



REFERENCE ONLY

UNIVERSITY OF LONDON THESIS

Degree phd

Year 2005

Name of Author BOTTALTA

COPYRIGHT

This is a thesis accepted for a Higher Degree of the University of London. It is an unpublished typescript and the copyright is held by the author. All persons consulting the thesis must read and abide by the Copyright Declaration below.

COPYRIGHT DECLARATION

I recognise that the copyright of the above-described thesis rests with the author and that no quotation from it or information derived from it may be published without the prior written consent of the author.

LOANS

Theses may not be lent to individuals, but the Senate House Library may lend a copy to approved libraries within the United Kingdom, for consultation solely on the premises of those libraries. Application should be made to: Inter-Library Loans, Senate House Library, Senate House, Malet Street, London WC1E 7HU.

REPRODUCTION

University of London theses may not be reproduced without explicit written permission from the Senate House Library. Enquiries should be addressed to the Theses Section of the Library. Regulations concerning reproduction vary according to the date of acceptance of the thesis and are listed below as guidelines.

- A. Before 1962. Permission granted only upon the prior written consent of the author. (The Senate House Library will provide addresses where possible).
- B. 1962 - 1974. In many cases the author has agreed to permit copying upon completion of a Copyright Declaration.
- C. 1975 - 1988. Most theses may be copied upon completion of a Copyright Declaration.
- D. 1989 onwards. Most theses may be copied.

This thesis comes within category D.

This copy has been deposited in the Library of UCL

This copy has been deposited in the Senate House Library, Senate House, Malet Street, London WC1E 7HU.

Essays on Household Labour Market Participation,
Housing, and Wealth Accumulation Decisions

Renata Bottazzi

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

Doctor of Philosophy
University College London

London
September 2005

UMI Number: U592646

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U592646

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

The chapters of this thesis analyse elements of households' consumption, labour supply and saving decisions within the life cycle framework. The focus is on three main issues arising in this context. The first considers modelling both a durable good (housing) and labour supply choices together with life cycle choices over consumption and savings. The importance of modelling these features together comes from the existence of explicit earnings-related borrowing constraints when taking out a mortgage. Empirical evidence is provided for the UK and the modelling exercise, although not aimed to reproduce the evidence of a single country, is calibrated to the UK.

The second issue concerns incorporating expectations into the model. Two approaches are followed in the different chapters. One approach infers expectations from past realisations. The alternative approach uses expectations elicited in survey interviews.

The third issue relates to whether individuals save enough for their retirement. This question is addressed by using expectations on retirement outcomes, elicited directly in accordance with the approach mentioned above, and collected for a representative sample of the Italian working population. Data of interest to our exercise is for the years before and after a series of major pension reforms.

Contents

1	Introduction	13
2	Labour Market Participation and Mortgage Related Borrowing Constraints	19
2.1	Introduction	19
2.2	The literature review	22
2.3	The theoretical framework	24
2.4	The data	27
2.4.1	Sample selection	28
2.4.2	Description of variables	30
2.4.3	Descriptive statistics	32
2.5	The empirical model	39
2.5.1	Estimation method	41
2.5.2	Empirical results	46
2.5.3	Endogeneity	52
2.5.4	Sensitivity analysis	57
2.6	Conclusion	63
2.7	Appendix to chapter 2	66
2.7.1	House price data at LAD level: missings	66
3	Explaining Life-Cycle Profiles of Home-Ownership and Labour Supply	67
3.1	Introduction	67

3.2	Related literature	69
3.3	Observed Patterns of Home Ownership	70
3.4	Model of housing and labour supply	75
3.5	Calibration	79
3.5.1	Exogenous parameters	80
3.5.2	Calibrated parameters	82
3.5.3	Non-Calibrated profiles	84
3.6	Labour Supply and Home-ownership	86
3.6.1	Baseline model	88
3.6.2	Sensitivity to utility function specification	91
3.7	Implications of Incomplete Markets	94
3.7.1	Varying downpayment requirements	95
3.7.2	Varying debt to income requirements	95
3.8	Welfare implications of wage uncertainty	98
3.8.1	Welfare criterion	98
3.8.2	Welfare with CRRA/Cobb-Douglas preferences	99
3.9	Conclusions	112
3.10	Appendix to chapter 3	114
3.10.1	Computational methods	114
3.10.2	Estimation of wage variance	115
3.10.3	Labour supply with binding income-related borrowing constraint	116
4	Retirement Expectations, Pension Reforms, and Their Impact on Private Wealth Accumulation	118
4.1	Introduction	118
4.2	The Italian pension system: a decade of reforms	122
4.3	Expectations of retirement outcomes	127
4.4	The effect of the pension reform on expectations	130
4.4.1	Descriptive analysis	131
4.4.2	Regression estimates	133
4.5	The offset between pension wealth and private wealth accumulation . .	137

4.5.1	Regression estimates	139
4.5.2	Implications	144
4.6	Conclusions	146
4.7	Appendix to chapter 4	148
4.7.1	The pension award formula before and after the reform	148
4.7.2	The adjustment factor for pension wealth	149
4.7.3	The Survey of Household Income and Wealth	151
4.7.4	Employment transition matrix	152
5	Conclusion	154

List of Figures

2.1	Female participation by tenure	33
2.2	Female participation by mortgage repayments	35
2.3	Female participation by “obligation ratio” (mortgage repayment/other hh income)	37
2.4	Female participation by obligation ratio and age of youngest child 0-2 .	38
2.5	Female participation by “obligation ratio” and age of youngest child 3-4	38
2.6	Female hours of work by tenure	59
2.7	Female hours of work by mortgage repayments	59
3.1	Low education: Home ownership rates by cohort	72
3.2	Low education: Home ownership rates by year	72
3.3	High education: Home ownership rates by cohort	73
3.4	Transition probabilities over the life-cycle	73
3.5	UK House price series	81
3.6	Simulated home-ownership rates	84
3.7	Simulated transition probabilities	85
3.8	Net liquid asset position by age	86
3.9	Consumption by age	87
3.10	Labour supply by age	87
3.11	Labour supply by home-ownership status	91
3.12	Labour supply: baseline vs. non-homeownership	92
3.13	Labour supply by home-ownership status, $\varepsilon_{c,l} = 1.25$	93
3.14	Labour supply by home-ownership status, $\varepsilon_{c,l} = 0.83$	94

3.15 Homeownership varying downpayment requirements	96
3.16 Labour supply varying downpayment requirements	96
3.17 Homeownership varying the maximum debt-to-income ratio	97
3.18 Labour supply varying the maximum debt-to-income ratio	97
3.19 Low education: Change in consumption	109
3.20 Low education: Change in labour supply	109
3.21 Low education: Change in home-ownership	110
3.22 High education: Change in consumption	110
3.23 High education: Change in labour supply	111
3.24 High education: Change in home-ownership	111

List of Tables

2.1	Descriptive statistics of variables used in the estimation (N=6255)	29
2.2	Female participation by age and tenure	34
2.3	Female participation by age and mortgage repayments	36
2.4	Female participation by age and level of obligation ratio	39
2.5	Female participation by age, level of obligation ratio and age of the youngest child	40
2.6	Conditional ML Logit	48
2.7	Probit model	50
2.8	Conditional ML Logit - Endogeneity: control function technique	57
2.9	Conditional ML Logit - Sensitivity analysis: definition of participation	58
2.10	Conditional ML Logit - Sensitivity analysis: house value	60
2.11	Net financial wealth by age and level of obligation ratio	62
2.12	Conditional ML Logit - Sensitivity analysis: mortgage 1995-2000(*) . .	64
3.1	Logit Regressions for Becoming a Home Owner	74
3.2	Parameter Values	79
3.3	Calibration Statistics	83
3.4	Simulated Labour Supply by Homeownership status	90
3.5	Compensating variations in consumption	102
3.6	Compensating variations in assets	104
3.7	Compensating variations in consumption, labour supply and home-ownership	

4.1 Retirement age and pension award formula after the Italian pension reforms	124
4.2 The statutory replacement rate before and after the pension reform . .	126
4.3 Expected and statutory replacement rate: descriptive statistics	132
4.4 The effect of the reform on the expected replacement rate: regression results	136
4.5 The offset between private wealth and pension wealth: OLS and LAD estimates	140
4.6 The offset between private wealth and pension wealth: instrumental variable estimation	143
4.7 Population and sample means of selected variables	152
4.8 Transition matrix among employment groups	153

Acknowledgements

I am deeply indebted to my supervisors, Richard Blundell and Ian Preston. Richard has made sure I was looked after at the IFS since the summer of 2000 when I wrote my M.Sc. dissertation. He pointed me towards my current research topics and discussions with him have been a constant source of inspiration; on the many occasions when I have needed to check where I was going he has always been able to provide insightful comments and to see and show me the positive aspects of my research outcomes. His unfailing optimism has been crucial in enabling me to pursue my research and a valuable lesson for life. Ian has been an equally constant source of support, always available to discuss ideas and patient enough to read any piece of work that I would send him. His rigorous and neat interpretation of economic problems has been a sure guide that has both helped my understanding and improved the quality of my work in innumerable ways.

I consider myself very lucky to have had the privilege to spend the years of my Ph.D. in such a stimulating and friendly environment as the IFS. I thank everyone who has given me support and advice, and in particular Laura Blow, Ian Crawford, Andrew Leicester and Zoe Oldfield, who hosted me while I was doing my M.Sc. dissertation and never got rid of me. The many answers Zoe provided about the BHPS data when I first started were of great help. James Banks has also contributed to this work through many informal and challenging discussions.

I am extremely grateful to the ESRC for the financial support it provided through the Centre for the Microeconomic Analysis of Fiscal Policy at the IFS (grant number RES-544-28-5001). The analysis in chapter three was also supported by the ESRC

funded project “Modelling Life-Cycle Labour Supply and Asset Holdings with Fixed Costs” (grant number RES-000-23-0283).

I would like to express my gratitude to Orazio Attanasio, Hamish Low and Lars Nesheim for engaging me in a project in which I had so much to learn and they had so much to teach. Working with them and with Matt has been a pleasure and a challenging experience. Orazio has taken care of my research as if he were my supervisor. I owe an enormous debt to Hamish for providing us with his Fortran code, from which I have learned an incredible amount. Himself and Lars were always available for questions and doubts and always provided valuable advice.

Tullio and Mario also made sure I was not left alone in analysing Italian pension reforms. Tullio made his enthusiasm for research known to me when I first went to Salerno and I have not noticed it vanish. Mario took me patiently and wisely through the Italian data. His hospitality always made working together in Naples a very pleasant experience.

Starting at the IFS with Elena gave us the chance to share the burden of the PhD courses. The lunches that followed in the IFS common room together with Emla have been a much better part of the story. Giovanni and Marco are two good friends and colleagues with whom I have had stimulating discussions.

Emla and Antonio were the flatmates and friends that made me feel at home during my first years in London. We no longer share a home but they still have a special place in my thoughts.

I thank Matt, who has been cheerful, witty, funny on all occasions, and supportive and helpful beyond limits. His calm and generosity can still surprise me. His thorough thinking makes him the first person I share my doubts with.

My family has been wonderfully encouraging throughout all these years spent abroad. Andrea, my brother, has always made sure I would not lose track of local politics, and my cousin Giulia that I would not miss the crucial moments of her degree in medicine. Vanessa, my niece, has not given up complaining that I do not attend her birthday parties.

This thesis is dedicated to my parents. My mum is a wonderful example of gen-

erosity and will, and her approach to a life that is sometimes pretty unpleasant for her, has been inspiring during all these years and will be for many more to come. My father has given his unconditional support and made sure there was always someone waiting for me when arriving in Bologna.

Declaration

1. No part of this thesis has been presented to any University for any degree.
2. Chapter 3, except for Section 3.8 and Section 3.10.2, was undertaken as joint work with Orazio Attanasio, Hamish Low, Lars Nesheim and Matthew Wakefield.
3. Chapter 4 was undertaken as joint work with Tullio Jappelli and Mario Padula.

Renata Bottazzi

Chapter 1

Introduction

The chapters of this thesis analyse elements of households' consumption, labour supply and saving decisions within the life cycle framework. The idea underlying this framework is of rational consumers aiming to maximise utility by allocating their resources optimally between consumption and leisure and over time. The life cycle framework has been developed and refined in a stream of research that has followed the seminal contributions of Modigliani and Brumberg (Life Cycle Hypothesis, 1954) and of Friedman (Permanent Income Hypothesis, 1957).¹ In standard versions of the model, it is assumed that the utility function is intertemporally separable, so that the utility derived from today's level of consumption, and from other current choices, does not affect or depend on choices made in other periods. Individuals maximise the discounted sum of the future stream of utility, which translates into keeping the marginal utility of money constant over time (and not necessarily into constant consumption over time, as assumed in Modigliani and Brumberg's original Life Cycle Hypothesis approach).

When outcomes are uncertain, the maximisation of inter-temporal utility will require the formation of expectations. While this has long been recognised, the theoretical advances that allowed for a systematic exploration of how uncertainty might affect decisions, occurred only in the mid-1980s. Before that time, models were often analysed under a set of assumptions that led to "certainty equivalence". Utility was

¹See Browning and Crossley (2001) and Carroll (2001) for critical summaries of the evolution of models within the life cycle framework.

assumed to be quadratic, and marginal utility of consumption linear, so that maximising expected lifetime utility would translate into a flow of consumption independent of the degree of uncertainty about future outcomes. In other words, only first moments of the stochastic processes over future outcomes would matter. As pointed out by Carroll (2001), it was only with Zeldes (1984) contribution that the “certainty equivalent” assumption was relaxed to allow for more realistic assumptions about how uncertainty could affect consumption decisions. Blanchard and Mankiw (1988) showed that consumption paths could depend on the uncertainty faced by individuals, rather than only on expected future outcomes. This led to the introduction of the idea of precautionary saving, whereby individuals reduce their current consumption when they face greater uncertainty about the future. Utility functions that would allow consumption (and other choice variables) to be sensitive to uncertainty, have since become common. A typical example is given by the Constant Relative Risk Aversion (CRRA) utility function.

If we think of the life of an individual as being made up of three stages, one in which she does not work but perhaps engages in education, one in which she works, and one in which she is retired, the life cycle model in general predicts that there should be borrowing in the first stage, accumulation of wealth in the second stage and wealth decumulation in the third stage. Browning and Crossley (2001) highlight a number of partially solved as well as unsolved puzzles, together with areas that have not been yet investigated in depth within the life cycle framework. Puzzles that remain unresolved include: little spending, in the first stage of the life cycle, by individuals with high expected lifetime wealth; and a lack of consumption smoothing in the transition from the second stage (work) into the third stage (retirement) of the life cycle (the so-called “retirement-consumption puzzle”). Partial consensus is instead reported over the view that the empirical correlation between consumption and income found during working life is due to a combination of precautionary savings and demographic changes over the life cycle. Areas that have as yet received little attention include modelling consumption goods with different properties, including different durability, as well as assets with different characteristics, such as different liquidity.

This thesis will focus on three main issues arising in the life cycle framework. The first concerns expanding the types of goods that are incorporated in the model, and in particular simultaneously adding both a durable good and the labour supply choice. The second involves exploring different methods that have been suggested to incorporate expectations in theoretical and empirical models. Finally, there is also an indirect contribution to the evidence on one of the unresolved puzzles identified above, namely that of the retirement-consumption puzzle. I will now briefly discuss each of these issues in turn, focussing on how they are addressed in the chapters of the thesis.

The first issue involves simultaneously adding both a durable good (housing) and labour supply choice to the life cycle framework. Consumption and labour supply are modelled alongside housing, which is an illiquid asset, a collateral good, and a good that directly affects household utility. This is the focus of chapters 2 and 3. The importance of modelling labour supply along with home ownership is thought to be particularly important in contexts where an explicit income-related borrowing constraint holds for households that want to take out a mortgage. In this respect, empirical evidence for the United Kingdom is presented in chapter 2, which investigates whether female labour supply is correlated with mortgage repayments. In this analysis, mortgage repayments are thought of as a dimension that picks up the existence of a binding liquidity constraint, particularly in the early stage of the life cycle when households purchase a home.

The empirical evidence of a positive correlation between female labour supply and mortgage commitments found in chapter 2 is the basis for investigating, in chapter 3, the relationship between home ownership, consumption, labour supply and saving in a structural life cycle model. Two institutional borrowing constraints associated with the home purchase are introduced - namely a down payment constraint and an earnings-related borrowing constraint - as well as transaction costs of buying and selling the home. The utility function is CRRA and therefore allows for prudent behaviour. The model is calibrated to home ownership profiles of the cohort born in the late 1940's and

produces labour supply profiles as well as profiles of consumption and asset accumulation. It is found that changes in home-ownership affect both consumption and labour supply. Moreover, the labour supply of home-owners is found to be higher than that of non-homeowners, with peaks at the time of house purchase. This is suggestive of constrained behaviour, particularly via the income-related borrowing constraint, but the effect is mostly small and smaller than that observed in the data. Part of the reason why the model does not seem able to generate the scale of effects seen in the data could be that in the model there is only one source of ex-ante heterogeneity, which is the education level of the household. Considering the potential importance of the earnings-related borrowing constraint, the relevance and welfare implications of different degrees of wage uncertainty are also measured. In particular, compensating variations in assets, subsequently translated into lifetime changes in consumption, labour supply and home-ownership, are obtained for scenarios with different variance in wages. Simulations show that labour supply and consumption choices in a scenario with compensating variations in assets and a higher wage variance are, respectively, lower and higher than in a scenario with lower variance and zero initial assets. Home-ownership is always higher at the beginning of life, since the compensating variation in assets relaxes the explicit borrowing constraints. However, it can be lower in the second part of the life cycle due to the implicit “no-bankruptcy” constraint being tighter in a scenario with higher wage variance, and in fact becoming binding.

The second issue concerns incorporating expectations into the model. A common way to model expectations within the life cycle framework is to infer expectations from past realisations. This however requires making a number of assumptions on expectation formation. In particular, the economist is assumed to know what information on realisations individuals have and how they translate this information into expectations. These two steps are implicit when, for example, agents in a model with stochastic earnings are assumed to form their expectations rationally by exploiting the properties of an econometrically estimated wage process. However, a strand of literature has now been developed on eliciting expectations in survey interviews (see Dominitz (1998,

2001) and Dominitz and Manski (1994)) in order to recover probability distributions over future outcomes without imposing any structure a priori. Chapter 3 will adopt the first approach with respect to three stochastic processes: wages, house prices and the interest rate. Chapter 4 will instead draw on the second approach in looking at point expectations over retirement outcomes.

The third contribution of the thesis is to provide some evidence, within the life cycle framework, that indirectly contributes to the debate around the retirement-consumption puzzle. Since pension income is the main component of income during the last stage of the life cycle, expectations of it are crucial in determining inter-temporal decisions of consumption and savings. Given that pension rules are normally set by governments, in order to form expectations correctly, individuals need to be informed about pension legislation, understand it, and adjust their expectations whenever changes in legislation take place. If any of these elements are missing, then current decisions over consumption and savings will be distorted by having inaccurate expectations. Chapter 4 analyses the expected retirement age and the expected ratio of pension benefits to pre-retirement income (the replacement rate) of working individuals in Italy, as collected in the Italian Survey of Household Income and Wealth. It then estimates the offset between expected pension wealth and private wealth accumulation. In some sense, this analysis could contribute to the debate on the drop in consumption around retirement. For example, if individuals do not adjust their expectations about the replacement rate after a change in legislation has reduced it, then they will make their inter-temporal consumption and savings plans by aiming to smooth marginal utility for the wrong amount of wealth. If they only realise their mistake at retirement, they will need to re-optimize their consumption plans and possibly be forced to reduce their consumption at that time. Eliciting expectations directly from individuals, rather than inferring them from current legislation, is a crucial feature of this analysis. A drawback of this approach is that little is yet known about the way in which people answer questions about expectations. On the other hand, it seems particularly relevant to obtain information on expectations in this way, as it is

what individuals believe that ultimately matters for their choices. It is found that individuals overestimate the replacement rate both before and after a series of changes to pension legislation. Moreover, the offset between private wealth and expected (state) pension wealth is found to be far from complete.

Chapter 2

Labour Market Participation and Mortgage Related Borrowing Constraints

2.1 Introduction

This chapter analyses the relationship between female labour market participation and mortgage commitments in a life-cycle setting.

Most of the literature concerning life-cycle models and borrowing constraints has focused on the impact of borrowing constraints on consumption, assuming labour supply to be exogenous. However, it is worth investigating whether this is a plausible assumption. In fact, labour supply may be thought of as a way to circumvent or to reduce the impact of borrowing constraints, and so assuming that it is exogenous would produce biased estimates of the impact of borrowing constraints on consumption.

The importance of mortgages in the context of borrowing constraints and labour supply decisions derives from two sources. The first is that housing is usually the most substantial component of assets for home owners. The second is that there is an explicit mortgage qualification constraint, based on household annual income and assets, when the mortgage is taken out.

The ways in which mortgage commitments can be interpreted, will now be analysed in turn.

When taking out a mortgage individuals normally face two types of constraints: a wealth constraint, typically in the form of a minimum downpayment, and an income constraint, typically in the form of a maximum ratio of mortgage payment to household earned income. In the UK, it is theoretically possible to borrow up to 100% of the property's value (although a higher lending fee, the so-called MIG -Mortgage Indemnity Guarantee-, normally applies for mortgages of more than around 90% of the property value) and the effective qualification constraint is the one related to the household income. The Financial Service Authority reports that "typically, the maximum mortgage a lender offers is three times the main earner's income plus one times any second earner's income, or two-and-a-half times the [households] joint income". Hence, there exists an explicit mortgage qualification constraint that is a function of household earned income.

Although it formally applies when the mortgage is taken out, it is possible to consider an income constraint as holding at every period during which the mortgage is being repaid as long as remortgaging is a possibility. Fortin (1995) justifies the presence of such a constraint at every period on the basis of the Canadian institutional setting, where people remortgage very frequently¹ or, alternatively, in the light of allocational inflexibilities introduced by mortgage payments in the household decision process. In other words, it is explicitly taken into account that mortgage payments might alter household consumption and leisure decisions due to the fact that they must be met at every period and that they normally constitute a substantial part of households' total income.

Moreover, a more general earnings-related borrowing constraint (see Alessie, Melenberg and Weber (1988) and Aldershof et al. (1997)) may apply to households that have taken out a mortgage. If capital markets are imperfect, people might not be able to borrow as much as they would like to, and one way to express the borrowing

¹In the selected sample used in this work, 20.5% of households take out additional mortgage on their home at least once during the observation period.

constraint is in terms of earned income. In particular, individuals who earn more can borrow up to a higher sum than people who earn less, on the basis that their higher earnings reflect higher human capital to be used as collateral. Households with mortgages might be particularly subject to this type of constraint if their home equity is still low and they do not have other assets to offer as collateral. However, it might be the case that individuals with higher earnings, together with a mortgage, hold assets of various forms, hence rendering this type of borrowing constraint binding only for households at the bottom of the income distribution.

An important issue that follows from the existence of a mortgage-related qualification constraint is that mortgage commitments are potentially endogenous in labour supply decisions. Most of the literature that addresses the relationship between mortgage commitments and labour supply decisions takes mortgage choice (whether to take out a mortgage, its size) as exogenous (see Fortin (1995), Aldershof et al (1997)). In other words, the borrowing constraint is considered to hold at every period but the amount of the mortgage payment is assumed to be exogenous. This will be the implicit assumption in the first part of this analysis as well. That is, it will be assumed that the direction of causality in the relationship between mortgage commitments and labour market participation runs from the former to the latter. The hypothesis of exogeneity is then tested in the second part of the chapter.

Most of the studies that have been carried out so far have analysed a single cross-section of the data set of interest and have thus relied on a static-level analysis. In this work, panel data from the British Household Panel Study will be used. Although a static model will be estimated, unlike static models based on a single cross-section, individual unobserved heterogeneity will be controlled for.

The chapter is organised as follows. Section 2.2 presents a survey of the related literature. Section 2.3 contains the theoretical framework that is used as a guide in the interpretation of the empirical results. Section 2.4 describes the data set as well as sample selection issues and the variables used in the analysis. It also contains a descriptive section on the relationship between participation and mortgage commitments. Section 2.5 presents the empirical model and the estimation methods, a description of

the empirical results and a discussion of endogeneity of mortgage choice. Section 2.6 concludes.

2.2 The literature review

The effect of mortgage market constraints has been analysed in relation to different types of household decisions such as tenure, savings and labour supply. Either the institutional borrowing constraints regarding the downpayment or the ratio of mortgage payments to household income, or both, have been considered. Some works have instead adopted a more general borrowing constraint.

Yoshikawa and Ohtake (1989) examine housing demand and female labour supply in the context of a three-period life-cycle model in which households choose the type of tenure (renting/owning), consumption of other goods and female labour supply, subject to a life-time budget constraint and to an additional constraint related to the downpayment for those who choose to own. They estimate the savings function and the female labour supply function for the two tenure types from a cross-section of the Japanese National Survey of Family Income and Expenditure, by using a two-step Heckman procedure. Tenure choice and savings are made to depend, among other variables, on permanent income of the household head, total net assets and price of land. It is found that changes in husbands' permanent income and/or the price of land affect the tenure decision and that this switching effect makes the price of housing affect the savings rate in two opposite directions. Whether the switching effect affects female labour supply as well, is not examined.

Fortin (1995) analyses the relationship between household labour supply and the mortgage qualification constraint expressed in terms of a maximum gross debt service ratio (mortgage payments/gross household income). The theoretical setting is that of a multiperiod life-cycle model with utility maximization over leisure and consumption, subject to the current period allocation of life-time wealth and to the additional mortgage qualification constraint based on earnings. It is assumed that this mortgage qualification constraint holds at every period on the basis of the Canadian institutional

setting, where people remortgage very frequently. Using 1986 data from the Canadian Family Expenditure Survey, the first estimate is of a reduced form model of female labour supply in relation to a set of housing variables that include the value of the house owned, the balance of mortgage and dummy variables for different levels of the ratio of mortgage charges to gross family income (exclusive of wife's labour income). Both a higher balance of mortgage and a higher ratio of mortgage charges to other family income positively affect female labour market participation and hours of work. Subsequently, two labour earnings equations are estimated in relation to the theoretical model. In order to test whether binding constraints for mortgage qualification influence wives' participation, one equation relates to the case in which the housing constraint applies to both the husband and the wife, and the other equation relates to the case in which it applies only to husband's earnings. It emerges that the housing constraint should apply to both spouses but the husband's earnings should be given a greater weight than the wife's.

Aldershof et al. (1997) analyse the relationship between female labour supply and housing consumption by developing a life-cycle consistent model in which households maximise expected lifetime utility over female leisure, housing consumption and non-housing consumption, subject to a housing production function, to the standard lifetime budget constraint and to an earnings-related liquidity constraint.² Hence, labour supply and housing consumption are jointly determined. Separability of preferences between these two choices is tested by estimating the female labour supply function conditional on housing consumption and it is not rejected. Estimation is based on the 1989 wave of the Dutch Socio-Economic Panel. Moreover, it is found that female labour supply is positively affected by the ratio of mortgage interest payments to total family income (exclusive of the woman's earnings), which is considered to be a proxy for the presence of binding liquidity constraints.

Del Boca and Lusardi (2001) examine whether imperfections in the credit market,

²It is through this earnings dependent liquidity constraint that mortgage commitments are expected to relate to female labour supply. It is argued that the constraint is more likely to be binding in cases where the household has high mortgage commitments.

and in particular in the mortgage market, spill over to the labour market by using the 1989 and 1993 cross-sections from the Italian Survey of Household Income and Wealth. Exogenous changes in the mortgage market, such as the reduction of the down payment and wider access to the mortgage market (due both to the financial liberalization brought by the European unification in 1992 and to the Amato Act of 1990), are considered to be an important source of variation to identify these effects. The decision to participate in the labour market and to obtain a mortgage are modelled in terms of a latent variables simultaneous equation model. Mortgage debt is introduced in the empirical estimation using a dummy and a continuous variable for the amount outstanding. It is found that it has a positive and significant effect on female labour participation. A dummy and the outstanding value owed are also used for family debt and for other types of debt (car, appliances) to check whether mortgage commitment is different from other types of debt. The direction of causality between borrowing constraints and labour supply is assessed by relying on variables proxying for the credit system³ and the changes in the mortgage market over time.

2.3 The theoretical framework

The theoretical framework that is adopted as a guide for the empirical specification is one of dynamic programming where individuals choose labour supply (participation) and consumption according to the value function

$$V_t(A_t) = \max_{P_t, c_t} [u_t(P_t, c_t, Z_t) + \beta E_t V_{t+1}(A_{t+1})] \quad (2.1)$$

³They include whether people have a checking account, how many banks they have been using and how many credit cards they use

subject to the standard asset accumulation rule and to a mortgage-related liquidity constraint,⁴ as follows:

$$A_{t+1} = (1 + r_{t+1})[A_t + (y_t + w_{ft}P_t - c_t - m_o)] \quad (2.2)$$

$$k(y_t + w_{ft}P_t) \geq m_o$$

where

A_t = assets at the beginning of period t

$u_t(\cdot)$ = intra-temporal utility function at period t

P_t = hours of work (participation)

c_t = consumption

Z_t = demographic variables

E_t = expectation operator based on information at time t

β = consumer's discount factor

r_t = rate of return between period $t - 1$ and period t

y_t = unearned (non-asset) income (husband's income)

w_{ft} = female wage (earned income if P_t represents participation)

m_o = mortgage payment at period t (assumed to be constant, i.e. non depending on the interest rate)

k = proportion of total household income to be allocated to mortgage payments

Utility is assumed to be intertemporally separable and the intra-temporal utility function, $u_t(\cdot)$, is assumed to be strictly concave and monotonically increasing in c and decreasing in P (there is no presumption regarding Z). Moreover, the amount of

⁴The mortgage-related borrowing constraint has been introduced following Fortin (1995). Also if thought of in terms of allocational inflexibility, it is meant to be relevant for homeowners with mortgage but not for renters. In fact, it could be argued that homeowners with mortgage who are incapable of keeping up with mortgage payments could always move to a smaller house in the same way as renters could move to a cheaper accommodation. However, as opposed to renters, in order to do that, they would be subject to a mortgage qualification constraint based on earnings. Alternatively, by staying in the same accommodation, they would be subject to allocational inflexibilities (captured by the mortgage-related borrowing constraint) in a way that renters would not be.

the mortgage is assumed to be exogenous so that the mortgage payment at period t is assumed as given.

First-order conditions are obtained from standard dynamic programming techniques and are represented by the following:

$$u_{ct} = \beta E_t[(1 + r_{t+1})\lambda_{t+1}] \quad (2.3)$$

$$u_{Pt} = \beta E_t[(1 + r_{t+1})\lambda_{t+1}]w_{ft} - \mu_t k w_{ft} \quad (2.4)$$

$$\lambda_t = \beta E_t[(1 + r_{t+1})\lambda_{t+1}] \quad (2.5)$$

$$\mu_t [k(y_t + w_{ft}P_t) - m_o] = 0, \quad \mu_t \geq 0 \quad (2.6)$$

where u_{ct} and u_{Pt} denote the first-order derivatives of the utility function with respect to consumption and hours of work, μ_t represents the Kuhn-Tucker multiplier associated to the mortgage-related constraint and λ_t is the marginal utility of wealth ($\partial V_t / \partial A_t$).

When constraint (2.6) is binding, $\mu_t > 0$ and $P_t = \frac{m_o/k - y_t}{w_{ft}}$. Hence,

$$\mu_t = f(m_o, y_t, w_{ft}) \quad (2.7)$$

The effect of the mortgage constraint, when binding, is that participation becomes a function of mortgage payment, other family income and female wage. By contrast, when it is not binding, $\mu_t = 0$ and participation does not depend on mortgage payment.

