixed costs, high education

substitution and wealth effects. Then we run the model with households expecting to pay the baseline fixed cost, but when transactions occur the household pays only the low fixed cost (column ‘b.’). There is therefore no substitution effect associated with the lower fixed cost.¹⁹ The results of this exercise are reported in table 2.10, 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 2.10 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

¹⁹The part of the wealth effect which would induce an ex-ante change in behaviour is also excluded from this experiment and so only the ex-post effect of higher wealth is calculated.

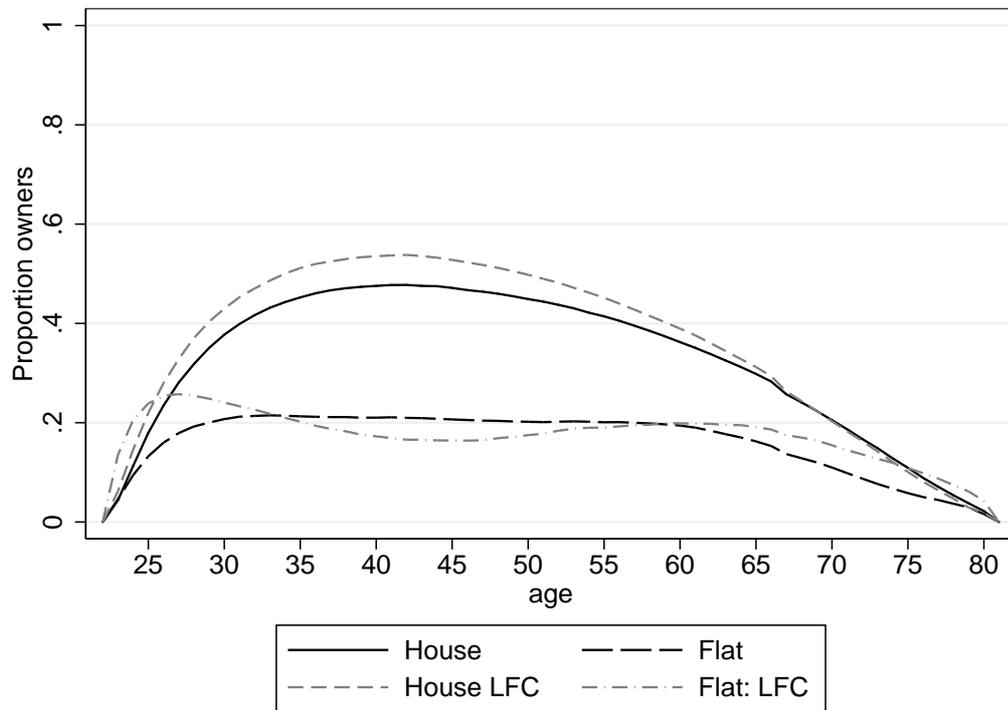


Figure 2.3: Response of ownership to change in fixed costs, low education

at the beginning of life (third from last row in each panel) and by a “consumption equivalent”, which is the proportional increase in consumption at every age that is needed to make agents indifferent (ex ante) between the two scenarios being compared. Since the utility function with housing terms, and the credit constraints, means that the model does not have a solution that is homothetic throughout, the consumption equivalent can not be calculated analytically. Instead it is found through an iterative procedure that relies on expected discounted lifetime utility as derived in the numerical solution to the model for comparisons between the scenarios reported in columns a. and c. For comparisons involving the scenario for column b. in which the low fixed cost is a “surprise”, there is no numerical solution for expected utility (rather the solution for the relevant policy functions is that with the higher fixed costs). For these cases the iterative procedure for finding the consumption equivalent relies on simulated expected

lifetime utility. Fuller details of the computational procedure for finding the welfare measures are reported in the appendix at the end of this chapter.

Table 2.10: Effects on consumption, homeownership, and welfare, due to changing fixed costs

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

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

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

The table shows that the halving of the fixed cost leads to an increase in consump-

tion, and also to an increase in the expected level of home ownership (or equivalently, in the expected number of periods of ownership), across the lifecycle. The results also show the extent to which the cut in fixed cost leads to an increase in the number of housing market transactions during the lifetime: with the lower fixed cost, the likelihood of moving directly from flat to house ownership (i.e. of ‘trading up’) at some age is more than doubled compared to in the high fixed cost case, while the likelihood of the converse change in ownership (‘downsizing’) approximately trebles. Also evident from the table is that consumption behaviour is almost the same in the scenarios of columns ‘b.’ and ‘c.’, while housing market behaviour is almost constant between scenarios ‘a.’ and ‘b.’. That is, the change in the consumption behaviour happens even when the change in the fixed cost is not expected, but the change in behaviour in the housing market only occurs when decisions are made on the basis of knowledge of the lower transactions cost.

The welfare measures in table 2.10 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 lower 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.

2.4 The Response to Shocks

In the previous section (subsection 2.3), we saw how responsive housing demand and consumption are to changes in house prices or income that affect the entire lifecycle. In this section, we contrast those results with the effect of a small shock to 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.²⁰ We consider the effect on ownership and consumption of house price and income shocks occurring at age 30, 40 or 50.

Figures 2.4 and 2.5 show, for high and low education respectively, the average effect on home ownership (panel ‘a.’) and consumption of a 1% upwards shock to the house price at age 30, 40 and 50.²¹ The results show larger percentage impacts for the low education group, as the fact that the house price is a bigger proportion of lifetime wealth in the low education group (on average) outweighs the fact that fewer low education individuals own.

Though the size (in percentage terms) of effects is different across the two education groups, patterns of results are rather similar. The effect of the shock on home-ownership rates is an initial decline, with the largest reduction in home-ownership (of around 1% for the high education group, and 1.2% to 1.5% for the low education group) experienced around five years after the shock. The peak decline does not occur in the year of the shock because an initial contraction in the number of buyers has a lasting effect on the stock of home owners, and this is compounded by the persistence in the house price process, meaning that the shock continues to discourage new purchases for several years after it is initially felt. In the longer term the home-ownership rate

²⁰Since income is stochastic, the experiment is that the realisation of income is 1% higher than a draw from the distribution would imply.

²¹Each of the lines plotted in figures (2.4) and (2.5) shows the percentage difference in a. the home ownership rate or b. average consumption in two simulation runs between which the only exogenous change is to the house price shock at the given age such that the level of the house price at that age is increased by one percent. The simulation runs are based on 5000 sequences of house price shocks, and each of the 5000 drawn prices at the relevant age are increased by exactly one percent.

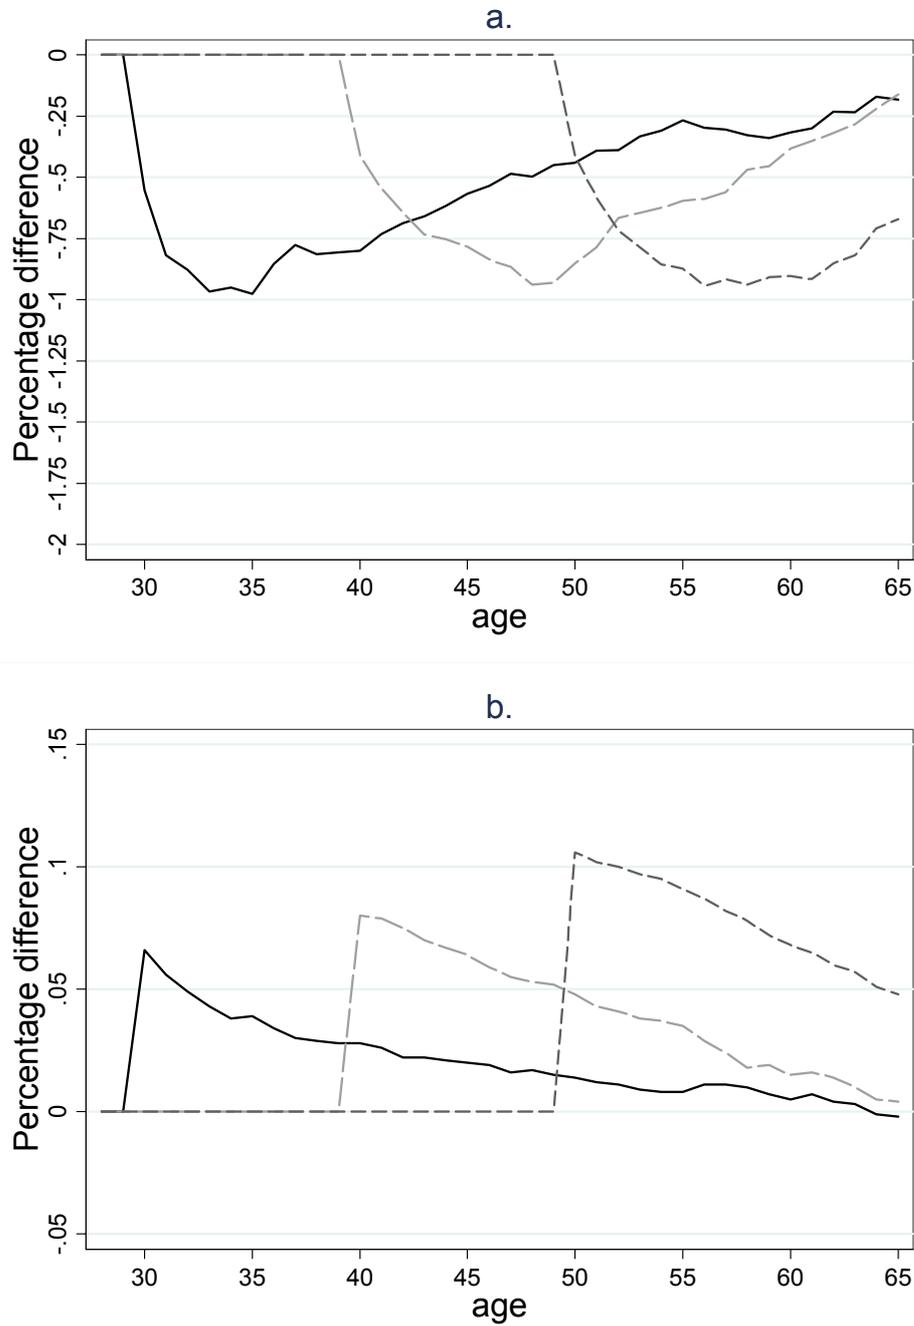


Figure 2.4: Effect of 1% house price shock at 30, 40 or 50, on: a. home ownership and b. consumption; High education.

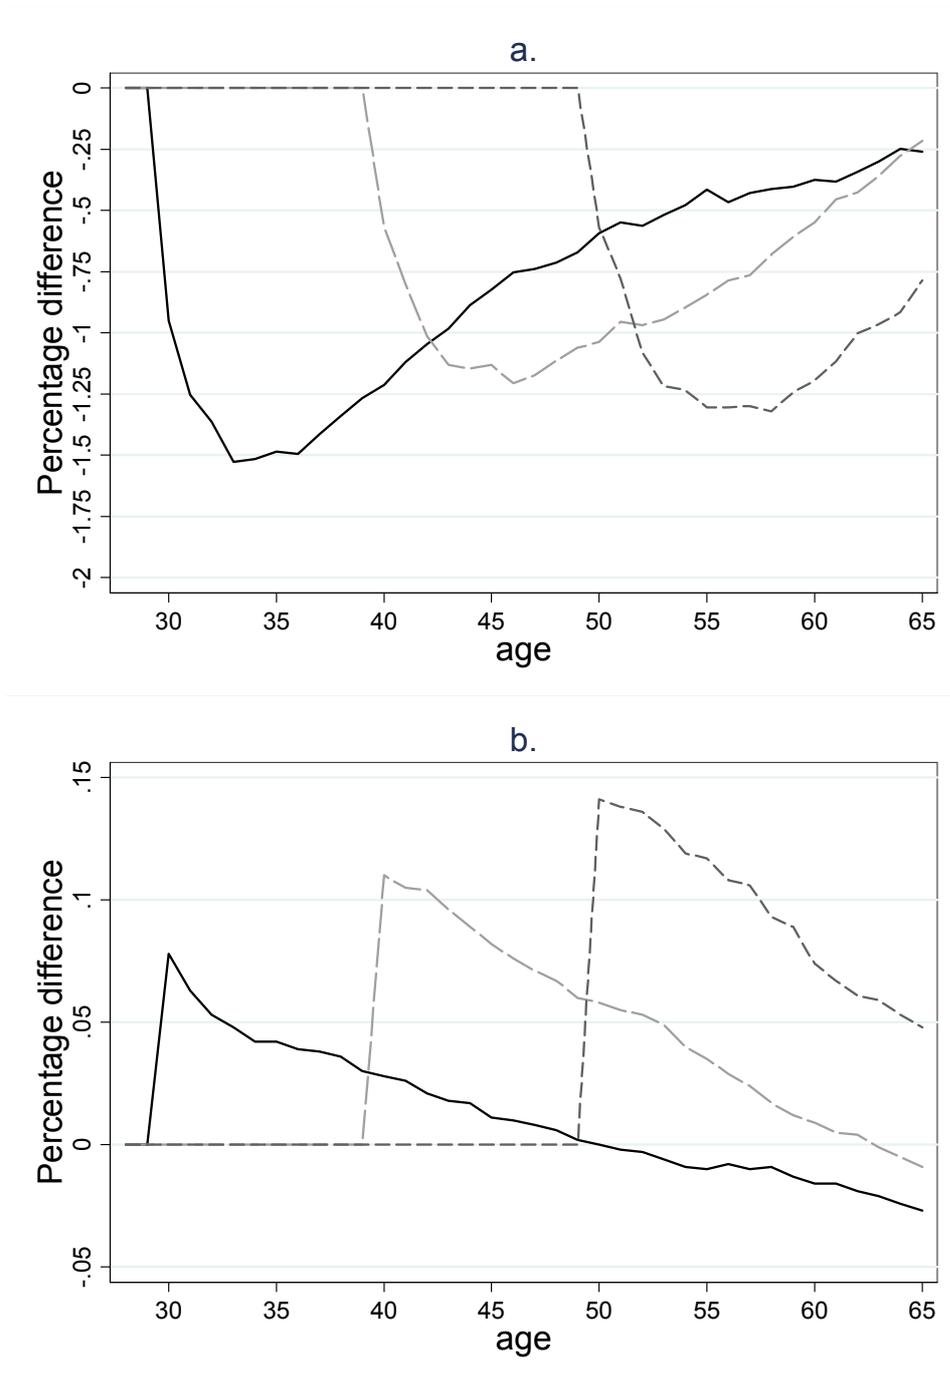


Figure 2.5: Effect of 1% house price shock at 30, 40 or 50, on: a. home ownership and b. consumption; Low education.

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 A2.1 shows the rate at which the effect of a shock dissipates from the expected level of the house price). It also reflects that over time agents can modify their saving behaviour to mitigate the shock.

Turning to the effect on consumption, we see that in the model the positive house price shock increases average consumption regardless of the age at which the shock is experienced. The increase in consumption is seen to be strongest for the eldest group who are most likely to have large equity in their home, and who have the shortest horizon (on average) over which to spread extra wealth. This pattern across age-groups confirms the intuition behind the empirical analyses of Attanasio and Weber (1994) and Attanasio et al. (2009, 2005b) investigating why house price shocks have been strongly correlated with consumption growth in the U.K. over the past 30 years; we shall return to that issue (and a fuller discussion of related literature) in the chapter 3.

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

Figures 2.6 and 2.7 show, for high and low education respectively, the effect on average on home ownership (panel ‘a.’) and consumption of a 1% upwards shock to income at age 30, 40 and 50.²² As with the house price shocks just discussed, the

²²Each of the lines plotted in figures (2.6) and (2.7) shows the percentage difference in a. the home ownership rate or b. average consumption in two simulation runs between which the only exogenous change is to the income shock at the given age, such that the level income at that age is increased by

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

Figures 2.6 and 2.7 show that income shocks at the given ages lead to increases in both home-ownership rates and non-durable consumption. The increase in consumption is immediate and (since income shocks are permanent²³) persistent. The proportionate increase in consumption is somewhat less than the proportionate increase in income. One reason for this is that some of the increase in lifetime wealth is spent on housing. A second reason also explains why older groups of individuals increase their consumption less (in proportionate terms) on average when the income shock occurs. These individuals enjoy the income shock for a smaller proportion of their 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 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.

exactly one percent for every individual. The simulation runs are based on 5000 individuals each with their own realisation of the (idiosyncratic) income process.

²³Income is a random walk.

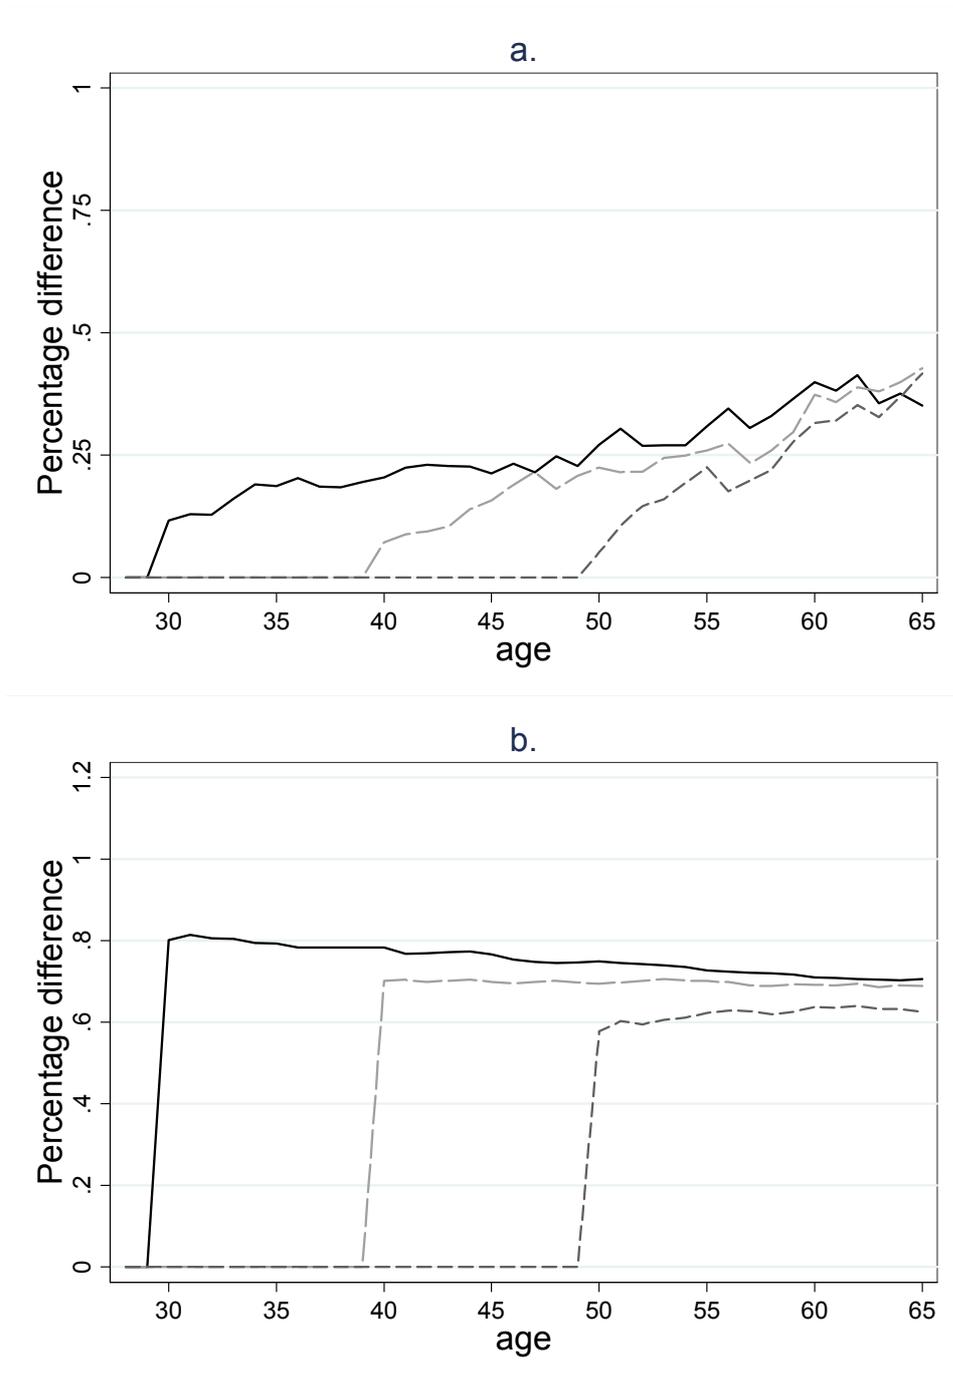


Figure 2.6: Effect of 1% income shock at 30, 40 or 50, on: a. home ownership and b. consumption; high education.

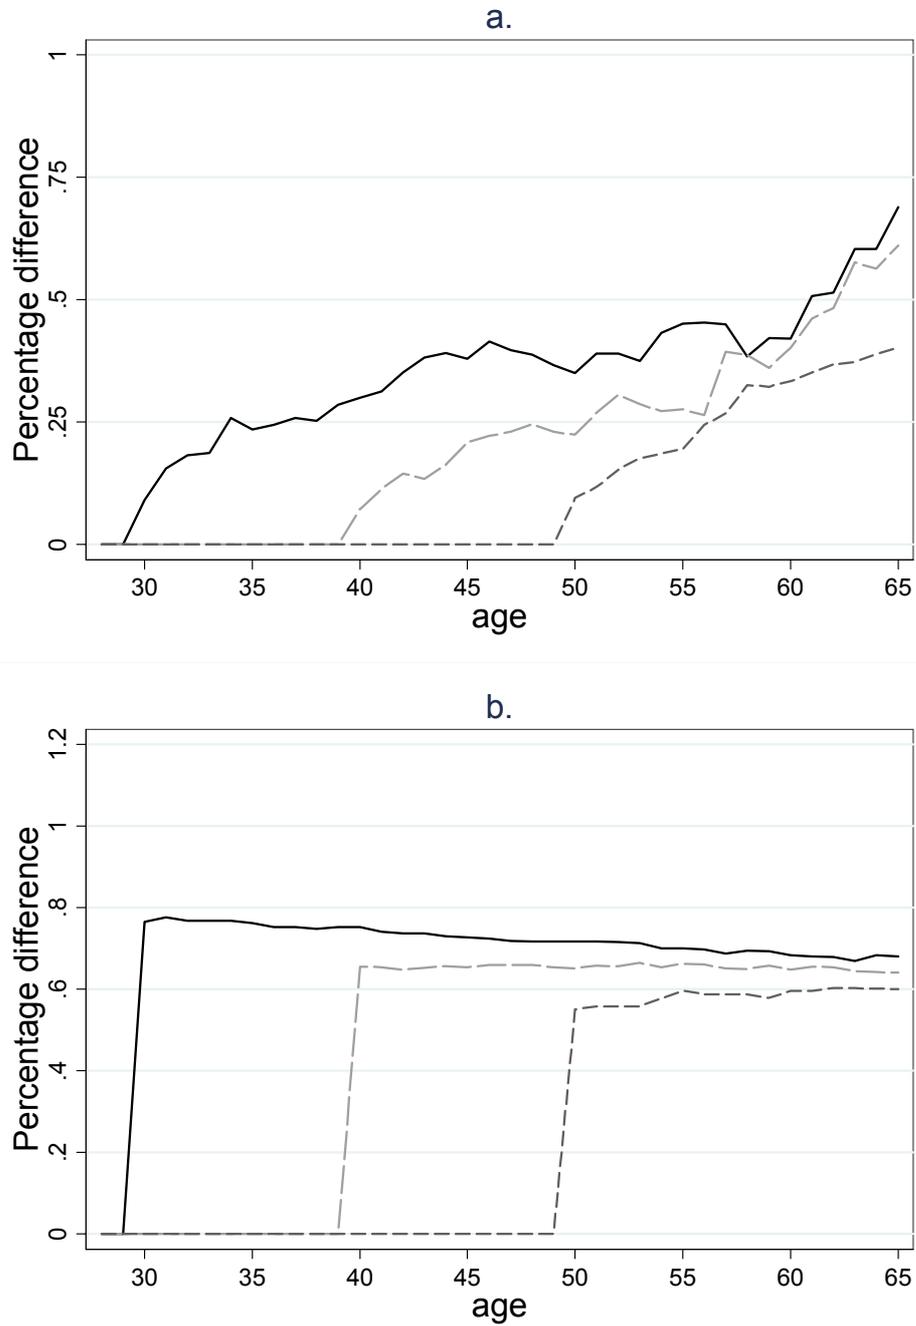


Figure 2.7: Effect of 1% income shock at 30, 40 or 50, on: a. home ownership and b. consumption; low education.

2.5 Conclusion

In this chapter we have seen a life-cycle model of households choosing how much to consume or save, and whether to own a flat, a house or no housing, in each period of life. The model was carefully constructed to realistically capture mortgage-related borrowing constraints, which are often represented in a stylised way in existing literature.

Within the model we have examined how sensitive households choices are to the parameters of the house price process, and the income process, and also to tightness of mortgage-related borrowing constraints. We found that the level of the house price and income have marked effects on behaviour and that when incomes go up, and when the price of housing (i.e of houses and flats goes up) this leads to an increase in demand for houses, but to a *fall* in the demand for flats.

We have also seen that for a given level of house prices the level of (proportional) fixed costs in the housing market can have a noticeable effect on consumption and welfare (primarily through a lifetime income effect) and on housing market behaviour (a substitution effect). In the final part of the chapter we saw that shocks to house prices or income during the lifecycle will have somewhat different effects on behaviour depending on the age at which they occur.

As mentioned in chapter 1, the exploration of the model in this chapter, as well as being interesting in itself, can be seen as a preparatory exercise for the work in the next two chapters.

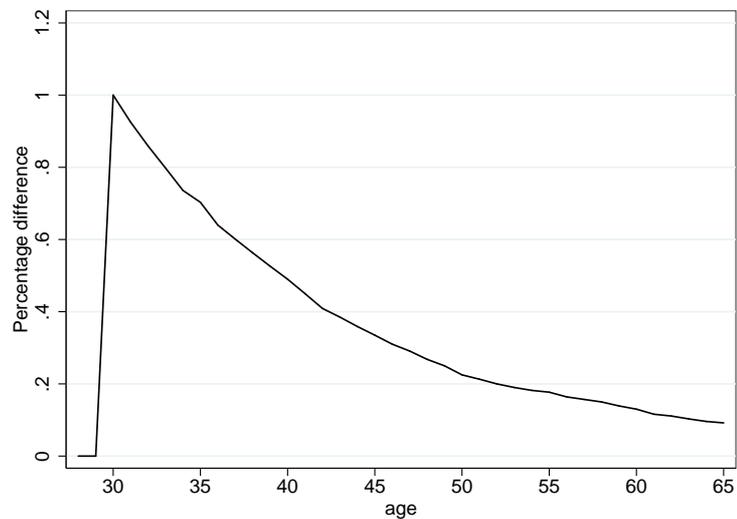
A2 Appendix to chapter 2

A2.1 Extra tables and pictures

Appendix Table A2.1: Ownership rates

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

Table A2.1 shows ownership rates by age group in the calibration run.



Appendix Figure A2.1 Effect of HP shock at age 30 on expected HP at older ages

Figure A2.1 shows the effect on the expected house price at ages 31 to 60 of a 1% house price shock at age 30.

A2.2 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.²⁵

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

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

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

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.

²⁵The development of these techniques was a key part of the research for the project RES-000-23-0283, funded by the U.K. Economic and Social Research Council (ESRC), (Attanasio et al., 2005c).

cesses. This is done as described in Tauchen (1986). We preferred Tauchen’s method of equally spaced nodes over the quadrature based method proposed by Tauchen and Hussey (1991), because this has been shown to be more robust to very high persistence in the modelled processes (Floden, 2007).