Using first-order conditions (2.3)-(2.5), optimal consumption and labour supply equations are defined as follows:

$$c_t = c_t(\lambda_t, w_{ft}, Z_t, \mu_t) \quad (2.8)$$

$$P_t = P_t(\lambda_t, w_{ft}, Z_t, \mu_t) \quad (2.9)$$

The equations derived for consumption and participation are the so-called *Frisch equations*.⁵ In this framework, λ is interpreted as capturing all the information from other periods that is required in order to obtain the optimal choice in the current period (e.g. it could be thought of as reflecting household permanent income).⁶

⁵Demand functions derived by holding the marginal utility of wealth constant.

⁶In an empirical specification, λ would then be modelled as an individual fixed effect.

In the theoretical framework presented so far, optimization is carried out over hours of work. However, the empirical specification that will be introduced later is expressed in terms of a participation rather than hours equation as a function of other family income, demographics, ratio of mortgage payment to other family income (also interacted with demographic variables). This discrepancy is purely due to the easier analytical tractability of a continuous variable rather than a discrete one, and the model is meant to be just an indication of the way in which the mortgage-related constraint affects labour supply decisions. Finding a positive effect of the mortgage variable on participation is taken as evidence that the mortgage borrowing constraint is binding, and hence that having a mortgage distorts households participation decisions.

2.4 The data

The data used for this work is the British Household Panel Study (BHPS), waves 3-10 (1993-2000).⁷ It reports information on British households on the basis of yearly interviews conducted on an original sample of approximately 5,000 households (circa 10,000 individuals).⁸ The panel nature of the data set means that the same individuals are interviewed each year.⁹

The data contains both household and individual level information. At the individual level, apart from demographic characteristics such as age, region, number of children and education, there is detailed information concerning current labour force status, labour force history, health and personal finances (including different sources of income and of investment). Moreover, detailed wealth data is collected every five years.

At the household level, there is detailed information regarding household composition, household expenditure and tenure type of the home lived in, and details on rent

⁷Waves 1 and 2 are not being used since wave 2 reports many missing observations on housing variables.

⁸Children are interviewed if above 16.

⁹In fact, the selected sample that is being analysed in this work reports on average 4.3 observations per individual over 8 years. This is partly due to attrition and partly to sample selection (see next section).

or mortgage payments and on loan repayments. In particular, housing information is provided with different degrees of detail according to whether households own the accommodation, are paying a mortgage or are renting. Each year, all respondents are asked about the type of accommodation and the number of rooms. Homeowners, including those with a mortgage, also provide their estimate of the value of the house (based on current prices). Further, owners with mortgage are asked to provide some information every year and some additional information the first time they are interviewed at their current address. Among the information which is provided only by people being interviewed for the first time (or who have moved address since previous interview), there is the following: the year in which the mortgage was taken out, the original cost of the house, the original amount of the mortgage (excluding later additions), the number of years the mortgage has still to run and the type of mortgage (whether a repayment mortgage, an endowment mortgage, a mixture of the two or some other type). Every year homeowners with a mortgage are asked whether any additional mortgage has been taken out, its amount, the destination use of it (home extension/ improvement, car purchase, other consumer goods, other). They are also asked to provide the last total monthly instalment on the mortgage or loan.¹⁰ Finally, there are three direct questions on liquidity constraints related to either rent or mortgage payments. In particular, respondents are asked whether there have been any difficulties paying for the accommodation in the last twelve months, and, in case of affirmative answer, whether this has resulted in borrowing money or in cutting back on other household spending. Respondents are also asked whether or not the household has been more than two months behind with the rent or mortgage payment.

2.4.1 Sample selection

Since the purpose of this work is to analyse whether mortgage-related borrowing constraints affect households' labour participation, the focus will be on female participation on the ground that males normally work full time and so their labour supply

¹⁰It includes life insurance payments if they are paid together with the mortgage, and it includes both the premium and the interests if it is an endowment mortgage.

behaviour is already constrained. Only people who are in couples (either married or cohabiting) and of different sex will be taken into account and, in line with Fortin (1995), it will be people who are in couples in which the husband works full time and so receives a regular salary. The number of observations drops from 90998 to 45082.

Self-employed are selected out on the basis of the assumption that their labour supply behaviour is different from employed people.¹¹ Moreover, renters are removed from the sample as it is assumed that the tenure decision is exogenous and the focus is on homeowners either with or without a mortgage. Finally, only women in the age range 25-45 are kept for the analysis. This is based on the assumption that above this range the mortgage qualification constraint is unlikely to be binding, moreover it is in this range that interactions with the presence of children are more likely to take place. The final working sample is of 12510 observations, that is, 6255 households. Summary statistics of the variables used in the estimation are reported in Table 2.1.

Table 2.1: Descriptive statistics of variables used in the estimation (N=6255)

variable	Mean	Std. Dev.	Min	Max
female participation	0.818	0.386	0	1
other hh income	1996.309	1123.378	235.177	21483.810
log(other hh income)	7.494	0.438	5.460	9.975
age (rescaled by 25)	9.741	5.694	0	20
obligation ratio	0.168	0.096	0	1
youngest child 0-2	0.155	0.362	0	1
youngest child 3-4	0.122	0.327	0	1
no. of children	1.174	1.060	0	7
no education	0.119	0.324	0	1
education O level	0.453	0.498	0	1
education A level	0.203	0.402	0	1
higher degree	0.221	0.415	0	1

Since we have an unbalanced panel, the number of households is in fact 1475, of which 281 are only observed once over 8 periods and 258 are observed over the whole

¹¹In particular, self-employed people do not receive a fixed wage as is assumed in this analysis.

period (8 times). A varying number of 120 and 200 households are observed 2-7 times. On average, people are observed 4.3 times out of 8 periods. The sample selection plays an important role in this apparent high attrition, both through selection of people in couples and through the choice of the age range.

2.4.2 Description of variables

Female participation

It is a dummy variable defined on the basis of whether the female did any paid work the week before the interview and, in case of negative answer, whether she had a job that she was away from in the same week.¹² Since this definition may assign women who did casual work and who were on maternity leave to the group of participants in the labour market,¹³ two alternative definitions of participation have been used to check the sensitivity of the results. One is based on whether the number of weeks of employment during the year before the interview was positive and the other is based on whether the declared current employment status is that of being in paid employment. None of these alternative definitions change the pattern of the results.

Demographic information

The BHPS includes all the standard demographic information such as age, education, number of children, age range of children and region of residence.

In the analysis that follows female's age has been rescaled, so that the youngest females in the selected sample (25 years old) are the reference group. Its square has also been included in order to take into account non-linearities in the relationship between participation and age.

Children possibly play a central role as the focus of the analysis is on female participation at early/middle stage of the life cycle. The number of children is controlled for, as is age of the youngest child using dummies which indicate whether the youngest

¹²“Any paid work” includes any number of hours, including Saturday jobs and casual work. The reasons of being away from work include those on maternity leave, on holiday, on strike and on sick leave.

¹³It turns out that only 8 observations register the female as participating while she is in fact on maternity leave.

child is either between 0 and 2 years of age or between 3 and 4 years of age. As long as this adds information to the analysis, a dummy for the youngest child being above 5 years of age will be included. Female's education and the region of residence are used as additional control variables whenever the estimation method does not eliminate the fixed effect. Education is summarised by four dummy variables: no education, O level, A level, degree or higher. Dummies are defined for each of nineteen regions.

Income measures

Since the focus of this analysis is on female labour participation, it is necessary to control for income effects related to unearned income. Hence, the log of other family income defined as all household income but female's earned income (i.e. husband's earned and non-earned income and the female non-earned income) is included as explanatory variable. In particular, monthly non-labour income is recovered from the sum of all non-labour income in the year prior the start of the interview period, whereas labour income is given by the usual gross pay per month.¹⁴

Mortgage and housing information

Of all the housing information described in the previous section, the main variable that has been used in this context is the monthly mortgage payment. The ratio between the monthly mortgage payment and the other family income is the so-called *obligation ratio* (*or*). As opposed to the ratio that enters the mortgage qualification constraint, this ratio has other family income rather than total family income in the denominator, the difference between the two being the female earned income.

The obligation ratio is expected to capture the effect on household labour market participation of the variation in the burden of mortgage commitments relative to household income. In other words, it is expected to capture the effect of mortgage-related borrowing constraints on labour market participation. Although the institutional mortgage qualification constraint should hold only when the mortgage is taken out, it is plausible to think that it may hold also in subsequent periods. In fact, a positive correlation between the obligation ratio and female participation may be taken

¹⁴It is gross payment in the month before interview as long as this is the usual payment; when only the net monthly income was available, BHPS imputed values have been used.

as evidence of this hypothesis. Fortin (1995) suggests that the mortgage qualification constraint may hold after the mortgage has been taken out as people may remortgage quite frequently. On the other hand, variation in interest rates may provide another reason for variation in participation related to mortgage commitments. For instance, assuming that the constraint binds, a household that obtained a mortgage on the basis of only male's income and then faces an increase in the obligation ratio due either to higher interest rates or to a decline in male's earnings, would probably have to increase its labour market participation in order to lessen the effect of the underlying mortgage-related borrowing constraint. Alternatively, it could re-mortgage but in that case it would face the institutional mortgage qualification constraint. Hence, either implicitly or explicitly, the earnings-related mortgage constraint would hold not only when the mortgage is taken out but also in subsequent periods.

2.4.3 Descriptive statistics

Table 2.2 and Figure 2.1 illustrate female labour market participation behaviour by different types of tenure and age as obtained from the pooled selected sample (years 1993-2000).¹⁵

Both document a strikingly higher participation rate for females in households that have a mortgage, particularly in the age range 30-35, when presumably labour supply decisions are strongly related to the presence and the age of children. Whereas participation for women without a mortgage shows a definite U-shaped pattern with its minimum at the age of 32-33, participation for those with a mortgage only decreases slightly from 84% to 80%. The participation rates of renters appear to be in between the participation rate of owners with no mortgage and owners with a mortgage for any age after 29 and it is the lowest before age 29.

Since households that own the house outright are less than 4% of the sample that includes renters,¹⁶ and since 80% of households in the same sample own with a mortgage, it is necessary to rule out the possibility that the above result is driven by the

¹⁵The figures have been obtained by running mean least-squares smoothing and confidence intervals have been constructed from pointwise standard errors of smoothed values of participation.

¹⁶4.5% of the sample with no renters.

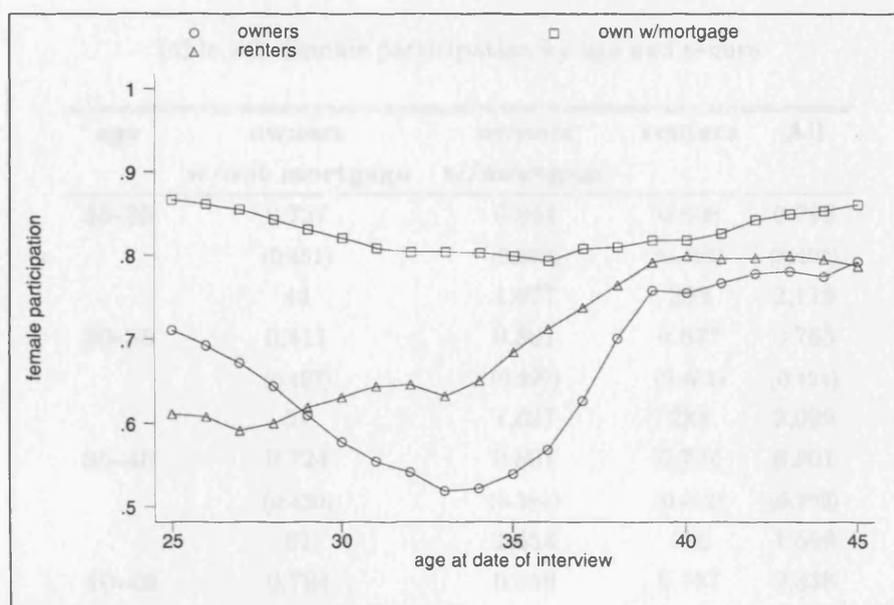


Figure 2.1: Female participation by tenure

small number of observations for the group of outright owners, particularly for age below 35. Rather than dividing the group of owners into those who own outright and those who have a mortgage, the two groups are defined according to whether the monthly mortgage payment is below or above the 33rd percentile of the distribution of the monthly mortgage payment (dummy for low/high mortgage payment), where the group with low mortgage payment includes those that own outright. As documented in Table 2.3 and in Figure 2.2, up to the age of 36 the participation rate of these two groups presents the same features as the groups of owners outright and owners with mortgage, with a definite U-shaped pattern for those with low monthly mortgage payment (although, as expected, the average participation rate for the group of owners with low mortgage is now higher than before). After the age of 36 the two patterns are very similar, possibly due to the fact that the mortgage-related constraint is no longer binding. The participation rate of renters, on the other hand, is now lower than the one of both groups of home owners.

In what follows it is then further explored the relationship between the participation

Table 2.2: Female participation by age and tenure

age	owners w/out mortgage	owners w/mortgage	renters	All
25-30	0.727 (0.451) 44	0.844 (0.362) 1,677	0.608 (0.489) 398	0.798 (0.402) 2,119
30-35	0.411 (0.497) 51	0.801 (0.399) 1,697	0.637 (0.482) 281	0.765 (0.424) 2,029
35-40	0.724 (0.450) 87	0.807 (0.394) 1,454	0.784 (0.412) 158	0.801 (0.399) 1,699
40-45	0.784 (0.414) 97	0.848 (0.359) 1,148	0.787 (0.411) 141	0.838 (0.369) 1,386
All	0.688 (0.464) 279	0.824 (0.381) 5,976	0.671 (0.470) 978	

Notes: Cells contain, top to bottom, the mean, the standard deviation (in brackets), and group size of female participation in the labour market.

rate of home owners (with and without a mortgage) by focusing on a sample that excludes renters. The obligation ratio is used instead of the mortgage monthly payment in order to net out the effect of household income (excluding female labour income) from a measure of allocational inflexibilities imposed by the mortgage.¹⁷ Figure 2.3 follows the same logic as Figure 2.2 but is based on the obligation ratio rather than the level of mortgage monthly payment. It illustrates the pattern of female participation according to whether the obligation ratio is below or above its 33rd percentile level.

¹⁷In fact, the institutional mortgage borrowing constraint imposes an upper bound on the mortgage level in terms of household income. Hence, if household income does not change significantly over time, and mortgage payments do not decrease (for instance due to decreasing interest rates) it is plausible to expect that mortgage payments are positively related to household income. Hence, the effect of the burden imposed by the mortgage on female labour supply must be measured in relative terms (relative to household income, excluding female labour income) rather than in absolute ones.

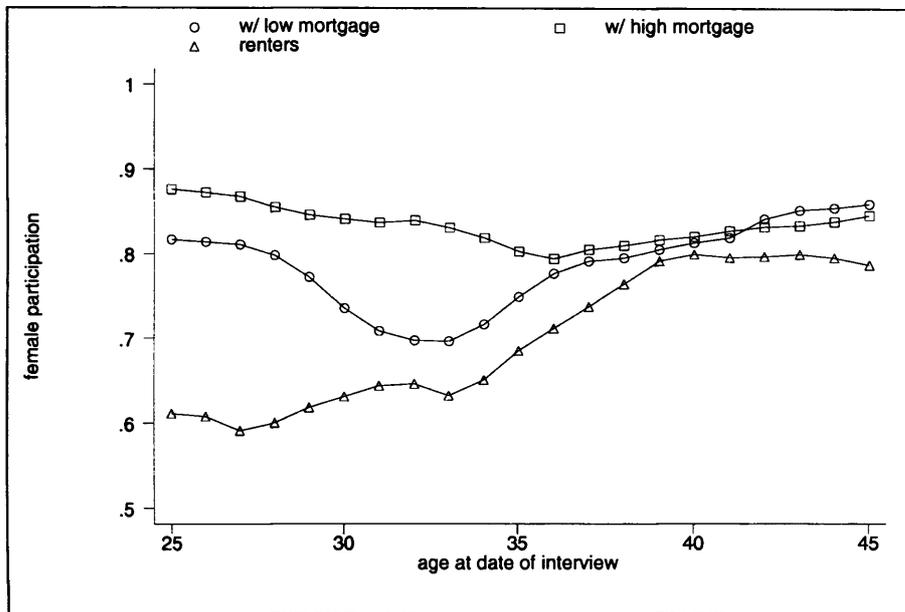


Figure 2.2: Female participation by mortgage repayments

As for the mortgage level, the pattern of female participation for the group with low obligation ratio is U-shaped, with a minimum at the age of 32-33, whereas participation for the group with high obligation ratio stays at a significant higher level and has only a slight decline at the age of 35. This indicates that other household income is positively related to monthly mortgage payment (otherwise the relative and the absolute mortgage payments would have different effects on labour participation) and that there is a positive correlation between mortgage-related allocational inflexibilities (as represented by the obligation ratio) and female labour market participation.

This latter feature is explored further in Figures 2.4 and 2.5 by controlling for the presence of young children in the household. In fact, it might be argued that mortgage decisions and fertility decisions are not separable. It turns out that even when the youngest child is between 0 and 2 year old a high obligation ratio significantly increases the probability that females work more than in cases where the obligation ratio is low, at least after the age of 32. When the youngest child in the household is between 3 and 4 year old, the same result holds in the age range 32-37.

Table 2.3: Female participation by age and mortgage repayments

age	monthly mortgage repayment	
	$\leq 33^{rd}pctile^*$	$> 33^{rd}pctile^*$
25-30	0.800	0.857
	(0.400)	(0.350)
	460	1,261
30-35	0.698	0.827
	(0.460)	(0.379)
	496	1,252
35-40	0.792	0.809
	(0.406)	(0.393)
	552	989
40-45	0.854	0.834
	(0.353)	(0.373)
	583	662
All	0.788	0.833
	(0.408)	(0.373)
	2,091	4,164

Notes: Cells contain, top to bottom, the mean, the standard deviation (in brackets), and group size of female participation in the labour market.

*The 33rd percentile of the distribution of monthly mortgage repayments is 213 GBP

Tables 2.4 and 2.5 report female labour market participation by 5-year-age groups and obligation ratio quartiles (the lower quartile is defined for $or \leq 0.104$ and the upper one for $or \geq 0.219$). In table 2.4 children are not controlled for. It emerges that female labour market participation is higher the higher the obligation ratio for any age group. Particularly in the age range 25-35, labour market participation for the group with the highest obligation ratio is about 20 percentage points higher than for the group with the lowest obligation ratio: female participation is 72% for those aged 25-30 with an obligation ratio below 10.4%, it is 93% if their obligation ratio is above 21.9%, and it is 78.7% and 83.8% if their obligation ratio is in between 10.4% and 21.9%. For females aged 30-35, participation goes from 66% to 86% when switching from the lowest to the highest obligation ratio.

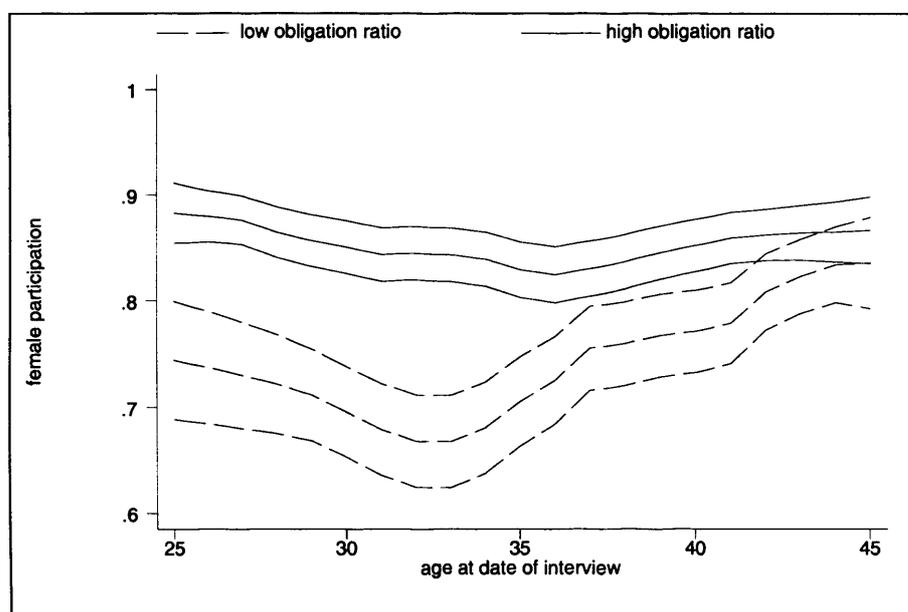


Figure 2.3: Female participation by “obligation ratio” (mortgage repayment/other hh income)

When controlling for the presence of the youngest child in the age range 0-2 (Table 2.5), it is still true that females with the highest obligation ratio have a higher participation rate than those with the lowest obligation ratio. However, for the age group 25-30, the participation pattern is no longer increasing for intermediate levels of the obligation ratio. Whereas those having an obligation ratio below 10.4% have a participation rate of 65%, those having an obligation ratio between 10.4% and 21.9% have a participation rate of, respectively, 59% and 61%. Also when controlling for the presence of the youngest child in the age range 3-4 some non-linearities are present for the age group 25-30. It is noteworthy that participation is around 50% for all the females between 25 and 40 year of age with an obligation ratio below 10.4% whereas it is well above 80% for those with an obligation ratio above 21.9%.

Of course these results are obtained ignoring the panel structure of the data set and hence they can only be taken as a rough indication of a positive relationship between mortgage-related allocational inflexibilities and female labour market participation.

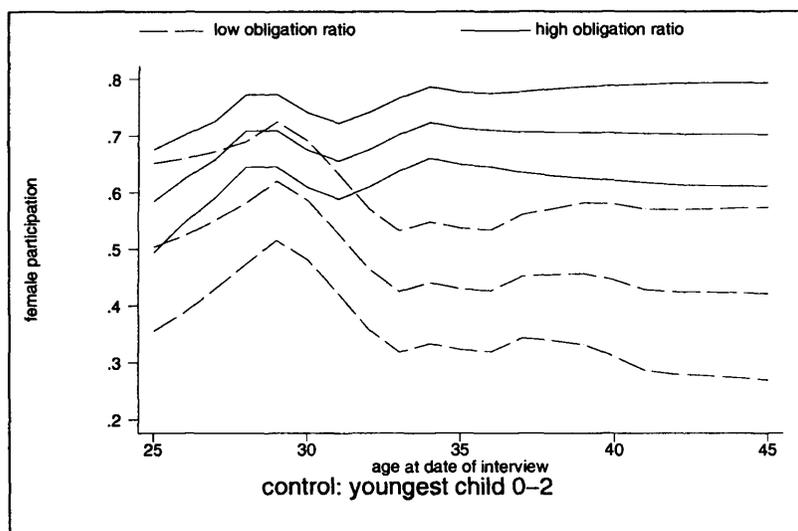


Figure 2.4: Female participation by obligation ratio and age of youngest child 0-2

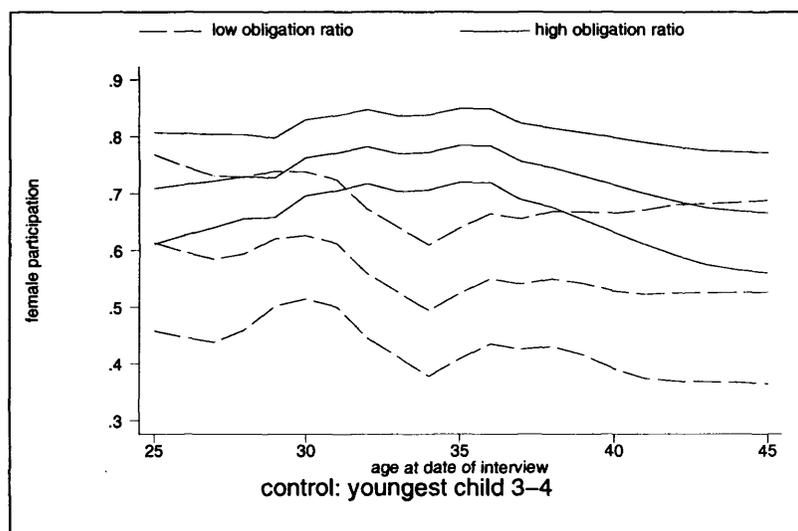


Figure 2.5: Female participation by “obligation ratio” and age of youngest child 3-4

The analysis of this relationship in a regression framework that accounts for unobserved heterogeneity in labour market participation will be carried out in the next sections using the panel structure.

Table 2.4: Female participation by age and level of obligation ratio

age	$or \leq 0.104$	$0.104 < or \leq 0.158$	$0.158 < or \leq 0.219$	$or > 0.219$
25-30	0.721	0.787	0.838	0.926
	(0.450)	(0.410)	(0.368)	(0.260)
	222	390	520	589
30-35	0.658	0.769	0.832	0.865
	(0.475)	(0.422)	(0.375)	(0.342)
	354	428	493	473
35-40	0.733	0.809	0.809	0.896
	(0.443)	(0.393)	(0.393)	(0.305)
	487	404	341	309
40-45	0.820	0.825	0.829	0.953
	(0.384)	(0.381)	(0.378)	(0.212)
	501	342	210	192
All	0.742	0.796	0.829	0.905
	(0.437)	(0.403)	(0.377)	(0.293)
	1,564	1,564	1,564	1,563

Notes: The obligation ratio is defined as the ratio between monthly mortgage payment and other family income (household income minus female's earned income).

Cells contain, top to bottom, the mean, the standard deviation (in brackets), and group size of female participation in the labour market.

2.5 The empirical model

A static female participation equation with unobserved heterogeneity is estimated. Specifically, the form of the estimated equation is:¹⁸

$$P_{it} = 1 \{ \beta \ln y_{it} + \gamma Z_{it} + \delta H_{it} + \theta_1 age_{it} * or_{it} + \theta_2 yoch02_{it} * or_{it} + \alpha_i + \varepsilon_{it} \geq 0 \} \quad (2.10)$$

¹⁸It is worth noticing that as long as mortgage monthly payment is small relative to household's other income, a specification including the the log of other household income and the level of obligation ratio (and no interactions) as explanatory variables represents an approximation (a first order Taylor expansion) of a specification containing only the log of *net* income (other family income net of mortgage payment). In that case, the effect of the mortgage variable would be interpreted in terms of income effect.

Table 2.5: Female participation by age, level of obligation ratio and age of the youngest child

Youngest child 0-2				
age	$or \leq 0.104$	$0.104 < or \leq 0.158$	$0.158 < or \leq 0.219$	$or > 0.219$
25-30	0.651	0.587	0.612	0.737
	(0.481)	(0.495)	(0.489)	(0.442)
30-35	63	92	121	114
	0.381	0.660	0.692	0.667
35-40	(0.490)	(0.476)	(0.463)	(0.474)
	63	94	130	111
40-45	0.385	0.615	0.680	0.775
	(0.493)	(0.496)	(0.471)	(0.423)
All	39	26	50	40
	0.571	0.750	0.571	1.000
All	(0.535)	(0.463)	(0.534)	-
	7	8	7	1
All	0.488	0.627	0.656	0.714
	(0.501)	(0.485)	(0.476)	(0.453)
All	172	220	308	266
	Youngest child 3-4			
age	$or \leq 0.104$	$0.104 < or \leq 0.158$	$0.158 < or \leq 0.219$	$or > 0.219$
25-30	0.513	0.688	0.696	0.886
	(0.507)	(0.467)	(0.464)	(0.321)
30-35	37	64	69	44
	0.515	0.632	0.795	0.829
35-40	(0.504)	(0.485)	(0.406)	(0.379)
	66	87	88	76
40-45	0.489	0.717	0.729	0.860
	(0.505)	(0.455)	(0.449)	(0.351)
All	47	46	48	50
	0.714	0.600	0.500	0.857
All	(0.488)	(0.507)	(0.527)	(0.378)
	7	15	10	7
All	0.515	0.665	0.735	0.853
	(0.501)	(0.473)	(0.442)	(0.355)
All	157	212	215	177

Notes: The obligation ratio (or) is defined as the ratio between monthly mortgage payment and other family income (household income minus female's earned income).

Cells contain, top to bottom, the mean, the standard deviation (in brackets), and group size of female participation in the labour market.

where P is a binary variable indicating whether the female participates in the labour market.

$\ln y$ is the log of other family income and Z is a vector of variables that capture demographic characteristics. Typically it includes (a polynomial in) age and the number of children as well as the age of youngest child. H is a vector of mortgage and housing variables. In the analysis that follows, it includes the obligation ratio (or), that is, the ratio of mortgage monthly payment to other family income and, possibly, the value of the house (in logs, to capture an income effect), the remaining mortgage life, the total mortgage outstanding.

The variable of interest is or and its interactions with age and with the number of children. Interactions with age are meant to capture a different effect of the mortgage-related borrowing constraint at different stages of the life-cycle. In fact, for people who take out a mortgage when they are 25, this interaction captures the effect related to the remaining life of the mortgage. The interaction with a dummy for the presence of young children in the household is meant to control for the possibility that the effect of mortgage commitments is different for people with and without young children.

α_i is the individual specific effect and ε_{it} is the time-varying error term.

2.5.1 Estimation method

The equation that is estimated in this work belongs to the class of non-linear panel data models with individual specific effect, and in particular to the class of discrete choice panel data models, as represented by the following:

$$y_{it} = 1 \{x_{it}\beta + \alpha_i + \varepsilon_{it} \geq 0\} \quad t = 1, 2, \dots, T \quad i = 1, 2, \dots, n$$

As a special case, if it is further assumed that ε_{it} 's are independent and logistically distributed conditional on $\alpha_i, x_{i1}, x_{i2}, \dots, x_{iT}$, it follows that

$$\Pr(y_{it} = 1 | x_{i1}, x_{i2}, \dots, x_{iT}, \alpha_i) = \frac{\exp(x_{it}\beta + \alpha_i)}{1 + \exp(x_{it}\beta + \alpha_i)}$$

as in the standard logit model,¹⁹ with the only difference being the individual

¹⁹Similarly, if ε_{it} 's are independent and normally distributed conditional on $\alpha_i, x_{i1}, x_{i2}, \dots, x_{iT}$, then

$$\Pr(y_{it} = 1 | x_{i1}, x_{i2}, \dots, x_{iT}, \alpha_i) = \Phi(x_{it}\beta + \alpha_i)$$

specific effect, α_i .

Estimating β requires dealing with the individual specific effect, α_i . There are basically two methods for doing this, and they will be surveyed in this section. Essentially, one approach, which defines the so-called *fixed effects* model, imposes no assumptions on the relationship between α_i and the explanatory variables and uses instead a method that “eliminates” the individual specific effect on the basis of the same idea that informs differencing in the linear panel data model with fixed effects. A second approach, which defines the so-called *random effects* model, assumes that both the individual specific effect and the idiosyncratic shock are independent of observable characteristics $(x_{i1}, x_{i2}, \dots, x_{iT})$. The distribution of α_i conditional on $x_{i1}, x_{i2}, \dots, x_{iT}$ is specified parametrically (semiparametrically) so that the individual specific effect is then integrated out.

In the *fixed effects* model the idea is to “eliminate” the individual specific effect by allowing it to be correlated in any form with the explanatory variables. A consistent estimator for β can be obtained by conditional maximum likelihood, where conditioning occurs with respect to the data $(x_{i1}, x_{i2}, \dots, x_{iT})$ and to a sufficient statistic for the fixed effect.²⁰ If the sufficient statistic depends on β , the parameter to be estimated, then the conditional distribution of the data given the sufficient statistic depends on β , and not on α_i , and so β can be estimated by maximum likelihood. The problem with this method is that there is no common sufficient statistic for the non-linear panel data models such that the conditional distribution of the data given the sufficient statistic depends on β . It follows that the method for constructing a likelihood function that does not depend upon the fixed effect is strictly related to the specific non-linear functional form that is chosen as a representation of the data.

One case in which the conditional maximum likelihood method can be successfully applied is the one outlined above, where ε_{it} 's are independent and logistically distributed conditional on $\alpha_i, x_{i1}, x_{i2}, \dots, x_{iT}$ (Conditional ML Logit). Here, the suffi-

where $\Phi(\bullet)$ is the standard normal cumulative distribution.

²⁰A *sufficient statistic* for α_i is a function of the data such that the distribution of the data given the sufficient statistic does not depend on α_i .

cient statistic that “eliminates” the individual specific effect and lets the conditional distribution depend on β is given by $\sum_{t=1}^T y_{it}$, i.e. the number of times that $y_{it} = 1$ for the individual. Hence, in this application a sufficient statistic is given by the number of times that each female participates in the labour market over the observation period (1993-2000).²¹ The drawback of “eliminating” the unobserved fixed effect in this fashion is that also observed fixed effects do not enter the conditional likelihood function and hence cannot be used as controls. In fact, identification requires that right hand side variables vary over time within individuals. Moreover, the explanatory variables must be strictly exogenous, that is, current shocks, ε_{it} , must be uncorrelated with past, present and future values of the the explanatory variable x : $E(\varepsilon_{it}|x_{i1}, x_{i2}, \dots, x_{iT}) = 0, t = 1, 2, \dots, T$.²²

This is a very strong identification assumption. In this work, a violation of strict exogeneity may occur if ε_{it} is correlated with or_{it} , the obligation ratio (i.e. if current mortgage payment is driven by a shock in the female’s participation, such as an unexpected lay-off that makes the household obtain a low mortgage, or if current mortgage payment is correlated with past participation and the “true” model is one with lagged participation but the estimated model ignores the lags). Moreover, assuming also children as strictly exogenous means that labour supply decisions do not affect fertility

²¹Since cases in which the female does not switch between participation and non-participation (i.e. cases in which she participates at every period and cases in which she never participates) do not contribute to the likelihood function, β is in fact estimated on the basis of females that switch status at least once between period 1 and period T . This means that the only relevant information for the conditional distribution is given by the cases in which $\sum_{t=1}^T y_{it} \neq 0, T$.

²²As pointed out by Honoré (2002), this assumption is probably unrealistic in most economic context in which t represents time and particularly in cases in which y_{it} is the outcome of an individual’s optimization problem so that it is expected that y_{it} enters as an explanatory variable in the equation for $y_{i,t+1}$. This is very likely to be the case for females labour force participation decisions.

In a model with lagged participation as explanatory variable, a weaker assumption than strict exogeneity would be predeterminedness. In other words, given a model like the following:

$$y_{it} = \gamma y_{it-1} + \beta_0 x_{it} + \beta_1 x_{it-1} + \alpha_i + \varepsilon_{it}$$

x would be predetermined if $E(\varepsilon_{it}|x_{i1}, x_{i2}, \dots, x_{it}, y_{i1}, y_{i2}, \dots, y_{it-1}) = 0$, i.e. if current shocks were uncorrelated with past values of y and with past and current values of x .

decisions.

As a further point, in these types of models the parameter(s) of interest are identified only up to scale²³ and conditional on the individual specific effect. As pointed out by Honoré (2002), by knowing the coefficient of the explanatory variable in a fixed effects logit model it is possible to judge the relative importance of different time-varying explanatory variables as well as to calculate the effect of the explanatory variables on the probability that the dependent variable takes the value 1 conditional on a particular value for the individual specific effect. However, it is not possible to calculate the average effect of the explanatory variable(s) on the same probability taken across the distribution of the individual specific effect in the population.

The *random effects* model “eliminates” the individual specific effect by assuming a parametric distribution for α_i conditional on $x_{i1}, x_{i2}, \dots, x_{iT}$ so that α_i can be integrated out of the conditional distribution of the data.

For the Probit model, the traditional assumption is that of independence between α_i and $x_{i1}, x_{i2}, \dots, x_{iT}$, although Chamberlain (1984) in fact allows for some correlation between them. Using Wooldridge (2002) definition, the “traditional random effects probit model” assumes that the distribution of the individual specific effect, α_i , conditional on the observables is as follows: $\alpha_i | x_{i1}, x_{i2}, \dots, x_{iT} \sim N(0, \sigma_\alpha^2)$, which implies that α_i and the vector $(x_{i1}, x_{i2}, \dots, x_{iT})$ are independent and that α_i is normally distributed.

As already mentioned, a particular case is that of Chamberlain (1984)’s random effects probit model, where the conditional distribution of the individual specific effect is allowed to depend linearly on the observables (assumption 2 below). More formally, Chamberlain’s random effects probit model is obtained under the following assumptions:

1. $(\varepsilon_{i1}, \varepsilon_{i2}, \dots, \varepsilon_{iT})$ is independent of α_i and of $(x_{i1}, x_{i2}, \dots, x_{iT})$, with a multivariate

²³Arellano (2000) recalls that in the logit case the scale normalization is imposed through the variance of the logistic distribution (and, in general, by the form of the cumulative distribution of $\varepsilon_{it} | x_{i1}, \dots, x_{iT}, \alpha_i$, if known).

normal distribution: $(\varepsilon_{i1}, \varepsilon_{i2}, \dots, \varepsilon_{iT}) \sim N(0, \Sigma)$, $\Sigma = \begin{bmatrix} \sigma_1^2 & & & 0 \\ & \sigma_2^2 & & \\ & & \dots & \\ 0 & & & \sigma_T^2 \end{bmatrix}$

so that $\Pr(y_{it} = 1|x_{i1}, x_{i2}, \dots, x_{iT}, \alpha_i) = \Phi\left(\frac{x_{it}\beta + \alpha_i}{\sigma_i}\right)$, where $\Phi(\cdot)$ is the standard normal distribution;

2. the distribution for the individual specific effect conditional on $(x_{i1}, x_{i2}, \dots, x_{iT})$ is linear in the x s and normally distributed:

$$\alpha_i = \lambda_0 + \lambda_1 x_{i1} + \lambda_2 x_{i2} + \dots + \lambda_T x_{iT} + \nu_i \text{ where } \nu_i \sim N(0, \sigma_\nu^2) \text{ and independent of the } x\text{s.}$$

Given assumptions 1. and 2., the distribution for y_{it} conditional on $(x_{i1}, x_{i2}, \dots, x_{iT})$ has a probit form:

$$\Pr(y_{it} = 1|x_{i1}, x_{i2}, \dots, x_{iT}) = \Phi\left(\frac{x_{it}\beta + \lambda_0 + \lambda_1 x_{i1} + \lambda_2 x_{i2} + \dots + \lambda_T x_{iT}}{(\sigma_i^2 + \sigma_\nu^2)^{1/2}}\right)$$

A more parsimonious version of Chamberlain’s model allows the individual specific effect to depend on the average of x_{it} , $t = 1, 2, \dots, T$, which we call \bar{x}_i , rather than on each single x_{it} , as follows:²⁴

$$\alpha_i = \lambda_0 + \lambda_1 \bar{x}_i + \nu_i \text{ , where } \nu_i \sim N(0, \sigma_\nu^2) \text{ and independent of the } x\text{s.}$$

The distribution for y_{it} conditional on $(x_{i1}, x_{i2}, \dots, x_{iT})$ then takes the following form:

$$\Pr(y_{it} = 1|x_{i1}, x_{i2}, \dots, x_{iT}) = \Phi\left(\frac{x_{it}\beta + \lambda_0 + \lambda_1 \bar{x}_i}{(\sigma_i^2 + \sigma_\nu^2)^{1/2}}\right)$$

Alternatively, a random effects logit model is defined under the assumption that the distribution of the time-varying disturbances conditional on the individual specific effect, $\varepsilon_{it}|\alpha_i$, are independently distributed according to a logistic *cdf* and that the individual specific effect, α_i , conditional on $(x_{i1}, x_{i2}, \dots, x_{iT})$ is normally distributed.

More generally, the joint distribution of the data conditional on observables is defined as follows:

$$\Pr(y_{i1}, y_{i2}, \dots, y_{iT}|x_{i1}, x_{i2}, \dots, x_{iT}) = \int \Pr(y_{i1}, y_{i2}, \dots, y_{iT}|x_{i1}, x_{i2}, \dots, x_{iT}, \alpha_i) dF(\alpha_i|x_{i1}, x_{i2}, \dots, x_{iT})$$

²⁴See Wooldridge (2002), Chapter 15.

and $F(\alpha_i|x_{i1}, x_{i2}, \dots, x_{iT})$, the *cdf* of the individual specific effect conditional on observables, has some specified parametric (or semiparametric) form.

This makes the model fully parametric and so, if the explanatory variables are strictly exogenous, maximum likelihood or methods of moments estimation can be applied.²⁵

In this work both maximum likelihood random-effects logit and maximum likelihood random-effects probit are estimated. The comparison between these two models, which rely on a different parametric specification of the individual specific effect, and of these two models with the conditional maximum likelihood estimator, should give us an indication of the nature of unobserved heterogeneity.

Finally, estimation of a logit model is performed on the pooled sample by ignoring the panel structure of the data. In other words, it is assumed that observations are i.i.d. and follow a logit distribution and so the fact that observations may be correlated over time within individuals due to the presence of unobserved individual heterogeneity is not taken into account. Hence, the individual specific effect is assumed to be zero. Moreover, the error term, ε_{it} , is assumed to be uncorrelated over time and with the explanatory variables, x_{it} .

$$y_{it} = x_{it}\beta + \varepsilon_{it}, \quad \varepsilon_{it} \sim iid$$

$$\Pr(y_{it} = 1|x_{it}) = \Lambda(x_{it}) = \frac{\exp(x_{it}\beta)}{1 + \exp(x_{it}\beta)}$$

The comparison between the pooled logit and the conditional logit estimators should give an indication of the importance of unobserved heterogeneity that can be taken into account by using the panel structure of the data.

2.5.2 Empirical results

As emphasised in the previous section, one of the main identifying assumptions underlying the estimates obtained here is that the explanatory variables are strictly ex-

²⁵Arellano and Honoré (2000) recall that there exists a practical issue in using maximum likelihood, in terms of the speed and the accuracy in the calculation of a multinomial normal cumulative distribution.

ogenous. Hence, both feedback effects from lagged participation to current and future values of the obligation ratio and to current and future fertility decisions (predetermination) and simultaneous decision about the mortgage and participation are assumed away.

The conditional ML logit, by allowing the individual specific effect to depend upon the explanatory variables, and by allowing this dependence to hold in an unspecified manner, is the least restrictive estimation method adopted here. Since only observations where at least one transition in participation has occurred over the observation period contribute to the likelihood function in a conditional logit, from the sample of 6255 observations, 4598, corresponding to 1186 individuals, are dropped and 1657 are used in the estimation.

Estimation results from a conditional logit are reported in Table 2.6, column 1. The variable of interest, the obligation ratio (*or*), has a positive and significant effect on participation for the reference group of 25 year olds. The interaction of *or* with age shows a negative sign, and significant at the 1% level, suggesting that the mortgage-related constraint has a decreasing impact on participation over the life cycle. It takes approximately 9 years (i.e. until the age of 34) to offset the positive effect of the obligation ratio on the participation of a female with no children.

As expected, children have a strong impact on female participation. In particular, having a youngest child aged either 0-2 or 3-4 has a negative and significant effect on participation, as does the presence of any additional child in the household. Whether it is the negative effect of children or the positive effect of mortgage commitments that dominates depends on the stage of the life cycle in which children enter. In fact, since the coefficient on the interaction between the dummy for having the youngest child in age 3-4 is not significant, the only relevant interaction between the obligation ratio and children is the one involving the dummy for having a youngest child aged 0-2. So, for instance, for a 25 year old female in a household with mortgage constraints and with the youngest child between 0 and 2 years of age the net effect on participation is positive, whereas for a female in the same situation but with no mortgage constraints the effect on participation is negative. However, when the youngest child in the household is

Table 2.6: Conditional ML Logit
 Dependent variable: female participation in the labour market

	Conditional ML	ML Random Effects		ML Pooled	
log(other hh income)	-1.547*** (0.391)	-0.943*** (0.205)	-1.237*** (0.222)	-0.557*** (0.084)	-0.576*** (0.092)
age	0.333*** (0.113)	0.229*** (0.072)	0.218*** (0.073)	0.094*** (0.032)	0.085*** (0.033)
age squared	-0.007 (0.004)	-0.008*** (0.003)	-0.007** (0.003)	-0.004*** (0.001)	-0.003** (0.001)
obligation ratio	5.843** (2.505)	8.023*** (2.232)	7.705*** (2.358)	6.856*** (1.199)	7.172*** (1.237)
age*obligation ratio	-0.679*** (0.221)	-0.335* (0.176)	-0.388** (0.189)	-0.117 (0.093)	-0.130 (0.095)
youngest child 0-2	-1.525*** (0.413)	-1.853*** (0.344)	-2.011*** (0.362)	-0.758*** (0.192)	-0.905*** (0.197)
youngest child 3-4	-1.437*** (0.225)	-1.584*** (0.197)	-1.686*** (0.193)	-0.885*** (0.103)	-0.949*** (0.105)
child 0-2*obl. ratio	-3.802* (2.090)	-3.437** (1.758)	-3.387* (1.864)	-3.419*** (1.050)	-2.996*** (1.073)
no. of children	-0.677*** (0.173)	-1.025*** (0.105)	-0.957*** (0.112)	-0.511*** (0.038)	-0.495*** (0.039)
education O level			0.990*** (0.327)		0.224* (0.117)
education A level			1.060*** (0.359)		0.232* (0.134)
higher degree			1.817*** (0.383)		0.460*** (0.139)
intercept		10.977** (1.722)	11.408*** (2.263)	5.534** (0.680)	3.638*** (1.054)
region		-	√	-	√
σ_α		3.249 (0.148)	3.497 (0.191)		
No. of observations	1657	6255	6255	6255	6255

Notes: Standard errors in parentheses.

*** significant at 1% level; ** significant at 5% level; * significant at 10% level

aged 0-2, the positive effect of mortgage commitments on participation dominates the negative one deriving from children only for women 28 or younger. In other words, a 30 year old female with the youngest child aged 0-2 has a higher probability of

participating if she has no mortgage commitments.

Other family income, in logs, has the expected negative sign.

The conditional logit estimation results are compared with estimates from a random effects logit model in Table 2.6, column 2. As opposed to the conditional (fixed effects) logit, all the observations are used (rather than just those in which there is a transition). Identification requires the *cdf* of the idiosyncratic shock conditional on the individual specific effect be logistic and the individual specific effect be normally distributed as well as independent of the explanatory variables. Moreover, as for the conditional logit model, strict exogeneity is assumed throughout. Both a specification that includes controls for education and region of residence and one that omits them are reported. In both specifications, all coefficients retain the same sign as in the conditional logit estimation.²⁶ However, the magnitude of both the obligation ratio and the age/obligation ratio interaction changes substantially, being, respectively, 5.843 and -0.679 according to the conditional logit and 8.023 and -0.335 according to the random effects logit. This casts doubts on the validity of the assumption underlying the random effects model, that the unobserved individual-specific effect be uncorrelated with the explanatory variables. In this case, it would mean assuming that preference towards work be uncorrelated with the obligation ratio. This appears to be very unlikely since the mortgage is given according to total family income, including female labour income (hence, participation).

The random effects probit model (see Table 2.7) produces substantially the same results as the random effects logit, both in terms of significance and in terms of magnitude, once the rescaling factor (of approximately 1.8) is taken into account. This result suggests that the estimation results obtained under the random effects model are not driven by the functional form (probit or logit) assumed for the individual specific effect. What seems to play the major role is the assumption of independence between

²⁶Most of the coefficients on the time-varying explanatory variables do not change noticeably according to whether these “fixed” effects are or are not included. The coefficients that show the biggest change are those for the log of other household income, the obligation ratio and the dummy for the youngest child being 0-2.

the individual specific effect and the explanatory variables that underlies the random effects model but does not need to hold for the fixed effects estimation.

Table 2.7: Probit model
Dependent variable: female participation in the labour market

	ML Random Effects		Chamberlain Random Effects
log(other hh income)	-0.531*** (0.112)	-0.685*** (0.121)	-0.617*** (0.123)
age	0.123*** (0.040)	0.118*** (0.041)	0.142*** (0.043)
age squared	-0.004*** (0.002)	-0.004** (0.002)	-0.004** (0.002)
obligation ratio	4.083*** (1.152)	4.128*** (1.170)	2.488** (1.246)
age*obligation ratio	-0.171* (0.093)	-0.201** (0.095)	-0.260** (0.102)
youngest child 0-2	-1.041*** (0.194)	-1.116*** (0.201)	-1.112*** (0.196)
youngest child 3-4	-0.895*** (0.109)	-0.950*** (0.109)	-0.893*** (0.106)
child 0-2*obl. ratio	-1.862* (0.982)	-1.985 (1.027)	-1.593 (1.012)
no. of children	-0.568*** (0.059)	-0.526*** (0.062)	-0.506*** (0.057)
education O level		0.536*** (0.186)	
education A level		0.607*** (0.204)	
higher degree		1.013*** (0.217)	
intercept	6.095** (0.892)	6.394*** (1.269)	5.782*** (0.930)
region	-	√	-
mean(obligation ratio)			6.620*** (1.324)
mean(oblig.ratio)*mean(age)			-0.008 (0.108)
σ_α	1.838 (0.082)	1.973 (0.129)	1.854 (0.094)
No. of observations	6255	6255	6255

Notes: Standard errors in parentheses.

*** significant at 1% level; ** significant at 5% level; * significant at 10% level

One can test whether or not the data support the assumption of independence between the individual specific effect and the explanatory variables by estimating a “Chamberlain random effects probit”. This involves defining the individual specific effect as a linear function of a vector of explanatory variables (or their average over time) and adding these variables into the probit model. The null hypothesis is that the coefficients on the variables that define the individual specific effect are jointly zero and this is tested against the alternative that there is some correlation taking the form of a conditional normal distribution with linear expectation and constant variance. In Table 2.7, column 3, we report the results obtained by adopting, as conditioning variables for the individual fixed effect, the individual average over time of both the obligation ratio and of the interaction between age and the obligation ratio.²⁷ The test on the joint significance of these two coefficients makes us reject the null hypothesis ($\chi^2(2) = 39.60$), so that the usual random effects probit is rejected in favour of a random effects probit that allows for some correlation between the individual specific effect and the explanatory variables. We take these results as further supporting the choice of a fixed effects logit since this estimator remains consistent whether or not there is any correlation (of whatever form) between the individual effect and the explanatory variables of the model.

Finally, estimation of a logit model is performed on the pooled sample and is reported in table 2.6, column 3. Comparison with the conditional logit estimates is expected to inform on the gain arising from acknowledging the panel structure of the data, i.e. from taking into account that observations may be correlated over time within individuals due to the presence of unobserved individual heterogeneity. As for the random effects models, education and region of residence are controlled for. The

²⁷This is to say that we run a standard random effects probit on our usual set of explanatory variables augmented with the individual mean of the obligation ratio over time and with the product of the individual means of age and of the obligation ratio over time. This corresponds to assuming that the conditional distribution of the individual specific effect has the following form:

$$\alpha_i | x_{i1}, x_{i2}, \dots, x_{iT} \sim N(\lambda_0 + \lambda_1 \bar{\sigma r}_i + \lambda_2 \bar{\sigma r}_i * \overline{age}_i, \sigma_v^2)$$

coefficients on all variables of interest retain the sign that was found by conditional logit estimation method and most of them are significant. However, the interaction between age and the obligation ratio is now insignificant. Moreover, the magnitude of the interaction term changes considerably relative to the conditional logit (from -0.679 to -0.130 , when education and region are controlled for, and -0.117 when they are not). This perhaps suggests once again that unobserved preference towards work is in fact relevant in modelling participation and that it is correlated with age and with mortgage commitments.

2.5.3 Endogeneity

The variable of interest in this analysis, the obligation ratio, is likely to be endogenous for a number of reasons. One way in which the error term of our model could be correlated with the obligation ratio is through reverse causality between the mortgage and female labour market participation. So far, we have assumed the mortgage choice is given and consequently we have analysed the relationship as running from mortgage choice to labour market participation. Due to the existence of the institutional mortgage qualification constraint (whether and how much one can borrow is a function of household labour earnings, hence also of female labour participation prior to taking out the mortgage), it is plausible to think of the causality as running from participation to the mortgage. As long as participation is a fixed individual effect over the period analysed here, this should not be a problem for our estimation. In fact, conditional logit estimation deals with the individual specific effect by allowing it to be correlated with the obligation ratio and the other explanatory variables in any unspecified way. Our estimation would be biased and inconsistent if, instead, today's participation in the labour market were a function of future mortgage payments in a way that is not "fixed".²⁸ This would be the case, for instance, if participation today were driven by changes in the expectation of future mortgage commitments.

²⁸Recall that conditional logit estimation requires that the explanatory variables are strictly exogenous. Focussing on the obligation ratio, this requirement translates into the following condition:

$$\Pr(P_{it} = 1 | or_{i1}, \dots, or_{it}, \dots, or_{iT}; \alpha_i) = \Pr(P_{it} = 1 | or_{it}; \alpha_i)$$

Another potential source of endogeneity lies in simultaneous decisions about the mortgage and labour market participation. Even after controlling for the individual specific effect, it could be the case that the idiosyncratic shock (ε_{it} in equation (2.10)) is correlated with the obligation ratio if, for instance, a common shock hit the obligation ratio and participation in the labour market simultaneously.

In order to test for endogeneity, we use house price data as an excluded variable in a control function framework. House prices are presumably correlated with the obligation ratio but uncorrelated with labour market participation, which makes them a suitable instrument. We use two different data sets for house prices, which we will briefly outline hereafter. A discussion of the method and results of the test will follow.

The data

We first use data on house prices that contains quarterly information on residential property transactions by house type (flat, detached, semi-detached, terraced) at the Postal Sector level between 1995 and 2000.²⁹ In order to match it with the BHPS, we have aggregated it at the Local Authority District level,³⁰ which is the minimum geographical area recorded for each individual. Then, we have taken annual average prices (ratio between annual volume of transactions and annual number of transactions), RPI adjusted, by house type and Local Authority District (LAD). Therefore, the vector of the mortgage variable, the obligation ratio for the years 1995-2000, is instrumented with a vector of house prices for the corresponding years, appropriate for the Local Authority District and the house type of the household. The BHPS sample includes years 1993 and 1994 but house prices are collected only from 1995 onwards. 1500 observations (of the 6255) are missing due to this. A further 555 observations are

²⁹Residential property transaction data were built by Experian, and made available through MIMAS, using information supplied by HM Land Registry.

³⁰Conversion has been done at the MIMAS webpage (<http://convert.mimas.ac.uk>), within the Updated Area Master Files project (based on the ONS All-Fields Postcode Directory (AFPD)). In some cases, the Local Authority Districts as defined in the BHPS did not match with the Census definition as of 1998, particularly for Scotland and Wales. As a consequence, the match is not always 1:1. If more Census districts form a BHPS district, the price index of the latter is the result of a weighted average of the prices of the contributing districts, each of which with equal weights.

dropped due to missing house prices mostly in Scottish LADs.³¹

Since it might be argued that current house prices are not suitable instruments for a mortgage that could have been taken out several years before, we also collect information on house prices at the time the mortgage was taken out (RPI adjusted). Unfortunately, we are not aware of any data set that collects house prices at Local Authority District level as far back in time as mortgages were taken out by households in our sample (the earliest dates back to 1968 although 95 percent of households took out the mortgage in 1980 or after). We then use house prices at regional level.³² Unlike the data at LAD level, house prices are now the average dwelling price for all dwellings. We had to sacrifice geographic and house type detail in order to find earlier data. Since the mortgage is taken out at one point in time, in order to capture the variability over time within individuals we interact the house price measure with current (annual) mortgage interest rates.³³ That is, the vector of obligation ratio between 1993 and 2000 is instrumented with a vector of the interactions between the average house price in the region of residence at the time the mortgage was taken out and mortgage interest rates between 1993 and 2000.³⁴

³¹See Appendix (section 2.7.1) for a detailed list of LADs and corresponding number of missing observations.

³²The geographic units are "Standard Statistical Regions", namely: North, North-West, Yorkshire and the Humberside, East Midlands, West Midlands, East Anglia, London, South-East, South-West, Wales, and Scotland. The source of the data is the Survey of Mortgage Lenders made available through the Office of the Deputy Prime Minister at www.odpm.gov.uk.

³³This is justifiable on the basis that most mortgages in the UK have variable interest rates.

³⁴We should also note that using prices at LAD level carried the cost of losing observations for the years 1993-1994, when prices were not available. This is not the case when using prices at regional level for the time the mortgage was taken out, although some observations are still missing due to not observing either the year the mortgage was taken out or the region of residence.

The test

The test of endogeneity for the mortgage variable (the obligation ratio) in our regression is performed within a control function approach. We write our binary model as follows:

$$\begin{aligned} P_{it} &= 1\{\mathbf{x}_{it}\beta_0 + \alpha_i + \varepsilon_{it} \geq 0\} \\ &= 1\{h(\mathbf{z}_{1it}, \mathbf{y}_{it}) + \alpha_i + \varepsilon_{it} \geq 0\} \end{aligned} \quad (2.11)$$

where α_i is the individual specific effect and ε_{it} is the idiosyncratic shock; $\mathbf{x}_{it} = (\mathbf{z}_{1it}, \mathbf{y}_{it})$, \mathbf{y}_{it} is the endogenous variable (the obligation ratio), and \mathbf{z}_{1it} is a vector of all the other (exogenous) explanatory variables.³⁵ \mathbf{y}_{it} is in turn determined by the exogenous variables \mathbf{z}_{1it} and an “excluded instrument”, \mathbf{z}_{2it} , given by house prices (or by the interaction between house prices and interest rates), as follows:

$$\mathbf{y}_{it} = \mathbf{z}_{it}\pi + \delta_i + \mathbf{u}_{it} \quad (2.12)$$

and

$$\mathbf{z}_{it} = (\mathbf{z}_{1it}, \mathbf{z}_{2it}), \quad (2.13)$$

δ_i and \mathbf{u}_{it} are, respectively, the individual specific effect and the idiosyncratic error term

As pointed out in Blundell and Powell (2001), the control function approach uses estimates of the reduced form error terms u_{it} as “control variables” for the endogeneity of the regressor \mathbf{y}_{it} in the original equation (2.11). Testing the significance of these “control variables” is therefore a test of endogeneity of the regressor \mathbf{y}_{it} .

The control function assumption is that

$$\varepsilon_{it} \perp \mathbf{y}_{it} | \mathbf{u}_{it}, \delta_i, \mathbf{z}_{it} \quad (2.14)$$

In order to integrate \mathbf{u}_{it} out, we therefore need to know the form of the distribution of ε_{it} conditional on \mathbf{u}_{it} . If the joint distribution of ε_{it} and \mathbf{u}_{it} were normal, as in Smith and Blundell (1986), one could write ε_{it} conditional on \mathbf{u}_{it} as linear: $\varepsilon_{it} = \mathbf{u}'_{it}\gamma + \eta_{it}$. In our context, where estimation is performed by conditional logit, we cannot assume joint

³⁵The function h is left generic to allow for interactions between our exogenous and endogenous variables (in particular, age and the obligation ratio).

normality of the two error terms and linearity of their conditional distribution. Instead, we say that ε_{it} is some function of \mathbf{u}_{it} plus an error term (η_{it}), and we approximate this with a second-order Taylor series expansion.

It follows that the conditional expectation of the binary variable P_{it} given the regressors \mathbf{x}_{it} , the fixed effect α_i and the reduced form error terms \mathbf{u}_{it} , now takes the form

$$\begin{aligned} E(P_{it}|\mathbf{x}_{it}, \alpha_i, \mathbf{u}_{it}) &\simeq \Pr(\mathbf{x}_{it}\beta_0 + \alpha_i + \gamma_1\mathbf{u}_{it} + \gamma_2\mathbf{u}_{it}^2 + \eta_{it} \geq 0) & (2.15) \\ &= \Lambda(\mathbf{x}_{it}\beta_0 + \alpha_i + \gamma_1\mathbf{u}_{it} + \gamma_2\mathbf{u}_{it}^2) \end{aligned}$$

and the test of endogeneity is a test of joint significance of the coefficients γ_1 and γ_2 . (2.15) is estimated by a two-stage procedure that allows us to replace \mathbf{u}_{it} and \mathbf{u}_{it}^2 with their estimated counterparts $\hat{\mathbf{u}}_{it}$ and $\hat{\mathbf{u}}_{it}^2$ obtained from the first stage estimation of (2.12).

In practice, the reduced form equation (2.12) is estimated by a within-groups regression of the obligation ratio on the set of exogenous variables (log of other income, quadratic in age, dummies for the youngest child aged 0-2 or 3-4, number of children) and the “excluded instrument”, i.e. the log of the current house prices at LAD level or the log of the interaction between the average house price in the local region at the time the mortgage was taken out and current mortgage interest rates. The results are reported in the top panel of Table 2.8; column 1 and column 2 report, respectively, the outcomes from using the two different sets of instruments. The t-ratio for our instruments are, respectively, 15.93 for the log of house price at LAD level and 16.17 for the log of the interaction of house prices at regional level and mortgage interest rates. A quadratic form of the estimated residuals from the first stage estimation is included in the conditional logit regression as indicated in equation (2.15). A $\chi^2(2)$ test of their joint significance takes the values of 1.6063 and 0.0084, respectively for the case where regional or LAD prices are used. Since they do not appear to be significant, we conclude that we cannot reject the null hypothesis of no endogeneity in our model.

Table 2.8: Conditional ML Logit - Endogeneity: control function technique
 Dependent variable: female participation in the labour market

	control variable 1	control variable 2
<i>First stage: Within-groups regression</i>		
log(other hh income)	-0.1534*** (0.0043)	-0.1587*** (0.0053)
age	0.0135*** (0.0011)	0.0044*** (0.0015)
age squared	-0.0002*** (4.50e-05)	-2.06e-04*** (6.13e-05)
youngest child 0-2	0.0042 (0.0029)	0.0049 (0.0038)
youngest child 3-4	-0.0006 (0.0028)	2.19e-04 (0.0035)
child	-0.0021 (0.0021)	-0.0041 (0.0028)
control variable	0.1412*** (0.0087)	0.0854*** (0.0054)
intercept	-0.6078*** (0.1174)	0.4055*** (0.0645)
No. of observations	5775	4125
<i>Second stage: test of residuals</i>		
1st-step est. residuals	7.511 (6.821)	0.660 (8.816)
(1st-step est. residuals) ²	-18.784 (26.301)	-1.954 (40.140)
$\chi^2(2)$	1.6063	0.0084
No. of observations	1524	867

Notes:

Control variable 1: log(REG house prices at t=mortgage taken out*current interest rates)

Control variable 2: log(LAD house prices, 1995-2000) (original panel: 1993-2000)

Bootstrapped St. Err. (2nd stage) in parentheses (500 bootstrap samples of size n)

*** significant at 1% level; ** significant at 5% level; * significant at 10% level

2.5.4 Sensitivity analysis

As mentioned in section 2.4.2, sensitivity analysis is performed with regard to the definition of participation. Results are reported in table 2.9 and bring the same conclusions

as the definition of participation adopted for the main analysis.