For the runs presented in this draft of the paper, we have used 105 nodes in each ‘conditional’ asset grid (we have separate grids underlying each conditional value function, since assets are limited by different borrowing constraints depending on the home-ownership choice). Points are more dense in the lower range of the asset grids, to make sure that non-convexities in the value function are not overlooked in the maximisation process. Income and the house price are each represented by a grid of fifteen nodes. Monte-Carlo experiments showed that these grid sizes were sufficient to capture the modelled processes to a high degree of accuracy. With this set up, the model solution and simulation takes around 3 hours on a desktop PC.

Method for deriving consumption equivalent measures

When constructing welfare measures in subsection 2.3.3 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:

$$\tau = (\overline{U}_0/\underline{U}_0)^{(1/(1-\gamma))} \quad (2.10)$$

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 2.10 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 2.10.
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 2.10 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 2.4), we had to set up the discrete approximation to the relevant price in such a way that an increase of exactly one percent could be implemented. This required altering the grids underlying our Markov chain approximations, and in particular a move away from the evenly spaced grids that we adopted for our baseline solution and simulation. In order to maintain the accuracy of the approximation, we increased the number of points in the relevant grid.

Chapter 3

The Coincident Cycles of House Prices and Consumption in the U.K.: Do House Prices Drive Consumption Growth?

3.1 Introduction

Over much of the past thirty-five years, the cycles of house prices and consumption have been relatively closely synchronised in the UK. Figure 3.1 shows the remarkable comovements of house price growth and consumption growth (notice however that the scale of house price movements is two and half times that of consumption growth) . The correlation between house prices and consumption makes changes in house prices an important indicator for those wishing to judge inflationary pressures within the economy. Indeed, this indicator is closely watched by the Bank of England's Monetary Policy Committee precisely because, in the words of Nickell (2004), "... The evidence suggests that house price inflation is significantly related to household consumption

growth and hence to aggregate demand growth and future consumer price inflation in the economy”. The perceived importance of this indicator is also illustrated by recent press discussion of what could happen if house prices were to decline or stagnate from current levels after an extended period of real growth (see Economist (2007), Giles (2008)).

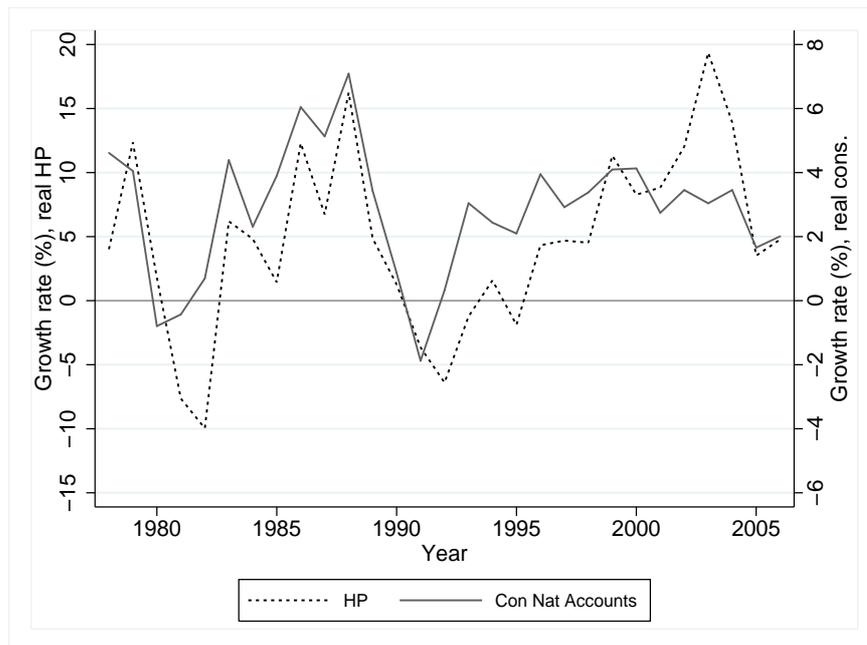


Figure 3.1: Growth in real consumption and house prices

To understand the implications of house price movements for consumption levels in the economy, and therefore the appropriate policy response to conditions in the housing market, it is necessary to understand what drives the link between house prices and consumption. Three main mechanisms to explain this co-movement have been proposed in the literature:

1. A “wealth effect”: increases in house prices raise households’ wealth, raising their desired level of expenditure.
2. A “credit constraints” channel: house price growth increases the collateral available to homeowners, thus loosening borrowing constraints and facilitating higher

consumption (see Muellbauer and Murphy (1997) and particularly Muellbauer and Murphy (1990) for a discussion of mechanisms 1. and 2)¹.

3. A “common causality” model: factors such as changes in expected income growth, tax changes or changes in credit market conditions lead to increases in both households’ expenditure and house prices (see King (1990) and Pagano (1990)).

It has been argued that these different mechanisms might have effects that are observationally equivalent if one looks only at macro-data, and so researchers must use data on the behaviour of individuals or households in order to disentangle the different effects. Attanasio and Weber (1994) (hereafter AW, 1994) used such micro-data in order to examine the consumption boom of the late nineteen-eighties, which coincided with a rise on real house prices of slightly more than 40% in the space of four years. In order to identify what features of the micro-data they should look for, these authors drew on insights from a life-cycle model. AW (1994) argued that if wealth effects were important, then these were likely to have the biggest effects on the consumption of older individuals who are most likely to have equity in any housing assets, and who have a relatively short time horizon over which to distribute the consumption of the extra wealth. They also simulated a simple lifecycle setup (with no housing asset) and showed that an upward revision of expectations of future income (productivity) could lead to the consumption of the young (with longer to enjoy the higher income stream) responding more strongly than that of the old. The finding that the eighties consumption boom was driven in a large part by strong consumption by the young, was therefore a key component of an argument that this boom, and the simultaneous boom in house prices, may have owed much to common causes. AW’s exercise was recently extended in a variety of ways by Attanasio et al. (2009, 2005b), (hereafter ABHL), who, by and large, confirmed the results in AW (1994).

The complicated nature of the relationship between housing assets and non-durable

¹Recent evidence relating to this issue can be found in Aron et al. (2007)

consumption means that it may be perilous to rely on a stylised model for insights. AW themselves acknowledged this when mentioning that the consideration of credit constraints, and modeling of housing decisions and house price booms, would be interesting extensions that were beyond the scope of their model. The purpose of this paper is to derive the implications of changes to house prices and earning innovations for the consumption of different groups in the population. We can therefore check in a rigorous fashion whether such a structural model confirms intuitions about how and why the consumption of different groups might correlate with house prices, intuitions which have formed the basis of past empirical tests. We will also explore whether our model can be used to add to our understanding of what has driven consumption growth during fluctuations over the last 35 years in the U.K. economy, and by performing counterfactual simulations attempt to quantify the effects of different factors that move consumption.

Our central contribution comes from our systematic use of a realistic structural life-cycle model of consumption, savings and housing choices, to inform the interpretation of empirical analyses intended to distinguish between the three mechanisms proposed to explain the correlation of house prices and consumption. The first part of the exercise that we undertake involves constructing and numerically solving the structural model which is the main tool of our analysis, and calibrating this model to match U.K. data. The model, which we describe in detail in section 3.3, includes some innovative features. We model the financial markets available to the agents in our model to be a realistic representation of the U.K. mortgage market. We calibrate the stochastic processes faced by our agents to include both idiosyncratic and aggregate components. The former are calibrated using micro data. The latter include *aggregate* shocks to house prices and incomes (which are experienced by everyone in the economy at the same point in time). We estimate the parameters of the time series processes for house prices and aggregate earnings using data covering the last 35 years in the U.K.. We are not aware of other studies that use this combination of aggregate and micro data

in the calibration of an individual-level model and yet it is, in our opinion, important given that we want to understand aggregate fluctuations by aggregating individual consumption.

We choose the parameters of our individual model (such as the preference for housing services) to reproduce some cohort level facts, such as the level of home ownership and its evolution over the life cycle. Having successfully matched these moments, we can simulate individual behaviour given the set of aggregate innovations to house prices and earnings that we estimate on our time series data. Aggregating up the simulated behaviour of heterogeneous individuals, we are able to check the extent to which our model is able to reproduce the features of aggregate consumption growth in the U.K.. Notice that the moments of aggregate data are *not* used to calibrate our individual-level model. Thus this type of comparison to aggregated data forms an extra check on the validity of our model (alongside comparison to micro data moments), and such a check might be used in other contexts where micro models should reproduce patterns of aggregated, as well as individual, behaviour.

The next step of our analysis is to simulate behaviour under a set of counterfactual scenarios in which the mechanisms that might drive the link between house prices and consumption are shut down in turn. The construction of these counterfactuals is, by definition, an exercise that cannot be undertaken using data and (as explained in section 3.4) forms our main means for exploring how we might disentangle the influence of the different mechanisms that might drive the link between house price shocks and consumption growth.

Our results (see section 3.4) provide a firmer theoretical grounding for reduced form empirical analyses of the type conducted by AW and ABHL. We show that in a model with credit constraints and simultaneous housing and consumption choices, a house price shock that drives consumption changes through a wealth effect will lead to the biggest consumption responses from older groups. As documented in AW and ABHL (whose results are described in more detail in section 3.2), this age pattern of responses

is the reverse of that observed in data. Our model also shows that the pattern in the data could instead be explained by a shock to aggregate incomes of the kind that has been suggested as a possible common cause of house price and consumption growth. In addition to these basic intuitions, we can quantify the size of these effects in the model, within a framework that is shown to fit the data in a number of dimensions.

The structure of the paper is as follows. Section 3.2 provides a fuller description of evidence which points towards the conclusion that factors other than wealth effects have driven the recent correlation between house price shocks and consumption growth in the U.K.. Section 3.3 describes the lifecycle model that we use, and how we calibrated this model to match recent U.K. history. Section 3.4 contains our main analyses of different scenarios within the model, and discusses how the analysis of these scenarios relates to empirical results. The final section concludes.

3.2 Existing evidence

Amongst recent papers that have used micro-data to assess the relationship between house prices and consumption in the U.K., there appears to be little consensus. Campbell and Cocco (2007), argue that their empirical results suggest that wealth effects are the most likely explanation of the correlation between house prices and consumption. They find that increases in consumption observed during recent house price booms, were mainly driven by increases in the consumption of home owners (rather than renters) and older consumers rather than younger ones. Campbell and Cocco (2007) also use a structural model to assess whether endogeneity of the home-ownership decision might bias their empirical results. By contrast, ABHL find the relationship between consumption growth and house prices to be strongest for younger individuals as argue that the common causality channel is most important. Both papers use micro-data from the UK Family Expenditure Survey / Expenditure and Food Survey

and yet reach very different conclusions.²

To be more precise about the findings, we summarise the exercise of ABHL. Their baseline specification is the following:

$$\ln X_t^{ch} = \alpha^c + f(\text{age}) + \gamma' z_t^{ch} + \varepsilon_t^c + u_t^{ch} \quad (3.1)$$

where c stands for cohort and h stands for household; X is non-housing expenditure, α^c a cohort-specific intercept, $f(\text{age})$ is a quintic in age, z_t^{ch} are a set of demographics, and ε_t^c and u_t^{ch} are cohort-specific and household-specific error terms.

The purpose of running this baseline consumption function, is to look at which groups deviate most markedly from predicted consumption levels in times of consumption boom or consumption bust. ABHL look at this question by calculating residuals from the regression, averaging these residuals by year and age-group, and then plotting the results. We have updated this analysis using data from 1978 to 2005/6 (ABHL had data running until 2001/2), and the results are shown in figure 3.2. The figure shows that in times of large consumption boom or bust - particularly the late 1980s boom and the mid 1990s slump - it is the consumption of the young that most deviates from expected levels. These are also times of unusually strong house price growth. Thus the figure is a good indication that it is the consumption of the young that is most strongly associated with house price growth. To anticipate section 3.4, a tool in our analysis of how our model relates to the data will be to recreate figures similar to 3.2, but using simulated data.

To look in more detail at the role that house price shocks have in explaining these patterns in residuals, ABHL extend the specification 3.1 to explicitly control for the

²It is unclear why this is the case, though there are several methodological differences between the papers. A comparative study by Cristini and Sevilla (2008) attempted to replicate both studies as closely as possible. The ABHL results were found to be robust, whereas the Campbell and Cocco results were sensitive to the specification.



Figure 3.2: Consumption function residuals, by age group

influence of (regional) house prices. The specification is extended as follows:

$$\ln X_t^{ch} = \alpha^c + f(\text{age}) + \gamma' z_t^{ch} + \theta'_{ag} g(\text{hp}) + \varepsilon_t^c + u_t^{ch} \quad (3.2)$$

where ag are three age groups: young (aged less than 35), middle-aged (aged between 35 and 60) and old (aged over 60).

Given the potential for complicated interactions in the relationships being analysed using equation 3.2, ABHL try a number of different specifications for the function $g(\text{hp})$ in an attempt to capture the relationship between house prices and consumption in a flexible fashion. Table 3.1 reports the coefficients of interest for two of their specifications.³ The ‘shocks’ specification reported in table 3.1 includes as regressors predicted (log) house prices⁴ and the difference between predicted and realised (log)

³The full set of independent variables in the regressions is: a quintic in age of head of household; cohort dummies; occupation dummies; region and month dummies; controls for family composition; educational attainment dummies.

⁴Where house prices are predicted from real interest rates and regional incomes.

house prices. The latter term is intended to capture the effect of unexpected house price growth for these three age groups. The house price terms in the ‘growth’ specification reported in table 3.1 measure the effect on consumption of proportionate growth in house prices. The strongest associations between unexpected and proportionate house price growth and consumption are seen to show up for young individuals. Since this group contains individuals who are unlikely to hold large amounts of housing wealth, and this is also the group for which (due to the planning horizon) wealth would have the smallest immediate effect on consumption, these results have been interpreted as suggesting that wealth effects have not been the main driver of the correlation between house prices and consumption.

Table 3.1: ABHL house price terms by age groups: predicted price specification

House Price Terms	Estimated Coefficients	
	<i>Shocks spec</i>	<i>Growth spec</i>
<i>Predicted log house price</i>		
Young	0.291*	
Mid-age	0.292*	N/A
Old	0.294*	
<i>Shocks to log house price</i>		
Young	0.188*	
Mid-age	0.088*	N/A
Old	-0.012	
<i>Proportional growth in house price</i>		
Young		0.209*
Mid-age	N/A	0.127*
Old		0.042

Notes: These results are taken from Attanasio et al. (2009).

The empirical evidence indicates that, in the recent past, it has been young individuals (those aged 35 and under) who have had the biggest deviations (on average) of consumption from expected levels in times of boom and bust. The young group are also the group for which swings in house prices are most able to explain such deviations of consumption from expected levels. Section 3.4 will use graphical analysis of residuals,

plus counterfactual simulations (i.e. without some of the aggregate shocks), to look at whether our model produces similar patterns and to analyse how we can understand what could be driving such patterns in the data. Before that, it is necessary to outline the structure of the model.

3.3 The model

As discussed in section 3.1, our central contribution lies in our systematic use of a realistic structural lifecycle model of consumption, savings and housing choices, to disentangle and distinguish the three mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. As such, our work is building on several recent contributions that have examined the relationship between housing (or durables) and consumption in structural frameworks.⁵ Such analyses have found that the nature of housing as consumption good, asset, source of collateral, and potential intergenerational heirloom, makes the relationship complex. Fernandez-Villaverde and Krueger (2005), for example, argue persuasively that a housing asset may have an important role in explaining the observed “hump shape” of profiles of consumption and durable consumption over the lifecycle. Flavin and Nakagawa (2004) discuss the fact that when the amount of housing can only be adjusted at cost, this may explain why both housing and non-durable consumption tend to be smooth over time apart from at infrequent periods of large adjustment. Similarly, the analysis of Li and Yao (2007) shows how a housing asset can result in consumption behaviour (including for housing consumption) that is very insensitive to income in some ranges, but very sensitive in others.

Our contribution to this literature comes in a large part from the way we apply our model to match features of aggregate data, as we simulate the behaviour of a

⁵As mentioned in the previous section, Campbell and Cocco (2007) makes some use of a structural model, while Cocco (2005) is a contribution to another growing literature on portfolio choice in the presence of housing.

series of cohorts designed to resemble cohorts of the U.K. economy. The model is constructed with some of the stochastic elements faced by the agents - the process for house prices, and a component of the income generating process - that are thought of as aggregate in that they affect all members of the population at the same time. As well as using data to estimate the parameters of the processes generating shocks to these aggregate features, we also have estimates of the path of aggregate shocks that have been experienced in the U.K. economy since the beginning of the 1970s, and we use this path as an input into our simulations.

More precisely, we take the data on actual growth rates in house prices and aggregate earnings in the economy each year, and input these into our simulations at the correct age for the particular cohort that we are simulating. We repeat this process for each five year cohort born between the 1910s and the 1970s, and create simulated data for a large number of different households (i.e. realisations of the idiosyncratic shocks) in each cohort given the house prices and the aggregate income that they actually faced. These simulated data are the basis for our comparisons to data on the U.K. economy, including to household survey data which incorporate the same set of cohorts.

It is important to stress that our simulations take as an input the actual path of prices and incomes in the U.K. economy. Since the relationships between choices, prices and income in our model are not linear, simulating for a specific path of prices and aggregate income as we have done gives different results from simulating behaviour on average (across different levels of the house price and aggregate income) for particular shocks. Our simulations based on actual shocks to house prices and incomes also form the “baseline” against which we can compare counterfactual scenarios in which the mechanisms that might drive the link between house prices and consumption are shut down in turn; these scenarios and how they compare to our calibrated baseline are the subject of section 3.4, below.

In order to carry out the exercise described in the previous paragraph, our model

had to be carefully constructed with certain novel features, and is therefore a contribution in its own right to the understanding of lifecycle choices in the face of uncertainty. The model is of a household's lifecycle choices concerning consumption and saving, and whether or not to own housing. Since we want to capture how aggregate shocks feed in to aggregate consumption, we are careful to model aggregate income and house price uncertainty in a way that reflects the data. We are particularly careful to model correlation in the shocks in the two processes since house prices and incomes both tend to go up (down) when the economy is performing strongly (poorly), and this must be reflected in agents (rational) expectations. To our knowledge, building these aggregate shocks, and their empirical correlation, into a model of the kind we construct has not been done previously. In addition, building on the models presented in Bottazzi et al. (2007) and in chapter 2 and Attanasio et al. (2007), we are careful to model mortgage-related borrowing constraints as realistically as possible since these are likely to have first-order effects on behaviour. We now describe the model, and its calibration to UK data, in detail.

3.3.1 The household maximisation problem

The households in our model are (ex-ante) heterogeneous in one dimension: their level of education. In practice this will imply the calibration of different earning processes for households with different level of education. In addition, each household experiences different idiosyncratic shocks to earnings. And at a given point in time, households of different ages are present.

A household lives $T = 59$ periods (ages 22-80). In every period $t \leq T$, the household maximises utility by choosing consumption, with $c_t \in R_+$, and whether to own a flat, a house, or no housing (which may be thought of as costless rental), with $h_t \in \{0, 1, 2\}$.⁶

⁶Recently we have been exploring the effect of adding a rental price (linked to the house price to avoid the addition of an extra state variable) to the budget constraint. This changes the properties of the calibration since ownership becomes financially more attractive and so the housing preference and discount parameters must be adjusted to match proportions of house and flat owners. However the change in parameters has no effect on the results regarding price shocks and the consumption of

The household value function in period t is given by:⁷

$$V_t(A_t, h_{t-1}, P_t, Y_t, Z_t) = \max_{\{c_t, h_t\}} u(c_t, h_t) + \beta EV_{t+1}(A_{t+1}, h_t, P_{t+1}, Y_{t+1}, Z_{t+1}) \quad (3.3)$$

subject to

$$A_{t+1} = R_{t+1} \begin{cases} (A_t + W_t - c_t - \kappa P_t(1 + F)I(h_t = 1) - P_t(1 + F)I(h_t = 2)) \\ \quad \text{if } h_{t-1} = 0 \\ (A_t + W_t - c_t + \kappa P_t(1 - F)I(h_t \neq 1) - P_t(1 + F)I(h_t = 2)) \\ \quad \text{if } h_{t-1} = 1 \\ (A_t + W_t - c_t - \kappa P_t(1 + F)I(h_t = 1) + P_t(1 - F)I(h_t \neq 2)) \\ \quad \text{if } h_{t-1} = 2 \end{cases} \quad (3.4)$$

where A_t is the start of period asset stock and $R_{t+1} = 1 + r_{t+1}$ and r_{t+1} is the (real) interest rate on the liquid asset; P_t is the price of housing which is realised at the start of period t ; F is the cost of selling or buying a house, which is proportional to the price; W_t is household income in period t , while Y_t and Z_t are the idiosyncratic and aggregate components of the income.⁸ The number of state variables in this problem (with four continuous states plus the current home ownership and time) means it is computationally demanding to solve; for more details on the solution method see appendix section A3.1.

We only allow for collateralised debt, i.e. households are only able to have negative financial assets when they are home owners, so that when they do not own a house

different age groups, which will be discussed in section 3.4.

⁷Throughout this description of the household's problem, I will stick to the notation of having only one index, t , that captures both time and age (for an individual born at a given date). In the simulations that produce the results reported from section 3.3.3, a key feature is that observed aggregate shocks are input for the date at which they hit each cohort. Thus for the simulations we have to be careful to keep track of the relationship between age and time for each simulated cohort.

⁸The household is assumed to be fully aware of the separate stochastic processes that generate these components.

($h_t = 0$) they are subject to the constraint

$$A_t \geq 0. \quad (3.5)$$

Home owners can borrow, and when they do so they are subject both to a terminal asset condition that translates into an implicit borrowing constraint, and to two explicit borrowing constraints. In particular, we impose the terminal condition $A_{T+1} = 0$. The specification of marginal utility becoming infinite at 0 consumption means this terminal condition prevents households borrowing more than they can repay with certainty. In addition to this implicit borrowing constraint, we model two explicit constraints. The first is a function of the value of the house and the second is a function of the household annual income. They determine how much a household is able to borrow at the time of purchase or when remortgaging, and translate into the following constraints in the period after the new mortgage is agreed:

$$A_{t+1} \geq -\lambda_h \kappa P_t (1+r), \quad \kappa = \begin{cases} 0 < \kappa < 1 & \text{if } h_t = 1 \\ 1 & \text{if } h_t = 2 \end{cases} \quad (3.6)$$

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

$$A_{t+1} \geq -\lambda_w W_t (1+r) \quad (3.7)$$

The explicit constraints on the downpayment and the debt to income ratio only apply when households buy the property or remortgage. Formulating the constraints in this way complicates the model solution (for details see section 2.2.1). In particular, it makes the model more complicated to solve numerically than, for example, having only an downpayment constraint that must be satisfied every period (as, for example, in Campbell and Cocco (2007), Li and Yao (2007) Cocco (2005), Campbell and Hercowitz (2004)). However it seems important to us to capture the institutional features of the

U.K. mortgage market, since these are likely to affect how house price shocks feed into consumption. The correlate of only applying the constraints in periods of buying or remortgaging is that when a household continues owning without remortgaging, they can keep their existing debt if they have negative financial assets:

$$A_{t+1} \geq A_t \quad (3.8)$$

Although there is no mortgage repayment schedule, the household does have to pay off mortgage interest each year in which it does not remortgage.

Utility and bequest functions

Households get utility from consumption, from home-ownership, and from leaving bequests.

The within period utility function is CRRA in non-durable consumption. This is augmented by a term reflecting the value of home-ownership:

$$u(c_t, h_t) = \exp(\theta\phi h_t) \frac{c_t^{1-\gamma}}{1-\gamma} \begin{cases} \theta, \phi \in \mathbb{R} \setminus \{0\} & \text{if } h_t = 0 \\ \theta \in \mathbb{R}, 0 < \phi < 1 & \text{if } h_t = 1 \\ \theta \in \mathbb{R}, \phi = 1 & \text{if } h_t = 2 \end{cases} \quad (3.9)$$

The parameter θ is a housing preference parameter which determines the utility that households obtain from owning a house rather renting it. ϕ determines the relative utility from owning a flat versus a house. These parameters are calibrated in our model; further discussion of the preferences just discussed can be found in chapter 2.⁹

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

$$b(\omega) = \tau * \frac{(\omega/\tau)^{1-\gamma}}{1-\gamma} \quad (3.10)$$

⁹Similar forms for preferences regarding housing are used by Ejarque and Leth-Petersen (2008) and Rios-Rull and Sanchez-Marcos (2008).

Where ω is the value of wealth (both financial and housing wealth, net of the fixed cost of selling a property) left over at the end of life T , after all shocks to resources (income and the house price) have been realised, and all consumption decisions have been made.¹⁰ The parameter τ is calibrated. A bequest motive, although not central to our analysis, is crucial to match certain features of the life cycle profile of home ownership, particularly at later ages.

3.3.2 The environment: exogenous stochastic processes

Households face three dimensions of uncertainty: shocks to house prices, which are aggregate (i.e. common across all properties in the economy); aggregate shocks to income; and idiosyncratic shocks to income. In the present version of the model, the interest rate on liquid assets and debt is fixed.¹¹

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

$$\ln W_t = d_t + y_t + z_t \quad (3.11)$$

where lower case has been used for logs, and d_t is a deterministic part to the income generating process, y_t is a persistent idiosyncratic stochastic element; and, z_t is the aggregate stochastic component.

The deterministic component d_t is hump-shaped over the working lifetime and is captured using (the log of) a polynomial. The coefficients of these polynomial are calibrated for the two education groups we consider. In practice, to fit the observed data we need a cubic specification (see Appendix A3).

¹⁰For more motivation of the modelling of bequests, see, for example, De Nardi (2004).

¹¹It is a relatively simple extension to add i.i.d shocks to the interest rate. Early experiments indicated that such a change has no qualitative impact on our results.

The idiosyncratic stochastic component y_t is modelled as an AR(1) process:

$$y_t = \rho_y y_{t-1} + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2) \quad (3.12)$$

The parameters of d_t and y_t are set based on a single estimation that is described in Appendix A3.

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

If we let HP stand for the house price and again use lower case letters to represent logs, this can be written as:¹²

$$\begin{bmatrix} z_t \\ hp_t \end{bmatrix} = \begin{bmatrix} \alpha_0^z \\ \alpha_0^h \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} z_{t-1} \\ hp_{t-1} \end{bmatrix} + \begin{bmatrix} u_t^z \\ u_t^h \end{bmatrix} \quad (3.13)$$

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

¹²Given the unit persistence in this equation, the constant terms α_0^z and α_0^h capture drift over time even though there is no time trend in equation 3.13, see Davidson and Mackinnon (2004), pp.606f. Since the process is in logs, the constant terms measure the trend growth rate.

distribution of the shocks is assumed to be normal:

$$\underline{u}_t \sim \mathbb{N}(0, \underline{\Omega}) \quad (3.14)$$

where

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

with π measuring the correlation between the shocks.

3.3.3 Calibration and estimation

The parameters of our model can be divided into three categories: those we take from other studies, those we estimate outside the model and those we calibrate to fit some moments of the micro data we consider. We estimate the parameters of the exogenous stochastic processes faced by the agents in our model using time series data. Our calibration exercise, instead, involves matching moments for life cycle levels of home ownership status in the U.K. for those aged 26-60 in 1990-2006. Since we simulate cohorts born between the 1910s and 1970s, we match the ownership levels for the cohorts in the relevant age range in the 1990s and early 2000s.

Parameters fixed or estimated outside of the model

For inputs into the calibrated model, we need to use data on earnings, the house price process and the interest rate on liquid assets. Values for parameters fixed or estimated outside the model are summarised in table 3.2.

Utility function

The preference parameter γ in the utility function is set to match the consumption elasticity of intertemporal substitution of 0.7 found in data (see Attanasio and Weber (1995)). This corresponds to a curvature $\gamma = 1.43$ for our within period utility function.