Table 2.9: Conditional ML Logit - Sensitivity analysis: definition of participation
Dependent variable: female participation in the labour market

	Definition 1	Definition 2
	(no. weeks worked) > 0	employment status: in employment
log(other hh income)	-1.871*** (0.459)	-1.129*** (0.347)
age	0.345*** (0.128)	0.347*** (0.097)
age squared	-0.006 (0.005)	-0.006* (0.004)
obligation ratio	5.888** (3.036)	5.561*** (2.091)
age*obligation ratio	-0.611** (0.267)	-0.571*** (0.188)
youngest child 0-2	-1.344*** (0.475)	-1.756*** (0.339)
youngest child 3-4	-1.871*** (0.260)	-1.112*** (0.198)
child 0-2*obl. ratio	-3.630 (2.528)	-3.013* (1.627)
no. of children	-1.044*** (0.200)	-0.505*** (0.154)
No. of observations	1388	2237

Notes: Standard Errors in parentheses

*** significant at 1% level; ** significant at 5% level; * significant at 10% level

Figures 2.6 and 2.7 also document the pattern of female labour market behaviour when hours of work are used rather than participation. Although a declining pattern in hours worked is observed in the age range 25-35 for both outright owners and owners with a mortgage (or for owners with low mortgage and owners with high mortgage), it is still true that a more pronounced dip is observed for the group of outright owners (alternatively, for those with low mortgage).

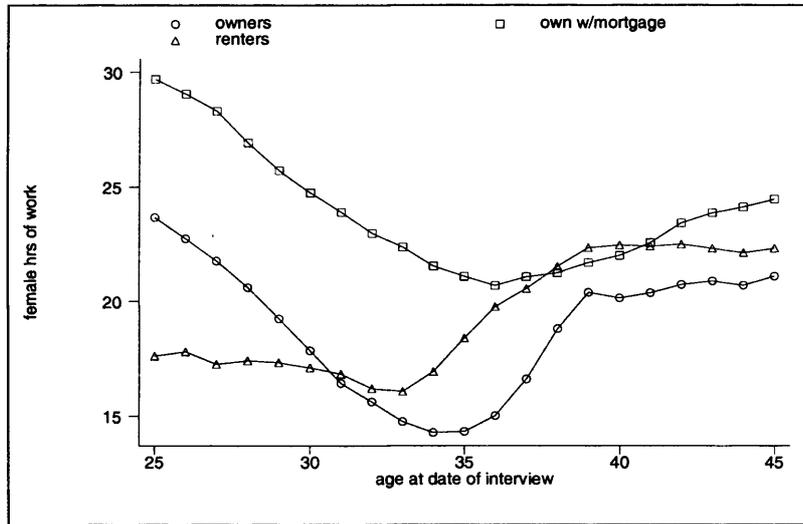


Figure 2.6: Female hours of work by tenure

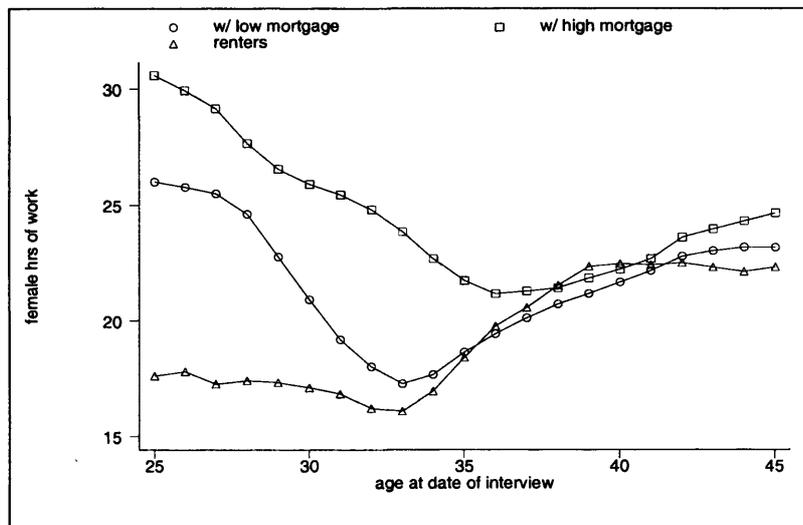


Figure 2.7: Female hours of work by mortgage repayments

Further sensitivity checks are performed by controlling for wealth. One concern regards real assets, and in particular whether it is necessary to control for the value of the house when analysing labour supply in relation to mortgages. In other words, we need to control for the possibility that some households have experienced an increase in

their house value that has relaxed their liquidity constraint. Including a self-reported measure of the value of the house (in logs), however, does not appear to change the results of our analysis.³⁶ As reported in table 2.10, the conditional logit estimates are almost identical to those of the basic model of table 2.6 (column 1) and the house value is not statistically significant.

Table 2.10: Conditional ML Logit - Sensitivity analysis: house value

Dependent variable: female participation in the labour market	
log(other hh income)	-1.501*** (0.396)
log(house value)	0.260 (0.388)
age	0.342*** (0.113)
age squared	-0.006 (0.004)
obligation ratio	5.962** (2.507)
age*obligation ratio	-0.661*** (0.222)
youngest child 0-2	-1.520*** (0.413)
youngest child 3-4	-1.442*** (0.226)
child 0-2*obl. ratio	-3.879* (2.091)
no. of children	-0.662*** (0.174)
No. of observations	1657

Notes: Standard Errors in parentheses

*** significant at 1% level; ** significant at 5% level; * significant at 10% level

Another concern relates to financial assets, in that it is necessary to rule out the possibility that those who appear to be more subject to liquidity constraints (in the

³⁶The same conclusion applies when including the ratio between the value of the house and total household income (excluding female's labour income). Results are not reported for brevity.

form of a higher obligation ratio), do not hold financial assets that could be used as collateral instead of human capital. If that were the case, claiming that having a mortgage makes the household work more, would not be correct as the liquidity constraint would not in fact be binding.

The BHPS collects data on household financial wealth every five years, namely in 1995 and 2000. Savings, investments and debt are reported separately by individuals, who also report whether they hold their assets jointly with someone else, so that a measure of net financial wealth at the tax-unit level can be constructed. Missing information or information for those who only provide bands for their assets are imputed according to the age of the head of the benefit unit, whether either of the adults in the benefit unit have completed any higher education and whether the head of the benefit unit is self-employed. The single components of net wealth are imputed separately.³⁷

As pointed out in Banks, Smith and Wakefield (2002), wealth information across the 1995 and 2000 waves is not fully comparable, due to the different definition of debt, which in 1995 does not include student loans and overdrafts, whereas it does in 2000. We then rely on the two single cross-sections of the data for our analysis. With only this data at hand it is not possible to perform a conditional logit estimation of our model with controls for wealth, which would allow comparability with the baseline model. We therefore investigate the issue at a descriptive level. Table 10 reports the net financial wealth in 1995 and in 2000 for increasing levels of the obligation ratio within each chosen age group. If increasing levels of the obligation ratio were to mean tighter liquidity constraints, we would want those households with higher levels of obligation ratio to hold lower net financial wealth. This is in fact what generally emerges in table 2.11, being violated only for the year 2000 for the top level of obligation ratio (note, however, that the difference between the third and the fourth column is small and that standard deviations are very large).

A final sensitivity check is performed with respect to the timing of taking out the mortgage. So far, all the observations of the selected sample have been used,

³⁷The wealth data used in this chapter has been derived and imputed by Banks, Smith and Wakefield (2002). For details, please refer to their paper.

Table 2.11: Net financial wealth by age and level of obligation ratio

Year 1995				
age	$or \leq 0.104$	$0.104 < or \leq 0.158$	$0.158 < or \leq 0.219$	$or > 0.219$
25-30	18,840 (37,890) 31	4,635 (11,695) 49	1,452 (9,208) 76	2,563 (10,148) 77
30-35	20,445 (50,673) 47	14,564 (47,771) 48	5,234 (17,205) 54	3,852 (13,653) 54
35-40	16,999 (33,245) 56	11,381 (30,278) 49	13,797 (32,122) 36	8,224 (17,226) 38
40-45	33,543 (151,211) 46	11,993 (25,522) 46	10,156 (30,770) 26	9,991 (44,620) 21
All	22,444 (83,860) 180	10,602 (31,519) 192	6,009 (21,228) 192	4,883 (19,238) 190
Year 2000				
age	$or \leq 0.104$	$0.104 < or \leq 0.158$	$0.158 < or \leq 0.219$	$or > 0.219$
25-30	13,563 (34,591) 27	1,520 (13,808) 39	1,376 (12,719) 58	2,285 (12,836) 80
30-35	18,346 (35,157) 31	9,895 (24,040) 53	5,226 (21,800) 66	7,360 (43,853) 68
35-40	9,721 (18,909) 61	11,900 (50,624) 52	8,886 (59,756) 48	9,464 (36,987) 48
40-45	24,346 (60,353) 68	18,181 (30,431) 46	2,976 (15,543) 34	10,169 (35,010) 31
All	17,024 (42,815) 187	10,731 (33,778) 190	4,623 (32,573) 190	6,400 (32,941) 227

Notes: Net financial wealth is defined as (savings+investments-debt and does not include housing. Savings, investments and wealth have been imputed separately when missing.

Cells contain, top to bottom, the mean, the standard deviation (in brackets), and group size of net financial wealth at individual level.

regardless of when the mortgage was taken out. The interaction term between the obligation ratio and age of course allows for the mortgage-related constraint to vary

over time. However, it is being investigated whether the qualitative results hold for households that have taken out a mortgage recently. In particular, also in relation to the use of house prices between 1995 and 2000 in order to test for endogeneity (see previous section), we perform our estimation on the sub-sample that took out the mortgage between 1995 and 2000 and for whom the house price is not missing, having a final sample of 1318 observations. Conditional logit estimation results are reported in table 2.12. Because of the small number of observations where a change in outcome is observed, the sample used in the estimation is made of 192 observations, which justifies obtaining very few significant coefficients. The qualitative results, however, are unchanged and both the obligation ratio (for 25 year old) and the interaction between the obligation ratio and having the youngest child between 0 and 2 remain significant and stronger in magnitude. Consistently with the nature of the sub-sample of those who have taken mortgages out recently, it would take longer to offset the positive effect of mortgage on the participation of a female with no children (16 years) relative to the baseline case (9 years).

2.6 Conclusion

This chapter contains an analysis of whether female labour supply is affected by mortgage commitments. It employs panel data techniques and uses the British Household Panel Study (waves 3-10). The sample used includes any woman aged 25-45, in a couple, and whose husband works full-time; it excludes self-employed individuals and renters. It is found that mortgage commitments, as captured by the ratio between monthly mortgage payment and household income excluding female's earned income, have a positive effect on female participation. However, the negative effect on female participation of having a young child is very strong and the combined effect of children and mortgage commitments on participation can stay negative.

As opposed to previous studies that have used cross-sectional data, the key advantage of the panel structure of this dataset is that it allows estimation of a static model that controls for unobserved heterogeneity. This is done by means of conditional logit

Table 2.12: Conditional ML Logit - Sensitivity analysis: mortgage 1995-2000^(*)

Dependent variable: female participation in the labour market	
log(other hh income)	-0.824 (1.028)
age	0.667 (0.429)
age squared	-0.020 (0.016)
obligation ratio	16.564* (9.937)
age*obligation ratio	-1.049 (0.898)
youngest child 0-2	0.734 (1.418)
youngest child 3-4	-0.574 (0.658)
child 0-2*obl. ratio	-10.363* (6.274)
no. of children	-1.143 (0.848)
No. of observations	192

Notes: ^(*) The sample is restricted to those who took out the mortgage between 1995 and 2000

*** significant at 1% level; ** significant at 5% level; * significant at 10% level

and random effects logit estimation. The conditional logit estimation method takes care of the individual specific component of the error term by allowing it to be correlated in any unspecified manner with the explanatory variables. Hence, preference towards work is allowed to be correlated with mortgage decisions.

Endogeneity of the mortgage variable is a potential issue; it could emerge due, for example, to correlation between mortgage decisions and transitory shocks or to reverse causality (from participation to the mortgage). A test of endogeneity is performed in a control function framework using house prices as control variables. We use, respectively, contemporaneous house prices by house type at Local Authority District level, and the interaction between house prices at regional level at the time the mortgage was taken out and current mortgage interest rates. The null hypothesis of no endogeneity

cannot be rejected.

Sensitivity checks are performed with respect to real and to financial assets. The value of the house, when controlled for, does not appear to be significant and does not change the results of the analysis. A first look at financial assets, which are measured in 1995 and in 2000, seems to rule out the possibility that many people with higher mortgage commitments could hold financial assets to use as collateral.

2.7 Appendix to chapter 2

2.7.1 House price data at LAD level: missings

Missing house prices at LAD level occur for the years 1993-1994 and mostly for Scottish LADs in subsequent years. In particular, the distribution of missing house prices for the years 1995-2000 is as follows:

Local Authority District	no. of observations in our sample for LADs where house prices are missing
E Yorkshire; Holderness	3
Thamesdown	67
Blaenau Gwent; Islywn	56
Edinburgh City	106
West Lothian	10
Clackmannan; Stirling	1
Falkirk	21
Annadale; Nithsdale; Stewarty; Wigtown	5
Dunfermline	6
Kirkcaldy; NE Fife	7
Aberdeen City	20
Banff & Buchan; Moray	6
Gordon; Kincardine & Deeside	7
Bearsden ; Clydebank; Strathkelvin	63
Cumbernauld & Kilsyth; Monklands	11
Clydesdale; Cumnock Doon ; Kyle Carrick	31
Cunninghame	46
Eastwood; Kilmarnock & Loudon	19
Glasgow City	7
Renfrew	2
Angus; Perth & Kinross	18
Dundee City	43

Chapter 3

Explaining Life-Cycle Profiles of Home-Ownership and Labour Supply

3.1 Introduction

For many households, housing wealth comprises a large fraction of total household wealth. For example, in the UK in 2000, housing wealth made up 80% of the non-pension wealth of households in the British Household Panel Survey. However, owning a home is often associated with large mortgage debt. By contrast, debt among non-home owners is small. Such differences in debt are likely to affect liquid savings, labour supply and non-durable consumption. The main aim of this chapter is to explore these interactions, and in particular to focus on how labour supply, home ownership and debt interact.

The importance of labour supply in this context has been a particularly neglected area of research. Labour supply is likely to matter for a number of reasons: banks and other mortgage lenders have explicit policies that relate debt to current household earnings, and so greater labour supply increases the availability of debt. Further, the ability to vary labour supply once a household has a mortgage gives extra protection

against interest rate uncertainty. At an empirical level, chapter 2 has studied some of these relationships, showing that women in households with greater mortgage commitments are more likely to participate in the labour force. Clearly labour supply, debt and home-ownership are jointly determined and we cannot treat the extent of the mortgage as exogenous. For example, individuals may have taken out greater debt knowing they were going to be working or they might have had to return to work because their debt payments were unexpectedly large. Our motivation in the current chapter is to try to understand these effects by building a structural life-cycle model of labour supply and housing choices. A further motivation is to understand the timing of house purchases over the life-cycle. This will again be affected by labour supply choices.

In the structural model, liquidity constraints lead households who have not yet purchased a house to work more to increase savings and to bring forward the date at which they can afford to buy a house. Further, working more can help relax any income related constraint on mortgage borrowing. We calibrate our model to match the level of home ownership in the data and then simulate the model to address a number of questions. First, we show that the labour supply of home owners is greater than that of non-owners. Second, we show that changes to the financial environment change home-ownership and labour supply patterns: reducing downpayment requirements leads to households buying their homes earlier in the life-cycle and also to lower labour supply; increasing the multiple of household income that households can borrow against leads to earlier home ownership, with an effect on labour supply at the time of house purchase.

We also analyse the welfare cost of wage uncertainty, given the importance of the endogenous borrowing constraint in determining home-ownership and labour supply choices. We first compute compensating variations in assets between scenarios with different wage variance and then translate them into changes in consumption, labour supply and home-ownership overall in the life-cycle.

The rest of the chapter is structured as follows. Section 3.2 reviews the main literature that is related to our contribution. Section 3.3 reports life-cycle patterns of

home ownership and labour supply in the data. Section 3.4 presents the structural model. Section 3.5 discusses the calibration strategy and shows the calibrated solution. Section 3.6 addresses the question of how labour supply and home ownership interact. Section 3.7 addresses the question of how changes to the capital market affect home ownership and labour supply and section 3.8 provides a welfare analysis of wage variance. Section 3.9 concludes.

3.2 Related literature

This chapter analyses life-cycle decisions concerning consumption and savings, labour supply and homeownership when households are subject to borrowing constraints and face transaction costs of trading the durable good (housing). It is related to several recent papers that have incorporated durable consumption into life-cycle models. One strand of research has analysed households decisions over portfolio allocations and/or savings levels when a durable good is added to a standard life cycle setup. Lustig and Nieuwerburgh (2002) study (non-durable) consumption growth and aggregate stock returns in a general equilibrium economy in which housing provides both utility services and collateral services, and in which its price is determined endogenously. Moving to a life-cycle framework with exogenous house prices, Flavin and Yamashita (2002) generalise the Grossman and Laroque (1990) model of optimal consumption and portfolio allocation when utility derives from a durable good the stock of which can be adjusted only at cost, by allowing for both a non-durable and a durable good as arguments of the utility function. They show the implications of this addition for the optimal portfolio composition over the life-cycle. In a similar set-up, Flavin and Nagakawa (2004) study the implications of the presence of housing for the life-cycle profile of non-durable consumption. Fernandez-Villaverde and Krueger (2002) study the life-cycle profiles of the durable and non-durable goods, as well as of asset allocation, when households can use the durable good as collateral, but contrary to Lustig and Nieuwerburgh (2002) the house price is exogenous.

Our contribution to this literature consists is twofold. First, we want to model household behaviour under standard explicit institutional borrowing constraints associated to the purchase of a house. Therefore, we consider an earnings-related borrowing constraint in addition to the downpayment constraint embedded in the previous models. Second, and in relation to the first point, we want to allow households choose labour supply, as well as non-durable consumption and savings, and home-ownership status. In the same way that durable consumption has been found to generate spill-over effects on non-durable consumption and asset allocation, one also would expect labour supply to be affected, particularly given the explicit earnings-related constraint. In this respect, there is empirical evidence, discussed in chapter 2, of a correlation between labour supply behaviour and mortgage commitments but, as far as we know, there is no attempt to model these features in a structural model of life-cycle behaviour.

Since the labour supply component is a key feature of our analysis, this chapter also relates to the literature on life-cycle choices of consumption, labour supply and savings, where durable consumption is not modelled explicitly. Low (2005) finds that flexibility in hours worked changes life-cycle profiles of assets, in that borrowing when young is greater and saving when middle-aged is greater than when labour supply is fixed. Moreover, he also finds that uncertainty causes individuals to work longer hours and to consume less when young. We will be able to analyse labour supply behaviour when the introduction of a durable good brings with it an explicit borrowing constraint that is a function of hours worked, as well as an additional dimension of uncertainty.

3.3 Observed Patterns of Home Ownership

The aim of this section is to highlight the main facts about home ownership that we want to understand using our model. We begin by showing how home ownership rates vary over the life-cycle and across education groups, and then consider at what age individuals are likely to transit from renting to home ownership.

In order to plot proportions of home owners at different ages, we use Family Expenditure Survey (FES) data on individuals born between 1910 and 1959, and construct

a pseudo-panel data set based on 5-year date of birth intervals. Figures 3.1, 3.2 and 3.3 split the cohorts between high and low education groups (where high education means at least some college education) and report how home ownership has varied with age within each cohort and education group. We report the proportion owning their home and the proportion either owning or in government housing. There appear to be substantial increases in home-ownership by age, particularly for the low educated. However, part of this increase is a year effect caused by the sale to tenants of some government owned housing that took place in the 1980s. Clear evidence of the effect this policy can be found in Figure 3.2, where we plot, for the low education group, home ownership over time for all the cohorts born in the 1930s, 1940s and 1950s. All the cohorts experienced a parallel increase in the proportion of home owners during the 1980s, suggesting that this increase was not just a life-cycle effect. The same graph for the high education group shows instead an almost flat profile, indicating that this group was not affected by the so-called “right-to-buy” phenomenon of the 1980s. A further check of the year effect can be found in Figure 3.1, where we include the total proportion either owning or in government housing. For the low educated, this combined proportion varies little by age or by cohort.¹ For the high educated (Figure 3.3) there remains some age effect: the proportion owning their homes (or owning and rent - not shown) increases with age up to about age 35. There is relatively little decline in home-ownership as households age, but this may mask transitions to houses of smaller size.

In our model we will want to capture not only the likelihood of homeownership at different ages, but also the fact that home ownership is a very persistent state. This persistence is highlighted in Figure 3.4 where we report estimated probabilities of moving from renting to owning, that are predicted from the regression reported in Table 3.1.² The logit regression is based on BHPS data for 1991-2002 and models the likelihood of transition to owning in period t for those renting in period $t - 1$.

¹The cohorts plotted in Figures 3.1 and 3.3 are, left to right, 1950-54, 1940-44, 1930-34, 1920-24, 1910-14

²These predictions are for an individual living in a couple, with no children and in a conurbation outside the South-East of England, in the year 2002.



Figure 3.1: Low education: Home ownership rates by cohort

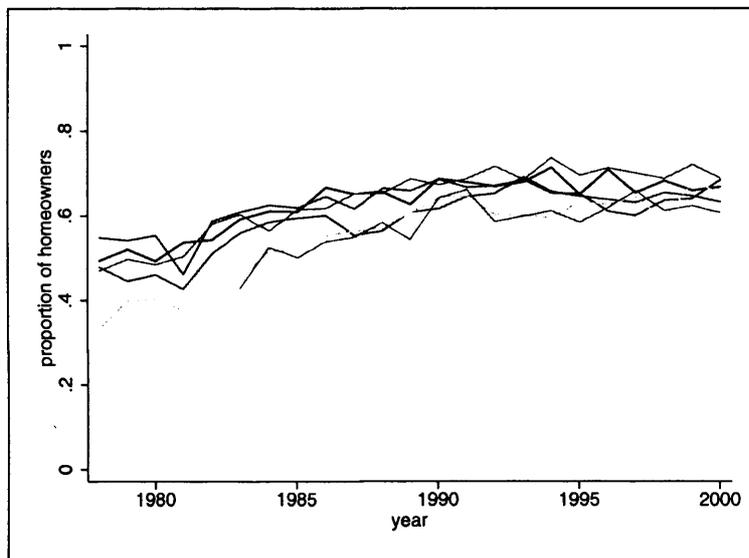


Figure 3.2: Low education: Home ownership rates by year

Transitions from renting to owning peak in the late twenties and the high education group shows a higher transition rate, which is a feature consistent with their higher home ownership rate.

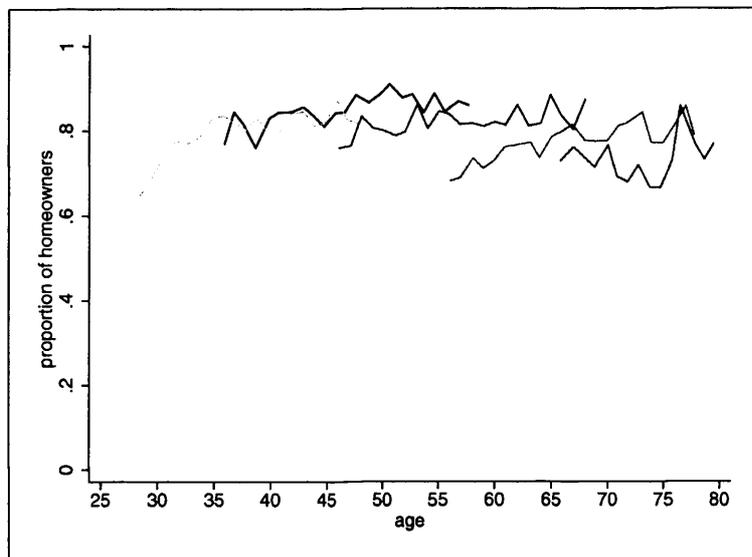


Figure 3.3: High education: Home ownership rates by cohort

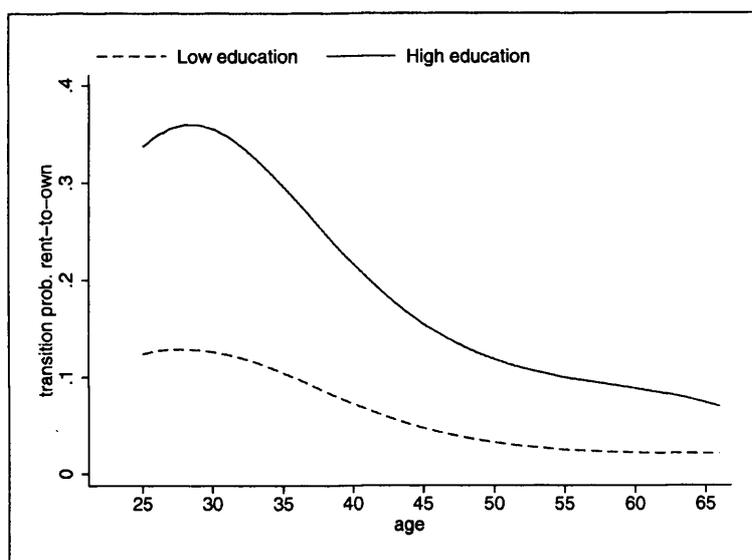


Figure 3.4: Transition probabilities over the life-cycle

We want to use this evidence on home ownership over the life-cycle to calibrate the model that we will use to analyse the interaction between housing and labour

Table 3.1: Logit Regressions for Becoming a Home Owner

	<i>High Education</i>		<i>Low Education</i>	
<i>Age</i>	0.22	(1.96)	0.059	(3.08)
<i>Age</i> ²	-0.0075	(-1.94)	-0.0019	(-3.13)
<i>Age</i> ³	$1.06 * 10^{-4}$	(1.85)	$2.52 * 10^{-5}$	(3.03)
<i>Age</i> ⁴	$-5.47 * 10^{-7}$	(-1.77)	$1.18 * 10^{-7}$	(-2.88)
Couple	0.069	(4.67)	0.023	(5.76)
Children	0.052	(0.31)	0.0093	(0.24)
No. of Children	-0.19	(-2.15)	-0.056	(-2.71)
Children * Age	-0.0033	(-0.75)	$5.57 * 10^{-4}$	(-0.53)
No. of Children * Age	0.0052	(2.18)	0.0012	(2.03)
Conurbation	-0.024	(-1.59)	-0.011	(-2.79)
London/SE	$-3.8 * 10^{-4}$	(-0.03)	0.011	(2.56)
Number of observations	2492		9963	
Mean Probability: Rent to Own	0.16		0.050	
<i>Pseudo R</i> ²	0.0695		0.0746	
<i>LR</i> - χ^2	152.33		296.17	

Notes: Marginal effects (z-statistic of underlying coefficient). Dependent variable = 1 if the household moves from renting to home-ownership, 0 otherwise. The sample is households renting in t-1. Year dummies are also included in the regression.

The measure of *Pseudo R*² is $1 - L_1/L_0$.

supply choices. The evidence reported in chapter 2 suggests that women in households which own their homes are more likely to participate, and that this effect is especially strong amongst those with greater mortgage debt. Participation is likely to be greater for those with greater debt firstly because of a wealth effect and secondly because mortgage repayments can be seen as committed expenditure. Of course, participation and debt are joint decisions and we cannot treat the extent of debt as exogenous. This is the reason we resort to calibration of a structural model of housing and labour

supply: it enables us to model explicitly the labour supply decision when households are making a decision about housing tenure.

3.4 Model of housing and labour supply

A household lives T periods. In every period $t \leq T$, the household maximises utility by choosing consumption $c_t \in R_+$, housing $h_t \in \{0, 1\}$ and the fraction of household time devoted to leisure³ $l_t \in [0, 1]$. The household value function in period t is given by

$$V_t(A_t, h_{t-1}, p_t, w_t, t) = \max_{\{c_t, h_t, l_t\}} u(c_t, h_t, l_t) + \beta EV_{t+1}(A_{t+1}, h_t, p_{t+1}, w_{t+1}, t+1) \quad (3.1)$$

subject to

$$A_{t+1} = \begin{cases} (1 + r_{t+1}) [A_t + p_t h_{t-1} (1 - F) (I\{\delta = 1\} + I\{h_t = 0, \delta = 0\}) + w_t (L - l_t) - c_t], & \text{if } h_{t-1} = 1 \\ (1 + r_{t+1}) [A_t - p_t h_t (1 + F) + w_t (L - l_t) - c_t], & \text{if } h_{t-1} = 0 \end{cases} \quad (3.2)$$

where A_t is the start of period asset stock and r_{t+1} is the interest rate on the liquid asset; δ is an indicator of having to sell the house, which occurs with an exogenous probability; 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; the household has a time endowment L which is allocated to leisure l_t or work at wage w_t per unit of time.

For households which are home-owners at the start of period t , liquid resources in a particular period are given by liquid assets at the start of the period plus the value of the house if the household has been forced to sell ($\delta = 1$) or the household chose to sell ($\delta = 0, h_t = 0$), plus earnings from working. For households which are not home-owners at the start of period t , liquid resources in a particular period is given by

³Our current model assumes a single household labor supply choice with one wage. This is clearly restrictive and was necessitated by computational feasibility.

liquid assets at the start of the period minus the price of a house plus the proportional cost of buying if the household chooses to buy in period t , plus earnings from working.

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

$$A_t \geq 0 \quad (3.3)$$

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 allow for two explicit constraints. The first is a constraint on the fraction of the value of a house that a household is able to borrow at the time of purchase or when remortgaging. This implies a constraint on end of period assets, \tilde{A}_t , in the period that the new mortgage is agreed, of:

$$\tilde{A}_t \geq -\lambda_h p_t h_t \quad (3.4)$$

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

The second constraint is on the debt to income ratio. Individuals in period t are not able to borrow more than a multiple λ_y of their earnings in that period. This constraint is imposed whether or not house status changes in period t .

$$\tilde{A}_t \geq -\lambda_y w_t (L - l_t) \quad (3.5)$$

These constraints are on the stock of debt rather than only on new borrowing and are taken to be exogenous. There is no mortgage repayment scheme in our model, and households only have to repay mortgage interest when they sell their house. This means that any repayment that takes place before selling the house is the result of an optimal saving decision and not of an exogenous constraint on the flow of debt. The main reason why we do not include a mortgage repayment scheme is computational

tractability, as this would introduce additional states to our model. On the other hand, a mortgage repayment scheme would be a potential source of constrained behaviour in labour supply, and therefore relevant to our model. We instead choose to impose the income-related borrowing constraint at every period in which a household is a home owner as an alternative, less computationally costly, way of relating household labour supply to home ownership in periods after the time of house purchase. The decision to create this link is partly supported by the empirical evidence found in chapter 2, where female labour market participation appears to be related to the existence of an income-related constraint from the mortgage, at least for part of the life-cycle. Moreover, Fortin (1995) considered the assumption that households were subject to income-related constraints even beyond the date of house purchase, to be an appropriate one in contexts in which households could remortgage frequently.

In order to complete the picture of the model that we have adopted, we will now describe the form of within-period utility and the type of stochastic processes assumed for our analysis.

Utility function

The within period utility function has a Cobb-Douglas form between consumption and leisure and this is nested within a CRRA utility function. This is augmented by a term reflecting the value of home-ownership:

$$u(c_t, h_t, l_t) = e^{\theta h_t} \frac{[c_t^\eta l_t^{1-\eta}]^{1-\gamma}}{1-\gamma}. \quad (3.6)$$

The parameter θ is a housing preference parameter which determines the utility that households obtain from owning a house versus renting it, and is calibrated in our model. Given a choice of $\gamma > 1$, if $\theta < 0$ home owners obtain additional utility, measured by $e^{\theta h_t}$, relative to renters. Moreover, the marginal utility of consumption (and of leisure) is lower when owning than when not owning.

We can now summarise the main properties of housing in our model. Housing is considered as a good that gives additional utility (under appropriate parameter specification) to home owners relative to renters. Moreover, housing serves as collateral,

i.e. it allows borrowing, although subject to borrowing constraints related to its price and to household income. As an asset, housing has higher expected return and variance than financial assets. However, it can only be bought as a discrete unit (a house) and there are transaction costs of buying and selling housing, as well as an exogenous probability of having to sell the house that one owns (to move to another area, for example).