Table 3.2: Estimated / Fixed Parameters

Parameter	Value	Source
<i>Estimated/Fixed Parameters</i>		
Utility Parameters		
γ	1.43	Attanasio and Weber (1995)
β	1.0358^{-1}	
Aggregate House Price and Income Process		
α_0^z	1.66%	FES
σ_z	0.033	FES
α_0^h	3.58%	DCLG
σ_h	0.091	DCLG
τ	0.645	FES / DCLG
κ	0.6	BHPS
Idiosyncratic Income Process		
Deterministic component: cubic in age		BHPS
	High Edu	Low edu
ρ_y	0.76	0.77
σ_ξ	0.39	0.41
$\frac{\text{Median}P_{22}}{\text{Median}Y_{22}}$	3.3	4.4
Credit market Institutions		
λ_y	3.0	
λ_h	0.9	
\bar{r}	0.03	B.o.E.

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

The parameter ϕ indicating the relative utility value of a flat to a house is set at the same level as the price ratio of flats to houses, κ , so that $\phi=0.6$. We also choose a baseline discount rate of 3.58%, which matches the expected rate of return on housing.

Aggregate Processes: The house price and aggregate income shocks

As described by equations (3.13) and (3.14), the aggregate house price and income processes are specified in logs as a vector autoregression of order one (a VAR(1)), with drift. To estimate this process we use: data on the national mix-adjusted house price

series for the U.K.¹³, which are the same data underlying figure 3.1 and are freely available from the Department for Communities and Local Government (DCLG); and, data on the UK Family Expenditure Survey / Expenditure and Food Survey (FES). Since our model is set up in terms of real prices, we deflate both series to prices for the latest year using the all item Retail Prices Index, and we use data for the years 1969-2006. As anticipated in section 3.3.2, estimation of this process does not reject unit persistence, and so we impose that aggregate shocks are permanent as in equation (3.13). The other parameters returned from this estimation are trend growth rates for the house prices (3.58%) and income (1.66%), and standard deviations of the shocks to the processes of 0.091 and 0.033 respectively for the house price and income. The correlation of the shocks for the two processes is approximately 0.65, and our model assumes that people know this correlation and build it in to their expectations.

It is important to realise that in our simulation exercises, in addition to inputting the parameters we have been discussing to solve the model, we use as inputs the estimated income and house-price shocks that are derived from the estimation of the aggregate time series processes in equation (3.13). These processes are treated as aggregate so that the same shocks hit all individuals in a given period, and thus affect different cohorts at different ages. We feel that it is important that our simulations reflect actual realizations of aggregate shocks to drivers of consumption, since the complicated nature of the relationship between house prices and consumption choices is likely to mean that how this shows up in data will depend on the actual path that prices follow and not just the processes generating prices. The growth rates in house prices and aggregate income that we input into our model as aggregate shocks are shown in figure 3.9 in the appendix.

¹³We use the series reporting average house prices for all dwellings.

Idiosyncratic income process

On top of the aggregate income process, there are two further elements to the household income generating process (see equation (3.11)). These are a deterministic process, which captures the hump-shaped profile of household incomes over the working life and a 60% replacement rate in retirement, plus a persistent stochastic component which is described by a first order autoregression (equation (3.12)). The deterministic component plus the process generating shocks to the stochastic component are both education group specific, and the realised shocks to the stochastic component are idiosyncratic at the household level.

Parameters for both processes are estimated together as described in Appendix 3.5. Since this involves estimation of a dynamic process at the household level, we require panel data and so use the British Household Panel Survey (BHPS), which is available for the years 1991-2005. The estimation yields that the deterministic component can be approximated as a cubic that shows a hump-shape over the working life which is slightly more pronounced (in particular with a steeper slope at the beginning of the working life) for the college educated group than for those with only compulsory level education. The process generating stochastic innovations is quite similar across groups with a persistence parameter of slightly more than 0.75 and a variance of the shock of around 0.16 (exact values are in Table 3.2 or Appendix A3).

Credit market institutions

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

calibration).

Calibrated parameters

We select the preference parameter for housing, θ ,¹⁴ plus the parameter τ that determines the bequest motive, and the fixed cost of housing, F , by matching average life-cycle home-ownership rates between ages 26 and 35, and 36 and 60, for our two education groups. We assume that the parameters are common across the two education groups. As reported in table 3.3 we obtain values such that owning a house raises utility by approximately 1.5% and owning a flat raises utility by approximately 1%, the fixed cost of buying and selling is 3% of the property price, and the parameter τ takes a value of 4. Table 3.4 compares home-ownership rates predicted by the model to those observed in the data¹⁵; we do a reasonably good job of matching the moments of interest, particularly for the high-education group.

Table 3.3: Calibrated Parameters

Parameter	Value
θ	-0.015
ϕ	$\frac{2}{3}$
τ	4
$F_b = F_s$	0.03

3.3.4 Model fit

Since we do not match any moments of consumption, the properties of the evolution of consumption provide a useful check on whether our calibrated model of overlapping cohorts is doing a good job of matching data on the U.K. economy. To the best of our knowledge, at least in the consumption literature, it is an innovative contribution to

¹⁴Since the estimate of the intertemporal elasticity is taken from a paper which does not condition on home-ownership, there is a possible bias.

¹⁵Data come from the years 1990-2006, as years prior to 1990 are affected by the large-scale selling off of local authority housing.

Table 3.4: Calibration Statistics

Statistic	High Education		Low Education	
	Data	Model	Data	Model
<i>Ownership rate</i>				
Age 26 - 35	0.558	0.584	0.474	0.423
Age 36 - 60	0.794	0.822	0.632	0.681

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

aggregate up simulations of individual behaviour based on observed aggregate shocks, in order to check whether the model of individual choices is consistent with aggregate evidence. In the present context, it is particularly sensible to check the properties of simulated consumption in this way, since the question we wish to address concerns the evolution of aggregate consumption.

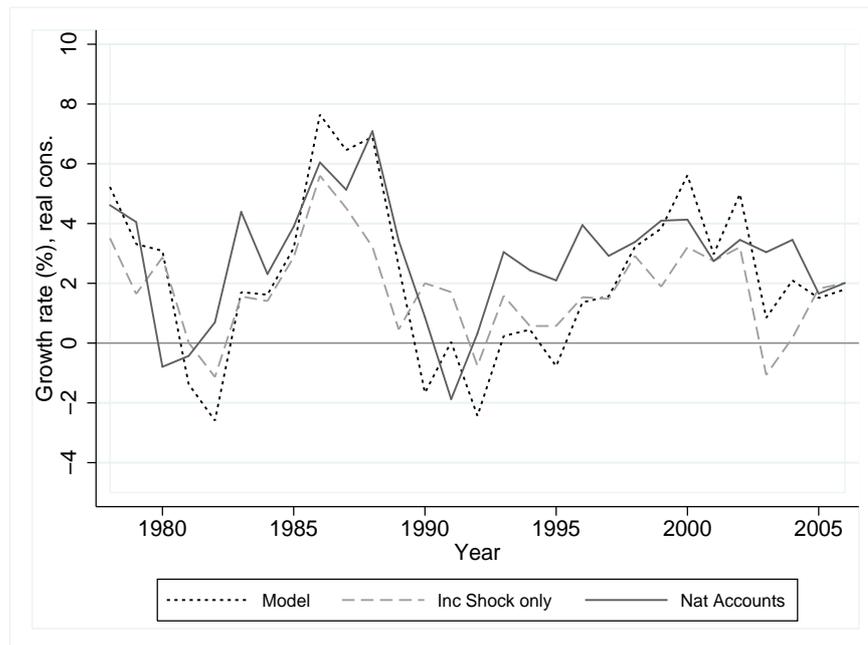


Figure 3.3: Growth in real consumption: Data versus model

In fact the model does a good job of matching the aggregate consumption growth rate. This is shown in figure 3.3 which compares the actual growth rate of aggregate consumption in the U.K. economy to that in our simulated data for the years 1975-

2006. The correlation between these two series is 0.75. To provide a benchmark of how good a match this is, the figure also includes simulated data for the growth rate of consumption based on a counterfactual scenario with no house price shocks (i.e. the innovation to the house price is set to zero every period). In this case the correlation to the data series is 0.58.

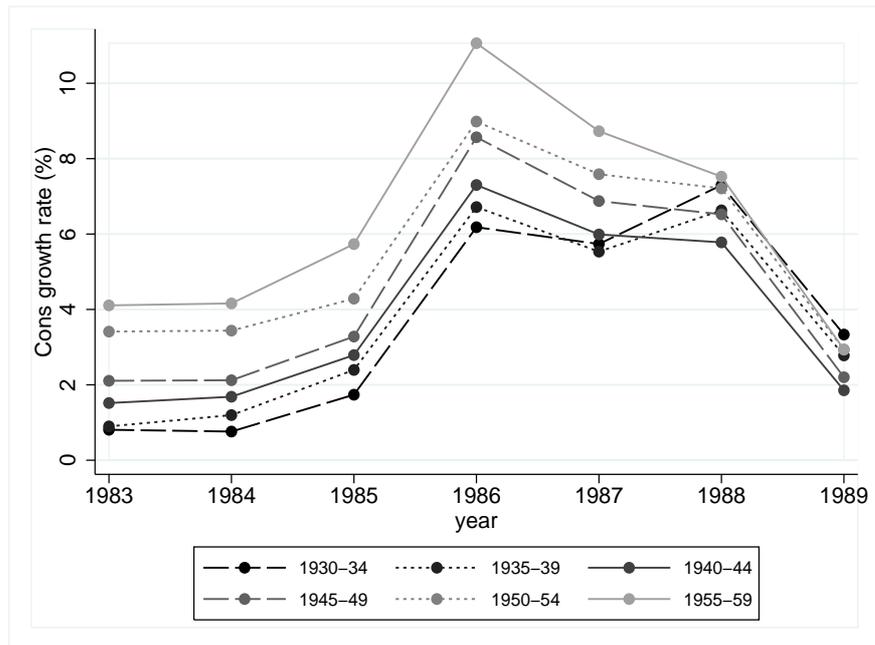


Figure 3.4: Model consumption growth by cohort: 1980s boom

This good match for the consumption growth rate is an indicator that our model might be useful for drawing quantitative, as well as qualitative conclusions. However, unfortunately the model does less well capturing patterns at a more disaggregated level, particularly by age group. For example, in the late 1980s consumption boom we know (from Attanasio and Weber (1994)) that the strongest consumption growth was from young groups. Figure 3.4 shows modelled consumption growth during this boom for different cohorts. During the middle years of the decade the model does indeed predict the fastest consumption growth for the young, but in 1988, the year of strongest consumption and house price growth, all cohorts are seen to have strong

consumption growth and the eldest cohort shown (aged in their late 50s - the black, long-dashed line) is predicted to have consumption growing just as rapidly as those cohorts in their 30s. This counterfactual pattern leads us to be cautious in drawing quantitative conclusions from our model about the extent to which house prices have caused shifts in aggregate consumption. Nonetheless the model is still useful for exploring the qualitative properties of consumption growth driven by increased housing wealth or by increased permanent income.

3.4 Booms and busts in house prices and consumption

The main motivation of our analysis using the calibrated life-cycle model is to explore how we can distinguish between the mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. The first part of this section looks at decomposing what is driving consumption growth in the model. The latter part of the section considers how patterns of consumption growth vary with age when house price or income shocks are more important. In each case the issues are explored through using our model to simulate counterfactual scenarios in which the possible mechanisms driving consumption growth are shut off in turn.

3.4.1 What is driving consumption growth in the model?

To disentangle what is driving consumption growth in our model, and why house-price and consumption growth are strongly correlated in the simulations, we compare a baseline run of our model to runs in which aggregate shocks to house prices and to incomes are switched off in turn. This method of decomposing the drivers of fluctuations is very much related to that outlined by Chiari et al. (2007) in a business cycles context.¹⁶

¹⁶Those authors emphasize that in computing the decomposition, the analyst wants to switch off shocks to variables in turn, but not to eliminate any effect of the level of other aggregate processes on expectations of shocks to the process whose influence is to be measured: one wants to isolate the effect of shocks each period, given rational expectations based on the observed state of the economy. In many cases this involves switching off the shocks, but having expectations based on the (multivariate) aggregate state of the economy that was actually observed. In our case with unit persistence (and zero

The aggregate shocks to income and house prices are the only drivers of changes in aggregate consumption growth in our model. We start by comparing aggregate consumption growth across simulations. This approach will give us a good idea of which factor is more important in driving swings in aggregate consumption growth in the model economy. We then examine the simulated data in a more disaggregate way to explore what micro-data on consumption can tell us about the relative importance of house prices in driving aggregate consumption growth.

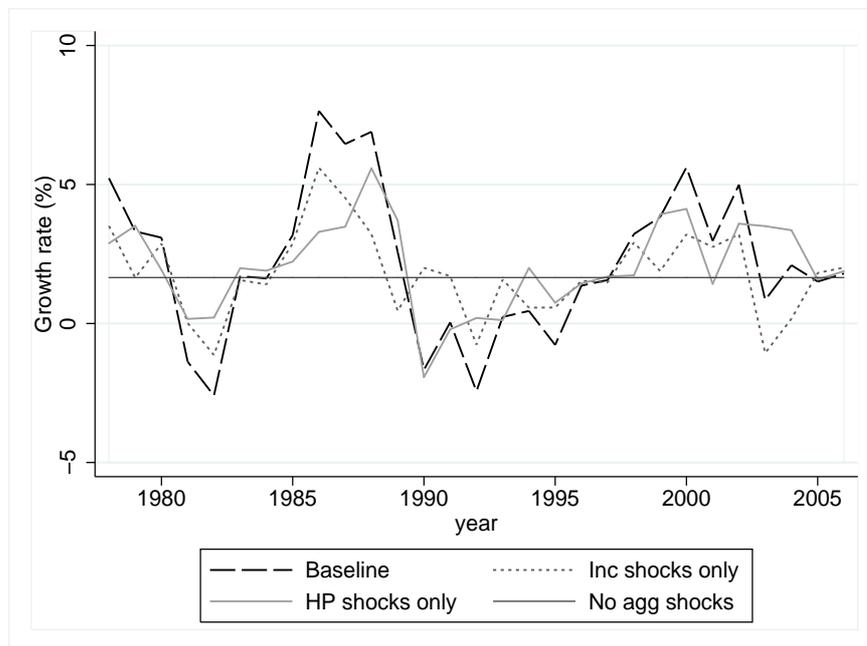


Figure 3.5: Model consumption growth: different scenarios

We saw in the previous section that our model appeared to generate a series for aggregate consumption growth that closely tracked that seen in aggregate data. Figure 3.5 shows the growth rate of aggregate consumption using different runs of the model in which the aggregate shocks are switched off. The “baseline” results are the growth rates we saw previously in figure 3.3, and the results with only the income shock were

cross terms) in the process generating aggregate states, the level of the house price or of aggregate income does not affect expectations about how the other aggregate state variable will move, and so our simulations do not need to keep track of an aggregate state measuring “where the economy would have been”. Thanks to Roman Sustek (Bank of England) for pointing out the parallel with this paper.

also reported in that figure. Figure 3.5 illustrates that when both house price and income shocks are shut down, consumption grows at the constant rate of approximately 1.7%. The other lines show consumption growth rates in the presence of only income shocks and only house price shocks. In the late 1970s and early 1980s, it is clear that income shocks are driving much of the movement in consumption growth: the growth rates when only house price shocks hit are typically closer to the constant level whilst growth rates when only income shocks hit are closer to those in the baseline (i.e. full shocks) economy. However, by the latter 1980s, the effect of changes in house prices on consumption growth becomes more substantial. By the house price boom of the late 1980s, more of the total consumption growth is driven by house prices than by incomes; similarly the house price bust of the early 1990s drives almost all of the negative consumption growth rates observed. Since then, house price shocks appear to have been the predominant factor driving aggregate consumption growth in the model. The relatively modest influence of income shocks in the later years modelled is due to the fact that aggregate income shocks, set to match the data shown in figure 3.9 in the appendix, have become rather weak since the late 1990s, compared to the volatile years of the 1970s and early 1980s.

One way to quantify “how much” of consumption growth is being driven by house price or income shocks, is to look at the proportion of the deviation of consumption growth in our baseline run from that in our run with no shocks, that is accounted for by the deviation in the modelled economy with only a single aggregate shock. Doing this over the whole period shown in figure 3.5, we find that house price shocks explain away about 55% of the deviation, while income shocks explain away about 35%.¹⁷ In line with the descriptive analysis of the previous paragraph, this difference is stronger after 1990 (for which period the proportions are 73% and 17% respectively), but is reversed (40% and 50%) for the 1978 to 1990.

¹⁷Though the income and house price shocks are the only aggregate shocks in the model, these two factors might also have a joint effect on consumption, and so there is no reason why these two numbers should sum to 100%.

3.4.2 How do the shocks affect different age groups?

The above analysis looked at aggregate trends in consumption growth under different versions of the model where house price and income shocks were present or absent. A more disaggregate analysis looking at these results for different year of age-groups will allow us to examine how the response to shocks varies for households at different points in the life cycle.

To carry out this exercise in a way that is directly comparable to the data, we use the kind of regression analysis that was described in section 3.2 and that generated figure 3.2, to generate similar pictures based on model data. The pictures show residuals from a consumption function, averaged by age at each point in time. The estimated consumption function controls for variation in demographic factors and skill levels.¹⁸

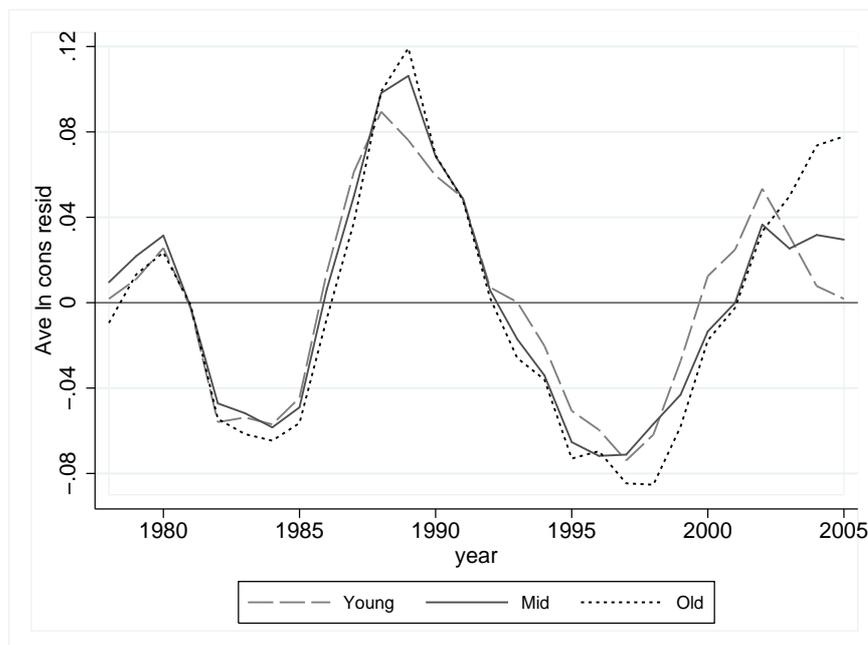


Figure 3.6: Residuals for log consumption: Model Results, Baseline

Figure 3.6 shows such averaged residuals for simulated data from the baseline run

¹⁸In the model we do not have all the demographic variables that can be measured in data, and the regressors are a subset of those mentioned in footnote 3. Specifically, we include a polynomial in age, plus measures of education and cohort.

of the model that featured both aggregate income and house price shocks. In contrast to the data picture (figure 3.2) that showed the young group experiencing the biggest swings in deviations from expected consumption, this figure shows three lines that generally move together. Indeed, if anything the biggest residuals at the top of peaks and the bottom of troughs, are for the oldest group. This is another example of the model predicting “too much” variation in the consumption of older groups of the population, a phenomenon discussed at the end of section 3.3.4. The results in figure 3.6 provide a useful baseline against which to compare the results for runs in which the only house price or income shocks are experienced.

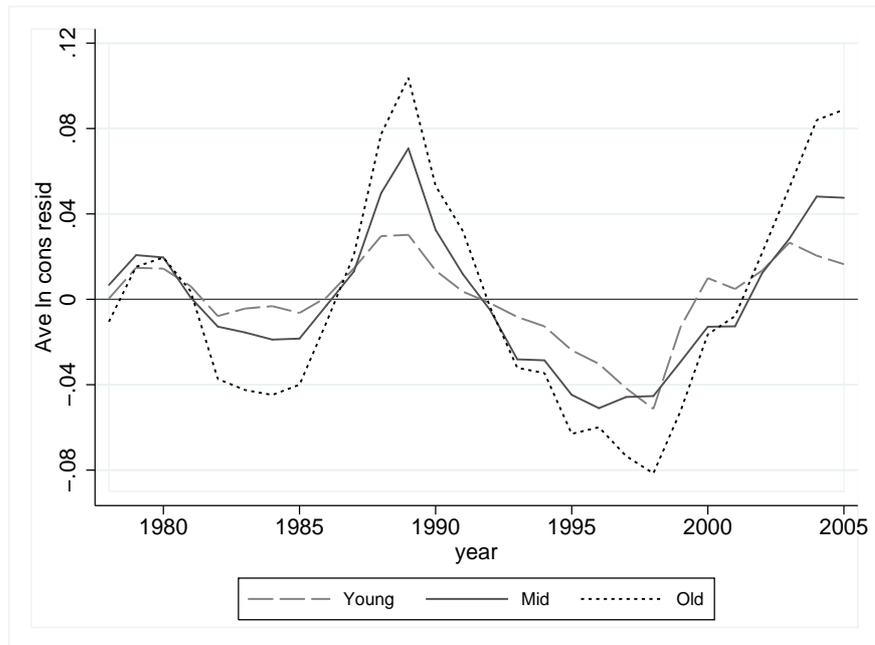


Figure 3.7: Residuals for log consumption: Model Results, House Price Shocks Only

Figure 3.7 shows the averaged residuals for simulated data based on an economy in which house-price movements are the only aggregate shocks. The results clearly show that, in the model economy, it is the consumption of the older group that responds most strongly to house price shocks. Individuals in this group tend to have high equity in housing, and have a short horizon over which to spread any new wealth. In times

of large shocks to house prices, the average consumption residual in the older group can be as much as two to three times larger than the same statistic for the young.

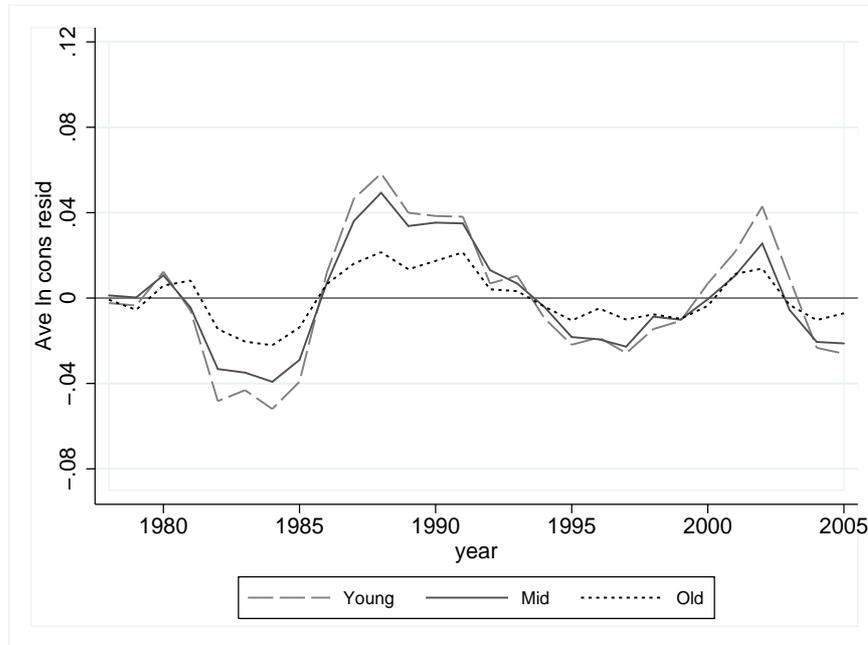


Figure 3.8: Residuals for log consumption: Model Results, Income Shocks Only

Figure 3.8 shows the averaged residuals based on a simulated economy in which income changes are the only aggregate shocks. In this case average residuals tend to be smaller than when house-price shocks operated, and the pattern is also reversed as the younger group have the largest deviations of consumption from expected levels. The contrast of this pattern to the case with house-price shocks, is important for the analyst who wants to understand why aggregate consumption and house prices tend to be positively correlated.

In sum, these age-group analyses have shown that different shocks translate into consumption responses that differ across groups. Younger individuals tend to respond more to aggregate-income shocks whilst older groups respond more to house price shocks. In a stylised lifecycle model, it is possible to see an intuition for why this occurs. A permanent income shock effectively acts as a shock to lifetime wealth, and

younger cohorts have a longer horizon over which to enjoy a positive shock (or suffer a negative shock) and so will adjust their consumption by more. Older cohorts respond more to house price shocks as they are more likely to have positive equity in their homes and so benefit from a wealth channel that allows them to adjust consumption. They would also respond more to this one off wealth shock simply because they expect fewer periods of life over which they can spread extra consumption.

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

What we can draw from our analyses is the importance of using micro-data to try to infer whether income shocks or house price shocks are responsible for consumption growth that we observe. If consumption shocks appear particularly strong for the young, that would suggest changes to aggregate incomes are most important in explaining consumption movements whereas if the shocks are strongest for the old, the housing wealth channel is likely to be most important. The earlier reduced form studies by Attanasio and Weber (1994) and Attanasio et al. (2009) and considered in section 3.2, indicated that excess consumption growth has been strongest for the young, which was taken to imply that a wealth effect was not the main driver of the observed correlation between house-price growth and consumption growth. Our results put this interpretation on a stronger theoretical footing.

3.5 Conclusions and further work

Our analysis has aimed to provide a deeper understanding of how we can distinguish between the mechanisms that have been proposed to explain the correlation of house price shocks and consumption growth. In particular, through our systematic use of a realistic lifecycle model, we have investigated whether structural modelling confirms the intuitions about how and why the consumption of different groups might correlate with house prices. These intuitions have formed the basis of past, reduced form, empirical tests. Our analysis has involved some methodological innovations in the way we have constructed our model and particularly in the way we have applied it to simulate the behaviour of a series of cohorts, designed to resemble cohorts of the U.K. economy.

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

The exercise we have constructed has returned interesting results, but can also be seen as a first step of a research agenda that aims at explaining the behaviour of house prices as an equilibrium process. We have, in effect, modelled the demand side of the housing market. In order to do this we have made a set of assumptions about the process generating house prices which have facilitated an empirical specification that could be approximated in our simulations. In doing this we have learnt that the precise shape and parameterisation of this process is an extremely important driver of the evolution of the joint decisions of housing demand and consumption expenditure, both at the household and the economy-wide level.

It is possible that building a set up with endogenous prices would help us to understand why consumption, and particularly the consumption of the old, seems to be less responsive to house-price shocks than is predicted by the current model. But more than that, it is a major research challenge to assess whether the volatile house price process we observe, and other features of aggregate consumption and housing market behaviour, can be obtained from a unified framework in which house prices are determined endogenously.

A3 Appendix to chapter 3

A3.1 Solution method

The setup of our model, with a discrete choice concerning home-ownership, coupled with fixed costs and the particular form of the borrowing constraints, mean that the functions of the household's optimisation problem will not be 'well behaved' and we cannot rely on the existence of smooth first-order conditions that could otherwise have been exploited to improve efficiency in solving the problem.¹⁹ Instead we rely on robust techniques developed in previous ESRC-funded research (ESRC grant RES-000-23-0283, led by Attanasio) which involve solving using iteration on the value function (rather than the first order condition), and finding different "conditional value functions" (one for each of the current choices of house ownership, flat ownership, and non ownership) which can be compared in order to determine the discrete choice.

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

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

We also use discrete approximations to the specified continuous processes for id-

¹⁹The 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.

idiosyncratic income, and for the house price and aggregate income. This involves modelling these processes using finite state Markov chains that mimic the underlying continuous-valued univariate or bivariate processes. This is done as described in Tauchen (1986). We preferred Tauchen's method of equally spaced nodes over the quadrature based method proposed by Tauchen and Hussey (1991), because this has been shown to be more robust to very high persistence in the modelled processes (Floden, 2007).

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

As explained in the paper, the profiles of behaviour that we simulate are obtained by simulating a series of cohorts designed to resemble cohorts of the U.K. population. Every member of a particular cohort experiences the same shocks to aggregate income and the house price at each age. Members of subsequent or earlier cohorts experience the same series of aggregate shocks, but at the relevant ages so that every cohort experiences the same shocks at the same point in time. The aggregate shocks that are input in our baseline run are designed to match actual shocks to house prices and aggregate incomes in the U.K., and these correspond to the growth rates shown in Figure 3.9.

The profiles of behaviour reported in the paper are obtained by simulating 13 cohorts (born between the 1910s and the 1970s) each containing 2000 individuals split

evenly between two education groups. Within education group, households differ according to their initial financial wealth and the idiosyncratic income shock that they face at each period.

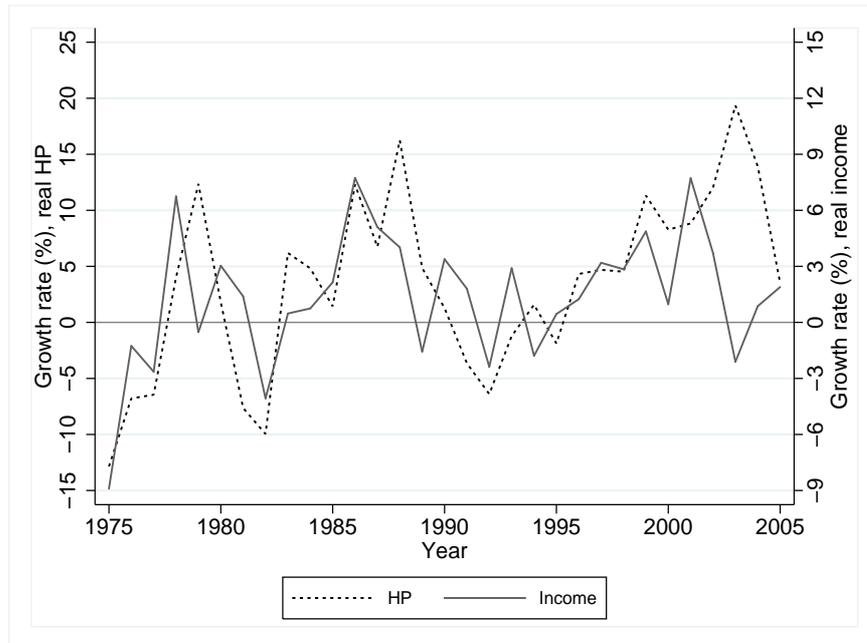


Figure 3.9: Growth in real house prices and "aggregate income"

A3.2 Estimating the idiosyncratic part of the income process

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

The estimation proceeded in two steps and was carried out separately for each education group. The first step was to regress income on a polynomial in age, and it turned out that in order to capture noticeable differences across education groups a cubic specification was required, but higher order terms did not enter significantly. Having fitted this regression, the persistent element of the process was then estimated as an AR(1) on the residuals from this regression. Results from this exercise are summarised in table 3.5.

Table 3.5: Estimated parameters of the idiosyncratic income process

	Low Education	High Education
age	-0.01068	0.01724
age squared	0.00184	0.00037
age cubed	-0.00005	-0.00002
ρ_y	0.76771	0.76365
σ_ξ^2	0.16580	0.15471

BHPS data were also used to fix the initial level of the house price relative to expected income (this was matched to the ratio of average income to the average house value for those aged 22-26). The 2000 BHPS dataset was our source for the distribution of financial assets for the two education groups at the beginning of adult life. This distribution was matched to that of individuals aged 22-26 in the relevant education group and we assumed that households have zero housing endowments at age 22.

Chapter 4

Modelling House Prices with Shocks to Income and Credit

4.1 Introduction

In the previous two chapters we have considered lifecycle models that incorporate a housing choice alongside the consumption/savings decision, with the process generating house prices taken to be exogenous to the model. Since the decisions of individuals about buying and selling property are key drivers of the demand to buy properties, and of their supply for sale, it seems desirable to endogenise property prices.¹ This chapter presents an overlapping generations model of heterogeneous agents in which households' decisions in the housing market determine the prices of houses and flats, and the volume of transactions each year.

The model does not feature any aggregate uncertainty, so an assumption of rational expectations amounts to perfect foresight regarding the housing market. The steady states that we describe involve no growth in the economy or the population, so that

¹Other factors that might affect demand and supply would be trends in the number of households in the economy, and the number of properties being built. Factors that alter house purchase and sale decisions of the existing population are perhaps the most likely to cause short term fluctuations in demand and supply, and therefore in prices.

house prices are constant. However, there are some dynamics in our analysis. We consider unanticipated shocks to incomes, and to the parameters of the mortgage market, and analyse transition paths to a new steady-state equilibrium after these (permanent) shocks. This analysis allows us to consider whether there are plausible mechanisms through which shocks to incomes, and to the mortgage market, might have contrasting impacts on prices and transactions in the housing market.

The interest in the exercise is not only to see whether these different shocks would have different effects, but also to see whether either could explain certain patterns in the dynamics observed in housing markets. In particular, house prices are observed to be substantially more volatile than average incomes (or GDP) in the economy: in the previous chapter we estimated the standard deviation of permanent shocks to the aggregate component of (log) incomes to be a third of the size of similar shocks to house prices. Additionally, housing-market transactions are observed to be positively correlated with house-price changes. Figure 4.1 shows quarterly data on housing market transactions and house price inflation since 1977. A regression based on these data of quarterly transactions on the house price inflation rate, plus a time trend and a constant, suggests that a one percentage point increase in the inflation rate is associated with an increase in the transaction volumes of about 3400 per quarter, or almost one-percent.²

A recent literature has emerged attempting to assess the hypothesis that, if individuals are forward looking but operate in an environment of partial insurance, then the dynamics of the housing market could be explained by the interaction between the institutions (fixed costs and mortgage-related borrowing constraints) of this market, with aggregate economic shocks. The idea was put forward by Ortalo-Magné and Rady (2006), who used an analytic model to show certain circumstances in which an income shock might lead to relatively large price swings, and transaction levels that covary positively with the price level. The exercise here might be thought of as build-

²And this is strongly statistically significant.

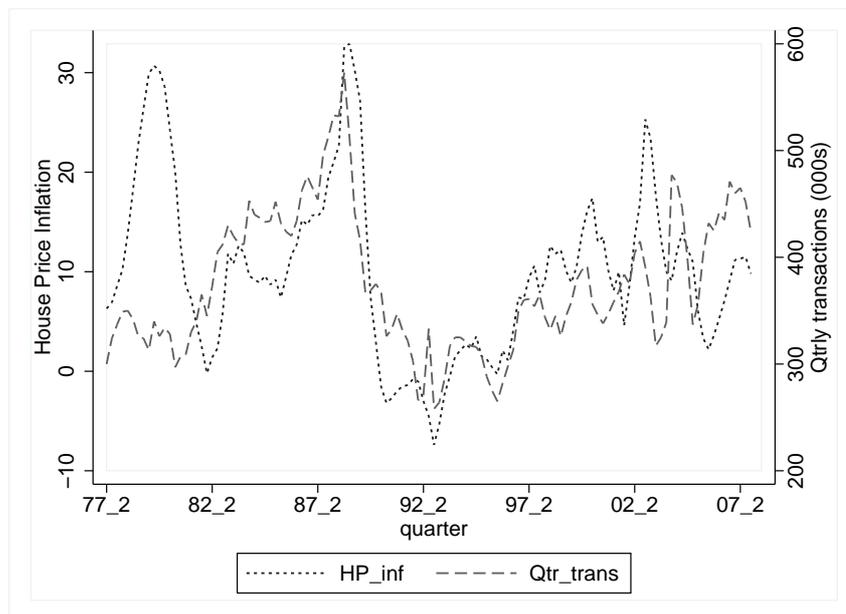


Figure 4.1: Housing transactions and house price inflation in the U.K

ing the thought experiment of Ortalo-Magné and Rady (2006) into a model that is more realistic in terms of the length of life of agents, the idiosyncratic uncertainty over labour incomes, and the nature of preferences and mortgage constraints. The model does though stop short of the model with aggregate uncertainty which is the aim of Rios-Rull and Sanchez-Marcos (In Progress).³ Nonetheless, an understanding of when and how large swings in property prices, and transactions levels that covary positively with prices, can occur in the present model, could still be a valuable contribution to understanding what might drive these patterns in economies including that of the U.K.

In the present version I will not try to be too precise in any quantitative assessment of whether plausibly sized shocks can lead to changes in housing prices and transaction volumes of a scale that is observed. Instead I will address whether one type of

³The theoretical and computational contributions mentioned in this paragraph present the relevant hypothesis and attempt to assess potential magnitudes of effects. Some empirical contributions have suggested that it might in fact be hard to explain housing market fluctuations from housing market institutions and individual behaviour alone (see Davis and Martin (2008), and articles discussed there), while there also exist theoretical contributions discussing alternative possible mechanisms (such as Kahn (2008) concerning the housing market and non-housing productivity shocks).

shock seems a more plausible candidate than the other for generating housing market dynamics. I will look at whether a given shock can lead to noticeable “overshooting” of house and flat prices beyond steady-state levels, and at whether such overshooting is associated with unusual transaction volumes. I will also try to disentangle what drives particular mechanisms in the model.

To anticipate, the results of the comparison between the effects of an income shock and of a mortgage-market shock, suggest that the shock to mortgage constraints is a more promising potential explanation for certain patterns in housing market fluctuations. It is only the mortgage shock that produces price swings substantially beyond the post-shock steady state during transition, and a transition path in which the house price and transaction levels are positively correlated.

The main force behind these patterns is relatively high saving (and low borrowing) in the economy with tight credit conditions, which results in a high amount of financial wealth in the economy immediately after a positive credit market shock. This in turn, and through the interaction with the newly generous mortgage constraints, leads to “excess demand” for both flats and houses, and so to price overshoots. The overshoots choke off demand but also encourage the supply of properties on to the market, hence the increase in transactions. This mechanism is different from the previously suggested (again see Ortalo-Magné and Rady (2006)) transmission of capital gains for the owners of starter homes into increased demand higher up the housing ladder, though this capital gains route will also operate in the economy modelled here. Thus the analysis suggests positive (negative) shocks that leave individuals feeling (ex-post) that they have built up “too much” (too little) financial wealth, as situations that might produce large housing market swings and transactions that move with prices.

The literature mentioned so far in this introductory section includes those papers that most closely relate to the exercise presented below. The work is also more broadly related to a larger set of papers concerned with building a housing or durable good/asset into a model of the macroeconomy. In this literature, Nakajima (2005) is

perhaps the most nearly related to the present exercise, as he focuses on house prices and particularly the relationship between earnings risk and house-price trends.⁴

The chapter is structured as follows. Section 4.2 describes the model in detail, including the method for solving for transition paths, while section 4.3 discusses the model results for income and mortgage market shocks. The final section is an extended conclusion. Since the current analysis is a preliminary exploration of the use of the model, this section discusses not only the findings of the chapter, but also considers at some length the options for further development of this research.

4.2 The Model

In this chapter we consider a model economy that is populated by N households. Each household lives for J periods and (so) there are n households from each of J overlapping generations (or cohorts).

Households maximize utility by choosing consumption and financial saving, and whether to own a flat, a house, or no housing, each period.

Households face idiosyncratic uncertainty over earnings. Shocks to earnings are mean zero and uncorrelated across families so that (in a large economy) there is no aggregate uncertainty over earnings.

Flat and house prices are determined endogenously. The level is such that the number of families wishing to live in each type of owner-occupied home is equal to the fixed stock of flats and houses in the economy. An equivalent way to say the same thing is that the number of families supplying flats and houses to the market for sale, is equal to the number generating demand to purchase each type of home in the current

⁴Other papers have considered how the housing/durable good influences the wealth distribution (Diaz and Luengo-Prado (Forthcoming); Gruber and Martin (2005)), the issue of trends and volatility in residential investment (Davis and Heathcote (2007); Fisher and Gervais (2007)), and whether downpayment constraints can explain a recent shift in U.S. home-ownership rates (Chambers et al. (2004)).

period.⁵ The assumption of a fixed housing stock is an extreme one, but does not seem unreasonable in the light of evidence that the volume of new houses being built has been relatively low and unresponsive to price fluctuations in recent years in the UK (see Barker (2006)).

Apart from the level of income, other important factors shaping households demand and supply decisions for flats and houses, are income-related and downpayment constraints on new mortgage borrowing, and (proportional) fixed costs relating to trading properties.

We assume perfect foresight so that the path of flat and house prices is known by all agents, and this is rationally built in to expectations. That is not to say that we do not consider dynamics.⁶ As well as steady states, we also consider how prices evolve after unanticipated shocks to incomes, or to credit market conditions. We study this by looking at transitions between steady states after these shocks. Of course, the combination of “perfect foresight”, and completely unanticipated shocks, is somewhat uncomfortable. Nonetheless, understanding *why* certain trends occur on such transition paths should still be useful for thinking about when similar patterns could occur in a more complex economy. As such, this exercise might be considered a useful compromise between an analytically tractable model (such as Ortalo-Magné and Rady (2006)), and the large computational burden of an equilibrium model with several dimensions of aggregate uncertainty.

4.2.1 Elements of the Model

In this subsection we first describe the household lifecycle problem that applies for each household in each cohort of the the model economy, including the specification

⁵We will try to stick to this terminology that demand and supply are counts of the numbers wishing to buy and sell (respectively) in the current period, while the number of properties in the economy is termed the “housing stock”.

⁶A related model of a perfect foresight economy with housing, is that of Fisher and Gervais (2007). Rather than modelling price setting, they have quantities move to equilibrate markets, and analyse residential investment volatility.

of the utility function and the structure of the income process. Second, we describe equilibrium in the model economy and the method for solving for flat and house prices, and in particular for a transition path for these prices after a shock to the economy.

The household optimisation problem, and constraints

As mentioned in the previous part of this section, the model is of J overlapping generations each containing n individuals, so that the total population size is $nJ = N$. In each of the J periods of life, households choose consumption and saving, and whether to own a flat, or a house, or no housing. They make these choices in order to maximise expected lifetime utility. Thus, the household value function for a household of age j in period t is given by:⁷

$$V_{j,t} \left(A_{j,t}, O_{j-1,t-1}, p_t^f, p_t^h, w_{j,t} \right) = \max_{\{c_{j,t}, O_{j,t}\}} u(c_{j,t}, O_{j,t}) + \beta E_t V_{j+1,t+1} \left(A_{j+1,t+1}, O_{j,t}, p_{t+1}^f, p_{t+1}^h, w_{j+1,t+1} \right) \quad (4.1)$$

subject to

$$A_{j+1,t+1} = R_{t+1} \left\{ \begin{array}{l} (A_{j,t} + w_{j,t} - c_{j,t} - p_t^f(1+F)I(O_{j,t}=1) - p_t^h(1+F)I(O_{j,t}=2)) \\ \quad \text{if } O_{j-1,t-1} = 0 \\ (A_{j,t} + w_{j,t} - c_{j,t} + p_t^f(1-F)I(O_{j,t} \neq 1) - p_t^h(1+F)I(O_{j,t}=2)) \\ \quad \text{if } O_{j-1,t-1} = 1 \\ (A_{j,t} + w_{j,t} - c_{j,t} - p_t^f(1+F)I(O_{j,t}=1) + p_t^h(1-F)I(O_{j,t} \neq 2)) \\ \quad \text{if } O_{j-1,t-1} = 2 \end{array} \right. \quad (4.2)$$

where $A_{j,t}$ is the start of period asset stock and $R_{t+1} = 1 + r_{t+1}$ and r_{t+1} is the interest rate on the liquid asset; p_t^f is the price of a flat in period t , and p_t^h the price of a house; F is the proportional fixed cost which is assumed to be the same for both

⁷In this chapter we have used ‘‘O’’ (for ownership) rather than ‘‘h’’ (for housing choice), in order to avoid notational confusion with the ‘‘h’’ used to indicate the larger property type (house) in p^h

buying and selling a house or flat, and constant; $w_{j,t}$ is household income; and j and t respectively index age and year. In this chapter we have added the j index since it will be necessary to explicitly take account of age and year separately when we define equilibrium; in this description of the household's problem (but not in the subsequent definition of equilibrium) an index for the household (within cohort) is still suppressed on all household specific variables, to simplify the notation.

Equation (4.2) is a standard intertemporal budget constraint, augmented by terms reflecting the house price and transaction costs that must be borne when trading housing.

Compared to the previous two chapters, a notable difference in the set up here is that there are separate prices for the flat and the house which enter the budget constraint and are both state variables of the household's problem: since these prices are now set to clear the house and flat markets (see below), we cannot here assume that the flat price is a fraction of the house price.⁸

The structure of mortgage-related borrowing constraints, and for interest repayments, is similar to that in the previous two chapters. For ease of exposition, we again (as in chapter 2) distinguish between beginning of period assets $A_{j,t}$ and end of period assets $s_{j,t}$.

We allow only for debt that is collateralised against a housing asset, and thus those who do not own must have end of period assets greater than zero ($s_{j,t} \geq 0$). Home owners can borrow, though they do so subject to the implicit constraint imposed by the terminal asset condition ($s_{J,t} \geq 0$)⁹ as well as formal constraints on new mortgage borrowing. These formal constraints are a downpayment constraint (equation (4.3),

⁸In the current perfect foresight setting, having separate state variables for these two prices is not computationally burdensome, at least for solving the household's problem, since future prices are non-stochastic. On the other hand, solving for the (known) path of prices after a shock has been the most challenging part of the model (see section 4.2.2, below).

⁹This constraint, together with the specification of marginal utility becoming infinite at zero consumption, prevents households from borrowing more than they can repay with certainty.

where $1 - \lambda_{h,t}$ is the downpayment requirement):

$$s_{j,t} \geq -\lambda_{h,t} p_t^x \quad (4.3)$$

where x could be f (for flat owners) or h (house owners); and a constraint that takes the form of a multiple of income (equation (4.4)):

$$s_{j,t} \geq -\lambda_{w,t} w_t \quad (4.4)$$

As previously, the mortgage-related borrowing constraints must be satisfied only when the household takes on new mortgage debt, either to buy a property or to increase the amount of debt being held. As explained in chapter 2, relative to applying the constraints in every period this structure increases the computational complexity of the problem because it means that the value function will not be universally concave. We use similar techniques to those described in that chapter to overcome this issue.¹⁰

A household that already owns and has an outstanding mortgage, can retain this debt without having to satisfy the mortgage-related borrowing constraints, as long as it pays back the new interest on this debt ($m_{j,t} = r_t s_{j-1,t-1}$). This means that such a household will have end of period debt that is no greater than at the end of the previous period.

Utility function

The within period utility function is a CRRA function in current consumption, augmented by a multiplicative term to capture the value of home ownership.¹¹

¹⁰Numerical solution of the household's problem in the current model is much faster than in the previous chapters because of the limited number of stochastic state variables, but this speed is required as the household's problem has to be solved many times when looking for a transition path between steady states.

¹¹Similar forms for preferences regarding housing are used by Ejarque and Leth-Petersen (2008) and Rios-Rull and Sanchez-Marcos (2008).

$$u(c_{j,t}, O_{j,t}) = \frac{c_{j,t}^{1-\gamma}}{1-\gamma} \exp(\theta\phi(O_{j,t})) \quad \begin{cases} \phi = 0 & \text{if } O_{j,t} = 0 \\ 0 < \phi < 1 & \text{if } O_{j,t} = 1 \\ \phi = 1 & \text{if } O_{j,t} = 2 \end{cases} \quad (4.5)$$

The multiplicative term is such that the utility from consumption is increased by a certain percentage for a house-owner, and a smaller percentage for a flat owner. These percentages are, respectively, 10 and 5 in the results that follow, as this means that the ratio of the flat price to average income ratio at the start of adult life approximately matches that from the data.

Income Process

During the working life,¹² income arrives exogenously and is assumed to be composed of a deterministic permanent component which captures the expected profile of income over the life cycle, and a persistent stochastic component which takes the form of an AR(1) process that is subject to idiosyncratic shocks:

$$\ln w_{j,t} = a_{j,t} + y_{j,t} \quad (4.6)$$

where:

$$y_{j,t} = \rho_w y_{j-1,t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathbb{N}(0, \sigma_\varepsilon^2)$$

and $a_{j,t} = a_{0,t} + a_1 j + a_2 j^2 + a_3 j^3$

During “retirement”, income is a given fraction of final working life income.¹³

Unlike in the models of the previous two chapters, in this chapter there is only one skill (education) group and all agents are ex ante identical, but differ in the idiosyncratic realisations of income. The parameters of the income process are estimated off panel data for Britain (from the British Household Panel Survey (BHPS)), using the

¹²These are the first 44 periods of life (ages 21-64); there are a further 16 periods of “retirement” and agents die after making their age 80 ($j = 60$) consumption.

¹³This fraction (the replacement rate) is set to 50%.

method described in the appendix to chapter 3 (section A3.2).¹⁴ In the results we present here, the estimated parameters are adjusted somewhat, for reasons described in the next subsection.

Fixed Parameters

The model structure just described requires that we fix several parameters in order to solve for choice functions and simulate behaviour. Table 4.1, shows the chosen values for the main such parameters in our baseline model.

Table 4.1: Parameter Values

Parameter	Value	Source
Income Process		
ρ_w	0.764	BHPS
σ_ε^2	0.155	BHPS
$a_{0,t}$	ln 1	Normalisation
a_1	0.00379	BHPS
a_2	0.00037	BHPS
a_3	-0.00002	BHPS
Preference Parameters		
γ	1.43	(Attanasio and Weber, 1995)
ϕ	0.5	HP/Income target
θ	-0.1	HP/Income target
β	1.006^{-1}	
Credit and Housing Market Parameters		
$\lambda_{y,t}$	4.0	
$\lambda_{h,t}$	0.8	
F	0.03	
r	0.006	

Notes: BHPS indicates that source is survey data from the British Household Panel Study; ODPM indicates data that were provided by the then Office for the Deputy Prime Minister, these data are now available from Communities and Local Government (<http://www.communities.gov.uk/index.asp?id=1165366>). The parameters with a t subscript are those that are the subject of experiments in the next subsection.

¹⁴In fact the parameters are those estimated for a separate piece of work (Attanasio and Wakefield, 2008 and forthcoming), and the cubic in age was used because this is better able than a quadratic to capture the pattern of income (on average) around the ages when young children are present in the household.

The parameters reported in table 4.1 show that a_1 , the linear term on wage in the wage process, is closer to zero than the equivalent parameters in the previous two chapters. This means that the income profile has low growth in the early twenties, and by the end of a working career income is expected to be rather lower than it was at age 21. The interest and discount rates are also set close to zero in this parameterisation. All of these choices were made in the light of the fact that house prices display zero growth in the steady state of the economy. The income profile has therefore been adjusted to ensure that expected house price to income ratios throughout a working life are matched by the model, and it is not the case that housing becomes more affordable as agents age, in a counterfactual way. Given the potential importance of constraints on mortgage borrowing in affecting housing demand, it would seem important to capture these price/income ratios in a realistic manner. The interest and discount rates were chosen so that housing did not seem hopelessly unattractive as an asset.

These parameter choices for income growth and the interest and discount rates will affect the utility parameters required to achieve a target house price, and so could be important in shaping outcomes of the model analysis. It should also be noted that in this case with utility that is not homothetic, it is not possible to normalise the model solution by dividing through by the permanent growth in the house prices. On the other hand, initial experiments with the parameters suggest that changing the income profile to a more standard hump-shape, for example, will not affect the nature of the results reported later in the chapter.¹⁵

The parameters of the mortgage market institutions in the model, were chosen to approximately match terms available on high street loans. The fixed cost (F) of 3% applies to both buying and selling.

Having fixed the other parameters, the parameters on utility from housing were

¹⁵The qualitative nature of the results regarding credit constraints is robust to having a more standard hump-shaped pattern for expected earnings, for example.

chosen to achieve a flat price of around 6.7,¹⁶ thus approximately matching to ratio of property prices to average incomes for young households in the data.

4.2.2 Solving for flat and house prices

The focus of the model for this chapter, is solving for equilibrium flat and house prices. As mentioned above, these prices must be such that the number of families wishing to be owner-occupiers of each type of home is equal to the fixed stock of flats and houses in the economy. An equivalent way to say the same thing is that the number of families supplying flats and houses to the market for sale, is equal to the number generating demand to purchase each type of home in the current period.

The analysis is concerned not only with steady states, but particularly with transition paths between steady states after unforeseen shocks to incomes, or the institutions (parameters) of the mortgage market. The broad strategy for finding a transition path is first to solve for the “before” and “after” steady states, and then to take a guess at a possible path between these initial and final positions. The guess must then be updated until the path is such that prices match expected prices along the path (with perfect foresight) and the flat and house markets are in equilibrium every period. Before describing this algorithm in more detail, a formal definition of equilibrium is required.

Equilibrium

The aggregate state variables that are subject to transition inducing shocks are the level of the the borrowing constraints (λ_h and λ_w) and average (or aggregate) income in the economy, a sufficient statistic for which is $a_{0,t}$.¹⁷ Notice that the model is of a

¹⁶The relatively high ratio of 6.7 reflects the fact that this number came from 2007, a time of high house prices.

¹⁷Changes in other common parameters of the income process, including the variance σ_ε^2 , would also shift the equilibrium of the economy but it has been assumed that these are fixed and so this is suppressed in the notation that follows. Similarly the exogenous fixed cost F will not be listed as a state of the economy.

small open economy, and the interest rate is assumed to be fixed at level r ; to simplify the notation we omit this from the list of aggregate states in what follows.

Given the sequences of the exogenous aggregate states $\{\lambda_{h,t}, \lambda_{w,t}, a_{0,t}\}_{t=1}^{\infty}$, and of income realisations $\{\{w_{i,j,t}\}_{i=1}^n\}_{j=1}^J\}_{t=1}^{\infty}$, an equilibrium for the economy is a sequence of variables $\{\{c_{i,j,t}, s_{i,j,t}, O_{i,j,t}\}_{i=1}^n\}_{j=1}^J, p_t^f, p_t^h\}_{t=1}^{\infty}$ such that:

1. households' choices of c and O optimally solve the household problem (4.1), taking prices as given and known with perfect foresight;
2. the markets for flats and houses clear, which is to say that at each date t the following are satisfied:

$$\sum_{j=1}^J \sum_{i=1}^n I(O_{i,j,t} = 1) = M^f$$

and

$$\sum_{j=1}^J \sum_{i=1}^n I(O_{i,j,t} = 2) = M^h$$

where M^h and M^f respectively are number of houses and flats in the housing stock, $I(\cdot)$ is an indicator function, i is a within cohort household identifier and j and t index age and time as before.

Steady states

As is customary, a steady-state equilibrium is an equilibrium of the economy with no aggregate shocks or growth, and such that the prices (p_t^h, p_t^f) are constant over time. The steady state also involves constant cross-sectional distributions (by age and income) of individuals over consumption, financial wealth holdings and home-ownership states, as each generation follows the same optimising choice functions.

Conceptually there are no major difficulties with computing the steady-state equilibrium for the model economy. In practice, to ensure that each step in the iterative procedure is a movement towards the equilibrium, one must be careful to take account of cross-price effects (particularly from the house price to the demand to own flats) as

well as own-price effects, when updating prices.