Stochastic processes

Households face three dimensions of uncertainty: over wages, over house prices and over the interest rate.

The house price follows an AR(1) process:

$$\ln p_t = d_t + \phi \ln p_{t-1} + \varepsilon_t \quad \varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \quad (3.7)$$

where d_t is the deterministic trend. We assume that house price risk is aggregate. The interest rate on liquid assets or debt is also an aggregate rate and follows an i.i.d process:

$$r_t = \bar{r} + \nu_t \quad \nu_t \sim N(0, \sigma_\nu^2) \quad (3.8)$$

The wage process is idiosyncratic and follows a random walk:

$$\ln w_t = a_t + v_t \quad \text{where } v_t = v_{t-1} + \xi_t, \quad \xi_t \sim N\left(-\frac{\sigma_\xi^2}{2}, \sigma_\xi^2\right) \quad (3.9)$$

where a_t is the deterministic growth in wages and has a hump shape over the life-cycle (see section 3.5.1).

We have made several simplifying assumptions in setting up our model and yet we have to deal with a problem that is computationally burdensome, involving 5 states, 3 of which are continuous variables, and 3 dimensions of uncertainty. We provide details of the computational method in Appendix 3.10.1.

3.5 Calibration

We begin this section by specifying exogenous parameters. We then report the results of the calibration and the simulated life-cycle profiles. We end the section by showing some implications of our baseline parameterisation.

Parameter values are summarised in table 3.2. For the inputs into the calibrated model, we need to use data on wages, the house price process and the interest rate on the liquid asset. For the calibration, we need life-cycle profiles of home-ownership status. To assess the calibration, we use data on labour supply choices over the life-cycle and labour supply conditional on being in debt, as well as transition probabilities between home-ownership states over the life-cycle.

Table 3.2: Parameter Values

Parameter	Value
<i>Calibrated Parameters</i>	
F	0.03
θ	-0.03
<i>Exogenous Parameters</i>	
λ_y	3.0
λ_h	0.9
γ	1.58
η	0.74
ϕ	0.94
σ_ε	0.089
σ_ξ	0.122
σ_ν	0.046
\bar{r}	0.015

3.5.1 Exogenous parameters

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 the UK institutional features. Households can borrow whichever amount is lower between three times household earnings ($\lambda_y = 3$) and 90% of the house price ($\lambda_h = 0.9$).

Utility function

The preference parameters η and γ in the utility function are set to match estimated elasticities in the data: the consumption elasticity of intertemporal substitution is set at 0.7 (from Attanasio and Weber, 1995) and the hours of work elasticity of intertemporal substitution is set at 0.3 (from Pistaferri, 2003). These numbers correspond to a curvature $\gamma = 1.58$ and to $\eta = 0.74$ for our within period utility function.

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 (Figure 3.5).⁴ We estimate an AR(1) process, with linear trend (equation 3.7), 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 (ϕ) of 0.94 and a variance of the shock (σ_ϵ) equal to 0.008. In the model we treat house price shocks as aggregate and we set particular realisations of the process to match realised house prices for the cohort born in the late 1940s. We assume households exist from age 22 to 66.⁵

⁴We use the series reporting average house prices for all dwellings.

⁵Having house price data only up to the year 2000, we can use actual realisations only up to age 53. For age 54-66 we draw shocks randomly from the distribution of the error term.

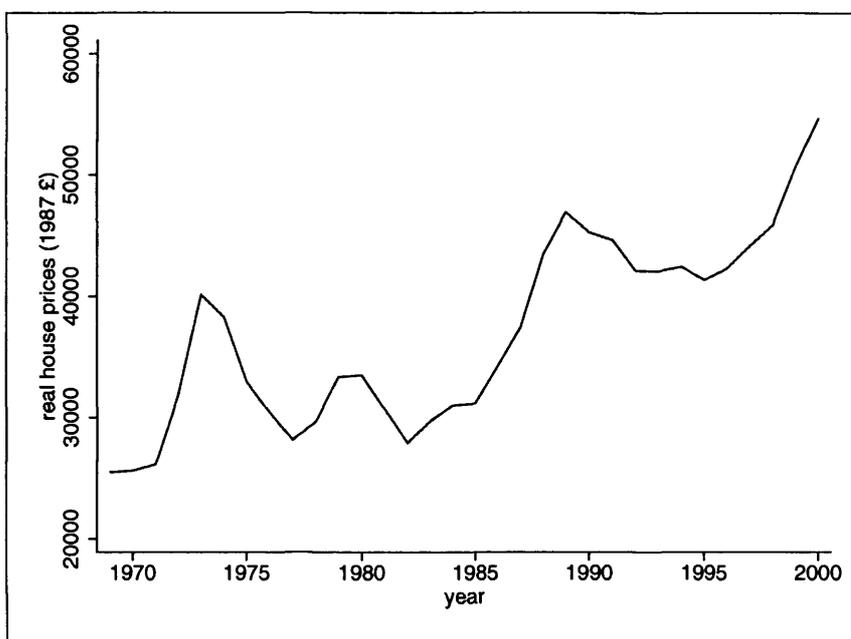


Figure 3.5: UK House price series

Wage process

For the wage process, we estimate the deterministic growth component from the Family Expenditure Survey (FES) and the wage rate uncertainty from the British Household Panel Survey (BHPS). The reason why we use two separate data sets is that the BHPS is only available starting in year 1991 and this would not allow us to obtain the wage growth for the 1940s cohort we calibrate our model to.

FES data are from the years 1978-2000, and are again organised into five year cohorts based on the date of birth intervals. We only base our estimates on males, working more than 10 hours a week, and obtain a measure of growth in median earnings by education group.⁶ In particular, we take median individual real earnings by cohort and education, and estimate, separately for each education group, a regression of the logarithm of median earnings on a quadratic term in the median age of the cell,

⁶Selecting males that work more than 10 hours per week means that average hours in our sample vary very little over the life-cycle and so this helps to mitigate the bias in estimating the growth of earnings instead of that of wages.

and cohort dummies. Defining the deterministic component a_t in equation (3.9) as $a_t = a_{1t}t + a_{2t}t^2$, we find that a_{1t} is 0.032 and 0.058, and a_{2t} is -0.0004 and -0.0008, respectively for the low and high education groups.

The variance of the log-wage process, σ_ξ , is estimated using BHPS data, years 1991-2002. We select male employees aged between 25 and 59, and we obtain a measure of real hourly wages as the ratio between the normal gross monthly pay (in real terms) and hours worked, converted from a weekly basis to a monthly one. Following the estimation procedure adopted in Blundell, Pistaferri and Preston (2004), we obtain $\sigma_\xi=0.018$ for both education groups (see Appendix 3.10.2).

Interest rate process

For interest rates we use the Bank of England base rates, years 1968-1996, and assume that interest rate shocks are aggregate in the economy. The average interest rate (\bar{r}) is estimated at 1.5 percent and the variance of the shock, σ_ν^2 , is 0.0021.

3.5.2 Calibrated parameters

We select the preference parameter for housing, θ ,⁷ and the fixed cost of housing, F , by calibration to match average life-cycle home-ownership rates between age 26 and 60. We assume that these parameters are common across the two education groups. We obtain $\theta = -0.03$ and $F = 0.03$ (i.e. the transaction cost for buying and selling is 3% of the house price). Figure 3.6 contains the simulated home-ownership profiles and Table 3.3 compares home-ownership rates predicted by the model to those observed in the data.⁸ The spike and subsequent dip in home-ownership around age 25 that can be seen in Figure 3.6, happen at times in which the aggregate house price is, respectively, at its lowest (year 1971) and highest (1973) levels of the 1970s. Since buying a house is relatively more costly, in terms of earnings, for the low education group, the biggest difference in home-ownership rate between the two education groups

⁷Since the estimates of the intertemporal elasticities are taken from papers which do not condition on home-ownership, there is a possible bias.

⁸Data come from the years 1991-2000, as years prior to 1991 are affected by the large-scale selling off of local authority housing.

appears to take place early in the life-cycle. Given that by age 35 most households have purchased their home, in Table 3.3 we produce the average rates by two age groups, 26-35 and 36-60, to assess the calibration for two phases of the life-cycle: one in which households are entering the housing market and one in which tenure decisions are largely completed. We do not consider the years of the life-cycle up to age 26 since we assume that households have zero financial asset and housing endowments at age 22, with the counterfactual implication that our model makes it particularly hard for households to become home-owners early in the life-cycle. Symmetrically, we do not consider the years of the life-cycle above age 60 since we do not model retirement behaviour or bequests, which has the implication that households run down all assets by age 66, leading to an overestimate of the amount of selling of homes towards the end of life. The model slightly over-predicts home-ownership rates early in the life-cycle for the high education group (top left cells in Table 3.3)

Table 3.3: Calibration Statistics

Statistic	High Education		Low Education	
	Data	Model	Data	Model
<i>Ownership rate</i>				
Age 26 - 35	0.74	0.80	0.58	0.53
Age 36 - 60	0.84	0.84	0.66	0.69
<i>Probability of Moving</i>				
Age 26 - 60	1.54	3.4	1.49	4.6
$\frac{\text{Median}p_t}{\text{Median}w_t(L-l_t)}$				
Age 22 - 26	2.9	2.8	3.5	3.5

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

Probabilities for moving from owning to renting are the proportion of home-owning households in the BHPS (1991-2000) that become non-owners between years of the panel survey.

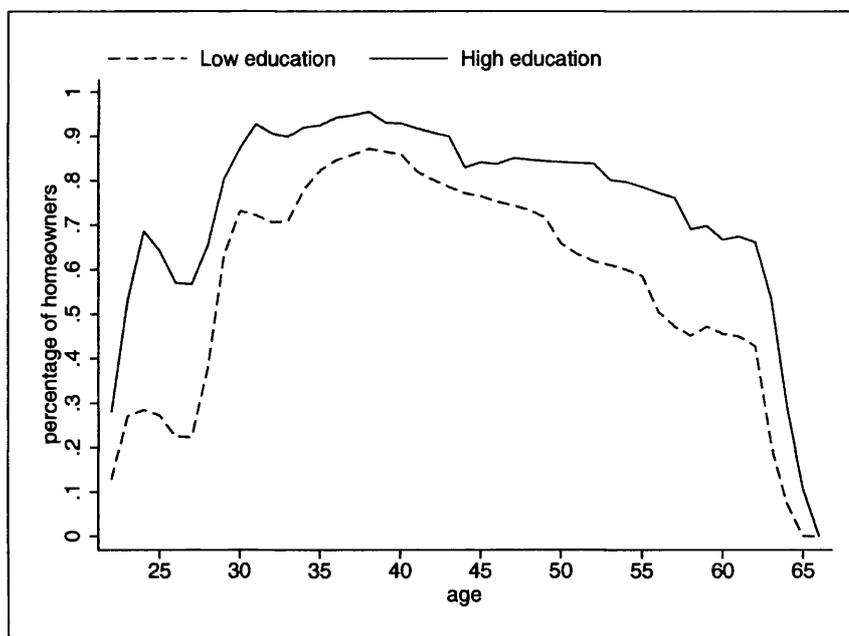


Figure 3.6: Simulated home-ownership rates

3.5.3 Non-Calibrated profiles

We turn now to showing predictions for and data on variables which have not been used in the calibration, to assess the validity of the model. We begin by showing in Figure 3.7 simulated and actual transition probabilities from renting to owning. The actual transitions are calculated from the BHPS as described in section 3.3. We match these transition paths fairly closely, although we under-predict the transition probability when young, particularly for the high education group. This is largely because of our assumption of zero initial housing and financial wealth, which means that households need some time to save up to meet the credit constraints to buy a house. We also over-predict the number of transitions back from owning to renting: as shown in Table 3.3 (“Probability of moving”), too many households sell their houses in our simulations. Many of these sales occur towards the end of the life-cycle. Since there are no bequests or retirement period in our model, households have a strong incentive to sell their housing to fund consumption in old age. It is therefore not surprising that we over-predict house sales late in life.

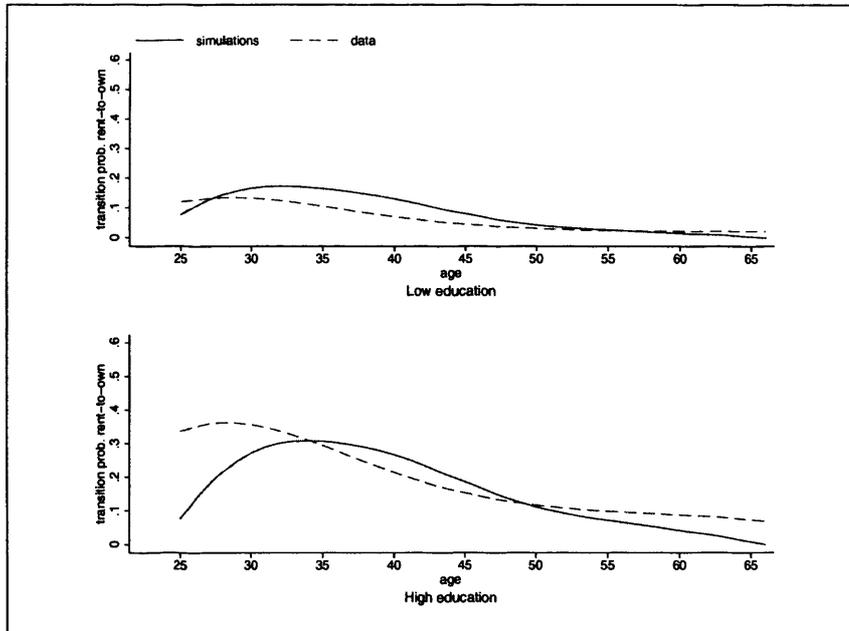


Figure 3.7: Simulated transition probabilities

We now turn to show other key predictions of the baseline model, namely the profiles on assets, consumption and labour supply.

Figure 3.8 reports simulated mean asset to mean income profiles. The definition of assets are all liquid assets which is liquid wealth minus outstanding debt. In our model, the rate of interest is the same on debt and liquid wealth and so we can concentrate on the net position. The large fall in the net position occurs when households buy their homes, the gradual subsequent reduction in debt is the voluntary repayment of the mortgage.

Figures 3.9 and 3.10 report simulated consumption and labour supply profiles for each education group. Consumption is higher for the higher education group throughout the whole life-cycle. The high education group supplies less labour in the first part of the life-cycle and more in the second part, which is consistent with their having a steeper wage profile and so a higher incentive for intertemporal substitution of labour

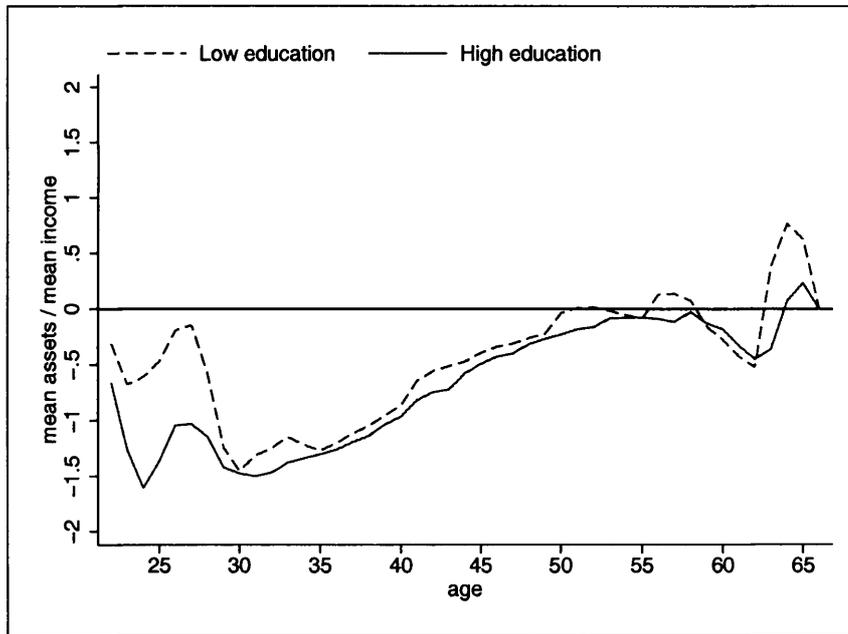


Figure 3.8: Net liquid asset position by age

than the low education group.

The early peak (age 26) and drop (age 28) in consumption, and corresponding drop and peak in labour supply, happen when home-ownership drops and subsequently picks up again, suggesting evidence of constrained behaviour.

3.6 Labour Supply and Home-ownership

One of the main aims of our model is to analyse the interaction between labour supply and housing choices. The key question is how labour supply varies by home-ownership. There are two reasons why we focus on the comparison between homeowners and non-homeowners, rather than between different levels of mortgage repayments. The first reason is that we only allow for collateralised debt, and therefore homeowners are those who can borrow and non-homeowners are those who cannot. The second reason is that mortgage payments are voluntary, and therefore distinguishing households by level of mortgage repayments, and comparing their labour supply behaviour, would not give us

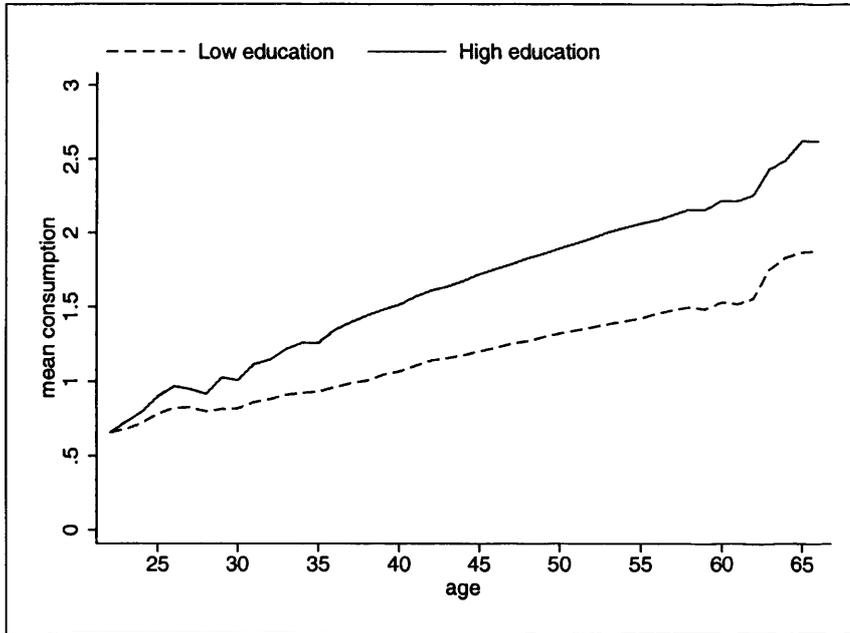


Figure 3.9: Consumption by age

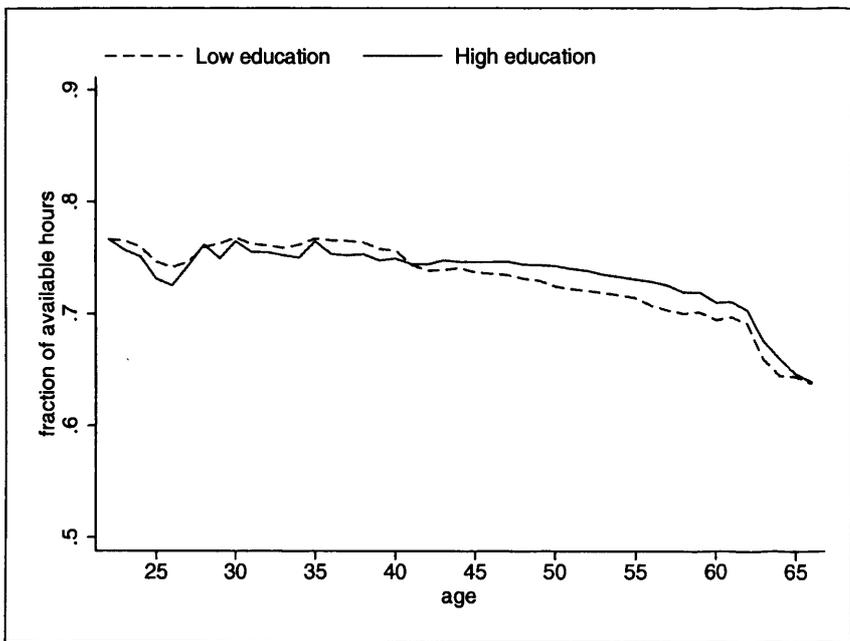


Figure 3.10: Labour supply by age

an indication of constrained behaviour. Those who repay higher amounts of mortgage will tend to be those with higher income draws, and we cannot expect their labour supply behaviour to be constrained.

The fact that only home-owners can borrow, and that the income-related borrowing constraint holds at every period, imply that our model predicts labour supply as being positively related to the proportion of income that households can borrow. In particular, given that the income-related borrowing constraint is a function of the current choice of leisure, the intra-period allocation between consumption and leisure, when the income-related constraint binds, is governed by the following:

$$\frac{\partial u(\cdot)/\partial l}{\partial u(\cdot)/\partial c} = w(1 + \lambda_y). \quad (3.10)$$

In other words, the constraint affects the price of leisure. We will now analyse how household labour supply is affected by this change in price by specifying the form of the utility function.

3.6.1 Baseline model

When utility (between consumption and leisure) is Cobb-Douglas, as in our baseline model, equation 3.10 implies the following relationship between labour supply and the proportion of income that households can borrow (see Appendix 3.10.3 for details):

$$ls_t = \eta L - \frac{\hat{A}_t(1 - \eta)}{w_t(1 + \lambda_y)} \text{ where } \hat{A}_t = \begin{cases} A_t - p_t h_t(1 + F) & \text{if buy house} \\ A_t & \text{if } h_t = h_{t-1} = 1 \end{cases}$$

and ls_t is labour supply at period t ($ls_t \equiv L - l_t$). Labour supply is therefore higher when the constraint is binding than when it is not (in which case the price of leisure is simply w_t).

Given this theoretical implication of the income constraint for labour supply, it is interesting to examine how much, and at what points in the life-cycle, this affects the behaviour of home-owners.

In Table 3.4 we report a comparison between the labour supply of homeowners and that of non-homeowners predicted by our baseline model. We find that homeowners

aged between 26 and 35 supply about 2.5% more hours than non-homeowners, and that those aged between 36 and 60 supply 5.4% and 3.5% more hours than non-homeowners in the low or high education groups, respectively. BHPS data for the sample of women aged between 25 and 45,⁹ in couples in which the partner works full-time, show that the difference in average hours supplied by the couple between homeowners and non-homeowners is 8.6% and 1.44%, respectively, for the low and high education group.¹⁰ For the sub-sample of women aged 25-35, this difference is more pronounced, amounting, respectively, to 12.2% and 2.9%. Our model, therefore, seems to underpredict the difference in labour supply behaviour between homeowners and non-homeowners for the low education group. However, some care needs to be adopted when making this comparison, since our model, by considering only the household unit, attributes the education level to the household, whereas the data figures that we have reported concern only the education level of the woman. Moreover, in the model a single wage is attributed to all the labour supplied in the period, whereas this is obviously not the case in the data. More generally, the model allows only one source of ex-ante heterogeneity, which is the education level of the household. By contrast, the data will reflect many other sources of heterogeneity, both observed and unobserved. The predictions from the model can be interpreted as an examination of how much diversity in behaviour can be explained on the basis of only limited heterogeneity in characteristics.

In Figure 3.11, we report labour supply by home-ownership status at each year of the life-cycle: home-owners work longer hours than those not owning in almost every period, and particularly at the time of buying, in a similar way to the evidence in chapter 2 and discussed above. This effect of greater labour supply among home-owners is driven primarily by greater debt holdings: greater debt means lower total expenditure in the current period, but when the home-owner ages and releases capital from the house, consumption and leisure both rise. In other words, the option of home-ownership raises expected lifetime utility, but it induces intertemporal substitution:

⁹These data are those used in chapter 2. We do not obtain statistics for age above 45 due to small cell sizes for the group of renters

¹⁰We assume that full-time hours for the partner correspond to 40 hours per week.

Table 3.4: Simulated Labour Supply by Homeownership status

	Low Education	High Education
<i>Age 26 -35</i>		
Homeowners	0.766	0.756
Non-homeowners	0.747	0.737
<i>Age 36 -60</i>		
Homeowners	0.742	0.743
Non-homeowners	0.704	0.718

Notes: reported fraction of time that the household allocates to labour supply.

households buy their homes, forfeiting current consumption and leisure, in order to profit from the fast house price growth and high consumption and leisure in the future.¹¹ Since the consumption elasticity intertemporal substitution is larger than the labour supply intertemporal elasticity, we would expect that more of the adjustment occurs through consumption. On the other hand, labour supply is directly affected by the borrowing constraint whereas consumption is not.

An alternative way of exploring the relation between labour supply and homeownership is to calculate labour supply profiles over the life-cycle for the case where households cannot purchase a house, to compare to the baseline case where homeownership is an option and labour supply is possibly constrained by home-ownership decisions and the associated credit constraints. In particular, we consider the average labour supply behaviour of home-owners in the baseline model and of the same group of individuals (i.e. with the same idiosyncratic wage shocks) in a model in which they cannot purchase a house. Figure 3.12 presents this comparison. It shows that in a model in which households can purchase a house, subject to borrowing constraints, the cross-sectional labour supply is higher than in the case in which they cannot purchase a

¹¹An additional motive to forfeit consumption and leisure in order to get into the housing market at an early age, may be that the housing asset provides insurance against future increases in its price - see Banks, Blundell, Smith and Smith (2004)

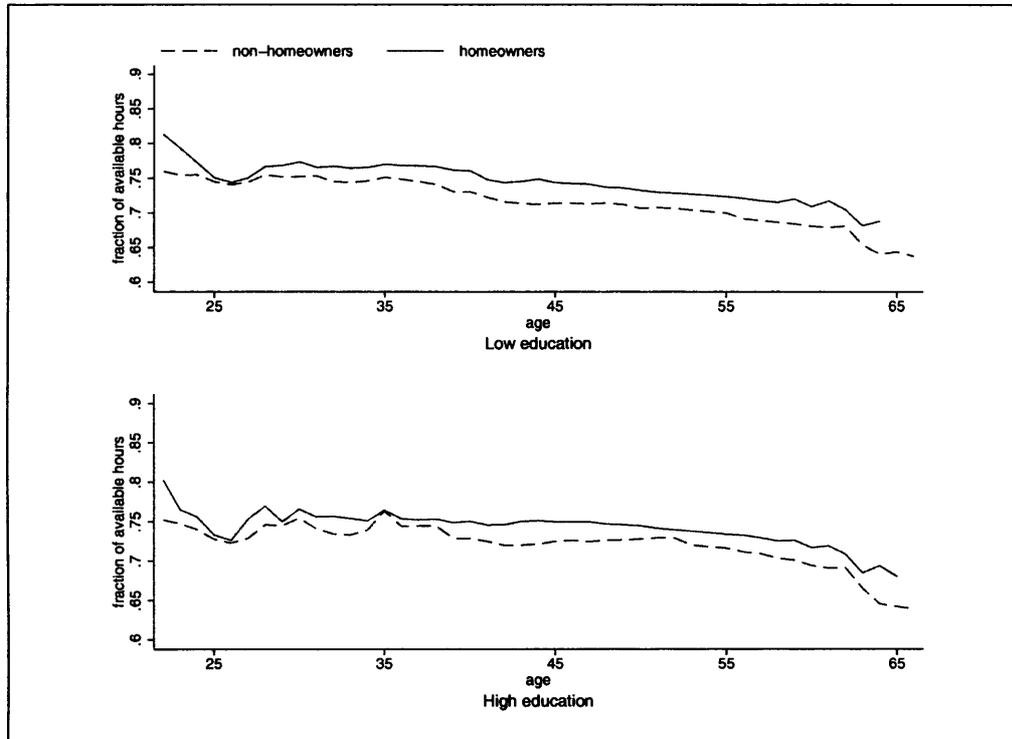


Figure 3.11: Labour supply by home-ownership status

house in the first half of the life-cycle, and lower afterwards. In other words, households substitute leisure inter-temporally since their expected utility from being home-owners and from future consumption is higher than when not being home-owners. There is, however, an additional effect driven by the income-related borrowing constraint. This is particularly evident when we compare labour supply of those households that become home-owners at the beginning of their life with their behaviour if they were not home-owners, and notice that labour supply is about 10 percent higher in the scenario in which they purchase a house. The remaining spikes in labour supply occur when a high proportion of households purchases a house and when house prices spike up.

3.6.2 Sensitivity to utility function specification

Choosing a Cobb-Douglas utility function as the form to embed in a CRRA, means choosing a particular case of CES function in which the elasticity of substitution be-

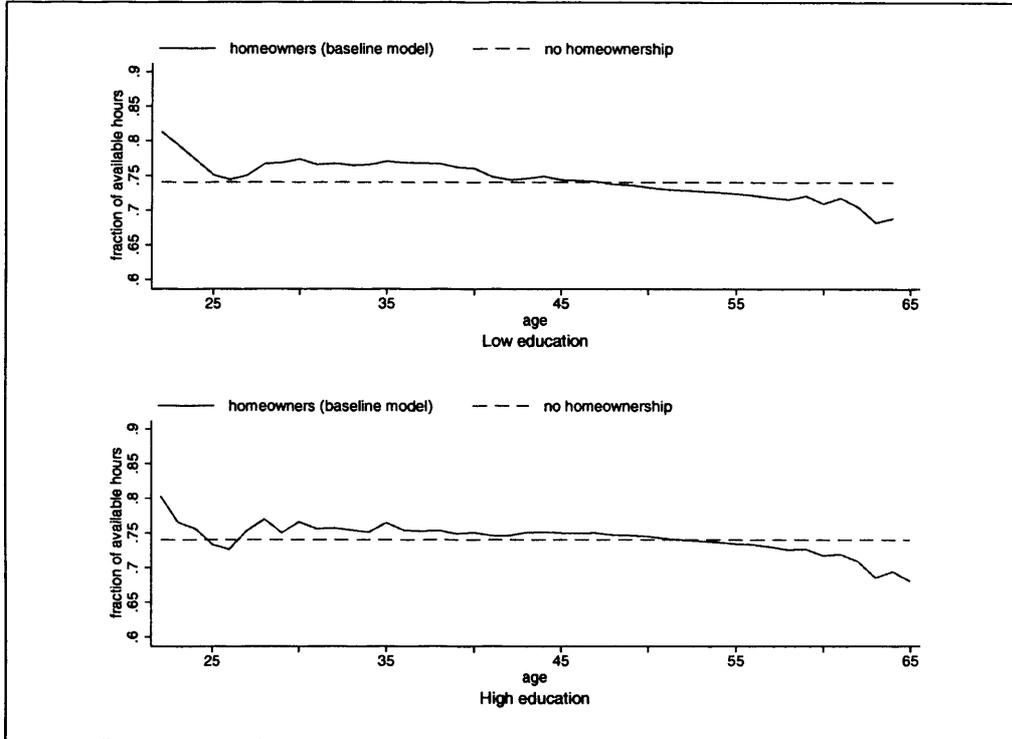


Figure 3.12: Labour supply: baseline vs. non-homeownership

tween consumption and leisure is equal to 1. We now want to check how sensitive our results are to this assumption. We therefore take a general CES utility function and analyse profiles of home-ownership and labour supply in two alternative cases in which substitutability between consumption and leisure is, respectively, higher and lower than in the Cobb-Douglas case. The within period utility function is now specified as

$$u(c_t, h_t, l_t) = e^{\theta h_t} \frac{(c^{-\rho} + \eta l^{-\rho})^{-\frac{(1-\gamma)}{\rho}}}{1 - \gamma}. \quad (3.11)$$

Figures 3.13 and 3.14 show labour supply profiles by home-ownership status for the cases in which $\rho = -0.2$ and $\rho = 0.2$. This means that the elasticity of substitution between consumption and leisure is, respectively, higher (1.25) and lower (0.83) than in the baseline model with a CRRA/Cobb-Douglas within period utility. Labour supply is still higher for home-owners corresponding to the stage in the life-cycle in

which a high fraction of the population purchases a house.¹² However, the difference in labour supply behaviour between home-owners and non-home-owners is bigger than in the baseline case when substitutability between consumption and leisure is higher (and smaller when substitutability is lower) than in the baseline model. With higher substitutability between consumption and leisure, household labour supply is more responsive to wages. Since homeowners tend to be those who have higher wages, this explains the pattern that we observe. The qualitative suggestion that the labour supply of home-owners tends to increase around the time when they buy still holds.

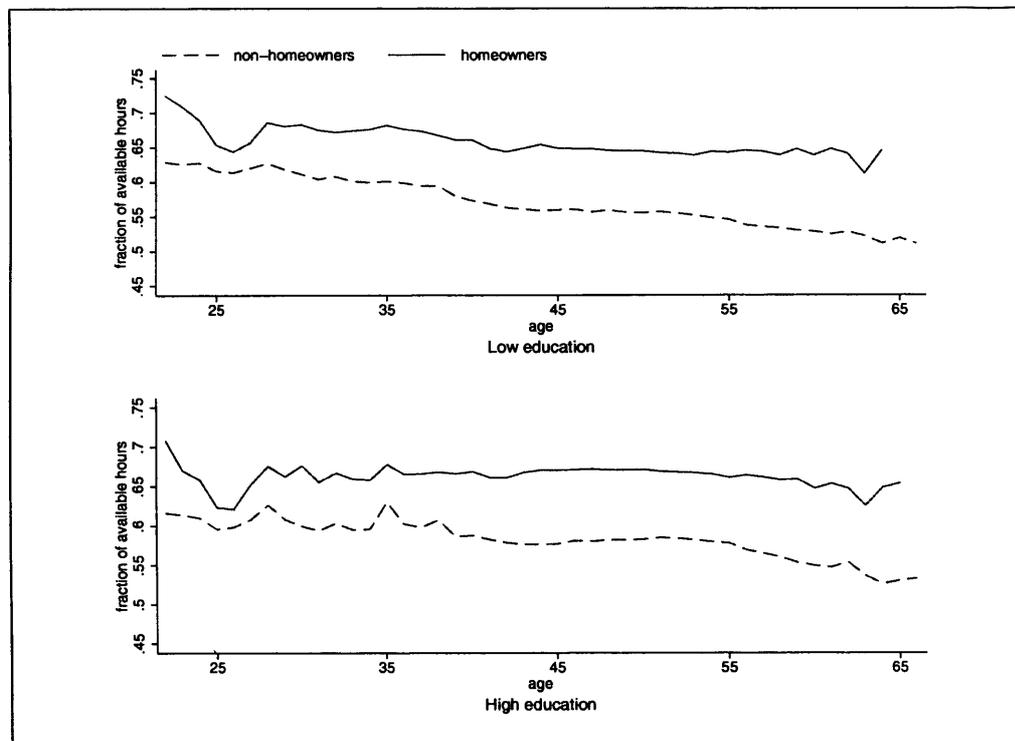


Figure 3.13: Labour supply by home-ownership status, $\varepsilon_{c,l} = 1.25$

¹²Home-ownership profiles are very similar to those shown for the baseline case, except for lower rates at the beginning of the cycle.

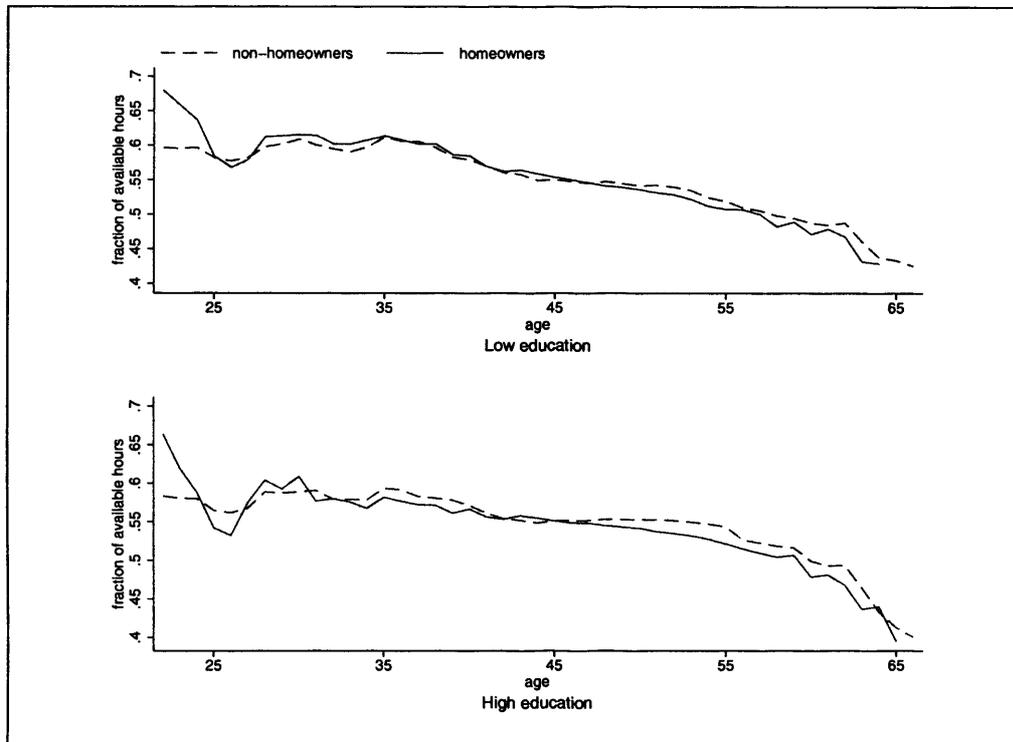


Figure 3.14: Labour supply by home-ownership status, $\epsilon_{c,l} = 0.83$

3.7 Implications of Incomplete Markets

In our baseline, the financial environment was set up to mimic the UK housing and capital market. In this section, we highlight the effect that changes in the capital market have on home-ownership patterns and on labour supply. We show first the effects of varying the size of downpayment needed to purchase a house, $(1 - \lambda_h)$. Second, we vary the multiple of household income that households can borrow against, λ_y .

3.7.1 Varying downpayment requirements

We assume households can borrow only a fraction λ_h of the house value. Figures 3.15 and 3.16¹³ show the effects on home-ownership and labour supply of varying λ_h while holding other parameters at their baseline values. Reducing downpayment requirements (increasing λ_h) leads to households buying their homes earlier in the life-cycle but makes very little difference later in the life-cycle. Although changes in λ_h have a small effect on labour supply, observing this can be informative about which constraint is binding. In particular, the low education group does not appear to change labour supply behaviour as the downpayment goes from zero to 50 percent and home-ownership changes substantially at the beginning of the life-cycle. This seems to indicate that the income-related constraint is not binding. On the other hand, the high education group increases both home-ownership and labour supply as the downpayment decreases, suggesting that the higher home-ownership associated to a lower downpayment makes the income-related constraint bind.

3.7.2 Varying debt to income requirements

We assume that there is a restriction on the debt to income ratio, λ_y , that households can hold. This is a common institutional feature in the UK and elsewhere (Survey of Mortgage Lenders). Figure 3.17 shows the effect on home-ownership of varying λ_y while holding other parameters at their baseline values. Like the case of decreasing the downpayment requirement, increasing the permissible debt to income ratio leads to earlier home-ownership, but in contrast to the previous case, it also has the effect of increasing the proportion of households owning their homes throughout the whole life-cycle. The effects on labour supply are again small but peaks in labour supply are still associated to peaks in home purchasing. The fact that this pattern is still observed even with $\lambda_y = 10$, suggests that peaks in labour supply at the time of purchase are at least partly due to smoothing consumption while meeting the downpayment requirement.

¹³This figure and figure 3.18 have been magnified to show the region where there are the biggest differences relative to the baseline model.

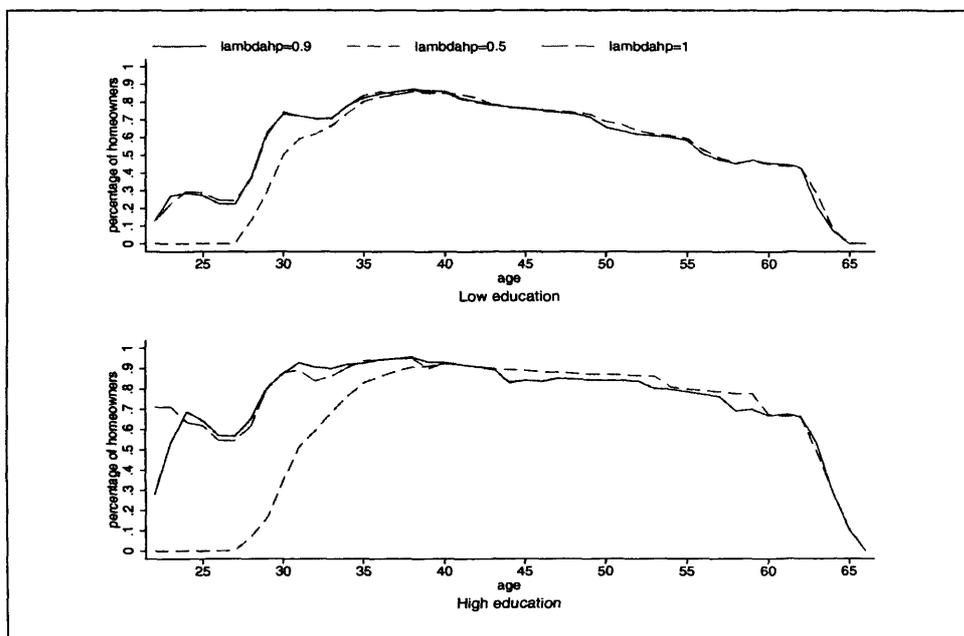


Figure 3.15: Homeownership varying downpayment requirements

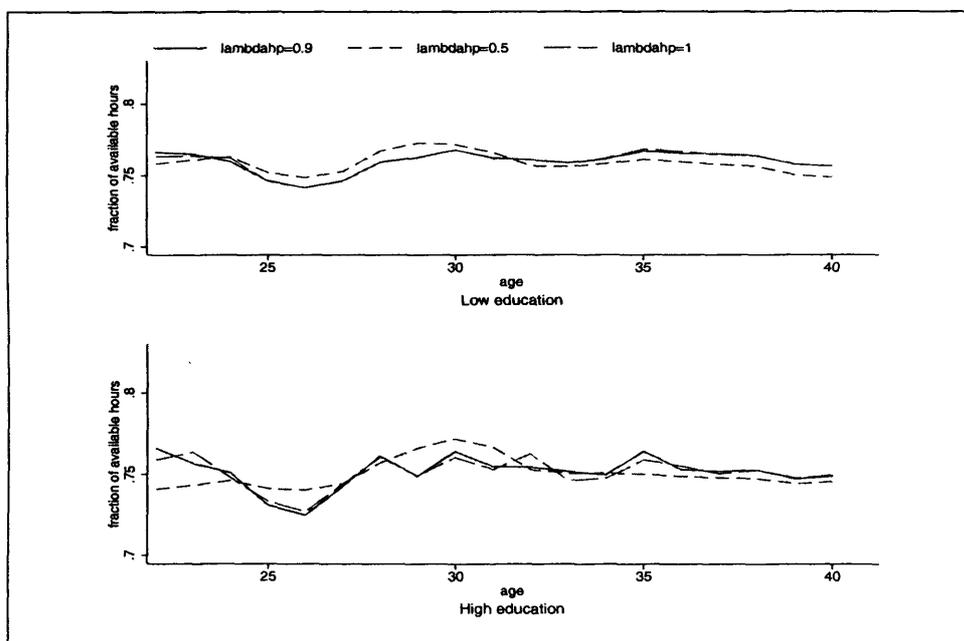


Figure 3.16: Labour supply varying downpayment requirements

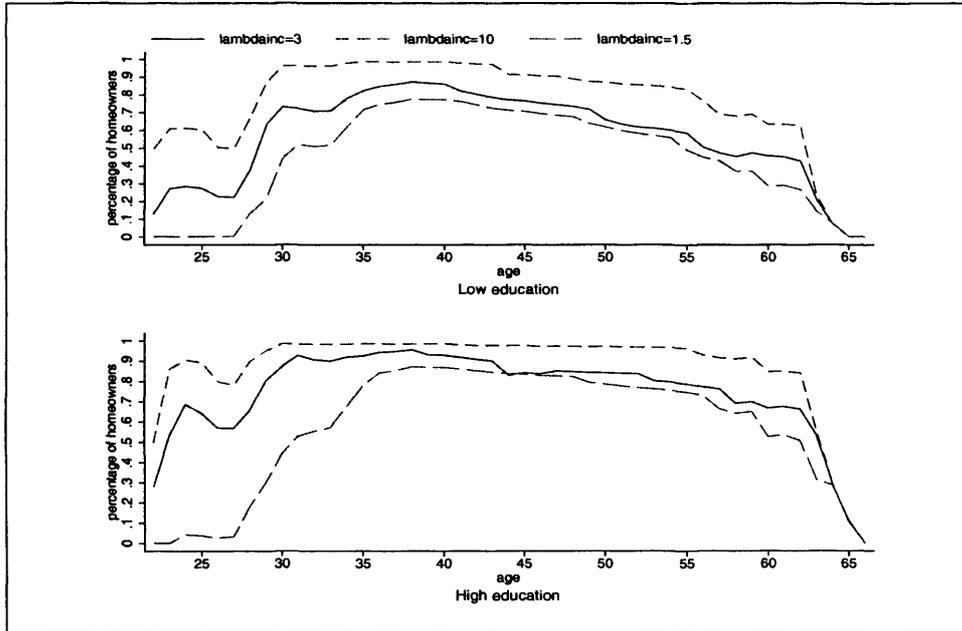


Figure 3.17: Homeownership varying the maximum debt-to-income ratio

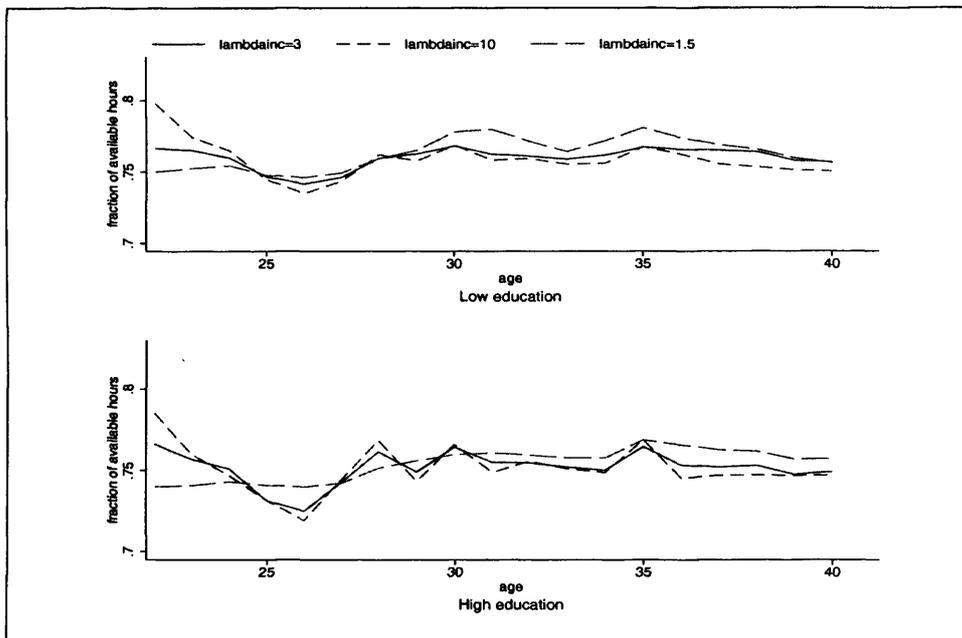


Figure 3.18: Labour supply varying the maximum debt-to-income ratio

3.8 Welfare implications of wage uncertainty

One important feature of the model introduced in this chapter is the existence of an earnings-related constraint for homeowners. We now want to analyse the implications of earnings uncertainty on household welfare. In particular, we want to ask how welfare changes in scenarios with different wage uncertainty, and how this translates in terms of consumption, labour supply, and home ownership choices. The first step for answering these questions is to illustrate which method to adopt for measuring welfare.

3.8.1 Welfare criterion

The welfare measure that we adopt is given by expected lifetime utility. That is, the welfare associated to a scenario with wage variance equal to σ_b , is given by:

$$\Omega_{\sigma_b} = \int_W \int_P V_1(A_0, H_0, p_1, w_1) dp dw$$

where $V_1(A_0 = 0, H_0 = 0, p_1, w_1)$ is the value function at the beginning of life, given zero initial endowments of financial assets and house, as assumed throughout this chapter, and p and w are the initial distributions of the house price and of wages, assumed to be independent one another. w is distributed log-normally with variance σ_b . As remarked by Aiyagari and McGrattan (1998), one way to interpret this function is in terms of the welfare of a typical household through the veil of uncertainty (*i.e.* before uncertainty is realised). In particular, our household knows whether it belongs to the high education group or to the low education group, and that its initial financial assets are zero and it has no house. However, it does not know what initial wage and house price it will draw. Since the value function at the beginning of life is the expected discounted lifetime utility, we can rewrite welfare as follows:

$$\Omega_{\sigma_b} = E_1 \left[\sum_{t=1}^T \beta^t u(c_t, h_t, l_t) \right]$$

where E_1 indicates expectations at the beginning of period 1, before uncertainty is realised.

In a model where utility is only a function of consumption, a way of measuring the welfare cost of different degrees of uncertainty is to compute the compensating variations in consumption, that is, the percentage change in consumption that has to occur at every period in order to keep welfare the same under two alternative scenarios. This is, for instance, the approach taken by Attanasio and Davis (1996). Huguet and Ventura (1999), as well as Heathcote et al. (2005), adopt this approach also in a model in which utility is a function of both consumption and labour supply. This is correct as long as it is possible to derive an intra-temporal allocation rule between consumption and leisure, so that the change in consumption takes account of the associated change in leisure. If we abstract for a moment from the additional choice of home-ownership that is present in our model, the condition that gives the compensating variation in consumption is as follows:

$$E_1 \left[\sum_{t=1}^T \beta^t u(c_t, h_t, l_t) \right] = E_1 \left[\sum_{t=1}^T \beta^t u((1 + \pi)\tilde{c}_t, \tilde{h}_t, \tilde{l}_t) \right] \quad (3.12)$$

where c , h and l are household's choice of consumption, home ownership, and leisure in the baseline scenario with wage variance σ_b , and \tilde{c} , \tilde{h} and \tilde{l} are household's choice of consumption, home-ownership, and leisure in the alternative scenario with wage variance σ_a . π is the compensating variation in consumption, i.e. the per-period percentage change in consumption, relative to the baseline scenario, that makes welfare the same under two alternative wage processes with different variance.

3.8.2 Welfare with CRRA/Cobb-Douglas preferences

Consumption-based welfare metric

When the felicity function is CRRA/Cobb-Douglas, welfare in the baseline scenario takes the following form:

$$\Omega_{\sigma_b} = E_1 \left[\sum_{t=1}^T \beta^t e^{\theta h_t} \frac{[c^\eta l^{1-\eta}]^{1-\gamma}}{1-\gamma} \right]$$

If we assume that households choose an interior allocation of leisure,¹⁴ we can use the intra-temporal optimal allocation rule between consumption and leisure, $l_t = \frac{c_t}{w_t} \frac{1-\eta}{\eta}$,

¹⁴We will deal later with the case in which they chose a corner allocation of leisure

so that equality between welfare in the two scenarios (equation (3.12)) requires the following:

$$E_1 \left[\sum_{t=1}^T \beta^t e^{\theta h_t} \frac{\left[c_t \left(\frac{1-\eta}{\eta w_t} \right)^{1-\eta} \right]^{1-\gamma}}{1-\gamma} \right] = E_1 \left[\sum_{t=1}^T \beta^t e^{\theta \bar{h}_t} \frac{\left[(1+\pi) \bar{c}_t \left(\frac{1-\eta}{\eta \bar{w}_t} \right)^{1-\eta} \right]^{1-\gamma}}{1-\gamma} \right] \quad (3.13)$$

or

$$\Omega_{\sigma_b} = (1+\pi)^{1-\gamma} \tilde{\Omega}_{\sigma_a}.$$

It follows that the consumption-based welfare metric, π , is obtained as a function of welfare in the two wage variance scenarios:

$$\pi = \left(\frac{\Omega_{\sigma_b}}{\tilde{\Omega}_{\sigma_a}} \right)^{(1-\gamma)^{-1}} - 1. \quad (3.14)$$

Note that this closed form solution for the consumption-based welfare metric is obtained under two assumptions: that households always choose the interior solution for labour supply, and that home ownership decisions are not affected by the change in consumption given by π . This is in fact not true in our model, where home ownership is non-separable from consumption and leisure in the utility function and where labour supply cannot become negative and is affected by the explicit income-related borrowing constraint for homeowners. We will address both these issues in the next section (“Compensating variations in assets”) but in this section we will maintain the two assumptions in order to obtain the consumption-based welfare metric, π .¹⁵ In this way we will be using a standard approach adopted in the literature (see Hugget and Ventura (1999), Krueger and Perri (2003) and Heathcote et al.(2005)) and we will be able to compare the predictions produced by this approach with the predictions obtained in the next section (where the assumptions are relaxed).

Table 3.5 reports the compensating variations in consumption, computed according to equation (3.14), in a model in which households are subject to an income-related

¹⁵The computation of welfare in each scenario will take into account the constraints specified for the model being analysed, and the non-separabilities between consumption, leisure and home-ownership. What does not take these features into account is the way in which the consumption-based welfare metric, π , is computed as a function of the welfare measures.

borrowing constraint (“Endogenous income constraint”) as well as in a model in which they are not subject to it. The first row in the table deals with a model in which households decisions are not subject to an endogenous income constraint. Labour supply is, however, still subject to a time allocation constraint. Assuming interior labour supply when computing π , is as if we were obtaining compensating variations in consumption for households that are never at a corner solution. Alternatively, we could compute the compensating variations in consumption for households assumed to be always at a corner solution given by the total time available, $\bar{l}_t = 1$, or equivalently (with Cobb-Douglas utility function) assumed not to change their labour supply behaviour across the two scenarios. This would give us an upper bound for the consumption-based welfare metric.¹⁶ Results for this case are reported in the Notes to Table 3.5.

The second row of Table 3.5 considers instead a model in which households make their decisions subject to an income-related borrowing constraint. The presence of the constraint affects households decisions through their expectations, since they are aware that they may become constrained at some point in time; therefore, it affects welfare and ultimately the consumption-based welfare metric, π . Given the way we compute π , the fact that the consumption-based welfare metric is different depending on whether or not there is an income-related borrowing constraint, means that welfare in the baseline scenario and in the alternative scenario have changed differently in the two models. The change is between 1.2 and 0.8 percentage points for the two education groups, which respectively amount to 8.5% and 5.5% increases.

Recalling that the wage process in the model presented in this chapter is of the

¹⁶In this case, equality between welfare in the two scenarios (equation (3.12)) requires the following:

$$E_1 \left[\sum_{t=1}^T \beta^t e^{\theta h_t} \frac{[c^\eta l^{1-\eta}]^{1-\gamma}}{1-\gamma} \right] = E_1 \left[\sum_{t=1}^T \beta^t e^{\theta \bar{h}_t} \frac{[(1+\bar{\pi})\bar{c}]^\eta \bar{l}^{1-\eta}]^{1-\gamma}}{1-\gamma} \right]$$

with

$$\bar{l} = \begin{cases} 1 & \text{if corner solution} \\ l & \text{if unvaried leisure across two scenarios} \end{cases}$$

which implies

$$\bar{\pi} = \left(\frac{\Omega_{\sigma_b}}{\bar{\Omega}_{\sigma_a}} \right)^{[\eta(1-\gamma)]^{-1}} - 1.$$

Under our parameter’s choice, $\eta = 0.74$ and $\gamma = 1.58$, $\bar{\pi} > \pi$.

Table 3.5: Compensating variations in consumption

Baseline scenario: deterministic wage. Alternative scenario: stochastic wage, $\sigma_\xi^2 = 0.018$

	Low education	High education
No income constraint	14.05%	14.34%
Endogenous income constraint	15.25%	15.12%

Notes: These are the compensating variations in consumption obtained according to equation (3.14). If, alternatively, we assumed that households were always at the corner solution $\bar{l}_t = 1$, or that they would not change their labour supply between the two scenarios, in a model with no income constraint the low education group and the high education group would need, respectively, a 19.45% and a 19.85% increase in consumption every period.

form

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

if we considered a scenario with $\sigma_\xi^2 = 0.018$ (the value estimated for the UK and for which the model is evaluated) versus a baseline scenario with deterministic wages (i.e. wages follow the deterministic trend a_t over time), we would obtain the compensating variations in consumption reported in Table 3.5. That is, to make a high education household in a scenario with stochastic wages (as given by σ_ξ^2) as well off as a household in a scenario with deterministic wages, the household facing stochastic wages would need to consume 14.34% more at every period when the household is not subject to an income-related constraint and 15.12% more at every period when it is subject to an endogenous income constraint.

Given that a crucial feature of our model is the endogeneity of labour supply choices in relation to home-ownership decisions, via the income-related constraint, it is actually important to include households at a constrained solution in leisure when obtaining measures of the welfare cost of wage uncertainty. Moreover, since our utility function is non-separable in home-ownership, consumption and labour supply, it is important to allow home-ownership decisions be different across two scenarios of wage variance.

We therefore turn to an alternative measure of welfare that is based on the change in initial asset allocations in the section below.

Compensating variations in assets

We now want to think what change in the financial asset endowment (at the beginning of life) would make a household as well off in two scenarios with different wage variance. The condition that sets the welfare equal in the two scenarios is the following:

$$\int_W \int_P V_1(A_0, H_0, p_1, w_1) dp dw = \int_W \int_P V_1(\tilde{A}_0, H_0, p_1, \tilde{w}_1) dp d\tilde{w} \quad (3.15)$$

We therefore want to find $\Delta A_0 = \tilde{A}_0 - A_0$, where A_0 is the level of initial assets under the baseline scenario with wage variance σ_b , assumed to be equal to zero, and \tilde{A}_0 is the unknown level of initial assets under the alternative scenario with wage variance σ_a . The solution is found numerically and the results are reported in table 3.6.

We ask what the change in initial assets should be in two different economies: one economy where households were subject to an income-related constraint when owning a house (“Endogenous income constraint”), and another economy in which this constraint did not exist. We perform this exercise under different scenarios with respect to the wage variance. First, we ask what the compensating variation in asset would be if $\sigma_\xi^2 = 0.018$ (i.e. the estimated UK log-wage variance) as compared to a scenario with deterministic wages (rows 1 and 5 for, respectively, low and high education). Recalling that initial wages are 1, and that “normal hours worked” are 0.74, we can translate these changes in assets in terms of “normal earnings” by dividing them by 0.74. For example, a high education household would need a change in assets at the beginning of life approximately equal to six and half times his earnings in period 1. We then ask what the compensating variation in assets would need to be if the log-wage variance were half its current value ($\sigma_\xi^2 = 0.009$) as compared to a case with no variance. Finally, we compare two scenarios with different log-wage variances: the current one in relation to half its value ($\sigma_\xi^2 = 0.018$ vs. $\sigma_\xi^2 = 0.009$), and twice the value of the current one in relation to the current one ($\sigma_\xi^2 = 0.036$ vs $\sigma_\xi^2 = 0.018$), as if to ask what

was the welfare loss in halving the log-wage variance and what would the welfare loss be were the current log-wage variance to double.

Table 3.6: Compensating variations in assets

	Endogenous income constraint	No income constraint	Difference
Low education			
a) $\sigma_{\xi}^2 = 0.018$ vs. <i>Deterministic wage</i>	4.24	4.02	0.22
b) $\sigma_{\xi}^2 = 0.009$ vs. <i>Deterministic wage</i>	2.23	2.05	0.18
c) $\sigma_{\xi}^2 = 0.018$ vs. $\sigma_{\xi}^2 = 0.009$	1.66	1.60	0.06
d) $\sigma_{\xi}^2 = 0.036$ vs. $\sigma_{\xi}^2 = 0.018$	2.42	2.43	-0.01
High education			
a) $\sigma_{\xi}^2 = 0.018$ vs. <i>Deterministic wage</i>	4.80	4.61	0.19
b) $\sigma_{\xi}^2 = 0.009$ vs. <i>Deterministic wage</i>	2.25	2.16	0.09
c) $\sigma_{\xi}^2 = 0.018$ vs. $\sigma_{\xi}^2 = 0.009$	1.89	1.78	0.11
d) $\sigma_{\xi}^2 = 0.036$ vs. $\sigma_{\xi}^2 = 0.018$	3.18	3.11	0.07

Notes: The table reports absolute changes in initial assets (ΔA_0) that make household welfare equal under two scenarios.

The compensating variations in assets contained in Table 3.6 show that the high education group needs higher compensations than the low education group in any scenario (as expected, given that the time trend in the wage profile is higher for the more educated). For the low education group, what seems to matter most towards welfare is whether there is any variance at all, rather than how much variance there is. Case d), which compares a scenario having log-wage variance $\sigma_{\xi}^2 = 0.036$ with a baseline scenario having log-wage variance $\sigma_{\xi}^2 = 0.018$, requires nearly half the compensation in assets that is required in case a), which compares a scenario having log-wage variance $\sigma_{\xi}^2 = 0.018$ with a baseline scenario having deterministic wages. This seems to be less the case for the high education group, for which the difference between case a) and case d) is much less marked. Both education groups require the highest compensation when a scenario with the current log-wage variance ($\sigma_{\xi}^2 = 0.018$) is compared to a baseline

scenario with no wage uncertainty. Also the difference between the compensation required in an economy where households are subject to an endogenous income-related constraint and the compensation required in an economy where there is not such a constraint is highest when comparing a scenario with $\sigma_{\xi}^2 = 0.018$ to a scenario with deterministic wages. This difference is positive, meaning that being subject to an income-related constraint makes households want a higher compensation than when not being subject to it. The pattern of this difference, however, does not seem to be monotonic with the level of the log-wage variance. For the low education group there is some evidence that the additional compensation required in an economy with an income-related constraint is decreasing with the level of the log-wage variance (see column 3, cases *b*) to *d*)) but this is not the case for the high education group.¹⁷

We now want to translate the compensating variations in assets in terms of consumption, labour supply and home-ownership choices. In order to do this, we simulate economies of 10,000 individuals and compare the average behaviour in a case with baseline wage variance and zero initial assets to a case with a different wage variance and the level of initial assets that makes welfare equal in these two scenarios (the “compensating variation in assets”). In other words, we simulate different cases corresponding to those reported in Table 3.6 and compare a case with higher variance and the level of initial asset reported in the table to a case with lower (or zero) variance and zero initial assets. For instance, in connection with row 1 of Table 3.6, we first simulate an economy in which households are subject to an income-related constraint (“Endogenous income constraint”) and obtain the average behaviour for low education households under two scenarios: one in which wages are deterministic and initial assets are zero (baseline scenario), and another in which the log-wage variance is $\sigma_{\xi}^2 = 0.018$ and initial assets are 4.24. We then simulate a different economy in which households are no longer subject to an income-related constraint (“No income constraint”) and obtain the average behaviour again under the baseline scenario and

¹⁷In any case, we are only evaluating three different scenarios and drawing inferences about intermediate cases would require assumptions such as that the relationship is monotonic.

under an alternative scenario in which the log-wage variance is $\sigma_{\xi}^2 = 0.018$ and initial assets are 4.02. The results of this exercise are reported in Table 3.7 (row 1, in this particular case) and Figures 3.19 - 3.24 (top panels of Figures 3.19 - 3.21 for this example). Figures 3.19 - 3.24 only illustrate the cases corresponding to rows 1, 3, 5 and 7 of Table 3.7 (log-wage variance $\sigma_{\xi}^2 = 0.018$ vs. deterministic wage and log-wage variance $\sigma_{\xi}^2 = 0.036$ vs. log-wage variance $\sigma_{\xi}^2 = 0.018$) since the other cases have very similar profiles and the same implications.