Algorithm for finding a transition path

Since the method for finding the transition path for flat and house prices is not entirely standard, it is worth outlining the algorithm. The shock hits the economy (which is in its steady-state equilibrium) at the beginning of period $t = 1$. The algorithm proceeds through the following steps:

1. Solve the household's problem, and simulate the initial steady state of the economy.
 - From this simulation store the distribution of households over financial asset holdings, home ownership status, and income draws, in the initial steady state, as this information will be used to simulate the economy immediately after the shock.
2. Take an initial guess at the path of prices $(\widehat{p}^h, \widehat{p}^f)$ between steady states.
3. Solve (over a long horizon) for each cohort back to the date of the shock, given this path of prices. Simulate to check whether the guess is a rational expectations equilibrium. If it is, stop.
4. If it is not a rational expectations equilibrium:
 - Update the guess for period $t = 1$ (leaving the path from period 2 forward unchanged), re-solve to period $t = 1$ and simulate demand and supply for homes in that year to check whether the new price achieves housing market equilibrium. Keep updating the guess for this year until equilibrium is achieved.¹⁸

¹⁸The current method for updating prices in a particular period involves exploiting information on own and cross price elasticities, where these elasticities are computed for a one-off price move, occurring in the relevant year.

- Once equilibrium in year one is achieved, repeat the procedure for year 2, but now aiming to achieve equilibrium in the housing market in years 1 and 2 simultaneously, given that the path of prices beyond year 2 remains fixed. Note that year one is solved with expectations reflecting the new guess for the year two price.
 - Keep repeating this procedure, but as equilibrium is achieved up to each given year, keep increasing the number of years in the search by one year at a time.
5. As the number of post-shock years being solved and simulated increases, so the prices should hopefully converge towards the “after” steady-state equilibrium. Once the prices seem to have settled at this steady state, return to step 3., and solve and simulate for a long horizon to check that the final steady state really has been reached. If the check is satisfied, stop. If not continue the search for a longer horizon.

Each step in the algorithm involves simultaneously solving for flat and house prices that achieve equilibrium in the housing market in each year of a set of years. As adjustments in prices within this set of years are made to balance markets, so these adjustments are built in to expectations for the earlier years in the set. Once equilibrium has been achieved across a given set of years, the number of years in the set is increased by one and the process is repeated. Following this procedure means that once we reach a point at which prices have converged to and settled at the after steady state, we can be confident that we have found a rational expectations equilibrium transition path.

This method is somewhat different from standard procedures for finding price paths in perfect foresight economies, such as that employed by Lee (2005) and that of Auerbach and Kotlikoff (1987). My understanding of the procedures employed by these authors is that after taking an initial guess at the price path, one proceeds by moving

forward along the transition path and achieving equilibrium in each year separately, given the initial guess at prices beyond that year. Thus, at each step in the algorithm one will find an equilibrium for each year, but this may not be a rational expectations equilibrium since prices forward may not coincide with the guess being followed. After completing this updating for a large number of years, the algorithms involve checking whether the paths follow rational expectations, and if not the year by year equilibrium search is repeated, but now starting from the set of (non-rational expectations) equilibrium prices that have just been found. The difference from my procedure lies in not adjusting prices and expectations earlier along the path, when solving for a given year.

I initially tried to implement an algorithm much like that described in the previous paragraph, before settling on my current method. However, I found that this algorithm did not reliably converge.¹⁹ Though my method is not standard, I am confident that it is finding correct transition paths. The final stage in the algorithm is always to plug in the final guess at the transition path and solve and simulate a long horizon, to check that the economy has converged to the final steady-state equilibrium. I have not considered analytically whether the transition path should be unique, but I have found that the algorithm converges to the same final solution from different initial guesses at the transition path.²⁰

Some further practical notes on the implementation of my algorithm are that:²¹ At each modelled date I populated the model economy with 60,000 individuals (households) drawn from 60 different cohorts; and, I assumed that there are 20,000 flats and 20,000 houses available for owner occupation so that the owner-occupation rate is

¹⁹It is possible that this had to do with the unusual features of the demand for flats, which were discussed in chapter 2.

²⁰To aid the speed of with which the model finds the transition path, my preferred initial guess is that prices go directly to the new steady state. The alternative that prices follow a linear trend over 25 years from the before to after steady states, would also be reliable.

²¹Another practical point is that in the current runs I have used only five nodes in the approximation for the stochastic part of idiosyncratic income. Experiments indicate that this low number of nodes may have affected the steady state solutions slightly, but the properties of the transition paths are not altered if the number of nodes is increased.

two-thirds, which is similar to the rate in the U.K. economy.

4.3 Income Shocks and Credit Constraint Shocks

In this section we look at two entirely unexpected shocks which hit our model economy. First we look at a shock to idiosyncratic incomes, such that at a point in time, everyone in the economy becomes 10 percent richer. Second we consider a shock to credit constraints, such that the multiples (λ_h and λ_y) that banks are prepared to lend change from 0.7 and 3.2, to 0.8 and 4.0. Table 4.2 summarises these different model parameters before and after the shocks. Note that the experiment of a change in incomes is described by a change in the exponent of the parameter a_0 (the intercept for the deterministic part of log income); since this is permanent, when it changes by a given proportion, this results in the same proportionate change for all incomes (including those of retired individuals) throughout the economy.

Table 4.2: Parameters before and after shocks

Parameter	Before	After
Income Shock		
exp a_0	1	1.1
λ_h		0.8
λ_y		4.0
Credit Market Shock		
\bar{w}_1		1.1
λ_h	0.7	0.8
λ_y	3.2	4.0

It is clear from table 4.2 that the “after” steady state is identical in the two cases that we simulate: the shocks leave the institutions of the economy (the parameters of the model) the same for both cases. This enables ease of comparison between the two simulated price transition paths. The comparison is also aided by the fact that the before steady states in both the “income shock” economy, and the “credit shock”

economy, are constructed to have the same flat price. Thus between steady states the adjustment of the flat price is the same in both of the modelled economies. This is not quite true for the house price: between steady states the house price adjusts slightly less in the credit shock economy than it does in the income shock case. The steady state prices are reported in table 4.3. We see that the change in the flat price between steady states is an increase of slightly more than 9%, while the change in the house price is slightly less than 9% in the income shock case, and around 5.4% in the credit shock case.

Table 4.3: House and flat prices before and after shocks

Price	Before	After
Income Shock		
p^f	6.67	7.28
p^h	11.43	12.45
Credit Market Shock		
p^f	6.67	7.28
p^h	11.81	12.45

Figures 4.2 and 4.3 show how home ownership at different ages in the steady states is affected by the change incomes or credit markets, and therefore in house prices. In the case of the income shock, the age-specific ownership rates for both flats and houses do not change by more than four percentage points at any age, and almost do not change at all below age 30. There is a small decrease in flat ownership between age 40 and the early 60s, offset by a small increase in early retirement. For house ownership the changes are in the opposite direction: a small increase in middle and late working age, with a small decrease either side of retirement age.

The changes between steady states are more substantial for the case of a credit market shock (though the gradient of the ownership profiles makes the larger differences during the twenties somewhat difficult to see). In this case, for both flats and houses the changes between steady states involve a shift of ownership towards a slightly earlier

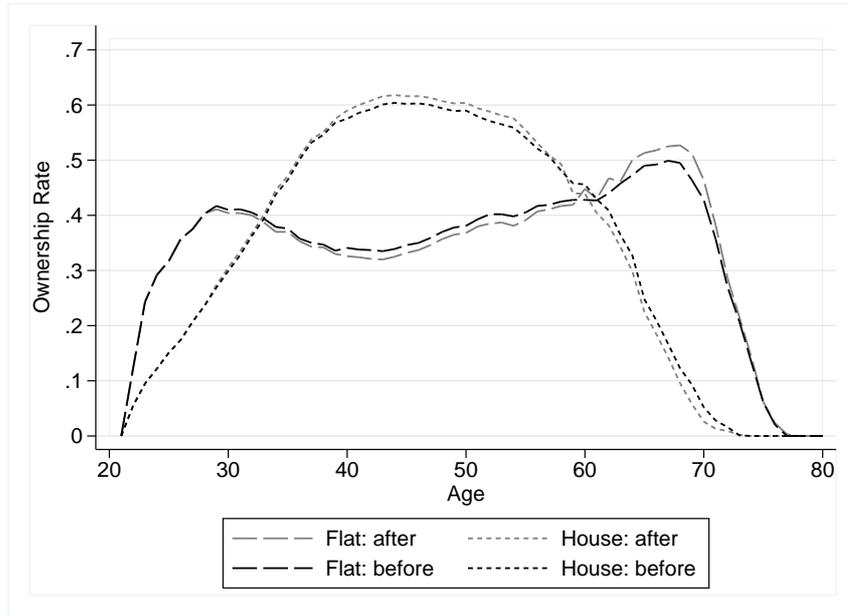


Figure 4.2: Ownership by Age: Steady states before and after income shock

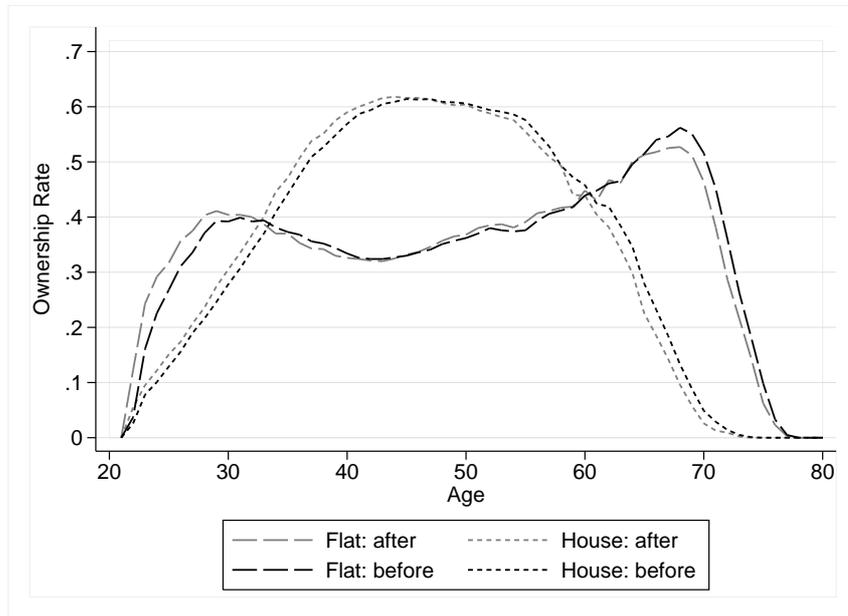


Figure 4.3: Ownership by Age: Steady states before and after credit market shock

part of the lifecycle, with increases of up to six percentage points for flat ownership for those aged less than thirty.

While table 4.3 shows prices in the before and after steady states, our real interest is in the evolution of the economy in transition between steady states. Figures 4.4 and 4.5 respectively show the paths of (log) flat and house prices, after the shocks to income (long dashes) and credit constraints (short dashes). We see that after the income shock, prices move very close to the after steady state levels immediately after the shock occurs, and do not deviate more than half a percentage point from these levels throughout the course of the transition to the post-shock steady state.²² In contrast, following the credit market shock, the flat price overshoots the post-shock steady state by slightly more than 0.04 log points (4 percent), and the house price overshoots by slightly less than 4 percent. After this initial overshoot, the prices gradually trend back down to the steady state, dipping very slightly below this level around 50 years after the shock, before settling at the steady state levels after around 65 years.

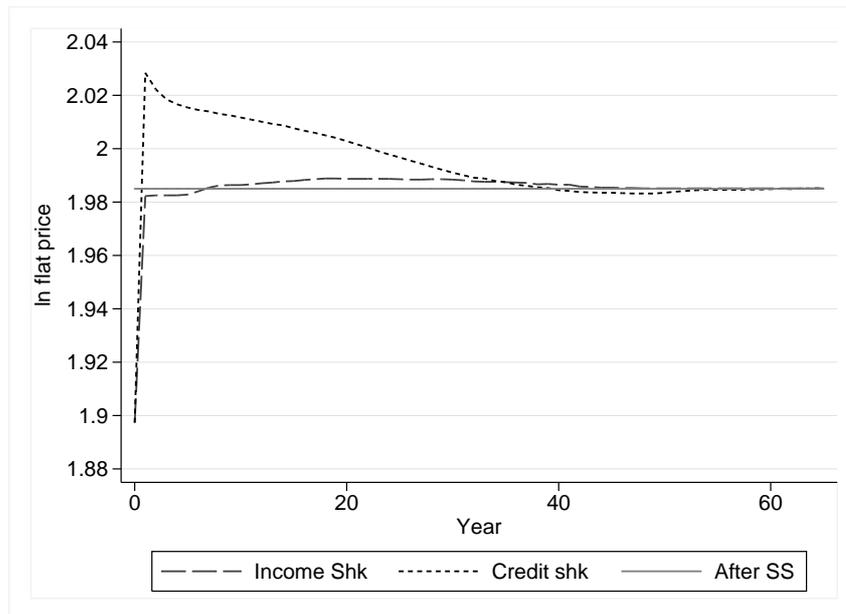


Figure 4.4: Log flat price following income and credit market shocks

As a supplement to the figures of the price paths in transition between steady states,

²²Though the prices go very close to the steady state levels, transition to the after steady state actually takes almost 50 periods (where being at the steady state is measured by having the housing market in equilibrium (repeatedly) and prices in the computations fixed at the steady state prices).

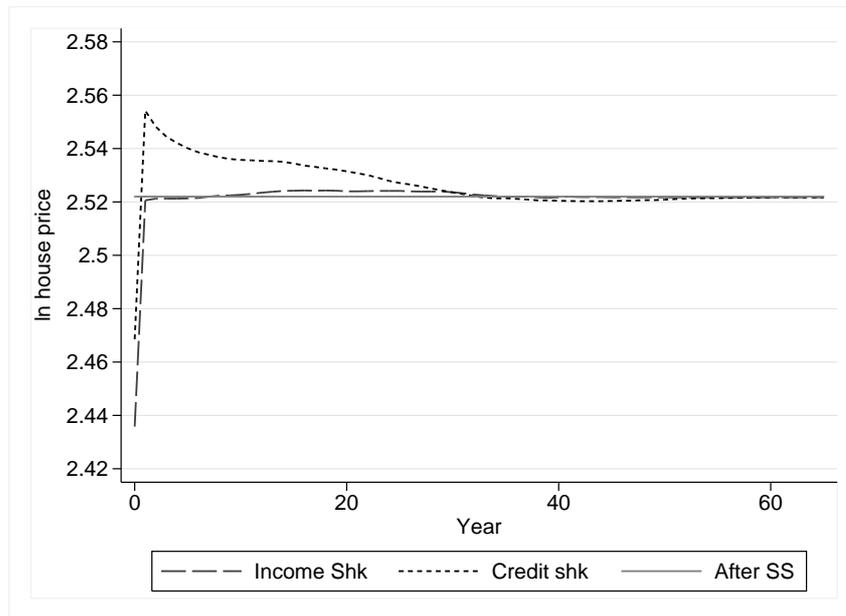


Figure 4.5: Log house price following income and credit market shocks

figures 4.6 and 4.7 show how the volume of flat and house transactions moves after the shocks. In the steady states, transactions occur as young individuals climb on to the property ladder while older individuals downsize in order to consume their wealth, and due to idiosyncratic earnings shocks. In the post-shock steady state, these factors lead to transaction rates of approximately 8% in the flat market (slightly fewer than 1600 of 20000 flats are traded each year), and about 3.5% in the house market (700 out of 20000).²³ In the figures, the number of transactions is expressed relative to these post-shock steady state values, so that a value greater than 1 indicates transactions exceed the steady state volumes. The marked contrast between the charts for the income shock and for the credit market shock, is that it is only in the case of the shock to the mortgage market that we see notable increases in the volume of transactions along the transition path.²⁴ In particular, in the early years after the shock to credit, flat

²³Summing these numbers and comparing to population size, we see that almost 4% of households are buying a property each period.

²⁴On the suggestion of both of my supervisors, as a sensitivity check, I repeated the income shocks experiment for a case in which both of the mortgage constraint parameters were much tighter than in the example reported here, even setting λ_h to 0.2 and λ_w to one. It turned out that this made very

transactions are above their steady state level by around 10%,²⁵ and house transactions are elevated by around 5%.

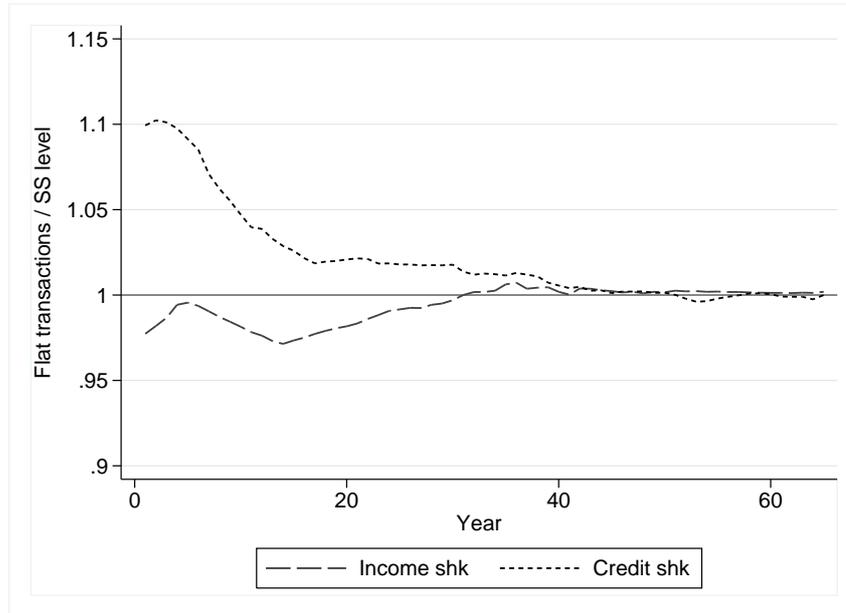


Figure 4.6: Flat transactions, relative to after SS level, following income and credit market shocks

A short summary of the results of these model runs is that house and flat prices substantially overshoot their new steady state levels in the aftermath of a mortgage market shock, but not following a shock to incomes in the economy. This overshooting of prices following the credit shock, is accompanied by a notable increase in property transactions when prices are high.

In the light of evidence that house prices tend to be more volatile than incomes, and that housing market transaction levels tend to be positively correlated with property prices, these findings, particularly for the credit market shock, are potentially interesting. However, our economy is stylised. To try to understand whether the results might generalise to a more complex economy, it is useful to try to disentangle what is

little difference to the pattern that prices and transaction levels do not deviate much from the steady state after this type of shock.

²⁵Transaction numbers in the simulated economy slightly above 1700 for flats, rather than slightly below 1600 in the steady state.

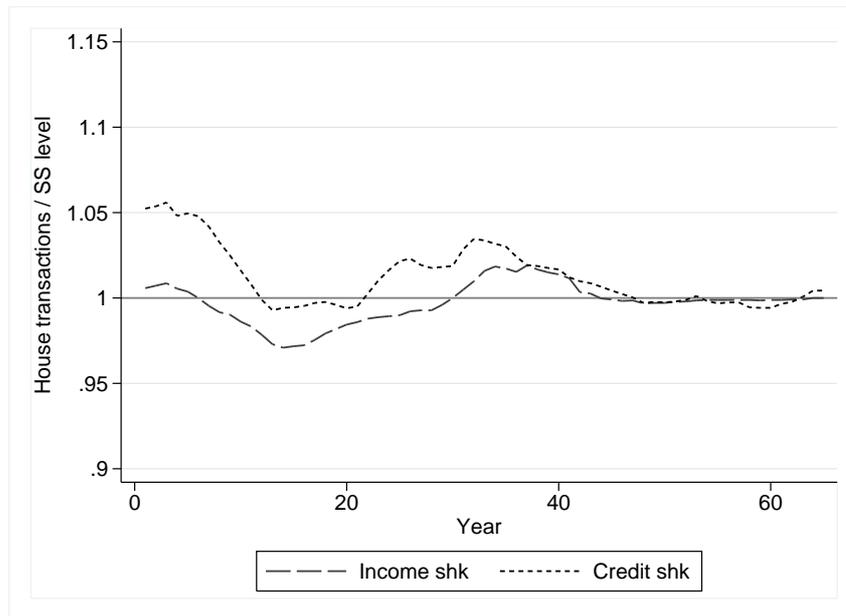


Figure 4.7: House transactions, relative to after SS level, following income and credit market shocks

driving the nature of these changes in prices and transaction levels.

To think about why prices and transactions move in these ways along the transition path, it is helpful to consider what factors shape the demand for and supply of houses in the model economy, and how these would be balanced if the economy jumped directly to the new steady state prices following a shock (and was expected to stay there). The structure of the model of household behaviour is such that whether a household will be willing to own a house or flat at given prices and expected prices, will be determined by:

- The amount of financial wealth held, $A_{j,t}$;
- Whether or not they already own, $O_{j-1,t-1}$;
- The current income level of the household, $w_{j,t}$;
- The parameters controlling mortgage borrowing ($\lambda_{h,t}$ and $\lambda_{y,t}$);
- The level of fixed costs.

In our experiments, once the shocks have occurred, levels of income, the tightness of mortgage constraints and levels of fixed costs, are all fixed at their post-shock steady-state levels. Thus the factors that prevent prices and transactions from settling immediately at their steady state levels must be those chosen by the households in the period before the shock: A_{t-1} and O_{t-1} . As we have seen in figures 4.3 and 4.2, ownership rates across ages are not dramatically changed between steady states. Table 4.4 shows that for the income shock, age-ownership rates almost do not change along the transition path for the economy. Thus for this case, a good starting point for thinking about what is driving the shape of the transition is to look at the distribution of financial assets around the shock.

Table 4.4: Ownership Rates Around Income Shock

Age-Group	Before SS	Year 1	Year 5	Year 10	Year 20	After SS
21-25	28.0	28.0	28.3	28.3	28.0	28.0
26-30	62.9	62.9	63.1	63.3	62.9	62.9
31-40	84.1	84.1	84.1	84.1	84.1	84.1
41-50	94.9	94.9	95.0	95.0	94.9	95.0
51-60	93.6	93.7	93.8	93.7	93.7	93.4
61-70	70.9	70.9	70.8	70.9	71.3	70.7
71-80	10.9	10.9	10.7	10.6	10.5	11.4

Rather than just considering the overall level of assets, it is more helpful to consider asset holding for those particular groups that most affect the demand for, and supply of, properties onto the market. The most active group of property buyers in the model economy are young families, either moving onto or up the property ladder: ownership rates increase from zero when families enter the labour market, to almost 80% (with slightly more flat owners than house owners) by age thirty. To think about whether more young people would (*ceteris paribus*) want to own properties immediately after the income shock than in the final steady state, we can look at the financial wealth of young individuals who do not own in the before steady state, and compare it to the wealth of similar individuals in the after steady state. Since ownership rates are

virtually identical in the two steady states in this age range, there is no composition problem in terms of the types of individuals in these before and after groups. Table 4.5 describes the wealth distribution for this group, in both steady states, and along the transition path (the units are such that a value of 1 is equal to mean income at the beginning of the economic life (age 21) in the before steady-state).

Table 4.5: Financial Asset Distribution: Non-owners aged 30 and under, around income shock

Percentile	Before SS	Year 1	Year 5	Year 10	Year 20	After SS
10th	0.00	0.00	0.00	0.00	0.00	0.00
25th	0.05	0.05	0.06	0.06	0.06	0.06
Median	0.41	0.41	0.45	0.43	0.42	0.43
75th	0.76	0.76	0.85	0.85	0.85	0.85
90th	1.20	1.20	1.30	1.32	1.32	1.32
Mean	0.50	0.50	0.54	0.55	0.55	0.55

Looking at the steady state columns first, we do not see dramatic differences in wealth levels or the shape of the distribution. Rather, wealth in the before steady state is slightly lower at all points in the distribution than is true in the higher-income steady state. In transition, wealth levels in general move smoothly to the new steady state, and are very close to it after ten years when all members of this age-group are “born” in the post-shock economy. Since wealth levels for this group are never substantially different from those in the final steady state, property demand will not be either. Since wealth is initially slightly below the steady-state level, demand will actually be slightly below the steady-state level, thus tending to depress prices.

While those aged under 30 who do not own are the most likely buyers of properties in the economy, sellers are drawn predominantly from those aged above sixty and owning either a flat or a house. Table 4.6 shows the distribution of financial wealth for this group, again for several different years after the income shock. As with the young non-owners, this table for the older owners shows that wealth holdings before the income shock are somewhat lower than in the post-shock steady state. At the extreme

of the distribution this could be driven by a slight shift in this age group towards flat owning (thus converting some resources into financial wealth), but the change is clear for a broad range of the distribution from the median upwards. This indicates that in the after steady state individuals do hold slightly more financial wealth (saved from their higher lifetime resources), even though their wealth has also increased due to the property price appreciation. The fact that this group has less financial wealth just after the shock than in the after steady state will again tend to decrease the number of households wishing to owner-occupy, relative to the steady state level: for given prices individuals with less wealth will be more inclined to exit the property market and instead use financial assets to smooth out consumption over the remaining years of life.

Table 4.6: Financial Asset Distribution: Owners aged over 60, around income shock

Percentile	Before SS	Year 1	Year 5	Year 10	Year 20	After SS
10th	-0.64	-0.64	-0.72	-0.69	-0.67	-0.69
25th	0.04	0.04	-0.03	-0.00	0.05	0.06
Median	0.97	0.97	0.97	1.01	1.07	1.10
75th	2.15	2.15	2.25	2.28	2.38	2.45
90th	3.29	3.29	3.37	3.47	3.64	3.90
Mean	1.31	1.31	1.34	1.44	1.57	1.68

This examination of wealth distributions for the young non-owners, and older owners, has shown that after the income shock wealth holdings are such that, if property prices jumped directly to the steady state level (and were expected to stay there), then there would be fewer young households wishing to buy properties, and more older households wishing to sell them, than would be true in the steady state. This must mean that in the absence of any expectation of large price rises, property prices immediately after the shock will fall short of the after-shock steady state level, otherwise a situation of excess supply would exist. Since there is no force in play to generate expectations of strongly increasing prices, this is indeed what we observe. The prices

do though immediately move very close to the after steady state level. This is not surprising given that wealth levels of the key groups of potential buyers and sellers are only slightly below the levels that will occur in the post-shock steady state. Given that the price immediately moves very close to the new steady state, and that patterns of ownership across generations are also close to the steady-state distribution, the incentives for younger generations are such that they will make consumption / saving and housing choices that closely match steady-state behaviour. This explains why, following the income shock, the economy converges to the post-shock steady state without large divergences of prices or transactions from the steady state levels.

Turning to the case of the shock to the mortgage market, we can again look at asset holdings of the group of likely buyers (young individuals who do not own) and potential sellers (older individuals who own) to see what might drive price and transaction movements after the shock. In this case we have to be a little more careful to allow for the composition (in terms of income levels) of these groups around the time of the shock, and in the steady state, since the proportion of owners at these ages changes slightly along the transition path (see table 4.7). For those aged less than thirty, ownership rates increase along the transition path, so the group of non-owners has slightly higher income immediately after the shock than in the post-shock steady state.²⁶ This difference reflects the fact that the credit shock shifts home ownership towards younger individuals, but the economy does not immediately move to the new steady state distributions. However, on its own it could not lead to overshooting of flat and house prices, as the higher income subset who are left as non-owners will not “over-demand” the steady state unless some factor other than income is at work. There must therefore be some other factor that is boosting the demand to buy properties, or depressing the supply of properties to the market, immediately after the shock.

One such factor is the asset levels of young potential buyers. The savings of these

²⁶ Average income in this group is 0.793 in the first year after the shock, and 0.784 once the post-shock steady state is reached.

Table 4.7: Ownership Rates Around Credit Market Shock

Age-Group	Before SS	Year 1	Year 5	Year 10	Year 20	After SS
21-25	20.5	24.1	27.3	27.8	27.9	28.0
26-30	57.8	57.7	59.8	62.4	62.8	62.9
31-40	81.9	81.4	81.3	82.2	83.7	84.1
41-50	94.4	94.0	94.3	94.2	94.5	95.0
51-60	93.9	93.3	92.4	92.4	92.4	93.4
61-70	76.4	75.9	74.8	73.5	72.0	70.7
71-80	14.4	14.5	13.7	12.8	12.2	11.4

Table 4.8: Financial Asset Distribution: Non-owners aged 30 and under, around credit market shock

Percentile	Before SS	Year 1	Year 5	Year 10	Year 20	After SS
10th	0.00	0.00	0.00	0.00	0.00	0.00
25th	0.11	0.10	0.06	0.06	0.06	0.06
Median	0.47	0.46	0.39	0.40	0.40	0.43
75th	1.06	1.12	0.92	0.85	0.85	0.85
90th	1.63	1.69	1.43	1.32	1.31	1.32
Mean	0.67	0.67	0.59	0.54	0.54	0.55

individuals are higher at all points in the distribution immediately after the shock, than in the post-shock steady state (see table 4.8). These individuals have been saving up to meet a tighter downpayment requirement, and now find they have “over-saved” relative to the optimum for the post-shock economy. Given the way that this extra wealth interacts with the downpayment requirement in the economy, this is an important source of extra demand to buy flats and houses immediately after the credit shock, compared to in the eventual steady state.

Similarly to the young potential buyers, so the older potential sellers of properties are seen to have higher wealth immediately after the credit shock, compared to in the eventual steady state. This group is smaller (more selected), and therefore has higher income on average, in the final steady state than in the immediate aftermath

Table 4.9: Financial Asset Distribution: Owners aged over 60, around credit shock

Percentile	Before SS	Year 1	Year 5	Year 10	Year 20	After SS
10th	-0.42	-0.42	-0.59	-0.59	-0.59	-0.69
25th	0.26	0.26	0.14	0.12	0.10	0.06
Median	1.35	1.34	1.23	1.19	1.18	1.10
75th	2.64	2.62	2.53	2.52	2.51	2.45
90th	4.11	4.04	3.97	4.21	4.09	3.91
Mean	1.94	1.94	1.88	1.91	1.86	1.69

of the shock.²⁷ Thus the higher wealth for this group certainly reflects higher saving in financial assets in the pre-shock steady state. For this group the higher financial wealth reflects the lower house price in the pre-shock steady state which leads to lower borrowing at the stage in the life-cycle when property is bought, and encourages higher financial saving as a substitute for housing wealth to fund retirement consumption. The key point of interest here is that the higher financial wealth will serve to discourage this group from putting their properties on the market: higher financial wealth means that there is less need to sell in order to have wealth available to smooth consumption. Thus, given the financial wealth distribution of this group, in order to induce selling that will reduce the proportion of owners in the group (as must happen between steady states), a price above the eventual steady state price will be required at some point in the transition.

Thus, the mechanism that drives price overshooting after the credit constraints shock is based on an “excess” of (financial) wealth in the economy immediately after the shock, compared to in the final steady state. This encourages high demand to buy houses from young households who can meet downpayment requirements, and also discourages supply to the market from existing owners who do not need to access their housing wealth so rapidly in order to smooth consumption. Together, these forces would generate excess demand with prices fixed at the post-shock steady-state levels,

²⁷Mean income increases from 0.294 to 0.305 between the first year after the shock, and the final steady state.

and so this generates overshooting. That these forces generate quite a large overshoot (around 50%) beyond final steady state prices, and also a surge in transaction volumes along the transition path, indicates that the movement of prices beyond the final steady state levels (and the associated creation of expectations of capital losses) more rapidly encourages older owners to sell, than it chokes off demand from potential young buyers. This partly reflects how the interaction between the downpayment constraint and wealth levels leads to a large increase in the home purchasing power of young households. The elevated level of transactions, with high numbers of young people buying as older households exit the property market, is the mechanism through which ownership shifts down the age distribution along the transition path.

In this section we have analysed two shocks to the model economy - one to income levels and one to mortgage-borrowing constraints - that lead to similar changes in steady state flat and house prices. However, we have seen that the path of the housing market between steady states is quite different for the two cases: for the income shock prices move almost directly to their new steady state levels, and transactions also stay around steady state levels. In contrast, the shock to mortgage-market constraints leads to prices overshooting the final steady state levels by almost fifty percent, and transaction levels are also high (relative to steady state) at the time when prices are high.

Ortalo-Magné and Rady (2006) proposed a mechanism through which house prices and transactions might move together, and substantially, following shocks to the economy. It is worth emphasizing that while the current model economy is similar to their analytical framework in terms of the perfect foresight and unanticipated shocks set up, the mechanism through which we get overshooting of house prices, and comovement of prices and transactions, is quite different. In their analysis the effect came from capital gains increasing the lifetime wealth of existing flat owners by more than income gains increased the lifetime wealth of successor generations of potential owners, and thus

these capital gains led to an overshoot of demand for houses along the transition path. In our case the mechanism is from a high amount of financial wealth in the economy immediately after the credit-market shock, which leads to “excess demand” for both flats and houses.²⁸

Of course the capital gains mechanism does operate in our economy but is not strong enough to generate any substantial overshoot of house prices following the income shock, while the credit constraint shock leads to a slightly larger overshoot for flat prices than for the price of the trade up property. In the economy we have studied it may be that very large flat price movements would be needed to generate capital gains that substantially outweigh the extra resources of generations that have increased income over several periods. Also, the increased wealth from capital gains will be offset in our economy through reduced saving in financial assets.

The final section of this chapter considers what extensions to this work would be needed to see whether the assets mechanism that we identify here could be important in generating house price fluctuations in the U.K. and other economies.

²⁸There is limited scope for this mechanism to ensue in the framework of Ortalo-Magné and Rady (2006), since utility that is linear in consumption means that all consumption is postponed until the end of life. In these circumstances there is no sense in which credit market parameters will lead to non-owners saving more than is consistent with a new steady state.

4.4 Conclusion and Next Steps

This chapter has been concerned with endogenising the prices of houses and flats. More precisely, a lifecycle model of consumption and saving, and decisions about owner-occupation, has been built into an overlapping generations framework in which house and flat prices are those that clear the property market each period.

The model has been constructed to capture certain realistic features of the housing good/asset, and the market for properties. In particular, properties are modelled as large (lumpy) objects that people like to own; ownership of housing can be adjusted only if transaction costs are borne, and so adjustments are infrequent; house purchases are often facilitated by mortgage borrowing, which is limited by downpayment and income-related constraints applied at the time when a mortgage is agreed; the size of the housing stock in the economy cannot be adjusted rapidly (indeed the model has a fixed stock of properties available to owner occupy).

The economy that is modelled contains households who are heterogenous in terms of the (stochastic) evolution of their earnings, but there is no aggregate income growth nor any aggregate uncertainty, so that the steady state has constant flat and house prices. Although we have not modelled aggregate uncertainty, we have considered the effects of unforeseen shocks to the economy, and analysed the transition path of the economy to a new steady state after such shocks. The transition path is one with rational expectations, and indeed perfect foresight, over property prices.

A key aim in studying shocks to this economy was to consider whether the interaction of such shocks with idiosyncratic risk in the economy, and particularly with mortgage-borrowing constraints, could plausibly lead to patterns of house price movements and housing transactions that are in line with those observed in data. Two types of (positive) shock were studied: a shock to income levels and a shock to mortgage-borrowing constraints. The shock to mortgage borrowing constraints seems more promising for explaining housing market fluctuations, as it produces larger price

swings (or, more precisely, bigger swings beyond the post shock steady state), and a transition path in which the house price and transaction levels are positively correlated.

The mechanism underlying this pattern in transition after the mortgage market shock is that the accumulation of financial wealth in the pre-shock steady state is higher than is consistent with the post-shock steady state. This means that immediately after the shock, the demand from young households to buy property is (*ceteris paribus*) higher than it will be once the steady state is reached, while the supply of properties onto the market from older household is lower than at the steady state. This pushes prices beyond the final steady state levels, and transactions are also high as high prices not only choke off demand but also increase the supply of properties to the market. Understanding the nature of this mechanism should help us to think about whether the results in our stylised economy, might generalise to a more realistic setting.

In the modelled economy, there is something uncomfortable in the combination of rational expectations and perfect foresight, with aggregate shocks that are unforeseen and so lie outside expectations. It would be preferable,²⁹ to analyse an economy in which agents form expectations over aggregate shocks and thus might be able to (incompletely) self-insure against them. In the light of the results of the present chapter, in such a model I would look for circumstances such that following a positive shock (i.e. one that is expected to increase house prices) to the economy people find (ex-post) that they have “oversaved”, as this seems a good candidate for a case in which prices might then rise substantially and transactions might increase at the same time. An expansion in credit might generate such a situation, particularly since before the expansion individuals who would like to borrow more cannot offset the unavailability of credit through their saving behaviour.³⁰

Before attempting to analyse such an ambitious model, I think that there are

²⁹This would also represent a move into supercomputing territory; cf. Rios-Rull and Sanchez-Marcos (2008) and Rios-Rull and Sanchez-Marcos (In Progress).

³⁰Of course this is speculation based on intuition: to properly explore the evolution of such a model economy will require getting the computer to tell us the numbers.

several exercises of interest that could be carried out within the current framework (or one very closely related to it). An obvious target for further research would be to check the robustness of the results reported here. It would, for example be worthwhile doing a more thorough job on matching the steady state of the modelled economy to features of the U.K. economy, and checking that the pattern of results described still applies with the modified preference (and other) parameters.³¹

An analysis that would involve some extension to the current model would be to introduce bequests into the economy. In section 4.3, the discussion of the mechanisms driving prices and transaction levels in transition focussed on the demand from young households who haven't yet got onto the property ladder, and the supply of properties to the market from older home-owners. This supply arises in the model because households get no utility from bequests, and so eventually choose to eat all their resources, including any house or flat that they own. In reality many individuals who become owner-occupiers while working, leave property as part of their estate. Capturing this would change the pattern of ownership, and the age distribution of buyers and sellers, in the model.

It is not, though, obvious that bequests would alter the patterns of results regarding shocks to the economy. It is (perhaps) tempting to think that having bequests, such that older owners will be much less inclined to sell, must change the results regarding movements in the number of transactions. However, what matters in the changes in the level of transactions is adjustments around the shock. In the economy with bequests the low number of older sellers would be a feature of both pre- and post-shock economies and as long as price increases reduce the incentive to buy, and increase the incentive to sell, the patterns of results we have observed could still go through. This would require exploration using the apparatus of a modified model.³²

³¹As mentioned in footnote 15, we have run the credit constraints case under a different assumption about the evolution of income over the lifecycle, and this delivered results qualitatively identical to those reported here.

³²A point to add is that it might not be straightforward to add the bequest to the model. It is not clear that the distribution of estates, and levels of ownership at death and at earlier ages, can easily

If the results do prove to be robust to checks and to bequests, it might be worthwhile thinking whether there are implications of the model that are testable in data. In the data there are formidable obstacles to separating the effects of income shocks from the effects of credit constraints on house prices, since these tend to move together with the economic cycle.³³ Looking within the model at how these different factors might influence the distributions of consumption and financial wealth, might suggest alternative means to separate the different possibilities (though this would need working out). Analysis of distributions might also suggest means for investigating the nature of the relationship between house prices and consumption growth (cf. chapter 3).

A finding of robust results would also support some additional analyses using the model. One line I intend to pursue is to use the framework that I have established to conduct some policy experiments. Interventions in the housing market are often discussed as a means to make housing more “affordable”, since policies such as subsidised mortgages, or cuts in the tax on housing transactions for first time buyers, seem well targeted at young aspiring owners. However, the effects of such policies on housing demand and house prices may undo welfare gains to these groups that seem plausible if prices are taken to be fixed, and in transition (particularly if prices overshoot) to a new equilibrium it may even be that individuals in the target group lose out. Such possibilities can be considered in the model presented.³⁴

Another possibility that I intend to investigate is to look at the effect of population growth (or more precisely, growth in the number of households in the economy) on house prices. In this instance the population growth could be considered as an

be captured with houses and financial assets that are fungible as a part of bequests. The model would also need to balance the wealth of those who die and the initial wealth of the generation receiving the bequests.

³³Or, looked at another way: it may be hard to separate a credit shock from an increase in borrowing because (expected) incomes have gone up; and, it may be hard to separate an income shock from a second-round effect on incomes when borrowing activity pushes up economic activity. There is a hotly debated literature concerned with identifying credit conditions and their impact (see Fernandez-Corugedo and Muellbauer (2006), Aron et al. (2007), and Besley et al. (2008)).

³⁴Indeed an initial analysis of a cut in the fixed cost (transactions tax) on flat purchases does suggest some overshooting in prices, though a welfare analysis has not been completed at the time of writing.

ticipated, thus avoiding the issue of unanticipated shocks. The apparatus used to compute a transition path for prices could then be used to compute a path of prices as the population grows. The aim of the analysis would be to use the model to assess how much of the (real) growth in house prices in the recent past can be attributed to population growth. If the model can deliver a plausible assessment of this, then it could be used to assess how expected future population growth might influence house prices, under various assumptions about the growth of the housing stock. Since the number of households in the U.K. is expected to grow considerably in the first half of the current century, this issue is of considerable importance for social welfare, and for policy makers.

Chapter 5

Tax reform and retirement

saving incentives:

Evidence from the ‘Stakeholder

Pension’ reforms in the U.K.

5.1 Introduction

Given demographic pressures on social security finances in most OECD countries, a common policy target is to encourage households to increase their own private retirement saving. In the United Kingdom (UK), many households already rely on private sources rather than social security for much of their retirement income (Banks et al., 2005) but, as in other countries, there has been well-publicised concern as to the extent of a ‘savings gap’ between how much working age individuals should save for retirement and what they are actually saving.¹ There is, however, little agreement

¹For the UK, Oliver, Wyman & Company (2001), Pensions Commission (2004, 2005) and Department for Work and Pensions (2002, 2006b) expressed concerns over the ‘adequacy’ of retirement saving. Banks et al. (2005) are more sanguine. Much the same debate has occurred over a long period in the United States: see, on the one hand, Bernheim (1992) and, on the other, Engen et al. (1999) and

in the evaluation literature as to what saving policies are effective. There has been a substantial debate around this question in the United States (as in, *inter alia*, Poterba (1994); *Journal of Economic Perspectives* (1996); Bernheim and Scholz (1993); Benjamin (2003); and, Chernozhukov and Hansen (2004)), but relatively little econometric evidence relating to this question for the UK despite the plethora of recent reforms affecting pensions and other savings instruments there.²

Greater financial incentives to encourage retirement saving are an obvious policy instrument. But it is difficult to target incentives on the marginal saver, so that more generous incentives may actually reduce private retirement saving for the intra-marginal saver through a wealth effect. The cost to the exchequer of providing incentives will also mitigate the impact of incentives on aggregate (i.e. public plus private) saving. Several changes to the UK's tax regime governing retirement saving have been implemented in the last two decades but, from an evaluation perspective, it is hard to disentangle the effects of these tax regime changes from other reforms taking place at the same time.

In this chapter, we consider a reform that embodied a differential change in tax incentives: the introduction of Stakeholder Pensions in the UK in April 2001. This reform, which was intended to encourage overall take-up of, and contributions to, private pensions, contained several provisions that are discussed briefly in the next section. An important feature of the change in the tax regime that accompanied the reform was that it only affected a sub-set of the population. We can therefore use a standard difference-in-differences evaluation technique to examine the impact of this component of the reform by comparing the behaviour over time of those who were potentially affected by the tax reform relative to those who were unaffected. The paper closest in spirit to the present one is Milligan (2003), which focuses on differential limit changes and Canadian retirement saving.

Scholz et al. (2006). The same issue has also arisen in countries such as Australia and New Zealand.

²See Disney et al. (2001), Attanasio and Rohwedder (2003) and Attanasio et al. (2005a) for some findings on the impact of earlier pension reforms on saving in the UK.

The format of the chapter is as follows. In Section 5.2, we briefly describe the overall Stakeholder Pension reform but focus in particular on the change in the tax regime for private pensions that occurred simultaneously with the introduction of Stakeholder Pensions. Section 5.3 describes the data and the estimator used to examine the impact of the change in the tax regime on take-up of private pensions. Section 5.4 discusses the empirical results of the effect of the reform on pension coverage and retirement saving. Section 5.5 provides a brief conclusion.

5.2 Stakeholder Pensions

Stakeholder Pensions were introduced by the UK government from April 2001 as a new tax-subsidised retirement saving instrument. Like existing personal pensions, and some occupational pension schemes,³ Stakeholder Pensions are ‘defined contribution’ schemes (i.e. an individual’s pension benefits depend on the accumulated value of her fund). Individual contributors to Stakeholder Pensions could choose to opt-out of part of the social security programme, and also to obtain tax relief on contributions to their Stakeholder Pension, up to an annual ceiling described in more detail shortly. Stakeholder Pensions differed from pre-existing Personal Pensions (introduced for employees in 1988) in having compulsory minimum standards, a different governance structure, guaranteed workplace access for those working for medium-scale or large employers, and a simpler and more uniform administrative cost structure. Further details on these aspects of the Stakeholder Pensions pension reform are given in Disney et al. (2007).

The government’s Green Paper (Department of Social Security, 1998) which proposed the introduction of Stakeholder Pensions identified middle income earners - defined as those earning between £9,000 and £18,500 per annum in 1998 prices - as a target group for the reform, although Stakeholder Pensions are available to everyone. High income earners, it was assumed, already had access to other retirement saving

³‘Occupational pensions’ are, in UK parlance, employer-provided pension plans.

instruments, and lower-earners were generally seen by the Green Paper to be better off accumulating rights in the public second tier pension (the State Earnings-Related Pension Scheme, SERPS, superseded in April 2002 by the more redistributive State Second Pension, S2P), rather than opting for a private pension arrangement.⁴ Furthermore other pension reforms introduced since 2001, most notably the introduction of the means-tested Pension Credit in October 2003, also make public provision more generous for people with low lifetime incomes, although there is a group of lower-earners - those with rich spouses - who would be less likely to gain from the introduction of the Pension Credit.⁵

The Green Paper also proposed a number of other changes to the pension regime, including a reform to the structure of tax reliefs that was also implemented in April 2001. This reform forms the main focus of the present paper. The broad features of the tax regime, before and after the advent of Stakeholder Pensions, are as follows. The UK has an individual-based income tax system. Tax privileged pension contributions, up to certain ceilings, can be made by the employer or the individual. Employer contributions receive up-front tax relief at the individual's marginal rate against income tax and National Insurance.⁶ Individual contributions receive 'relief' (deferral) at the higher-rate of income tax for higher-rate taxpayers and at the basic-rate of income tax for all other individuals, i.e. including those with individual incomes below the threshold at which income tax becomes payable. Returns are broadly tax-exempt and

⁴Under the United Kingdom's social security programme, individuals can choose to opt-out of the second tier of the public programme. In the case of a personal pension, the DWP pays part of the social security contribution (which is a proportion of earnings within a given range) into the opted-out pension account in return for the individual forgoing that part of second-tier pension benefit that would have accrued had they remained 'contracted-in'. The employee can then contribute further amounts to their account, accruing the tax reliefs described in the text. The individual is not required to opt out of the social security programme to open a defined contribution pension plan. For further details on opting-out incentives, see Chung et al. (2008).

⁵The need to target middle-earners had been queried at the time, since this group already had high rates of pension coverage (see Disney et al. (1999) and Table 5.2 in the text). For the 30% who were not covered, unstable incomes and less accessible savings made pension saving less attractive (Banks et al., 2002a).

⁶'National Insurance Contributions' (NICs) is the name given to social security contributions in the UK.

pensions are then subject to income tax at withdrawal, except for up to 25% of the fund that can be withdrawn tax-free.

Until April 2001, as depicted in Table 5.1, the ceiling on individual contributions to pension plans was proportional to earnings and more generous for older individuals. Individuals without earnings could not gain tax relief on pension contributions.

Table 5.1: Pre-2001 tax reliefs for defined contribution pension plans: Maximum contributions as a % of earnings by age

Age at start of tax year	Maximum contributions as % of earnings
35 or under	17.5%
36 to 45	20.0%
46 to 50	25.0%
51 to 55	30.0%
56 to 60	35.0%
61 to 74	40.0%

Notes: Contributions were subject to an overall earnings cap. In 2005-06, this was set at £105,600. Maximum contributions include contributions by both the employer and employee.

An important difference between the post-Stakeholder Pension tax regime (which applied to all personal pensions including Stakeholder Pensions) and the previous tax regime was that every individual, irrespective of any earnings, was able to make gross contributions of up to £3,600 a year to his or her private pension (which, for an individual receiving tax relief at the basic rate, would require a net contribution of £2,808). In the new regime, individuals were then allowed higher contributions in line with their earnings as in the previous regime in Table 5.1. Overall, the effect of this change was to raise contributions limits significantly for lower-earners, especially for younger age groups (since maximum contributions as a proportion of earnings are lower). Figure 5.1 depicts the effect of the change post-2001 on the maximum gross contribution limits by gross relevant earnings for the various age groups in Table 5.1. Note that individuals with no earnings could also contribute up to the £3,600 maximum. It is worth re-emphasising that the UK's tax system is individual-based so

that each individual in a couple could contribute up to this maximum.⁷

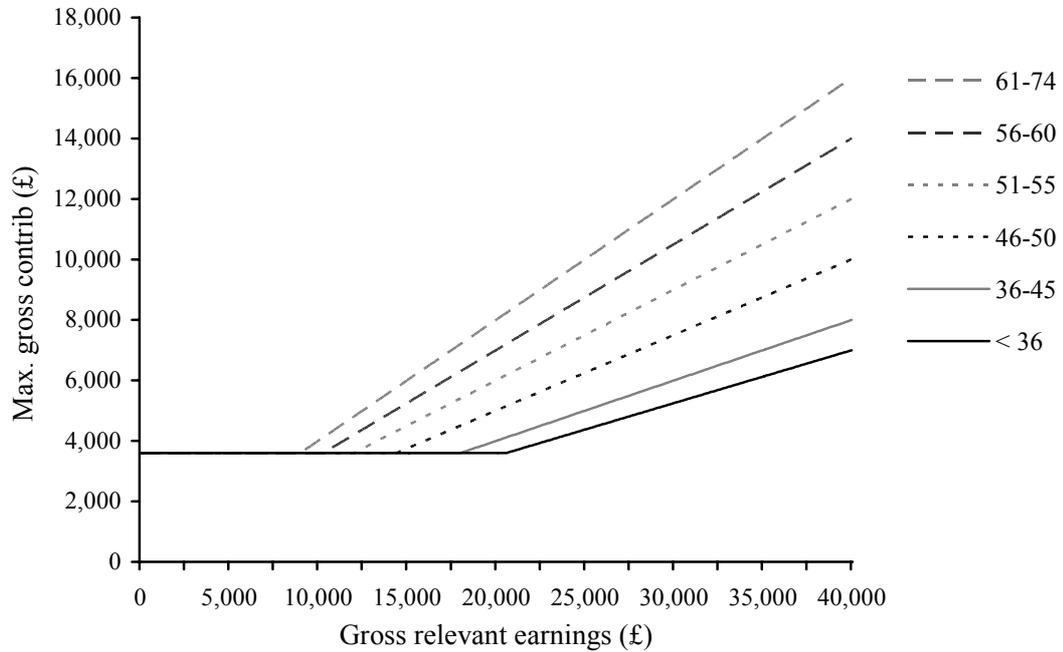


Figure 5.1: Maximum annual gross contribution limit, by annual gross relevant earnings and age, personal pension tax regime from 2001 to 2006

Although this change in tax regime was highlighted less than the ‘targeting’ of middle earners in the 1998 Green Paper (Department of Social Security, 1998) and subsequent discussion of the legislation, an implication of it was described in the Green Paper:

“The changes will also make it easier for partners to contribute to each other’s pensions, again within the overall contribution limits, should they choose to do so.” (p.63)

⁷Clark and Emmerson (2003) discuss other features of the tax treatment of Stakeholder Pensions. An even more sweeping reform to the ceilings on pension contributions was introduced in April 2006. Under these provisions designed to unify the tax regime for all types of private pensions - whether of the defined benefit or defined contribution form - there is an annual limit on contributions of 100% of earnings up to a ceiling of £215,000 (with the floor of £3,600 remaining) and a new lifetime limit on the value of the pension fund of £1.5m, rising over time. A further change which affected all potential contributors, was the abolition of the ‘carry-over’ provision for unused limits. Milligan (2003) examines the implications for intertemporal behaviour of a similar reform in Canada.

More precisely, Emmerson and Tanner (1999) noted that:

“The proposals may be of most benefit to high earners with non-working spouses who have already used up their own tax-free contribution limits or who want to maximise the value of their joint personal allowances in retirement.” (p.12)

In general, the consensus in the pensions industry and among analysts has been that the Stakeholder Pension reform was a failure, with little impact on either the take-up of private pensions or the real value of retirement saving. This apparent failure has contributed to the decision to enact yet another change in the UK’s retirement saving regime, with legislation in process to enrol employees whose employers do not offer an occupational pension arrangement automatically in a new system of Personal Accounts, unless the employee explicitly chooses otherwise. The rationale for this change towards greater prescription of retirement saving has rested in part on the presumption that the Stakeholder Pension episode illustrated the limited effectiveness of retirement saving policies that relied wholly on individual responses to perceived incentives.

Nevertheless, the introduction of Stakeholder Pensions provides a policy ‘experiment’ with, in effect, both a visible targeting of a new retirement saving instrument on a specified group and a change in the tax regime for pensions for another group; the latter reform was less publicised at the time but for those affected had a direct impact on financial incentives to save in a pension. We briefly show in the next section (and in more detail in Chung et al. (2008)) that the former aspect (i.e. the headline ‘targeting’ of the reform on middle earners) had no effect on take-up of private pensions. This result accords with the general view on the limitations of the reform.

We then focus the rest of the paper on the impact of the change in contribution limits on pension take-up among the affected group. In this latter respect, our line of argument follows closely the literature in North America that has attempted to identify behaviour in relation to retirement saving off differential changes in contribution

limits across sub-groups of the population (as in Venti and Wise (1987); Gale and Scholz (1994); Milligan (2003)). In analysing this change in tax reliefs, we show that the reform did indeed induce changes in behaviour, with significant increases in participation rates among sub-groups of the population who benefited from the increase in the ceiling on contributions eligible for tax relief. The consensus that the reform had little or no impact, and the inferences drawn from that conclusion, may thereby have been premature.

5.3 Empirical Strategy

As mentioned in the introduction, our empirical strategy is to use a standard difference-in-differences evaluation technique to examine the impact of the reform to pension contribution ceilings. In this section we describe the data and methods that are the basis of our estimations. The dichotomous nature of our outcome variable (the decision of whether or not to save in a pension) necessitates the somewhat involved estimation technique that we describe. Nonetheless, for interpretation of the results presented in section 5.4 it is important to keep in mind the underlying principle of difference-in-differences, which is to evaluate the reform by comparing the behaviour over time of those who were potentially affected by the tax reform to the behaviour of those who were unaffected.

5.3.1 Data sources and descriptive analysis

We investigate the determinants of the household decision to take-out a private pension using information from the Family Resources Survey (FRS). The FRS is a large-scale repeated cross section survey used to construct the UK's official statistics on income inequality and income poverty and so it elicits a rich set of information on each household's demographic characteristics, incomes (by detailed component) and other economic circumstances. The FRS asks individual respondents who are in work or

who have ever worked (below age 65) whether they or their employer contributes to a pension scheme. The pension arrangements that are explicitly delineated are a ‘personal’ pension, a company-run pension scheme, or a stakeholder pension. In addition respondents are asked whether the scheme is contributory or non-contributory, when they joined it and if it is ‘portable’, as well as more detailed questions about own contributions, contracted-out rebates paid into a Personal or Stakeholder Pensions (since individuals can have such schemes without making any additional contributions) and, in the case of a Stakeholder Pension, whether it was organised by the employer or the respondent.⁸

Table 5.2 Panel A provides data from the Family Resources Survey for the (tax) years 1999-2000 to 2002-03 on pension holdings by type for all employees. According to the table, overall coverage by private pensions has declined slightly over the period. Coverage by employer-provided plans has been constant, and a decline in coverage by Personal Pensions has been not quite offset by the introduction of Stakeholder Pensions and by a slight rise in the number of people with multiple plans.