Table 3.7: Compensating variations in consumption, labour supply and home-ownership

	Endogenous income constraint			No income constraint		
	Δc	Δls	Δho	Δc	Δls	Δho
Low education						
$\sigma_{\xi}^2 = 0.018$ vs. <i>Deterministic wage</i>	8.37%	-4.48%	10.62%	7.26%	-4.23%	-5.18%
$\sigma_{\xi}^2 = 0.009$ vs. <i>Deterministic wage</i>	4.56%	-2.31%	8.50%	3.66%	-1.99%	-3.94%
$\sigma_{\xi}^2 = 0.018$ vs. $\sigma_{\xi}^2 = 0.009$	3.21%	-1.90%	21.21%	2.31%	-1.65%	3.21%
$\sigma_{\xi}^2 = 0.036$ vs. $\sigma_{\xi}^2 = 0.018$	4.93%	-4.69%	29.79%	3.59%	-1.20%	8.90%
High education						
$\sigma_{\xi}^2 = 0.018$ vs. <i>Deterministic wage</i>	6.70%	-3.85%	-13.81%	6.89%	-4.13%	-8.78%
$\sigma_{\xi}^2 = 0.009$ vs. <i>Deterministic wage</i>	3.83%	-1.92%	-9.58%	4.03%	-2.12%	-5.93%
$\sigma_{\xi}^2 = 0.018$ vs. $\sigma_{\xi}^2 = 0.009$	2.21%	-1.70%	-1.38%	1.84%	-1.52%	0.09%
$\sigma_{\xi}^2 = 0.036$ vs. $\sigma_{\xi}^2 = 0.018$	3.72%	-3.42%	7.30%	3.48%	-3.36%	0.74%

Notes: The table reports percentage changes in consumption (Δc), labour supply (Δls) and home ownership (Δho) that correspond to the increases in initial assets reported in Table 3.6 that allow households in a high wage variance scenario to achieve the same welfare as under a lower wage variance scenario with zero initial assets. The changes in behaviour are expressed in relation to a baseline scenario with zero initial assets and lower wage variance.

Table 3.7 displays the percentage change in consumption, labour supply and home-

ownership over the life cycle under the alternative wage variance scenario with compensating variation in initial assets relative to the baseline wage variance scenario with zero initial assets, for each economy type. For instance, we find that in an economy in which households are subject to an income-related constraint, providing low education households with compensating variations in initial assets in a scenario with log-wage variance $\sigma_{\xi}^2 = 0.018$ allows them to choose on average, and over the lifetime, 8.37% more consumption, 4.48% less labour supply and 10.62% more home-ownership than in the baseline scenario with deterministic wages and zero initial assets. In general, we find that in the alternative scenario with compensating variation in assets and higher variance than in the baseline scenario, consumption is always higher and labour supply is always lower. Figures 3.20 and 3.23 show that the biggest drops in labour supply in a scenario with compensating variations in assets takes place at the beginning of life. This type of inter-temporal substitution is not unexpected, given that the wage trend is concave, and that the initial lump-sum in assets allows them to purchase the house without making the income-related constraint bind.

As for home-ownership, Table 3.7 shows that it can be either higher or lower in the scenario with compensating variation in assets, relative to the baseline scenario with zero assets and lower (or zero) variance. However, this is the overall change across the life cycle and it can be difficult to understand the underlying mechanism that generates this effect. Figures 3.21 and 3.24 shed some light on this. Homeownership appears to be no lower, and generally higher, in the first part of the life cycle in the scenarios with high variance and compensating variations in assets than in the scenario with low (zero) variance and zero assets. This behaviour early in life is consistent with the fact that the compensating variation in initial assets relaxes the explicit borrowing constraints (downpayment and income-related borrowing constraint) that in the baseline scenario bind and initially prevent households from buying. In the second part of the life cycle home-ownership can be lower in the scenario with compensating variations in assets. This happens predominantly in those cases in which the baseline scenario is one with deterministic wages and home-ownership is almost 100% until the end of the

life cycle. This kind of home-ownership behaviour in the second part of the life cycle in the two scenarios is related to the existence of an implicit “no-bankruptcy” constraint in our model, whereby the extent of a household’s debt must be less than the minimum possible flow of income plus the minimum house price. Towards the end of life, the minimum flow of future income is much lower in the scenario with stochastic wages than in the one with deterministic wages, so that with stochastic wages the constraint binds for some households, forcing them to sell the house.

A comparison between Table 3.7 and Table 3.5 gives us an idea of the importance of measuring welfare by allowing all of our choice variables to vary across different scenarios and by taking into account the income-related constraint for home-owners. In an economy where households are subject to an income-related constraint, when comparing a scenario with log-wage variance $\sigma_{\xi}^2 = 0.018$ to a scenario with deterministic wages, the compensating variation in consumption throughout the life cycle for, say, the high education group is 6.70% instead of 15.12% when we take into account changes in labour supply and home-ownership as well as potential constraints on labour supply due to the endogenous income constraint. Home ownership has a two-part role in determining these differences. Since owning a house directly increases both utility and (through its high return) wealth, the fact that the compensating variation in assets allows households to buy a house earlier, is an important determinant of the reduction of the compensating variation in consumption.

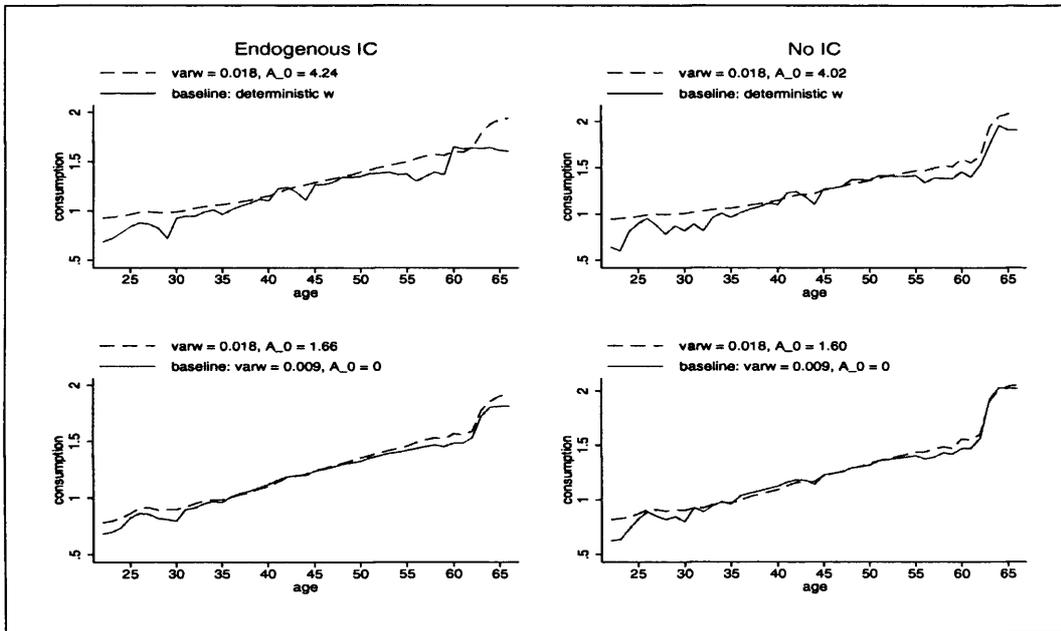


Figure 3.19: Low education: Change in consumption

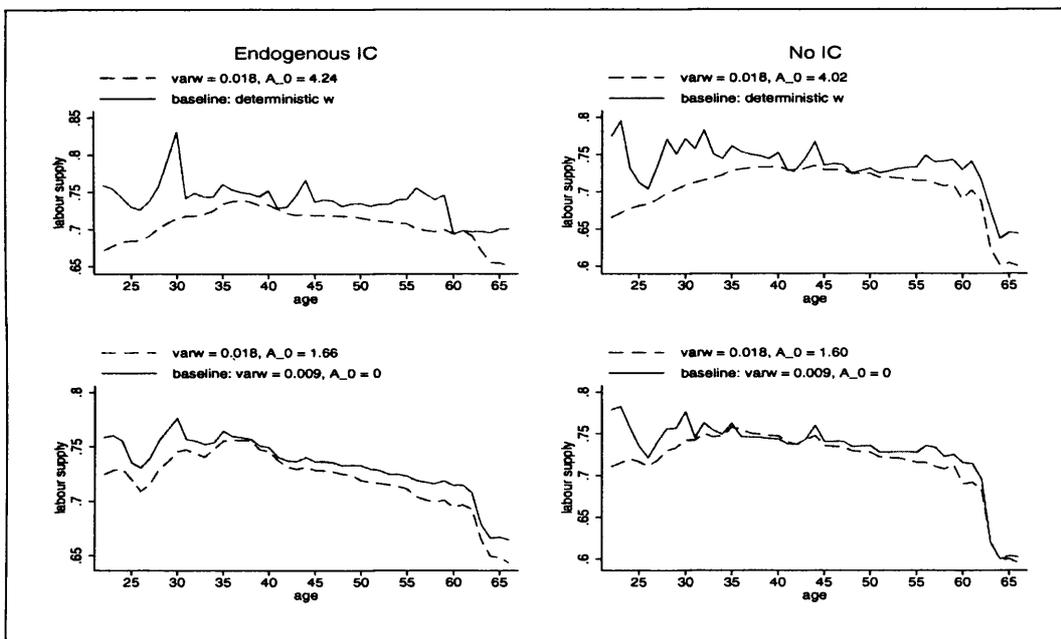


Figure 3.20: Low education: Change in labour supply

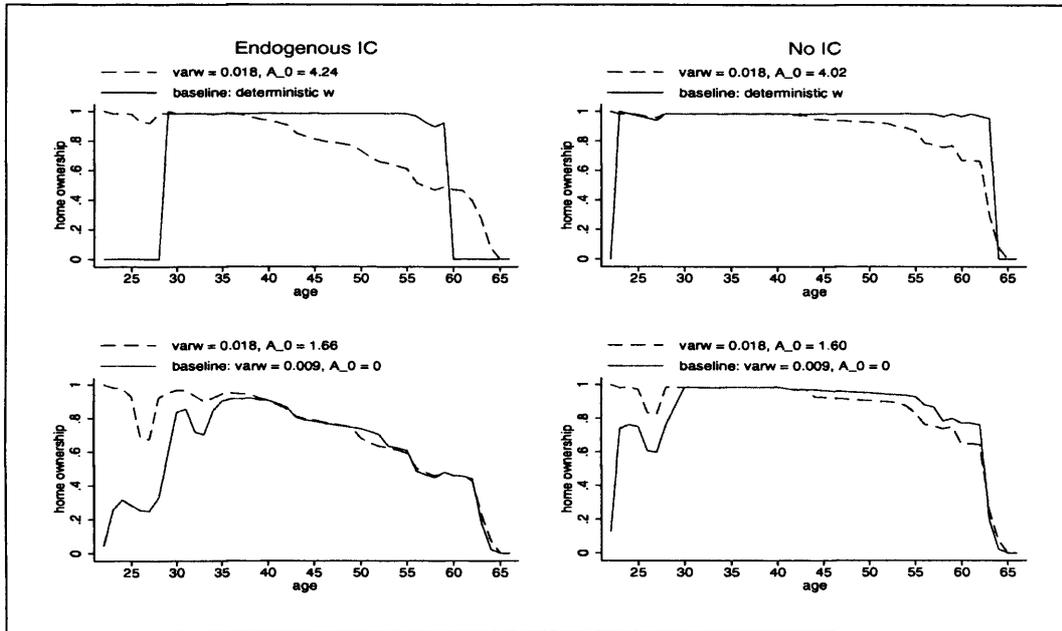


Figure 3.21: Low education: Change in home-ownership

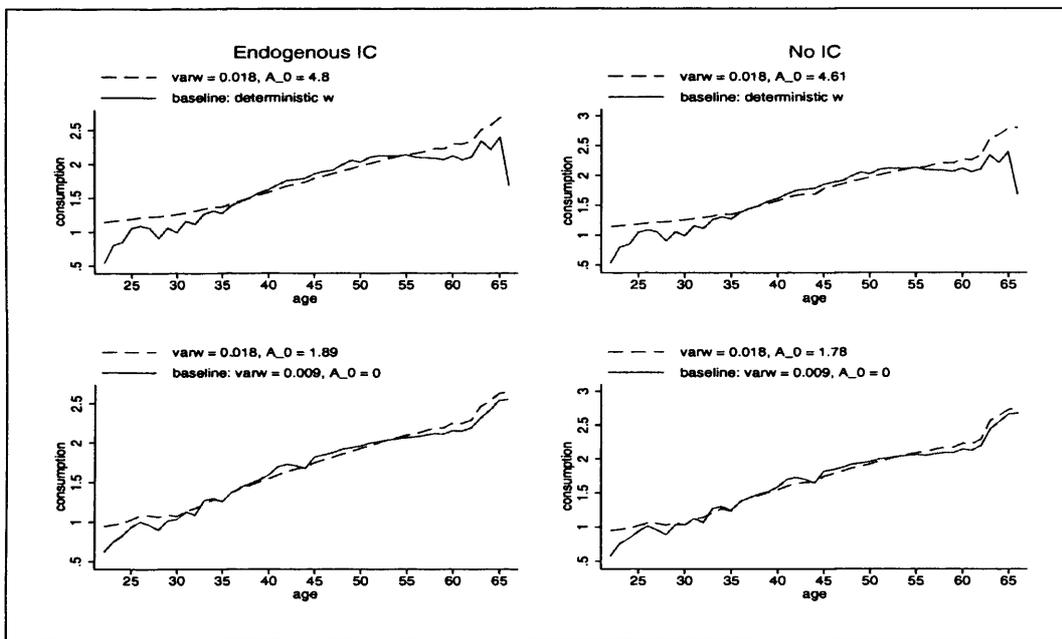


Figure 3.22: High education: Change in consumption

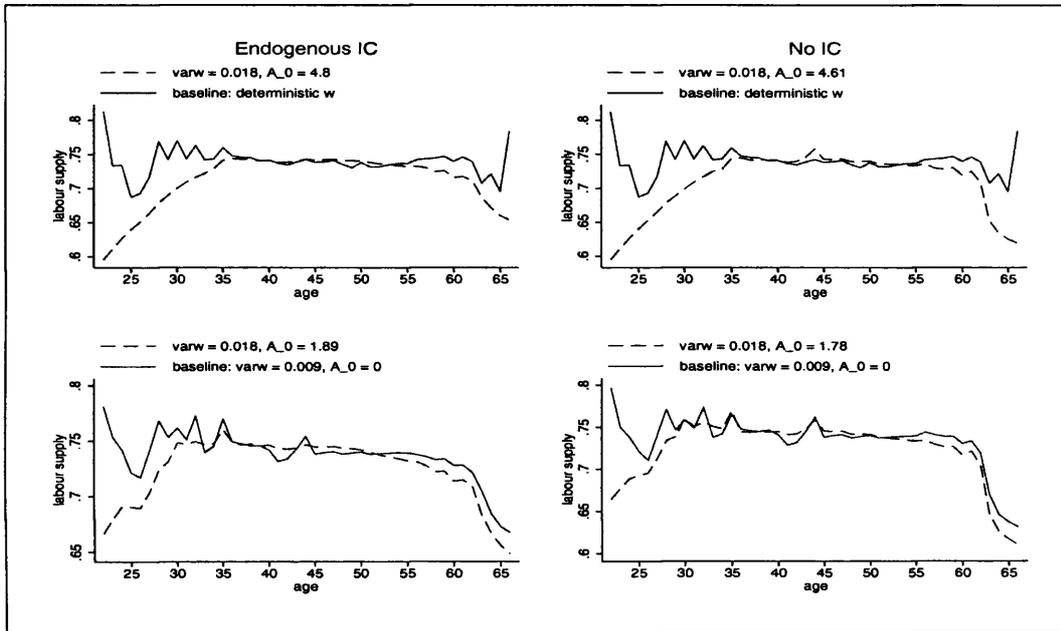


Figure 3.23: High education: Change in labour supply

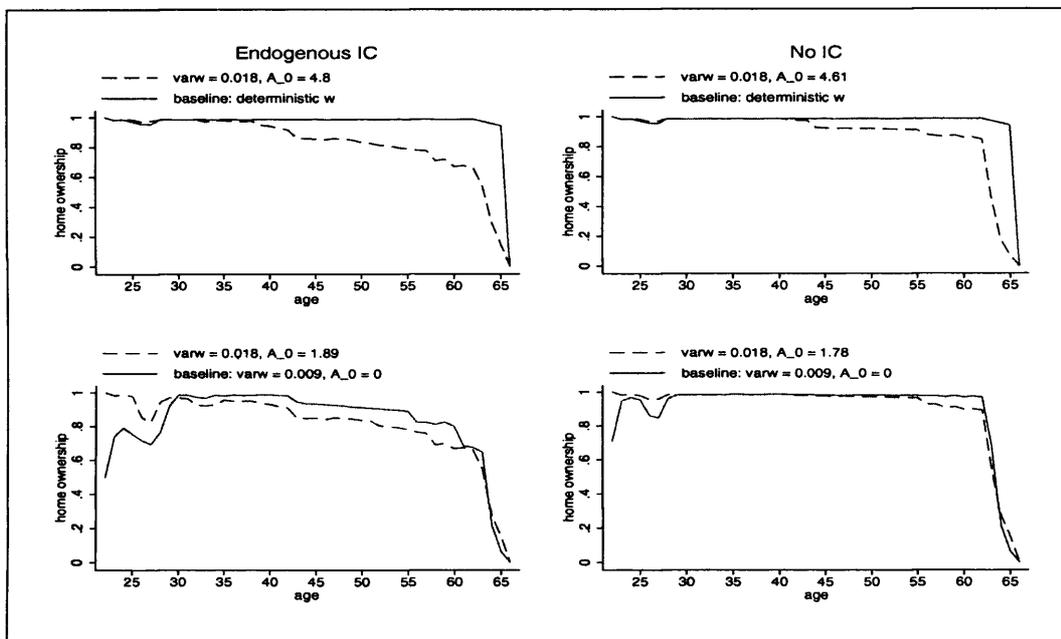


Figure 3.24: High education: Change in home-ownership

3.9 Conclusions

In this chapter we have built a structural life-cycle model of consumption and saving, labour supply and home-ownership decisions. We were motivated by two main goals. First, given empirical evidence suggesting a positive relation between labour supply and mortgage debt, as well as an institutional earnings-related borrowing constraint when taking out a mortgage to purchase a house, we wanted to analyse household behaviour within a set-up that would allow us to conduct experiments for analysing responses to changes in the economic environment. In particular, the existence of the institutional earnings-related constraint, makes labour supply and home-ownership decisions jointly determined. We therefore wanted to try and disentangle these choices by analysing labour supply behaviour of a home-owner (with mortgage-debt) in the counterfactual world in which he were not such. The second, connected, goal was to analyse home-ownership decisions over the life-cycle.

The model investigates the behaviour of households that differ ex-ante in the education level, and ex-post also in wage realisations. Their decision to purchase a house is subject to common institutional features, such as a downpayment constraint, an earnings-related constraint and transaction costs of buying and selling the house. Borrowing is only allowed for home-owners since the house serves as collateral. We calibrate the model to average life-cycle home-ownership profiles in the UK and obtain predictions for labour supply behaviour. We find that labour supply is higher for home-owners (borrowers), although the effect on labour supply is lower than data suggests, and that labour supply peaks at times of house purchase, when behaviour is constrained.

Investigating the consequences of changes in the institutional borrowing constraints produced the expected effects on home-ownership. However, whereas increasing the downpayment constraint only delays home purchase, decreasing the debt-to-income ratio reduces home-ownership at all ages. The change in labour supply is again small but consistent with the findings in the rest of the chapter.

Welfare implications of wage uncertainty have also been analysed. In particular, we wanted to examine to what extent a change in earnings uncertainty would affect house-

hold welfare not only through labour supply but also through home-ownership. Taking compensating variations in assets as a measure of welfare cost of wage uncertainty, we find that, particularly for the low education group, a higher compensation needs to be introduced in a case in which a scenario with wage variance is compared to a scenario with deterministic wages relative to a case in which the wage variance is doubled. Moreover, low education households subject to an income-related borrowing constraint would translate the compensating variation in assets into higher lifetime consumption, higher home-ownership and lower labour supply. This would not necessarily be true when the same households are not subject to an income-related borrowing constraint. In particular, home-ownership could be lower, despite the compensating variation in assets, in a scenario with stochastic wages than in a scenario with deterministic wages. This is a consequence of the extremely high home-ownership profile that results when wages are deterministic and there are no endogenous borrowing constraints.

The computational-intensive nature of the analysis associated to the model presented in this chapter forced us to make a number of simplifying assumptions, which are to be relaxed in future research. In particular, the assumption of a single earner in the household makes it hard to compare simulated profiles of labour supply to those observed in data, such as that analysed in chapter 2. A not-too-costly addition to our model would be to include a full-time worker as well as a second earner who decides how much labour to supply.

A second strong limitation was assuming the absence of a mortgage repayment schedule. Although we imposed that the income-related borrowing constraint would hold at every period, interpreting this as a proxy for a mortgage repayment scheme, we cannot be sure that labour supply responses in a model with such a scheme would be the same as those that we have examined.

A further extension would allow for more than one type (size) of housing. This would give an opportunity to investigate what parameterisations lead to households initially purchasing relatively small homes in order to circumvent constraints on labour supply, and trading up later in the life-cycle, even when fixed costs are present. However, this extension would add substantially to the computational burden.

3.10 Appendix to chapter 3

3.10.1 Computational methods

We solve our model via value function iteration, the main reason for this being the non-convexity of the value function.¹⁸ The solution for consumption, labour supply and home-ownership is found recursively from the last period of life, T , backwards. Since there is no bequest motive in our model, and since households are subject to the constraint of zero or positive assets at the end of life, they spend all that is available to them in the last period and allocate this optimally between consumption and labour supply, given the optimal choice of home-ownership. Given the optimal choices at $t+1$, $t < T$, they then need to choose home ownership, consumption, leisure and saving that maximise period t 's value function, subject to the borrowing constraints.

In order to compute the solution, we discretise the grids for the wage, the house price and the interest rate processes. Moreover, we choose a finite number of points for the asset grid. The wage and the house price processes are modeled as finite-state Markov chains that mimic the underlying continuous-valued AR(1) processes, as described in Tauchen (1986). Each of them is represented by 8 nodes. The interest rate is represented by 3 nodes, which are chosen by Gauss-quadrature rule. The asset grid contains 100 nodes for each home-ownership status at previous period (i.e. we have two conditional asset grids, particularly in relation to the fact that we only allow for collateralised borrowing and so only those owning a house at period $t-1$ can enter period t with negative financial assets). 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.

We store optimal decisions and value functions at grid points but households choices are not restricted to coincide with these points. We do linear interpolation in all the cases in which choices lie between points.

The profiles of behaviour reported in the chapter are obtained by simulating, for

¹⁸Non-convexities in the value function arise from having transaction costs associated to buying and selling the house.

each education group, an economy with 10,000 individuals. Within education group, households differ according to the wage shock that they face at each period, given that both the house price and the interest rate are common across households.

3.10.2 Estimation of wage variance

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

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

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

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

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

$$\text{var}(\varsigma_{a,t}) = \text{cov}(\Delta y_{a,t}, \Delta y_{a-1,t-1} + \Delta y_{a,t} + \Delta y_{a+1,t+1}) \quad (3.16)$$

Recalling that we only allow for permanent shocks in our wage process (equation (3.9)), and assuming that in our data the transitory shocks are uncorrelated, we base the estimation of the variance of permanent shocks on equation (3.16). In particular, we need to obtain real wages net of predictable individual components and use their

lags to construct the covariance between a one-period lag and a three-period lag, as follows:

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

where \tilde{w} is the logarithm of real wages net of predictable individual components, obtained as the residual from a regression, for each education group, of individual log-real wages on year, cohort and region dummies, a dummy for being married, household size, and number of children in the household.

3.10.3 Labour supply with binding income-related borrowing constraint

Let us consider our model (equations (3.1)-(3.5)) in the case in which a household owns a house in period t ($h_t = 1$) and the income-related borrowing constraint (3.5) binds. The constraints (3.2)-(3.5) boil down to the following constraint:

$$\tilde{A}_t = -\lambda_y w_t (L - l_t)$$

or

$$\hat{A}_t + w_t(l - l_t) - c_t = -\lambda_y w_t (L - l_t) \quad (3.17)$$

where $\hat{A}_t = \begin{cases} A_t - p_t h_t (1 + F) & \text{if buy house} \\ A_t & \text{if } h_t = h_{t-1} = 1. \end{cases}$

The Lagrangian for this problem is the following

$$\mathcal{L} = [u(c_t, h_t, l_t) + \beta EV_{t+1}(A_{t+1}, h_t, p_{t+1}, w_{t+1})] + \mu[\tilde{A}_t + w_t(L - l_t) - c_t + \lambda_y w_t(L - l_t)]$$

and the first-order conditions for consumption and leisure are therefore

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial c_t} = 0 &\Rightarrow \frac{\partial u(c_t, h_t, l_t)}{\partial c_t} = \mu \\ \frac{\partial \mathcal{L}}{\partial l_t} \geq 0 &\Rightarrow \frac{\partial u(c_t, h_t, l_t)}{\partial l_t} \geq \mu(1 + \lambda)w_t \end{aligned}$$

From the first order conditions we then obtain the intra-temporal allocation rule between consumption and leisure (when leisure is not at the corner solution $l_t = L$):

$$\frac{\partial u(c_t, h_t, l_t)/\partial l_t}{\partial u(c_t, h_t, l_t)/\partial c_t} = (1 + \lambda)w_t. \quad (3.18)$$

When the within-period utility function between consumption and leisure is Cobb-Douglas, as in our baseline model, the intra-temporal allocation rule (3.18) takes the form:

$$\frac{(1 - \eta)c_t}{\eta l_t} = (1 + \lambda)w_t$$

giving

$$c_t = \frac{\eta}{(1 - \eta)}(1 + \lambda)w_t l_t. \quad (3.19)$$

Combining (3.19) with constraint (3.17), we then obtain the labour supply function for constrained households:

$$L - l_t = \eta L - \frac{(1 - \eta)\tilde{A}_t}{(1 + \lambda)w_t} \quad (3.20)$$

where $L - l_t \equiv l_{s_t}$.

Chapter 4

Retirement Expectations, Pension Reforms, and Their Impact on Private Wealth Accumulation

4.1 Introduction

In all industrialized countries pension benefits represent a major component of retirement income, and therefore social security arrangements can have important effects on households' intertemporal choices. One of the most important issues in this area is to what extent individuals perceive and react to changes in pension legislation. Do people increase their saving and labor supply in response to a reduction in pension benefits? Is private wealth a good substitute for mandated accumulation in the form of social security contributions?

Answers to these questions usually proceed in two steps. In a first step, researchers estimate expected pension wealth, that is, the expected present discounted value of future benefits that workers are entitled to. In a second step, expected pension wealth is related to discretionary wealth and/or labor supply behavior. Difficult methodological

problems are encountered at each of these steps. The first step requires a model of the way in which individuals form expectations about future pension legislation. The second step requires to control for the possible endogeneity of expected pension wealth, and specifically of labor supply and retirement decisions, with respect to discretionary wealth accumulation decisions.

Even in the simplest scenarios, estimating future pension benefits is a difficult task. For the working population, expected pension wealth depends, among other variables, on the age at which workers expect to retire and on the expected ratio of pension benefits to pre-retirement earnings (the replacement rate). The standard approach taken in the literature is to estimate these variables from current and projected legislation on pension eligibility rules, accrual rates of contributions, productivity growth and mortality projections. The estimated pension wealth is then used for simulation analysis, to project the future path of social expenditures, or for estimating the impact of pension wealth on retirement decisions and private accumulation. Feldstein (1974) pioneered the analysis of the displacement effect of pension wealth on national saving using U.S. time series data. Since then, a growing literature has used individual level data to provide evidence on the degree of substitution between discretionary accumulation and pension wealth in the U.S. and other countries imputing pension wealth from legislation. Recent attempts include Gale (1998), Attanasio and Brugiavini (2003), and Attanasio and Rohwedder (2003) who use, respectively, U.S., Italian and U.K. microeconomic data and find that pension wealth crowds out discretionary saving, but at a rate of considerably less than one-for-one.¹

A different approach to analyzing the impact of social security on individual decisions relies on subjective expectations of retirement ages and benefits (Bernheim, 1990; Gustman and Steinmeier, 2001). This literature has been concerned with a rather different set of issues, which are, to a large extent, preliminary to the analysis of the effect of pension wealth on private wealth accumulation. Specifically, it analyses the degree of workers' information about the retirement benefits they are entitled to, the relation

¹Gruber and Wise (1999) use estimates of pension wealth to calculate the effects of pension arrangements on the retirement decision and on the labor force participation of the elderly.

between planned and actual retirement age, and the determinants of the probability distribution of expected retirement age (Disney and Tanner, 1999; Dominitz, Manski and Heinz, 2002).

The Survey of Household Income and Wealth (SHIW), a large representative survey of the Italian population carried out by the Bank of Italy, elicits retirement age and replacement rate expectations from 1989 to 2002. This is not the only survey eliciting such expectations but, to our knowledge, it is the only survey in which this information is available for an extended period spanning a set of intense pension reforms. During the period, the Italian government enacted three reforms (in 1992, 1995, and 1997), whose ultimate effect was to reduce the replacement rate of young workers relative to older cohorts. This chapter attempts to estimate the impact of these reforms on people's perceptions about their future replacement rate – a convenient summary index of the generosity of the pension system.

The analysis then focuses on the relation between expected pension wealth and private wealth. The reforms differently various cohorts and employment groups, providing an exogenous source of variability in pension wealth and an ideal instrument to estimate the offset between private and pension wealth.

Our framework calls attention to the fact that the effect of pension reforms on individual decisions depends on the extent to which people understand and react to the changes implied by the reform. The standard life-cycle hypothesis posits that a reduction in expected pension benefits should increase private wealth during the working life one-for-one. This offset is what Feldstein calls the substitution effect - pension wealth crowds out discretionary wealth. There are several potential counter-effects to a complete crowding out. Bequest motives, short-sightedness, liquidity constraints, risk associated with future reforms, and non-marketable future benefits are among the most cited reasons to explain why the offset between private and pension wealth might well be less than one-for-one. But there is another element that is potentially important: when pension reforms represent a permanent shift, individuals might not change their behavior, or adjust only partially to the new economic environment, because they are not informed, do not understand how the reform will affect their benefits or because

changes in expectations occur slowly. This is one of the elements that we investigate here.

In doing this we answer two relevant policy questions. First, to what extent do pension reforms affect workers' expectations? Second, provided that expectations are revised, how do these revisions affect retirement decisions and discretionary wealth accumulation? Previous literature does not distinguish between these two questions, and looks directly at the link between pension arrangements and saving decisions. Answering the first question is quite important in understanding to what extent people offset reductions in pension wealth after major pension reforms. As we shall see, answers to the first question also provide important empirical tools to address the second question.