⁸As a cross check, we examined responses from the General Household Survey (GHS), which asks somewhat different questions, primarily about coverage and membership, and also looks at aggregate data on pension scheme membership and contributions from Inland Revenue sources. The FRS provides detailed information on contributions, unlike the GHS. Both household surveys give significantly lower numbers for pension coverage and (more significantly, in the case of the FRS) for contributions than aggregate data from the Inland Revenue, perhaps reflecting under sampling in household surveys of contributors who make large contributions (i.e. the rich) and of other groups who may be contributing but are not asked about their contributions in the survey. However, it can be noted that aggregate data on total pension saving has been heavily revised downwards in recent years (although this applies more to data reported by the Office of National Statistics).

Table 5.2: Pension coverage by type of pension and earnings band: 1999–2000 to 2002–03

<i>Panel A: Employees Only</i>					
Year	1999-00	2000-01	2001-02	2002-03	Increase
Type of Pension	%	%	%	%	% point
Personal Pension	11.9	10.8	10.3	8.9	-3.0
Occupational Pension	0.0	0.0	0.8	1.3	+1.3
Stakeholder Pension	50.3	50.3	50.3	50.2	+0.4
Combination	2.0	2.2	2.3	2.5	+0.4
Agg. Coverage (%)	64.3	63.3	63.6	62.9	-1.4
Sample size	19,549	18,711	20,418	21,648	80,326

<i>Panel B: All aged 22 to state pension age</i>					
Year	1999-00	2000-01	2001-02	2002-03	Increase
by earnings band	%	%	%	%	% point
Zero	3.8	4.0	3.9	3.9	+0.1
Low	33.7	34.0	35.4	34.2	+0.5
Medium	61.9	60.8	61.2	60.6	-1.3
High	84.5	84.5	83.3	82.2	-2.3
Agg. Coverage (%)	47.2	46.8	47.4	46.9	-0.2
Sample size	27,259	25,887	28,026	29,657	110,829

<i>Panel C: All aged 22 to state pension age</i>					
Year	1999-00	2000-01	2001-02	2002-03	Increase
by limit increase	%	%	%	%	% point
Zero earnings	3.8	4.0	3.9	3.9	+0.1
Limit increase	46.8	46.1	46.9	46.5	-0.2
No limit increase	81.3	80.9	80.3	79.2	-2.1
Agg. Coverage (%)	47.2	46.8	47.4	46.9	-0.2
Sample size	27,259	25,887	28,026	29,657	110,829

Notes: The sample includes individuals aged 22 and over up to the state pension age, although a few individuals have to be excluded due to missing data. The sample in Panel B is that used for the regressions reported in later sections. Rounding explains why figures in the right-hand column may be slightly different from the difference between the 1999-2000 and 2002-03 columns.

Source: own calculations, Family Resources Survey 1999-2000 to 2002-03.

Panel B reveals the striking finding that coverage has fallen among the high and medium earnings groups over the period (these are the bands delineated by the Green Paper, of £18,500+ and £9,000 to £18,500 respectively) for those aged over 21 and

below state pension age (the sampling frame we subsequently use).⁹ Coverage has risen among lower-earners and even (marginally) among those reporting no earnings who are below state pension age. Finally, Panel C splits the data by the fraction of the sample that might potentially be affected by the new limit increase, as illustrated in Figure 5.1. While pension coverage was relatively stable among those who did, or would have, received an increase in their pension contribution limit, it was falling among those with higher earnings.

At first sight, some of findings in Table 5.2 are paradoxical given the stated aims of the policy. The data suggest that the introduction of Stakeholder Pensions has had no effect on overall coverage and indeed coverage by any kind of pension has fallen among the initial ‘target’ group of middle earners. Nor can these declines be explained by a decline in employer-provided occupational pension provision, since this remains constant. Finally, despite the Green Paper suggesting that lower-earners might be better off in the second pillar of the social security programme rather than opting for a personal pension, this is the only group to see a noticeable increase in private pension coverage.

In Chung et al. (2008) we formally test whether the reform affected coverage among the ‘target group’ of middle earners using a variant of the methodology outlined in the next section. We show that the only earnings group for which there was a significant positive effect on pension coverage were the group delineated as ‘low’ earners by the Green Paper, as suggested in Table 5.2, Panel B. We calculate that pension coverage amongst this group increased by 3.6 percentage points as a result of the reform. These results are therefore a more precise estimate of the differential trends that can be observed in Table 5.2.

The discussion in Section 5.2 combined with these results lead us to think that it

⁹We gross up weekly earnings data to provide these annual earnings bands. This inevitably produces measurement error - for example some people will wrongly be attributed ‘zero’ earnings for the year based on current zero earnings. In addition, the Green Paper sometimes refers to ‘£20,000’ and sometimes to ‘£18,500’ as the highest income of ‘middle earners’. In general we work with the latter definition in the FRS data, revalued over time in line with average earnings growth in the sample.

might be the change in the contribution limits, rather than the targeting of particular earnings bands, that had an impact on pension coverage. This possibility was highlighted in the two quotations in section 5.2 which pointed to the intra-household incentives to contribute to spouses' pensions. Accordingly, we examine the effect of the change in the contribution limits in the remainder of this paper. This is done by comparing the pension coverage over time of those who were potentially affected by the tax reform to coverage over time amongst those who were unaffected; the data underlying our analyses are described in Panel C of Table 5.2.

5.3.2 Modelling the pension take-up decision

Our test is whether the change in contribution limits enacted as part of the Stakeholder Pension reform in April 2001 affected pension saving decisions within households. This reform can be considered as the policy 'treatment'. Our identification strategy is straightforward, insofar as we can broadly identify the individuals (i.e. those with earnings below the limits at the time of the reform) who were, and who were not, affected by the tax change. Hence, our basic method of analysis is a difference-in-differences approach. However, our implementation strategy requires further explanation.

To formalise the test, write a general model of retirement saving in which Y_{it}^* is the outcome of the retirement saving decision of individual i at time t , related to a set of individual and household characteristics (X_{it}), an appropriate measure of earnings (Z_{it}), and, to capture trends over time flexibly, to a vector of time dummies (d_t):

$$Y_{it}^* = \beta' X_{it} + \gamma Z_{it} + \tau' d_t + \epsilon_{it} \quad (5.1)$$

Our outcome variable is a dichotomous indicator of whether or not an individual saves in a private pension at a particular point in time. Denoting this dichotomous outcome as Y_{it} , then we think of the (continuous) Y_{it}^* as a latent variable that measures whether or not an individual gains positive utility from saving in a private pension.

With a normally distributed error term, this set-up can be analysed using a ‘probit’ model.¹⁰

The hypothesis tested here is that the probability of purchasing a private pension changed differently for those who were potentially affected by the contribution limit increase compared to those who were not. The required counterfactual assumption is that in the absence of the reform the purchase probabilities for those who were and were not affected by the policy change would have followed a common trend. To implement the difference-in-differences exercise, we need to define a ‘treatment’ variable and a ‘post-reform’ variable. The latter is defined as the indicator I_t that measures whether the individual is observed after the beginning of April 2001. The ‘treatment’ variable is another 0-1 indicator, L_{it} , that measures whether, given an individual’s age and annual earnings, the post-April 2001 contribution limit rules would have been more generous to him/her than the pre-April 2001 contribution limit rules.¹¹ If we could estimate the linear relationship (5.1), then the difference-in-differences model would be estimated by equation (5.2) in which the extent of any difference-in-differences would be measured by the coefficient α on the interaction between ‘had a limit increase’ (L_{it}) and the indicator (I_t) for the post April-2001 period:¹²

$$Y_{it}^* = \beta' X_{it} + \gamma Z_{it} + \tau' d_t + \phi L_{it} + \alpha L_{it} I_t + \varepsilon_{it} \quad (5.2)$$

However, in a non-linear model such as the probit (used here), calculated ‘marginal effects’ on interaction terms cannot give a difference-in-differences measure analogous

¹⁰It is important that our modelling strategy allows for the discrete nature of our outcome variable: a linear probability specification may very likely lead to the prediction that those with zero or very low earnings have a negative ‘probability’ of saving in a pension.

¹¹For example, for an individual aged 35 or younger, this variable takes the value 1 if gross earnings are less than £20,571. This number is derived from the fact that before April 2001 individuals in this age range could contribute no more than 17.5% of their earnings to a pension, but after April 2001 this limit became the maximum of 17.5% of earnings or £3,600. £3,600 is (to the nearest pound) 17.5% of £20,571. Similar values are constructed for individuals in other age bands.

¹²The variable I_t is not entered independently in this regression since the time dummies capture this variation.

to the coefficients from a linear model. With the discrete outcome set-up, the common trends assumption may not hold for the expectations of Y_{it} (the saving probabilities), since even when affected by the same factors different baseline probabilities are likely to move at different rates simply because they are different distances from the bounds on probabilities of zero and one. Rather, common trends may hold for a transformation of the distribution of the outcome variable and specifically for the inverse probability function, which is assumed to be known and for the probit is $\Phi^{-1}(\cdot)$.¹³ In other words, the assumption of common trends is made for the index rather than for the probability itself. Following Blundell et al. (2004a) this can be written formally as that in the absence of any ‘treatment’ the following would hold¹⁴:

$$\begin{aligned} & \Phi^{-1}[E(Y_{it}|X_{it}; L_{it} = 1, I_t = 1)] - \Phi^{-1}[E(Y_{it}|X_{it}; L_{it} = 1, I_t = 0)] = \\ & = \Phi^{-1}[E(Y_{it}|X_{it}; L_{it} = 0, I_t = 1)] - \Phi^{-1}[E(Y_{it}|X_{it}; L_{it} = 0, I_t = 0)] \end{aligned} \quad (5.3)$$

where variables are defined as above.

The right hand side of this equality can be estimated from observations of the ‘control group’ (those not affected by the limit increase) before and after April 2001. Using the common trends assumption as now formulated, this information can in turn be used to construct a counterfactual of how the index would have evolved for each treatment group individual had the change in pension contribution limits not occurred.

¹³In addition to the points raised in the text about the common trends assumption, it is also the case that the ‘marginal effects’ on interaction terms in non-linear models that are automatically generated by software packages (in our case by STATA version 9.2) often do not give a true measure of ‘interaction effects’. For more details see Ai and Norton (2003).

¹⁴Time dummies are no longer included since separate probits are run for those observed pre-reform and those observed post-reform. Similarly analysis is done separately by whether or not the individuals would have received an increase in their private pension contribution limit and since this depends on earnings these are also excluded from this specification. Partners earnings, where relevant, is included.

The impact of the policy can then be evaluated as¹⁵:

$$\begin{aligned}
 I(X) = & E(Y_{it}|X_{it}, Z_{it}, d_t; L_{it} = 1, I_t = 1) - \\
 & -\Phi\langle\Phi^{-1}[E(Y_{it}|X_{it}, Z_{it}, d_t; L_{it} = 1, I_t = 0)] + \Phi^{-1}[E(Y_{it}|X_{it}, Z_{it}, d_t; L_{it} = 0, I_t = 1)] - \\
 & -\Phi^{-1}[E(Y_{it}|X_{it}, Z_{it}, d_t; L_{it} = 0, I_t = 0)]\rangle
 \end{aligned}
 \tag{5.4}$$

Blundell et al (2004) propose a method for implementing this 'difference-of-differences' estimator of the effect of the policy. A different relationship between the outcome and the observed characteristics is estimated for each group of agents defined according to the various interactions of whether or not the reform would have increased their contribution limit (which depends on their age and earnings) and whether they were observed before or after the reform was implemented. These relationships encapsulate the behavioural patterns of each group and the impact of the reform once it had been enacted. By predicting the outcomes for the 'treated, after' group (i.e. individuals characterised by $L_{it} = 1$ and $I_t = 1$ who were observed after 2001 and whose age and earnings were such that they were affected by the contribution limit increase) using the behavioural equations for the pre- and post- reform 'control' groups, one obtains an estimate of how the underlying index would have changed for individuals in the treated group in the absence of the reform. This can be used in combination with the behavioural equations for the treated group to construct the estimated effect (5.4). The final estimate of the effect uses predictions made for the 'treated, after' group and weighted according to characteristics (X_{it}) in this group. It can therefore be thought of as representing the average impact of treatment on the treated.

¹⁵Despite the similarity to the linear case, the nonlinear assumption exploited here entails two additional restrictions on the nature of the error terms: only group effects are allowed for and the groups being compared are assumed to have the same residual variance. See Blundell et al. (2004a), p.580.

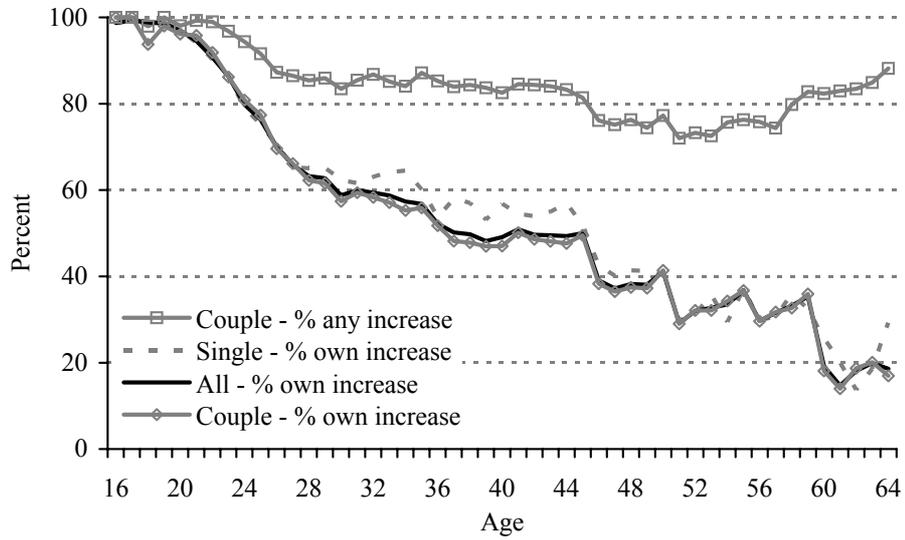
5.3.3 Implementing the estimator: data considerations

Before examining the results of the difference-in-differences exercise, we return to a caveat mentioned in footnote 10. Weekly earnings from the FRS are grossed up to obtain annual earnings, so there is measurement error in the indicator of whether the individual is affected by the change in contribution limits. This error would be particularly pertinent if we had sought to measure the impact of the reform solely by identifying individuals who were on their maximum contribution limits pre-April 2001 and who then benefited from the increase in limits in April 2001, since these individuals could only be identified with error. In any event, such a limited definition of the 'treated' group could lead to an underestimate of the overall effect of the policy since it is not true that all of those who might respond to the policy would be observed exactly on their contribution limit in the pre-reform regime. Some lower-earners might, for example, have been induced to take up a private pension rather than doing no pension saving because the large increase in their contribution limit meant that the benefits of making contributions could now outweigh the threshold fixed costs of purchasing a pension.

In contrast to some studies of limit increases in the United States, therefore, we do not see potentially 'treated' individuals simply as those who individuals who we calculate to be at the pre-reform contribution. Instead, we assume that all individuals or households who received a limit increase as a result of the reform, irrespective of whether their calculated contributions would have been at the pre-reform ceiling, are potentially 'treated'. From this more general viewpoint, the change in the limits had a large potential coverage. This is illustrated in Figure 5.2, which shows the fraction of respondents treated by age given our more general measure.

Figure 5.2 should be read in conjunction with Table 5.1 and Figure 5.1, which illustrate that the impact of the limit increase is driven by both age and earnings levels. At the ages of 36, 46, 51, 56 and 61 there are step changes in the likelihood of being affected by the limit increase because of the higher earnings proportions that

Panel A: Positive earners only



Panel B: All individuals

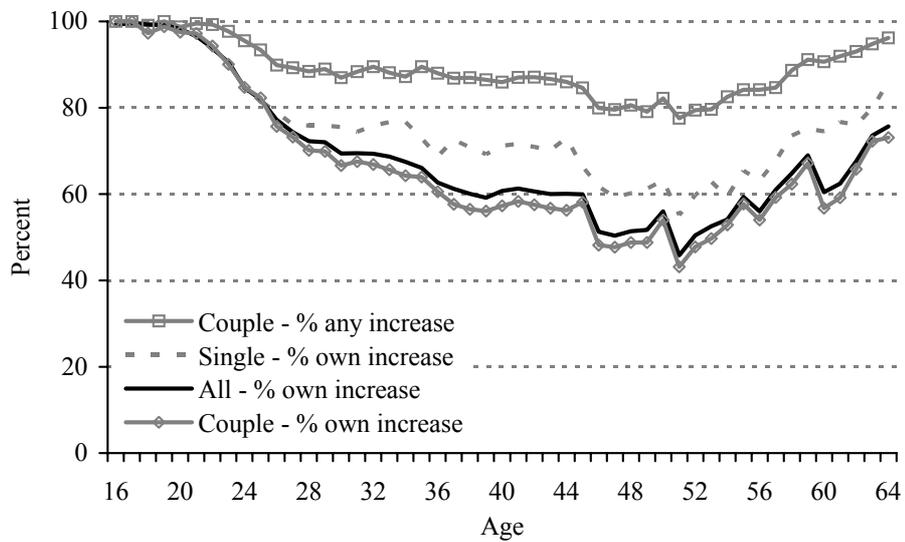


Figure 5.2: Percentage of individuals receiving limit increase by age and household type

Notes: As for Table 5.2.

Source: As for Table 5.2.

can be contributed at each of those ages. These steps are observable in the figures, especially at the three intermediate age ranges.

Eligibility for the limit increase is otherwise driven by two factors: the increase in earnings with age, and the proportion of households with zero earners in each age band. Panel A, which focuses on individuals with positive earnings, shows how rising earnings over the life cycle and the age dependent changes (described in Table 5.1 and Figure 5.1), in the proportion of earnings that can be contributed tax-relieved to a pension, combine to reduce the likelihood of being affected by the limit increase when only own earnings are considered. This also implies that, when we consider only positive earners, the 'control' group are on average older than the 'treated' group; an issue to which we return later. Panel A also shows that almost 100% of individuals aged 21 or under are affected by the limit increase, and so we exclude them entirely from the regression estimates on coverage.

Figure 5.2 Panel B, which incorporates zero earners within households, gives a somewhat different picture of the impact of the changes in limits by age. Essentially because of cohort differences in participation rates, a greater fraction of older households have at least one non-earner who would be eligible for the limit increase under the reform regime even if the primary earner is above the limit threshold.¹⁶ Thus the (negative) association of the probability of being treated by age is much less marked.

¹⁶This raises another small measurement issue that applies to those with no earnings who, pre-2001, should not have been contributing to a private pension and receiving tax relief - see the regime described in Table 1. Table 2 nevertheless suggests that we observe a few individuals who are contributing pre-2001, which arises (we surmise) because they had some earnings during the year even though we observe no current earnings at the time the individual was surveyed for the FRS. Under certain assumptions (notably concerning the volatility of earnings), this measured proportion of take-up among those with no earnings pre-2001 may be taken as an upper bound on the measurement error involved in grossing up weekly earnings to obtain annual earnings, both before and after the reform.

5.4 Did the reform increase private pension coverage?

5.4.1 Empirical Results

We estimate the difference-in-differences model described by equation (3). This involves estimating separate probits for the treated and controls, pre- and post-treatment date (2001), using the Family Resources Survey for the four years 1999-2000 and 2000-01 ('pre-treatment') and 2001-02 and 2002-03 ('post-treatment'), in order to calculate the estimated treatment effect, as described in the previous section.¹⁷ The regressors comprise whether the individual is single or in a couple, age dummies of both the individual and, where relevant, their partner (structured such that the bands coincide with the age ranges for the contribution ceiling bands described in Table 5.1), sex, age left school, partner's education, dummies for each period within the pre- and post-treatment regime, earnings and a full set of age-education interactions. We consider those aged over 21 but under the State Pension Age (65 for men and 60 for women). Rather than present a full set of probit estimates, we provide key calculated treatment effects in Table 5.3 along with statistical significance (which is estimated using bootstrapping) and sample sizes.

The results in Table 5.3 confirm that there is a positive effect of the change in the contribution limits on take-up of private pensions among the treated group. The table further analyses the sub-groups which are disproportionately affected.

The first row of Table 5.3 suggests that the overall result of the reform is an increase in private pension coverage of 2.1 percentage points among those affected by the limit increase. This result averages across positive and zero earners, and indirectly explains the result in Chung et al. (2008), insofar as this effect is concentrated among affected positive earners (+3.3 percentage points, statistically significant) rather than zero earners (+0.4 percentage points, not statistically significant). We did not find any evidence that the reform had a statistically significant impact on take-up among

¹⁷Data limitations rule out using earlier years before 1999-2000 in calculating the 'pre-treatment' probits.

sub-groups of zero earning respondents (e.g. by sex or whether or not in a couple) and so do not consider single zero earners further.

Among single respondents with positive earners, the total effect on private pension coverage is in an increase of 3.9 percentage points, which is statistically significant at the 5% confidence level. Strikingly, however, the effect among single women is large and statistically significant (+5.4 percentage points) whereas among single men, the effect is again positive but not statistically significant. Among couples, the overall effect is also positive and statistically significant (+3.0 percentage points); again, the effect is stronger (and statistically significant) in couples where the woman is the 'treated' member (i.e. below the contribution limit), whereas the impact on treated men is much smaller.

Finally, we disaggregate couples further by the earnings of the partner. In one specification, we use current earnings of the partner. However since current earnings are subject to transitory shocks we also use partner's education as a proxy for lifetime earnings, where more years of schooling are assumed to be associated with greater earnings and/or participation. A necessary caveat to this interpretation is that average years of schooling have increased cohort-by-cohort. The 'treated' group of positive earners is relatively young on average (see Figure 5.2.A) and is therefore more likely to have greater schooling than the 'control' group, *ceteris paribus*.

Table 5.3: Results: Impact of reform of contribution limits on take-up of private pensions for selected groups, using Blundell et al. (2004) procedure

<i>By zero or positive earnings: respondent</i>						
	Zero earners only		Positive earners only		All	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
All	3.9%	+0.4ppt	46.7%	+3.3ppt	28.5%	+2.1ppt

<i>Positive earnings only (respondent)</i>						
	Men		Women		All	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
Singles	39.2%	+3.1ppt	40.7%	+5.4ppt	40.1%	+3.9ppt
Couples	47.4%	+0.3ppt	50.3%	+4.3ppt	49.5%	+3.0ppt

<i>Couples, positive earnings only, by partner's earnings/partner's education</i>						
	Zero earner		Low-mid earner		High earner	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
Couples	41.3%	+5.7ppt	49.5%	+2.4ppt	55.4%	+2.9ppt

	Low education		Medium education		High education	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
Couples	47.5%	+3.7ppt	54.4%	+7.5ppt	51.6%	-2.2ppt

Notes: Statistical significance estimated using bootstrapping with 1,000 repetitions. Significant treatment effects at 5% level in **bold**. Controls for age, sex, school leaving age, and full age-schooling interactions, whether single or in a couple, and, where relevant, partner's age, partner's school leaving age and partner's earnings. 'Low-mid' earner is defined as in 1998 Green Paper 'low-middle' earner. Education: 'low' = leaving school age 15/16; 'medium' = age 18, 'high' ≥ 18 (tertiary). For sample sizes see Appendix Table A5.1.

Using partner's earnings and education, we note that the largest and sole statistically significant effect by partner's earnings group is on positive earners with a zero earning partner. We do not find clear-cut evidence that lower-earners with high earning partners had a particularly strong response to the reform (inasmuch as the coefficient for this group is positive but not statistically significant at the 5% level);¹⁸ that is, we cannot find conclusive evidence for the suggestion by Department of Social Security (1998) and Emmerson and Tanner (1999) - see section 5.2 above - that the reform allowed couples jointly to utilise the increased limits by having pension coverage for the lower-earning partner. When we use education as a proxy for long-term earnings of the partner, we get strong but slightly hard to interpret results. Treated respondents whose partners are educated up to 18 exhibit the largest positive responses to the reform (+7.5 percentage points) while a smaller estimated impact is found among those whose partner left school at (or before) the school leaving age. The impact on those partners have the highest education levels is actually negative, albeit not statistically significant.

Overall, the results show statistically significant effects, and for some groups the proportionate effects on pension coverage are substantial. In particular, we observe large effects on pension coverage for lower-earning women (both single and in couples). These results contrast with the conventional wisdom that the overall Stakeholder Pension reform had little or no impact on pension coverage. They also suggest that looking at aggregate trends and at 'target' groups such as middle earners has led analysts and policy-makers to overlook significant effects elsewhere in the earnings distribution. The lack of effects on middle earners likely arose because the effective policy 'treatment' was the reform of tax relief, which affected lower-earners. The analysis of the next subsection is designed to increase our confidence that the observed effects really were due to this policy reform.

¹⁸With the bootstrapped 5% confidence interval, we can only confirm that the effect lies between +1 and +7 percentage points.

5.4.2 Alternative explanations and sensitivity analysis

This section considers alternative explanations for our results. In particular, it considers the plausibility in this context of the 'common trends' assumption that lies at the heart of the 'difference-in-differences' model, and undertakes some sensitivity analysis which focuses on possible implications of pre and post-treatment heterogeneity of the population.

The validity of the common trends assumption

Our results assume that the trend in pension coverage (or more accurately in the index of the underlying propensity to contribute to a private pension) among those who received an increase in the contribution limit as a result of the reform would, in the absence of the reform, have been the same as among those who were unaffected by the reform to contribution limits. We are overstating (understating) the effect of the tax relief change to the extent that the positive change in coverage among the below-limit group would have been higher (lower) than that of the above-limit group in the absence of the 2001 reform. How plausible is this common trends assumption? Undoubtedly, the overall trend in pension coverage was affected by the change in the financial climate from the beginning of 1998 to the end of 2002, during which the FTSE 100 index of UK equities fell by 31% whereas, on average, UK house prices rose by almost exactly 50% (using the Halifax plc index). This might have induced savers to switch away from saving through private pensions (which were at this stage largely in equity-dominated portfolios) to invest in housing. So if above-limit households either exhibited greater substitutability in their asset portfolios, for example through economies of scale or greater financial acumen, or had different asset portfolios (i.e. more equity-dominated) this might explain the disparate trends.

There is mixed evidence for this alternative 'story' in our data. In comparisons of treated earners against controls (for example, of single people) it is certainly true that the 'typical' control is older and more likely to have greater housing equity, given

the fall in potential treatment rates with age illustrated in Figure 5.2. This would imply that treated households, with lower housing equity (and perhaps less financial acumen, although younger households tend to have greater formal education to offset this factor) might be less likely to reduce their pension saving in response to house price changes. On the other hand, it is hard to utilise this reasoning to explain away the differential results by the sex of the respondent and for couples by spouses' income and education.

If, instead, we focus on changes that might affect the 'treated' group rather than the control group, there is one reason for thinking that we might understate rather than overstate the magnitude of the effect. Here there is an important change in the benefit regime that coincides with the introduction of Stakeholder Pensions in 2001. In 2002, the State-Earnings Related Pension (SERPS), the second tier pension, was replaced by the State Second Pension (S2P); a change that had also been announced in the 1998 Green Paper. S2P is more explicitly redistributive towards low lifetime earners in its design. In addition the means-tested benefit for pensioners was formally indexed to earnings rather than prices from April 1999 (unlike the rest of the pension programme), so increasing its real value for those with low lifetime incomes and reinforcing the disincentive to save for retirement.

Analysis therefore suggests that, whereas replacement rates cohort-by-cohort for the public pension programme have already peaked for average earners, lower-earners are likely to see increasingly generous replacement rates from the public programme for several decades yet (Disney and Emmerson, 2005). To the extent that single women, or couples that include a zero earner or have low education (to take some of the groups where we have found a significant effect) are disproportionately likely to gain from these reforms to the public pension programme relative to the control group, we might expect the treated group's take-up of private pensions to have fallen faster than that of higher income groups, in direct contrast to the results found here.

Changing heterogeneity among the treated and control groups

Another possible alternative explanation for the observed 'treatment effect' lies in changes in the composition of households between the treated and untreated group. Household heterogeneity can arise from measured characteristics and from unobserved characteristics, such as individual preferences for saving. If it is assumed that such characteristics are fixed over time then our difference-in-differences estimator should be robust to their presence, although since our data have no panel dimension we cannot directly control for unobserved individual effects and it is for this reason that we are careful not to generalise our empirical measure of the 'treatment effect' to the untreated group. On the other hand, compositional effects of whatever kind could potentially bias our results if the composition of households within the treated group and the controls changed after the reform.

This bias would be most pertinent if observed earnings are endogenous to the reform. As demonstrated in Table 5.1, prior to 2001, an individual required some earnings to be eligible for tax relief on retirement saving; as Figure 5.1 shows, after 2001 there is an initial lump sum tax relief independent of earnings. It is in theory possible that some people with accumulated cash who wished to engage in retirement saving pre-2001 had to work, at least part-time, in order to obtain eligibility for tax relief; after 2001 this was no longer necessary. In this case observed earnings would be endogenous to the treatment. This possibility, although unlikely, might be pertinent for, say, older women who had low retirement savings and who were not regular workers. If the identification of the treatment is endogenous to the outcome variable, this clearly poses a potential problem for the analysis. Other possible scenarios, such as greater volatility of intra-household earnings over time, might also have similar effects on the composition of the treated relative to the controls.

We test this possible proposition using the FRS in a manner designed to test for common and differential trends in composition across earnings groups. We pool each 'treatment' group (low/zero earners below the contribution limit) with the 'control'

(those who would not have received an increase in their pension contribution limit) group. Now write the model:

$$J_{it} = \theta L_{i=1} + \gamma_1' X_{it} + \alpha_1' X_{it} L_{i=1} + v_{it} \quad (5.5)$$

Where J_{it} is an indicator variable of whether the observation of the individual occurred during the period in which Stakeholder Pensions were available, L is an indicator of whether or not the reform increased the pension contribution limit of the individual, X_i is a vector of explanatory variables such as age, schooling and partner's characteristics. This model tests two possibilities. First, the significance (or otherwise) of the vector of coefficients $[\gamma_1]$ tell us whether the characteristics of those observed among the pooled groups before the reform occurred are different from the characteristics of those observed after the reform was implemented. Second, the (lack of) significance of the vector of coefficients $[\alpha_1]$ tells us whether any changes in characteristics over time occurred differently between the control group and the treatment group. It is this second test which is important for our analysis - our results may not imply the policy effects we wish to identify if changes in the characteristics of individuals we observe are occurring differentially across the treatment and control groups. In fact the absence of differential changes in group composition is a more stringent test than is required to support the common trends assumption,¹⁹ and so passing the test would be reassuring for our analysis.

Equation (4) can be estimated first using data from the control group and the treated group who did have some earnings, then for the treated group who did not have any earnings and the (whole) control group. In both of these models the results (available on request) show that our sample is, on average, slightly older and slightly better educated in the period after stakeholder pensions were implemented. There is

¹⁹The common trends assumption could still be valid even if characteristics changed differentially over time across the control and treatment groups, but these specific characteristics did not affect the take-up of private pensions, or if differential changes in characteristics just happened to cancel out in their net impact on average take-up among the control and treatment groups.

little evidence of a differential change in characteristics between those in the control group and those in the treated group, especially among the treatment group with some earnings. For positive earners, the coefficients on the interaction terms [i.e. vector α_1] are not jointly different from zero at conventional levels of statistical significance ($\chi^2(18) = 25.08, Prob > \chi^2 = 0.13$). Looking at the individual interaction coefficients, the only significant coefficient is on male, which is negative, suggesting slightly fewer men in the control group relative to the treatment group after the reform. Since men are more likely to contribute to a private pension, this might slightly offset our treatment effect. On the other hand, since men are unlikely to be marginal workers, it provides no support for the argument that our treated group pre-treatment contained a preponderance of earners who were only working to obtain earnings on which to obtain tax relief for their pension contributions.²⁰

The composition of the control and treatment groups

Our final sensitivity test attempts to handle the implication of the fact that, for positive earners, our controls are on average older than our treated group (see Figure 5.2, Panel A). Although we interpret our measured effects as a ‘treatment of the treated’ rather than a consequence of differential limits across age groups, it is interesting to see whether the calculated effects are affected by choosing comparable controls facing a common rate regime. To do this, we focus only on the 36-45 age group. We choose this group because it is a relatively broad (10- as opposed to 5-year) age band that is faced with a single pension contribution ceiling of 20% of earnings in the pre-Stakeholder Pension regime and, conveniently, it contains around 50% of treated individuals and

²⁰For zero earners, we have evidence that fewer low educated people are likely to be observed after the reform. The coefficients on the interaction terms are jointly different from zero at conventional levels of statistical significance ($\chi^2(18) = 34.70, Prob > \chi^2 = 0.01$). Inspection suggests that our control group has a higher proportion of less educated people after the reform - that is, the overall fall in the fraction with low education was disproportionately concentrated among the treated group. Since people with more schooling are more likely to contribute to a pension, this facet tends to raise the measured treatment effect but in fact we observe no significant treatment effect when looking at the impact of the policy change on those without any earnings, as reported in Table 3.

50% controls (Figure 5.2, Panel A).

The significant results for this age group can be summarised as follows. First, as before, the results are only statistically significant for positive earners and, for single people, only among women. The latter coefficient is particularly large (+14.6 percentage points) and statistically significant. Among couples, the largest effects are found, as in Table 5.3, amongst those whose partner is a zero earner (+5.0 percentage points) and among those whose partner is ‘middle’ educated - that is, educated to age 18 (+8.0 percentage points). These results therefore lead us to believe that the magnitudes in Table 5.3 for the whole sample are broadly correct.

On balance, therefore, we do not believe that our key results in this section are driven by compositional changes or by the heterogeneity of our control and treatment groups.

Reflections on the results

Our analysis of the composition of our control and treatment groups, coupled with careful consideration of factors that could have disturbed the crucial ‘common trends’ assumption, reinforces our confidence that the results reported in subsection 5.4.1 can be attributed as causal effects of the stakeholder pension reform package. The effects that we found are for lower-earners and, as argued above, the lack of effects on middle earners likely arose because the effective ‘treatment’ in the policy reform package was the change in incentives arising from the reform of tax relief on lower-earners, rather than the highlighted targeting of middle earners.

Over all individuals affected by the change in contribution limits, the effect on pension coverage has more or less exactly offset what would otherwise have been a falling trend in pension coverage (see Table 5.2 above). For some groups of the population, the effects on pension coverage are substantial in proportional terms: we found effects of up to 5 - and in one case 7.5 - percentage points, for coverage in groups where around one-half of individuals have private pensions. An important pattern

is that we observe particularly strong increases in take-up of private pensions among lower-earning women.

To supplement these results on pension coverage, it would be nice to have results on how the reform has affected the level of pension saving. Unfortunately our data provides only limited information that could be used to explore this issue, as it only records information on individual (and not employer) contributions in to what are reported to be personal or stakeholder pensions. Using these data, in Disney et al. (2007) we tentatively concluded that an aggregate £160 million²¹ decline in reported pension contributions of this kind might have been 30% larger in the absence of the policy reform. This tentative figure supports the strong evidence on coverage rates presented in the current paper. The evidence shows that, contrary to the conventional wisdom, the stakeholder pension reform package did have important effects on the pension saving of some groups of the population.

²¹This amount is for contributions reported to the FRS, then grossed up using survey weights to give a national figure.

5.5 Conclusions

Our starting point was the policy debate concerning the best ways of encouraging people to save for their retirement. Stakeholder Pensions, introduced in 2001, were targeted by the government on middle earners as a means of filling a perceived gap in retirement saving products. The introduction of Stakeholder Pensions was also associated with a change in pension contribution limits which allowed lower-earners to make larger tax-relieved contributions to private pension schemes. Our analysis represents the first systematic attempt, to our knowledge, to examine the impact of these policy reforms on the probability of households engaging in retirement saving.

Aggregate data suggest that the introduction of Stakeholder Pensions had little impact on the overall propensity to save for retirement. The number covered by private pensions was static and there was a downward trend in pension saving over the period 1999 to 2002, either side of the reform in 2001. This apparent failure of the Stakeholder Pension reform was one factor that encouraged the government to introduce further pension reforms on 6th April 2006²² and to announce yet more reforms to the retirement saving regime on 22nd May 2006 (see Department for Work and Pensions (2006a,b)) with greater emphasis on a ‘default option’ that encourages individuals to invest a minimum amount in a private pension plan.

Our results suggest that analysts have been too quick in assuming that the 2001 reform had no effect, and also in assuming that individuals failed to respond to the changes in tax incentives. Exploiting a difference-in-differences estimator that allows for the dichotomous nature of the saving decision, we show that these aggregate trends conceal a more complex picture. In particular, our results show that a trend fall in coverage was partially counteracted by the introduction of Stakeholder Pensions, primarily through the associated change in contribution ceilings that disproportionately

²²Notably this reform again changed the limits on the amount of tax-relieved contributions that individuals can make in such a way that the vast majority of individuals will now be able to make tax-relieved contributions equivalent to their full year’s earnings each year.

benefited low and zero earners. We provide some evidence that women, both single and in couples, have benefited from the increase in the joint contribution limits within households, which was an additional intention of the policy. To put this in context, there was an underlying decline in private retirement saving in the early part of the decade (for reasons that we briefly discuss in the text) that would have been greater had it not been for the tax changes associated with the Stakeholder Pension reform in 2001.

In general, our results also suggest that individuals respond to tax incentives in making retirement saving decisions - a result incidentally confirming much of the US literature on the impact of contribution limits on saving in Individual Retirement Accounts (see again *Journal of Economic Perspectives* (1996) and the literature cited therein). The results also highlight that it is sometimes important to know the details of a given policy reform, rather than just the 'headline' target, in order to understand how the policy might work in practice. Since it is common for large impacts of policy reforms to be highlighted that turn out to be illusory on subsequent closer analysis, it is ironic that the introduction of Stakeholder Pensions in 2001 was quickly written off as having had little impact on retirement saving: our evidence suggests that some of the changes that were part of the Stakeholder Pension reform package had non-trivial effects on particular groups of the population and thus on aggregate retirement saving.

A5 Appendix to chapter 5

	Treatment		Control	
	Before	After	Before	After
All	33,987	36,619	19,159	21,062
Positive earners	19,101	21,002	19,159	21,062
Zero earners	14,886	15,617	19,159	21,062
Positive earners, singles				
All	5,353	6,154	3,956	4,303
Single women	3,184	3,748	1,638	1,980
Single men	2,169	2,406	2,318	2,323
Positive earners, couples				
All	13,748	14,848	15,203	16,759
Couples women	9,666	10,327	3,969	4,562
Couples men	4,082	4,521	11,234	12,197
Couples – zero earning partner	3,163	3,405	4,105	4,453
Couples – low/mid earning partner	6,085	6,712	7,250	7,916
Couples – high earning partner	4,500	4,731	3,848	4,390
Couples – low education partner	9,007	9,386	7,869	8,368
Couples – mid education partner	2,553	2,950	3,502	3,786
Couples – high education partner	2,156	2,473	3,777	4,540

Appendix Table A5.1: Sample sizes for analysis in Table 5.3

Chapter 6

Conclusion: Topics for Further Work

The conclusions of each chapter of this thesis have been discussed within the relevant chapters, and were also outlined in chapter 1 (the Introduction). Rather than repeat these points, this conclusion suggests some topics for further work that follow from analyses presented above.

One topic that has arisen throughout the work modelling housing demand in a lifecycle context, is the potential importance of bequest motives. If households choose not to downsize as they age because they have a strong preference to bequeath the family home (or indeed for some other reason), this could well influence how house price shocks influence consumption choices. It may also affect the evolution of house prices through an influence on the number of houses that are supplied to the market and how this responds to price signals. Properly understanding the operation of bequests in this context would not be simple, since incentive issues regarding the taxation of estates and insurance against the need for long-term health care in old age, are likely to interact with the preference to bequeath the house. However, as a better understanding of bequest motives and housing could significantly improve our understanding of several issues addressed in this thesis, further work here would be worthwhile. Subject to

concerns about selection, an empirical investigation comparing those with and without children could be a useful supplement to structural modelling work on this topic.

A particularly fruitful area for further research could be the analysis of the equilibrium model for the process generating house prices, and the final section of chapter 4 has already considered several directions in which I would like to expand my research. One focus would be to adapt the framework presented in the chapter, to compute a path of prices as the number of households in the population grows. This would facilitate an analysis of how much of the recent growth in real house prices might be due to population growth, and also of how expected future population growth could be expected to influence house prices. A more long term ambition for the equilibrium model would be to build a framework that explicitly allows for aggregate shocks, the likelihood of which would be built into expectations.

Regarding the analysis of the stakeholder pensions reforms, an obvious topic for further research would be an analysis of further reforms to pension contribution limits which were implemented in 2006 (see Department for Work and Pensions (2002)). Having a better understanding of how contribution limits, and other financial incentives, affect pension saving, would be very useful in an environment in which public policy aims to increase private saving as a means to supplement state pensions. The recent reforms also give more scope to test certain predictions of the lifecycle model, than was possible with the stakeholder pensions reforms. In particular, one could test whether the 2006 increase in contribution limits led any younger households to reduce current pension contributions because the reform allowed increased scope to exercise the option to wait before making contributions (see Milligan (2003) for a full description of this mechanism).

The combination of work on housing, and work on saving, suggests a further topic for future research, namely the analysis of a structural model that simultaneously considers housing choices, private pension saving and consumption / liquid saving decisions. Just as the housing good / asset adds complication to the lifecycle framework,

adding a pension also complicates the dynamic framework since wealth saved in this asset is locked away until the time of retirement. However, there is great value in studying the relationship between housing wealth and pension saving, due to its policy relevance. Recent assessments of the number of individuals in the U.K. who are likely to have inadequate resources for retirement have been shown to be sensitive to assumptions about how housing wealth is counted, (see Pensions Commission (2004) and Banks et al. (2005)). It is only with a theoretical understanding of how individuals substitute housing and pension wealth, that it would be possible to know what weight should be given to housing wealth when considering this adequacy question. The model would also be a sensible unification of the thus far separate research streams on housing, and on pension saving.

Bibliography

- AI, C. AND E. NORTON (2003): “Interaction terms in logit and probit models,” *Economic Letters*.
- ARON, J., J. MUELLBAUER, AND A. MURPHY (2007): “Housing Wealth, Credit Conditions and Consumption,” mimeo, Oxford University.
- ATTANASIO, O. P., J. BANKS, AND M. WAKEFIELD (2005a): “The Effectiveness of Tax Incentives to Boost (Retirement) Saving: Theoretical Motivation and Empirical Evidence,” *OECD Economic Studies*, 39, 145–172.
- ATTANASIO, O. P., L. BLOW, R. HAMILTON, AND A. LEICESTER (2005b): “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., R. BOTTAZZI, H. W. LOW, L. NESHEIM, AND M. WAKEFIELD (2007): “Explaining life-cycle patterns of home ownership,” mimeo, IFS.
- ATTANASIO, O. P., H. W. LOW, AND L. NESHEIM (2005c): “Modelling Life-Cycle Labour Supply and Asset Holdings with Fixed Costs,” ESRC specific grant.
- ATTANASIO, O. P. AND S. ROHWEDDER (2003): “Pension Wealth and Household Saving: Evidence from Pension Reforms in the U.K.” *American Economic Review*, 93, 1499–1521.
- ATTANASIO, O. P. AND M. WAKEFIELD (2008 and forthcoming): “The Effects on Consumption and Saving of Taxing Asset Returns,” paper prepared for the report of a commission on reforming the tax system for the 21st century, chaired by sir james mirrlees, Institute for Fiscal Studies, (<http://www.ifs.org.uk/mirrleesreview/reports/consumption.pdf>).

- ATTANASIO, O. P. AND G. WEBER (1994): "The UK Consumption boom of the late 1980's: aggregate implications of microeconomic evidence," *The Economic Journal*, 104, 1269–1302.
- (1995): "Is consumption growth consistent with intertemporal optimization? Evidence for the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121–1157.
- AUERBACH, A. AND L. KOTLIKOFF (1987): *Dynamic Fiscal Policy*, Cambridge University Press.
- BANKS, J., R. BLUNDELL, R. DISNEY, AND C. EMMERSON (2002a): "Retirement, Pensions and the Adequacy of Saving," IFS Briefing Note 29, Institute for Fiscal Studies, (www.ifs.org.uk/pensions/bn29.pdf).
- 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 (2002b): "The distribution of financial wealth in the UK: evidence from 2000 BHPS data," IFS Working Paper 02/21, The Institute for Fiscal Studies.
- BARKER, K. (2006): "Barker Review of Land Use and Planning: Interim Report Analysis," Tech. rep., HM Treasury, (http://www.hm-treasury.gov.uk/media/D/2/barker2_interim050706.pdf).
- BENJAMIN, D. J. (2003): "Does 401(k) Eligibility Increase Saving? Evidence from Propensity Score Subclassification," *Journal of Public Economics*, 87, 1259–1290.
- BERNHEIM, B. D. (1992): "Is the baby boom generation preparing adequately for retirement?" Technical report (unpublished), New York: Merrill Lynch.
- BERNHEIM, B. D. AND J. K. SCHOLZ (1993): "Private saving and public policy," in *Tax Policy and the Economy*, (vol. 7), ed. by J. Poterba, NBER, Cambridge: MIT Press.
- BESLEY, T., N. MEADS, AND P. SURICO (2008): "Household External Finance and Consumption," External MPC Unit Discussion Papers 25, Bank of England, (www.bankofengland.co.uk/publications/externalmpcpapers/extmpcpaper0025.pdf).

- BLUNDELL, R., M. COSTA DIAS, C. MEGHIR, AND J. VAN REENEN (2004a): "Evaluating the employment impact of a mandatory job search program," *Journal of the European Economic Association*, 2, 569–606.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2004b): "Consumption inequality and partial insurance," IFS Working Paper 04/28, The Institute for Fiscal Studies.
- BOTTAZZI, R., H. W. LOW, AND M. WAKEFIELD (2007): "Why Do Home Owners Work Longer Hours?" IFS Working Paper WP10/07, IFS, (<http://www.ifs.org.uk/wps/wp1007.pdf>).
- CAMPBELL, J. R. AND Z. HERCOWITZ (2004): "The Dynamics of Work and Debt," NBER Working Paper 10201, NBER.
- CAMPBELL, J. Y. AND J. COCCO (2007): "How Do House Prices Affect Consumption? Evidence From Micro Data," *Journal of Monetary Economics*.
- CHAMBERS, M., C. GARRIGA, AND D. E. SCHLAGENHAUF (2004): "Accounting for Changes in the Homeownership Rate," mimeo, Florida State University, (<http://mailer.fsu.edu/dschlage/garnet-dschlage/ownership-03.pdf>).
- CHERNOZHUKOV, V. AND P. HANSEN (2004): "The effect of 401k participation on the wealth distribution: An Instrumental Quantile Regression Analysis," *Review of Economics and Statistics*, 86, 735–751.
- CHIARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2007): "Business Cycle Accounting," *Econometrica*, 75, 781–836.
- CHUNG, W., R. DISNEY, C. EMMERSON, AND M. WAKEFIELD (2008): *Public policy and retirement saving incentives in the UK*, Cambridge: MIT Press, chap. 6, 169–210.
- CLARK, T. AND C. EMMERSON (2003): "The tax and benefit system and the decision to invest in a stakeholder pension," IFS Briefing Note 28, Institute for Fiscal Studies.
- COCCO, J. (2005): "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*.
- CRISTINI, A. AND A. SEVILLA (2008): "Do house prices affect consumption and why? A replication and comparison exercise," mimeo, University of Oxford, (http://www.eea-esem.com/files/papers/EEA-ESEM/2008/2656/comparison_ccabhl_140208.pdf).

- DAVIDSON, R. AND J. G. MACKINNON (2004): *Econometric Theory and Methods*, Oxford: Oxford University Press.
- DAVIS, M. A. AND J. HEATHCOTE (2007): “The Price and Quantity of Residential Land in the United States,” *Journal of Monetary Economics*, 54, 2594–2620.
- DAVIS, M. A. AND R. F. MARTIN (2008): “Housing, Home Production, and the Equity and Value Premium Puzzles,” mimeo, University of Wisconsin Madison, and Federal Reserve Board, (<http://morris.marginalq.com/2008-03-Davis-Martin.paper.pdf>).
- DE NARDI, M. (2004): “Wealth Inequality and Intergenerational Links,” *Review of Economic Studies*, 71, 743–68.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–48.
- DEPARTMENT FOR WORK AND PENSIONS (2002): “Simplicity, Security and Choice: Working and Saving for Retirement,” Tech. Rep. Cm. 5677, (<http://www.dwp.gov.uk/consultations/consult/2002/pensions/gp.pdf>).
- (2006a): “Personal accounts: a new way to save,” Tech. Rep. Cm. 6975, (http://www.dwp.gov.uk/pensionsreform/new_way.asp).
- (2006b): “Security in retirement: towards a new pensions system,” Tech. Rep. Cm 6841, (<http://www.dwp.gov.uk/pensionsreform/whitepaper.asp>).
- DEPARTMENT OF SOCIAL SECURITY (1998): “A new contract for welfare: Partnership in pensions,” Tech. Rep. Cm 4179, (<http://www.dwp.gov.uk/publications/dss/1998/pengp/pdfs/pensions.pdf>).
- DIAZ, A. AND M. LUENGO-PRADO (Forthcoming): “The Wealth Distribution with Durable Goods,” *International Economic Review*.
- DISNEY, R. AND C. EMMERSON (2005): “Public pension reform in the United Kingdom: what effect on the financial well being of current and future pensioners?” *Fiscal Studies*, 26, 55–82.
- DISNEY, R., C. EMMERSON, AND S. TANNER (1999): *Partnership in Pensions: A Commentary*, IFS Commentary 78, Institute for Fiscal Studies, London.
- DISNEY, R., C. EMMERSON, AND M. WAKEFIELD (2001): “Pension reform and saving in Britain,” *Oxford Review of Economic Policy*, 17, 70–94.

- (2007): “Tax reform and retirement saving incentives: Evidence from the introduction of Stakeholder Pensions in the United Kingdom,” IFS Working Paper WP19/07, Institute for Fiscal Studies, (<http://www.ifs.org.uk/wps/wp1907.pdf>).
- (Forthcoming (2009)): “Tax reform and retirement saving incentives: take-up of Stakeholder Pensions in the UK,” *Economica*.
- ECONOMIST (2007): “The Trouble with the Housing Market,” *The Economist* issue dated 24th March 2007.
- EJARQUE, J. AND S. LETH-PETERSEN (2008): “Consumption and Savings of First Time House Owners: How Do They Deal with Adverse Income Shocks?” CAM Working Papers 2008-08, University of Copenhagen. Department of Economics. Centre for Applied Microeconometrics.
- EMMERSON, C. AND S. TANNER (1999): “Stakeholder pensions,” IFS Briefing Note 1, Institute for Fiscal Studies, (http://www.ifs.org.uk/publications.php?publication_id=1761).
- ENGEN, E., G. GALE, AND C. UCCELLO (1999): “The adequacy of household saving,” *Brookings Papers on Economic Activity*, 2, 65–187.
- FERNANDEZ-CORUGEDO, E. AND J. MUELLBAUER (2006): “Consumer Credit Conditions in the U.K.” Bank of England Working Paper 316.
- FERNANDEZ-VILLAVERDE, J. AND D. KRUEGER (2005): “Consumption and Saving over the Life Cycle: How Important are Consumer Durables,” Working paper, University of Pennsylvania, (<http://www.econ.upenn.edu/~dkrueger/research/durables12192005sec.pdf>).
- FISHER, J. AND M. GERVAIS (2007): “First-Time Home Buyers and Residential Investment Volatility,” FRBC Working Paper 2007-15, Federal Reserve Bank of Chicago, (http://www.chicagofed.org/publications/workingpapers/wp2007_15.pdf).
- FLAVIN, M. AND S. NAKAGAWA (2004): “A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence,” Working Paper 10458, NBER.
- 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>).

- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*, Princeton University Press, Princeton, NJ.
- GALE, W. AND J. SCHOLZ (1994): “IRAs and household saving,” *American Economic Review*, 84, 1233–60.
- GILES, C. (2008): “Britains Property Market Faces its Moment of Truth,” *Financial Times* issue dated 12th April 2008, (<http://www.ft.com/cms/s/0/4d6b08c4-082a-11dd-a922-0000779fd2ac.html>).
- GRUBER, J. W. AND R. F. MARTIN (2005): “Durable Goods and the Distribution of Wealth,” in *New Directions and Applications in Control Theory*, Springer.
- JOURNAL OF ECONOMIC PERSPECTIVES (1996): “Symposium on ‘Government incentives for saving’,” *Journal of Economic Perspectives*, 10, 73–138.
- KAHN, J. A. (2008): “What Drives Housing Prices?” Working paper, Federal Reserve Bank of New York, (http://www.newyorkfed.org/research/staff_reports/sr345.pdf).
- KING, M. (1990): “Discussion (of “Is the U.K. Balance of Payments Sustainable?”),” *Economic Policy*, 5, 383–387.
- LEE, D. (2005): “An Estimable Dynamic General Equilibrium Model of Work, Schooling and Occupational Choice,” *International Economic Review*, 46, 1–34.
- 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.
- MILLIGAN, K. (2003): “How do contribution limits affect contributions to tax-preferred saving accounts?” *Journal of Public Economics*, 67.
- MODIGLIANI, F. AND R. BRUMBERG (1954): “Utility analysis and the consumption function: an interpretation of cross-section data,” in *Post-Keynesian Economics*, ed. by K. K. Kurihara, Rutgers University Press, New Brunswick, NJ, 128–197.
- MUELLBAUER, J. AND A. MURPHY (1990): “Is the UK Balance of Payments Sustainable?” *Economic Policy*, 5, 348–395.
- (1997): “Booms and Busts in the UK Housing Market,” *The Economic Journal*, 107, 1701–1727.

- NAKAJIMA, M. (2005): "Rising Earnings Instability, Portfolio Choice, and Housing Prices," mimeo, University of Illinois, Urbana-Champaign, (<http://www.compmacro.com/makoto/paper/050703earnhouse-paper.pdf>).
- NICHOLS, J. (2005): "Mortgage Contracts and The Definition of Demand For Housing Wealth," Ph.D. thesis, University of Maryland.
- OLIVER, WYMAN & COMPANY (2001): "The Future Regulation of UK Savings and Investment," Report prepared by oliver, wyman & company, Oliver, Wyman & Company, Executive summary is available at (<http://www.abi.org.uk/oliverwymanreport.pdf>).
- 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.
- PAGANO, M. (1990): "Discussion (of "Is the U.K. Balance of Payments Sustainable?")," *Economic Policy*, 5, 387–390.
- PENSIONS COMMISSION (2004): "Pensions: Challenges and Choices. The First Report of the Pensions Commission," Tech. rep., (<http://www.pensionscommission.org.uk/publications/2004/annrep/index.asp>).
- (2005): "A New Pension Settlement for the Twenty-First Century: The Second Report of the Pensions Commission," Tech. rep., (<http://www.pensionscommission.org.uk/publications/2005/annrep/annrep-index.asp>).
- POTERBA, J., ed. (1994): *Public policies and household saving*, Chicago: Chicago UP for National Bureau of Economic Research.
- 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.
- SCHOLZ, J., A. SESHADRI, AND S. KHITATRAKUN (2006): "Are Americans saving 'optimally' for retirement?" *Journal of Political Economy*, 114, 607–643.
- 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.
- VENTI, S. AND D. WISE (1987): "The determinants of IRA contributions and the effect of limit changes," in *Pensions in the US Economy*, ed. by Z. Bodie, J. Shoven, and D. Wise, Chicago: Chicago UP for National Bureau of Economic Research.
- 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.
- ZELDES, S. P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, 104, 275–98.