The chapter is organized as follows. Section 4.2 illustrates the Italian pension reforms of the last decade and discusses previous evidence. Section 4.3 presents the data on expectations on retirement outcomes available in the 1989-2002 Survey of Household Income and Wealth, providing the ground for our instrumental variables approach. Section 4.4 estimates the impact of pension reforms on the expected replacement rate exploiting the ample variability in the effects of the legislation on different demographic and economic groups. The main finding is that workers have revised expectations in the direction suggested by the reform, but the adjustment is far from complete. Section 4.5 relates discretionary wealth to expected pension wealth, using the variations in the effects of the reforms over time and across demographic groups, to construct an instrument for pension wealth. The empirical estimates suggest an offset between private wealth and expected pension wealth of about 50 percent.

Although the estimated substitution coefficient is on the high side of current estimates, we find that so far the Italian pension reforms of the 1990s had limited impact on private wealth accumulation, because people have revised only in part their expected pension wealth after the reform. Section 4.6 concludes by drawing attention to the crucial role of financial information and suggesting that in the coming decades a problem of inadequate savings could emerge for the cohorts most affected by the reforms.

4.2 The Italian pension system: a decade of reforms

Until recently, the Italian social security system featured high replacement rates, earnings-based benefits, indexation of pensions to real earnings and cost of living, generous provisions for early retirement, and a large number of social pensions (i.e. old-age income assistance). These features of the social security system were gradually implemented and extended during the post-war period, and especially between 1967 and 1975. The result was that the ratio of pension benefits to GNP reached almost 16 percent in 1992, the highest value among industrialized countries.

The late eighties and early nineties saw increasing alarm over the growing imbalance of the social security system expressed not only by economists and in official government documents, but also in the media. In the second half of 1992 the Amato government presented a fiscal package containing a major reform of the social security system. In 1995 Italy underwent a second major reform of the social security system, known as Dini reform. Social security legislation was further refined in December of 1997 by the Prodi government.

The main features of the reforms were an increase in the retirement age and minimum years of contribution for pension eligibility, abolishment of seniority pensions for all those who started working after 1995, a gradual reduction in pension benefits, and indexation of pension benefits to prices rather than to wages. The three reforms maintained the generous provisions of the pre-1992 regime for the relatively old workers, who in 1995 had at least 18 years of contributions, and different rules for private employees, public sector employees and self-employed.

Although the current regime combines some features of each of the three reforms, we do not detail their specific aspects.² In fact, we compare pension regimes and individual expectations *before the 1992 reform* and *after the 1997 reform*, omitting the transitional years between the Amato and Prodi reforms (1992-1997). Our data set allows us to observe workers in two regimes, one with generous provisions (before the Amato reform, or simply the pre-reform period) and one – ten years later – with

²Brugiavini (1999) provides ample details on the specific features of the sequence of the three Italian pension reforms.

much lower benefits (after the Prodi reform, or the post-reform period), at least for some categories of workers. We regard the availability of low frequency microeconomic data as a major improvement with respect to previous evidence.

The top panel of Table 4.1 compares statutory retirement ages in the pre-1992 regime with the post-1997 regime. For brevity we refer to workers with more than 18 years of contributions in 1995 as the “old”, to those with less than 18 years of contributions in 1995 as the “middle aged”, and to those who started working after 1995 as the “young”. In the new regime the young are entitled to a flexible retirement age (from 57 to 65), subject to incentives. For those already working in 1995 (the old and the middle-aged), the reform raises minimum retirement age for old age pensions of private sector employees (65 for men and 60 for women), but not for public employees and self-employed. For the old and middle aged, the reform raises minimum years of contributions for both seniority pensions and old age pensions; for the young, whose pension award formula is entirely contribution based (see below) the minimum years of contributions is just 5 years.

The shift to the new regime dramatically altered the pension award formula for new cohorts, but retained the main features of the pre-1992 formula for older workers. As indicated in the lower panel of Table 4.1, for the young the reform introduced contributions-based pension benefits. Specifically, in the new regime the pension is proportional to contributions, capitalized on the basis of a 5 years moving average of GDP growth. Since the contribution rate is 33 percent for private and public employees and 20 percent for the self-employed, in the new regime the self-employed will receive substantially lower pensions than employees. Actuarial equilibrium of the system is guaranteed by multiplying the sum of the contributions by a coefficient that takes into account life expectancy at retirement. The contributions-based model has identical minimum retirement age for males and females, in both old age and seniority pensions. However, the new regime applies only to the young cohorts, who entered the labor market after 1995, and will presumably start to retire after the year 2030.

For older workers, pensions are still computed using the earnings model. For the private sector, for instance, the pension is obtained as the number of years of con-

Table 4.1: Retirement age and pension award formula after the Italian pension reforms

RETIREMENT AGE				
Old age pensions				
	<i>Minimum retirement age</i>			<i>Minimum years of contributions</i>
	Private sector	Public sector	Self-employed	
Pre-1992 regime				
All workers	60(55)	65(60)	65(60)	15
Post-1997 regime				
Old	Progressively rising to 65(60)	65(60)	65(60)	Progressively rising to 20
Middle-aged	Progressively rising to 65(60)	65(60)	65(60)	Progressively rising to 20
Young	Subject to incentives: 57-65			5
Seniority pensions				
	<i>Minimum years of contributions</i>			
Pre-1992 regime				
All workers	35	20	35	
Post-1997 regime				
Old	40 before age 57 35 after age 57		40 before age 58 35 after age 58	
Middle-aged	40 before age 57 35 after age 57		40 before age 58 35 after age 58	
Young	Abolished			
PENSION AWARD FORMULA				
	Private sector	Public sector	Self-employed	
Pre-1992 regime				
All workers (Earnings model)	2% of $N_{contrib} * \bar{y}_{-5}$	2.33% of $N_{contrib} * \bar{y}_{-1}$	2% of $N_{contrib} * \bar{y}_{-10}$	
Post-1997 regime				
Old (Earnings model)	Gradually to 2% of $N_{contrib} * \bar{y}_{-10}$	Gradually to 2% of $N_{contrib} * \bar{y}_{-10}$	Gradually to 2% of $N_{contrib} * \bar{y}_{-15}$	
Middle-aged (Pro-rata model)	Earnings model before 1995, contribution model after 1995			
Young (Contribution model)	Contributions capitalised at 5-yr moving average of GDP growth. Capitalised sum multiplied by α_{RetAge}			

Notes: Old, middle-aged and young refer, respectively, to workers with more than 18 years of contributions in 1995, less than 18 years of contributions in 1995, or started working after 1995.

In the top panel, female retirement age is in brackets when different from males.

$N_{contrib}$ = number of years of contribution. \bar{y}_{-t} = average of the final t years of earnings. α_{RetAge} = coefficient that is a function of retirement age, to take into account life expectancy.

tributions, times 2 percent of the average of the last 10 years of salaries. For the middle-aged (less than 18 years of contributions as of 1995), pensions are computed according to a “pro-rata model”: earnings-related for working years before 1995, and contribution-related afterwards. The Appendix (section 4.7.1) provides further details on pension award formula before and after the reform, and of the specific provisions for public and private employees and self-employed.

While Table 4.1 provides a qualitative assessment of the pension reform, in order to show the magnitudes involved, in Table 4.2 we compute statutory replacement rates of before and after the reform of a worker retiring at 62 years, after 37 years of contributions. The example posits that the growth rate of individual earnings is 2 percent, and that the aggregate GDP growth rate is 1.5 percent. We distinguish between three categories of workers (private, public, self-employed), three cohorts (old, middle-age, young) and two periods (before and after the reform). The replacement rate is defined as the ratio of the first year’s pension to the last year’s earnings.³

In the pre-reform regime the replacement rates were the same for old, middle-aged and young workers, because the earnings model applied to all. However, in that regime replacement rates did differ considerably across occupational groups: 71.1 percent for private employees, 86.2 percent for public employees and 67.8 percent for the self-employed. The higher rates for public employees reflect more generous accrual rates (see Table 4.1) and pension award formulas (pensionable earnings were just the last salary).

After the reform workers are distinguished according to the number of years of contributions in 1995. In our example we still posit that each worker plans to retire after 37 years of work, but distinguish between an old worker with 27 years of contributions in 1995, a middle-aged with 10 years of contributions in 1995, and a young person who starts working in 1996. After the reform, the replacement rates of old private employees and self-employed are practically unaffected (-1 and -0.9 percent), while that of the old public employees falls by 5 percentage points. This differential

³We do not distinguish here between males and females, because the same pension accrual formula applies to both groups.

Table 4.2: The statutory replacement rate before and after the pension reform

	Pre-reform	Post-reform	Difference
<i>Private employees</i>			
Old	71.1	70.1	-1.0
Middle-aged	71.1	64.0	-7.1
Young	71.1	61.7	-9.4
<i>Public employees</i>			
Old	86.2	81.2	-5.0
Middle-aged	86.2	66.9	-19.3
Young	86.2	61.7	-24.5
<i>Self-employed</i>			
Old	67.8	66.9	-0.9
Middle-aged	67.8	45.2	-22.6
Young	67.8	37.4	-30.4

Notes: In the post-reform period, for middle-aged and young the replacement rate is based on the assumptions that the growth rate of earnings is 2 percent per year and the growth rate of aggregate GDP is 1.5 percent. In all cases the retirement age is 62, and each worker contributes for 37 year before retiring. In the post-reform regime the example considers an old worker who contributes 27 years before 1995 and 10 years after, and a middle-aged worker who contributes 10 years before 1995 and 27 years after.

effect is largely due to the reduced accrual rate of public employees (from 2.33 to 2 percent).

In contrast, middle-aged and especially young workers experience a much more dramatic reduction in replacement rates due to the reform. For private employees the change is -7.1 points for the middle-aged and -9.4 for the young; for public employees, -19.3 and -24.5 percent, respectively; and for the self-employed -22.6 and -30.4.

In summary, Table 4.2 shows that the reform has reduced pension benefits for the middle aged and the young, and for all cohorts of public employees. The implied magnitudes of change are substantial, because for some of the categories involved the replacement rate falls by over 20 percentage points. On the other hand, old private employees and old self-employed workers were basically unaffected by the reform. The

Italian pension reforms therefore provide a quasi-experimental framework to analyze the impact of reforms on individual expectations. Since the reform affects some population groups (the middle-aged, the young, and public employees) more dramatically than others (old private employees and old self-employed), we can study the impact of the reform by comparing the changes in the behavior of different groups of individuals before and after the reform.

4.3 Expectations of retirement outcomes

A recent strand of literature has analyzed the role of expectations in determining retirement outcomes in the U.S. and Europe. In general, the literature finds that expectations are reasonably informative about retirement outcomes, but also uncovers substantial heterogeneity across the population and reveals that many workers lack knowledge about the details of their pension plans. The earliest paper is Bernheim (1990), who compares retirement expectations and realizations in the U.S. Retirement History Survey and finds that men and wealthier individuals make more accurate plans. Disney and Tanner (1999) analyze expectations of retirement age in the U.K. Retirement Survey, and find that marital status and education have a significant effect in explaining systematic deviations of expectations from outcomes. The focus of the paper is on the distribution of actual retirement age, conditional on a given expected retirement age, rather than on the overall distribution of expectations and realizations. Gustman and Steinmeier (2001) use data from the U.S. Health and Retirement Study to analyze the degree of information about social security and private pensions. They find only a weak relation between expected retirement benefits and benefits estimated on the basis of social security earnings records and employers' descriptions of pension plans.

In this chapter, like most of the recent studies, we use point expectations.⁴ How-

⁴Some studies focus on the subjective probability distribution of retirement outcomes, rather than on point expectations of retirement age and benefits. Hurd and McGarry (1995) analyze the subjective probability distribution of the chance of working full-time past age 62 and of living to age 75 in the U.S. Health and Retirement Study. Dominitz, Manski and Heinz (2002) use the Survey of Economic

ever, while the focus of previous literature is mainly on expected retirement age, we focus on expected replacement rate, defined as the expected ratio of the first pension to the last salary. For any given expected retirement age, the replacement rate is a convenient summary measure of the generosity of the pension system and therefore a good proxy for expected pension wealth. Our data are repeated cross-sections, as opposed to the longitudinal data provided by the Health and Retirement Surveys in the U.S. or the U.K. The main advantage of the data used in the present study is that the sample spans a period set of intense pension reforms, which deeply changed the social security system.

The survey – the SHIW – is a large random sample of the Italian population drawn by the Bank of Italy every two years. Sample design, interviewing procedure, response rates and a comparison between sample and population means are reported in the Appendix (section 4.7.3). The survey covers several important topics related to retirement and pensions and collects data on the subjective assessment of expected retirement age and replacement rate. All workers (public employees, private employees, self-employed) are asked the following questions:

- *When do you expect to retire?*
- *Think about when you will retire, and consider only the public pension (that is, exclude private pensions, if you have one). At the time of retirement, what fraction of labor income will your public pension be?*

In Italy only about 5 percent of the workers are covered by occupational pension schemes, so for the overwhelming majority the social security replacement rate coincides with the overall replacement rate. The first question is posed in each survey year from 1989 to 2002; the second question only in 1989, 1991, 2000 and 2002. Since we are interested in studying workers' expectations about retirement income, we focus on the group aged 20 to 50 years. This implies that we include in our sample individuals born between 1939 (who were 50 years old in 1989) and 1982 (20 years old in 2002).

Expectations, which elicits the subjective probability distributions of eligibility for social security benefits and of the level of benefits. They report a high degree of uncertainty about future benefits even for people only ten years from retirement.

The composition of the sample of older workers is likely to reflect self-selection into higher expected retirement ages, and so these workers are dropped from the analysis. A small number of individuals younger than 20 are also excluded (less than 1 percent of the sample). We focus on how expectations change after the reform and therefore drop workers that are interviewed in the transitional years.

We define as the pre-reform period the pooled 1989-91 sample, while the post-reform period is the pooled 2000-02 sample. Finally, we consider only workers who are employed or self-employed in the survey year, excluding the unemployed, retirees and other individuals not in the labor force. Overall, we have valid responses on expected replacement rate for 9,766 males and 5,955 females.

As explained in Section 4.2, the pension reform has different effects depending on whether workers had contributed for more or less than 18 years at the end of 1995, and different again for those who started working after 1995. The SHIW records the age at which individuals started working. This allows us to compute the years of contribution at the end of 1995 for each worker and to define our groups accordingly.⁵

As a preliminary step, we compare the replacement rate that people expect with the rate that they should expect, given the relevant pension legislation at the time of the interview and the declared expected retirement age (henceforth “statutory replacement rate”). Pooling all observations, we find that for 75 percent of the sample the expected replacement rate is higher than the statutory rate. Expectation errors are higher for private employees, for the better educated and for females (for brevity, these regressions are not reported). These findings are in line with previous research, which generally concludes that there is considerable heterogeneity in expectations, and that many workers lack precise knowledge about their public pensions.⁶ What is most interesting,

⁵Our imputation procedure assumes no unemployment spells during the working life and is therefore subject to a certain amount of measurement error. As a sensitivity check, we assume that each individual starts working and contributing at age 20 (or 22) and define years of contribution as current age less 20 (or 22). These alternative definitions do not affect any of our results.

⁶Other surveys confirm that predictions of pension-related variables are not accurate. Boeri, Brsch-Supan and Tabellini (2001) analyze the results of a recent European survey on 1,000 households showing that only two thirds of individuals give the correct answer when asked about the social security contribution rate.

however, is to compare the replacement rate of specific employment groups and cohorts before and after the reform.

4.4 The effect of the pension reform on expectations

We use a difference-in-difference framework to study how the expected replacement rate has been affected by the three pension reforms. As with other studies that use a quasi-experimental framework, our tests rely on the assumptions that the pension reform is exogenous with respect to individual decisions – in particular, with respect to retirement age – and changes in sample composition. As far as the first assumption is concerned, we believe that the possible endogeneity of the reform can be safely ruled out. The reform was not implemented in order to offset different paths of retirement ages by different cohorts or employment groups. Rather, the 1992 reform was part of a major deficit-reduction package, prompted by a severe political crisis coupled with the dramatic devaluation of the lira; and it was followed shortly by the deepest recession of the post-war era. The 1995 and 1997 reforms were prompted by the huge projected deficits of the social security system and the attempt to meet the Maastricht criteria.

The second assumption posits that shifts in sample composition are exogenous with respect to pension expectations. Cohorts and gender are obviously determined at birth. As far as employment groups are concerned, we require that mobility across various sectors (for instance, from public to private employment or self-employment) are independent of pension expectations, i.e. that workers did not switch jobs as a result of the pension reform itself. Since the SHIW has a rotating panel component, we can check the validity of this assumption by computing the transition rates across the three employment groups between each pair of adjacent survey years from 1989 to 2002; the Appendix (section 4.7.4) reports the transition rates for 1989-91 and 2000-02. We find that, in each period, the probability of not changing sector is about 90 percent for each of the three groups. Furthermore, we do not reject the hypothesis that the degree of sector mobility is the same before and after the reform for each of the estimated transition matrices.⁷ Although we cannot test directly the hypothesis that workers

⁷As an example, consider the Shorrocks mobility index in 1989-91 (12.5 percent) and 2000-02 (13

did not change sector as a consequence of the reform, we take this as indirect evidence that the pension reform has not changed the overall pattern of workers' mobility across sectors.

4.4.1 Descriptive analysis

Table 4.3 reports the expected replacement rate of males and females in the various employment groups and cohorts considered. On average, the rate is high for all groups, for both males and females, reflecting the generous provisions of the Italian social security system. The expected rate ranges from 65.3 to 81.8 percent before the reform, and from 57.3 to 79.9 percent afterwards, and attains the highest value for public employees (between 74 and 82 percent for males, and between 69 and 74 percent for females). On the other hand, the self-employed report the lowest replacement rates.

The expected replacement rates decrease after the reform, for both males and females, and for all employment groups. For males, the reduction of the middle-aged is stronger than for the old, particularly for private employees (-8.4 percentage points) and self-employed (-12.2 points). Replacement rates also fall for females, but the difference between the old and the middle-aged is not as large as that for males.⁸

Qualitatively, the reduction in the expected replacement rate is consistent with several features of the reform. However, for most groups the magnitude of adjustment is not as large as implied by the reform. This can be seen by comparing the expected replacement rates with the statutory rates after the reform. For this purpose, we cannot use the example of Table 4.2, where we keep retirement age and years of contributions fixed. The reason is that the pension accrual formula relates the replacement rate to years of contributions, and therefore workers could offset part of the reduction in pension benefits by raising retirement age after the reform. As a result, the statutory

percent). The statistical test does not reject the hypothesis that mobility is the same in the two periods (the associated statistic is 0.12 and is normally distributed with mean zero and variance equal to one). The same test can be performed for each of the transition matrices (1991-93, 1993-95, 1995-98, 1998-2000), and in all cases we do not reject the hypothesis that mobility is constant.

⁸In Table 3, no comparison is possible for young workers because this group is not observed before the reform.

Table 4.3: Expected and statutory replacement rate: descriptive statistics

	Expected			Statutory		
	Pre-reform	Post-reform	Difference	Pre-reform	Post-reform	Difference
Males						
<i>Private employees</i>						
Old	79.2	74.9	-4.3	71.2	71.5	0.3
Middle-aged	79.2	70.8	-8.4	71.2	66.8	-4.4
Young		67.9			67.0	
<i>Public employees</i>						
Old	81.8	79.9	-1.9	86.2	80	-6.2
Middle-aged	80.6	76.6	-4	83.9	69.2	-14.7
Young		73.9			66.2	
<i>Self-employed</i>						
Old	69.0	61.2	-7.8	71.5	69.9	-1.5
Middle-aged	71.2	59.0	-12.2	69.6	49.2	-20.4
Young		61.8			41.1	
Females						
<i>Private employees</i>						
Old	77.2	71.1	-6.1	67.3	71.5	4.2
Middle-aged	76.9	70.0	-6.9	67.3	62.1	-5.2
Young		67.2			61.2	
<i>Public employees</i>						
Old	78.3	74.0	-4.3	79.2	79.5	0.3
Middle-aged	78.3	72.8	-5.5	76.9	65.4	-11.5
Young		69.1			55.3	
<i>Self-employed</i>						
Old	65.3	57.3	-8	67.8	69.9	2.1
Middle-aged	69.7	58.9	-10.8	62.3	45.2	-17.1
Young		57.4			35.5	

Notes: Data are drawn from the 1989-2002 SHIW. The pre-reform period is 1989-91, the post-reform period is 2000-02.

rates after the reform would reflect not only differences in pension rules across groups and pension regimes, but also the increase in retirement age.

To compute the statutory replacement rate, we therefore need information on retirement age after the reform. We indeed find that expected retirement age increases

after the reform (about 2 years for males and 3 years for females), which in turn entails higher replacement rates. It follows that the reduction in the statutory rates in Table 4.3 is lower than in Table 4.2, although the pattern is similar: the largest reductions in the statutory rates are for middle-aged public employees and self-employed.

We now compare the expected and the statutory rates in Table 4.3, focusing on the middle-aged, the group that is most affected by the reform. For males, the two groups that face the largest reductions in the statutory rates (public employees and self-employed) are also those whose expected replacement rates after the reform are furthest away from the statutory rates. For example, although self-employed have revised their expectations down by 12.2 percentage points after the reform, they should have further reduced them by 9.8 percentage points to reach the statutory level of 49.2. Similarly, the post-reform expected replacement rate is above the statutory one by 7.4 and 4 percentage points for, respectively, public employees and private employees. For females, the difference between post-reform expected and statutory replacement rates is similar for private and public employees (7.9 and 7.4 percentage points, respectively) and larger for the self-employed (13.7). Overall, the comparison indicates that expectations move in the direction suggested by the reform, but that the magnitude of the revision in expectations is not as large as implied by the reform.

4.4.2 Regression estimates

The drawback of looking at differences in the expected replacement rate over time is that this depends not only on the pension reforms, but also on other economy-wide phenomena. To control for other factors potentially affecting the expected replacement rate, we turn to a difference-in-difference framework. We can identify the effect of the reform on the expected replacement rate because there is one group of individuals (old private employees) that was unaffected by the reform, while other groups (public employees, self-employed, the young and the middle-aged) were affected and should have revised their expectations downward. Therefore to disentangle the effect of the reforms on expectations from other effects, such as common trends in determinants of labor supply and business cycle effects, we compare the difference over time in the

replacement rate of the middle-aged with the same difference for the old.

It is important to notice that our approach does not require panel data. What we need to observe is a representative sample of the various groups in each of the two periods and therefore rely on repeated cross-sectional data. The young cannot be used to evaluate the effect of the reform because they entered the labor market after 1995. Since they were sampled only after the reform was in place, they are dropped from the analysis.

We pool all data from pre- and post-reform periods and specify a reduced form for the expected replacement rate σ . We assume that before the reform σ is a linear function of socio-demographic variables X , employment status (private, public, self-employed) and depends on whether the years of contributions of 1995 are more or less than 18:

$$\sigma_i = X_i\beta + \alpha_0 + PUB_i\alpha_1 + SELF_i\alpha_2 + M_i\delta_1 + M_i * PUB_i\delta_2 + M_i * SELF_i\delta_3 + \varepsilon_i \quad (4.1)$$

The reference group in the regression equation is the group of old, private employees; the dummy variable M equals 1 for the middle-aged (less than 18 years of contributions as of 1995). The α coefficients capture the different rules applying to public employees (PUB) and self-employed ($SELF$) relative to private employees, whereas the δ coefficients measure the potential differences between middle-aged and old of the three employment groups. After the reform σ potentially shifts for all groups, so we augment the previous equation with terms that interact the cohort (M), the post-reform period ($POST$, where $POST$ equals one for the post-reform period) and the employment status ($SELF$, PUB):

$$\begin{aligned} \sigma_i = & X_i\beta + \alpha_0 + PUB_i\alpha_1 + SELF_i\alpha_2 + M_i\delta_1 + M_i * PUB_i\delta_2 + M_i * SELF_i\delta_3 + \\ & + POST_i\phi_1 + POST_i * PUB_i\phi_2 + POST_i * SELF_i\phi_3 + \\ & + POST_i * M_i * PUB_i\gamma_1 + POST_i * M_i * SELF_i\gamma_2 + POST_i * M_i * PRIV_i\gamma_3 + \varepsilon_i \end{aligned} \quad (4.2)$$

The ϕ coefficients capture the change in σ after the reform for the three employment groups: ϕ_1 measure the change for old private employees and ϕ_2 and ϕ_3 the additional

effects for public employees and self-employed. The γ coefficients measure the change in σ for the middle aged due to the reform, and are our main parameters of interest. We expect the reform to reduce the replacement rate ($\gamma_1 < 0, \gamma_2 < 0, \gamma_3 < 0$), and that this reduction is smallest for private employees and largest for the self-employed, as shown in Table 3 ($\gamma_3 > \gamma_1 > \gamma_2$).

The model is estimated separately for males and females, omitting the transitional 1993-1997 period. Table 4.4 reports the results. In the first specification we drop the control vector X , and regress σ on a set of group dummies. The results confirm the descriptive analysis. The coefficient estimates indicate that after the reform there is a reduction in the replacement rate of private employees by 4.2 percentage points (the estimated γ_3). The coefficients γ_1 and γ_2 for public employees and self-employed are both negative, but only γ_2 differs statistically from zero.

To benchmark the estimated γ 's, recall from Table 4.3 that σ should change by -4.4 percentage points for middle-aged private employees, -14.7 for public employees, and -20.4 for self-employed. Subtracting from these numbers the corresponding differences for the old, the appropriate benchmark for the difference-in-difference estimates is -4.7 for private employees, -8.5 for public employees, and -18.9 for the self-employed. According to the estimates in the first regression of Table 4.4, the difference-in-difference estimates are -4.2 for private employees, -2.0 for public employees and -4.3 for the self-employed. The coefficient for private employees is close to the benchmark, but the other two coefficients imply considerable underestimation of the effect of the reform.

The second regression adds to the basic specification regional and educational dummies and annual earnings (in thousand Euro). Working in the South and the level of income are positively related to σ . The effect of the education dummies is positive for high school, negative for university degree but never statistically different from zero. The γ 's are qualitatively unchanged, confirming partial adjustment of the expected replacement rate in the new pension regime.

The regressions for females uncover an across-the-board reduction in σ after the reform by 6.1 percentage points, but the employment dummies interacted with M and $POST$ signal no differential effect by employment groups or cohort after the reform.

Table 4.4: The effect of the reform on the expected replacement rate: regression results

	Males		Females	
Private, middle-aged, after	-4.202 (0.960)**	-4.552 (0.957)**	-0.718 (1.427)	-0.796 (1.411)
Public, middle-aged, after	-2.046 (1.357)	-2.585 (1.353)	-1.182 (1.336)	-1.476 (1.324)
Self-e., middle-aged, after	-4.291 (1.511)**	-4.298 (1.504)**	-3.118 (2.375)	-2.868 (2.350)
Public employee	2.689 (0.696)**	2.501 (0.699)**	1.077 (0.960)	-0.483 (0.974)
Self employed	-10.046 (0.839)**	-5.956 (0.965)**	-12.020 (1.405)**	-6.733 (1.508)**
Middle-aged	0.109 (0.684)	0.916 (0.689)	-0.421 (1.027)	-0.251 (1.022)
Public, middle-aged	-1.357 (1.148)	-1.603 (1.144)	0.414 (1.385)	0.695 (1.371)
Self-e., middle-aged	1.968 (1.341)	1.096 (1.339)	4.959 (2.074)*	4.267 (2.053)*
Post-reform	-4.289 (0.716)**	-3.961 (0.715)**	-6.110 (1.136)**	-5.735 (1.124)**
Public, after	2.325 (1.226)	2.453 (1.221)*	1.797 (1.532)	2.317 (1.518)
Self-e., after	-3.632 (1.317)**	-3.874 (1.313)**	-1.735 (2.110)	-2.188 (2.089)
Central Italy		0.724 (0.431)		0.362 (0.526)
Southern Italy		1.249 (0.370)**		2.569 (0.517)**
Earnings		0.235 (0.027)**		0.413 (0.045)**
High school degree		0.622 (0.356)		1.587 (0.486)**
University degree		-0.429 (0.586)		0.775 (0.708)
Constant	79.182 (0.425)**	74.249 (0.673)**	77.215 (0.737)**	70.683 (0.956)**
Observations	9766	9766	5955	5955
R-squared	0.16	0.17	0.11	0.13

Notes: All explanatory variables are dummy variables, except earnings (expressed in thousand euro). Young workers, who started working after 1995, are excluded. Standard errors robust to unknown form of heteroskedasticity are reported in parentheses. Two stars indicate statistical significance at the 1% confidence level, one star at the 5% level.

The corresponding coefficients are not statistically different from zero, while Table 4.3 implies a large reduction in σ for each employment group.

In summary, the regressions of Table 4.4 suggest that most groups have revised their expectations in the direction and magnitude implied by the reform. But the revision to the new pension rules has been far from complete. Two interpretations of the results are possible: an anticipation effect, or lack of information. If the reform had been anticipated, people would have adjusted downward the expected replacement rate even before the reform. This explanation clashes with the fact that, on average, the 1989-91 expected replacement rates were quite close to - or even overestimated - the statutory rates. Furthermore, the anticipation effect should be stronger in the years immediately before the reform; however, dropping 1991 and defining the pre-reform period as just 1989 does not change the results with respect to estimates in Table 4.4. Therefore, the most likely explanation for our findings is that, as of 2002, many workers did not fully understand the implications of the new pension regime and had not yet updated their pension expectations accordingly.

4.5 The offset between pension wealth and private wealth accumulation

So far our analysis suggests that people reacted to the pension reform by raising expectations of retirement age and reducing perceived replacement rates and pension wealth. However, the magnitude of the expectation revision is considerably lower than the actual magnitudes implied by the reform. This is an important first step in evaluating the effect of pension reforms on individual behavior. The next important step relates perceived pension wealth to private accumulation. Since the reform provides an exogenous source of variation in pension wealth across socio-economic and demographic groups, we are in a good position to assess the extent to which the revision in retirement age and replacement rate leads to changes in private wealth. In this section we therefore estimate the offset between pension wealth and private wealth using the reform as an instrument for pension wealth.

Our empirical specification relates private wealth to pension wealth, and to a set of observable variables potentially correlated with private wealth. More specifically, we estimate the following equation:

$$WY_{it} = \alpha + SSWY_{it}\delta + X_{it}\gamma + \theta_t + \varepsilon_{it} \quad (4.3)$$

where WY_{it} is private wealth of household i at time t , scaled by household disposable income, $SSWY_{it}$ is the ratio of expected pension wealth at retirement to earnings (evaluated at time t), X_{it} is a vector that includes age of the household head,⁹ disposable income, employment group dummies, education dummies, and region dummies;¹⁰ θ_t represents time effects. Age, income, education and employment sector are proxies for lifetime earnings, while year dummies capture macroeconomic effects. Sensitivity analysis is performed to check, among other aspects, for the inclusion of an additional vector of demographic characteristics such as family size and number of family workers.

Total net worth is defined as financial assets plus real assets (real estates and businesses) minus financial debt. As for the ratio of expected pension wealth at retirement to earnings (evaluated at time t), in order to keep its computation as simple and as tailored as possible to elicited expectations on the replacement rate and the retirement age, we use the following proxy for each worker's pension wealth:

$$SSWY_t = \left[P(N_t|t) \left(\frac{1+g_u}{1+r} \right)^{N_t-t} \sigma_t \right] \sum_{\tau=N_t}^T P(\tau|N_t) \left(\frac{1+g_N}{1+r} \right)^{\tau-N_t} \quad (4.4)$$

where σ_t is the expected replacement rate and N_t the expected retirement age elicited at time t ,¹¹ T the maximum length of life, $p(\tau|N)$ the probability of surviving to age τ , conditional on being alive at age N , g_u the growth rate of earnings for group u , r the real interest rate, and g_N the growth rate of pension benefits during retirement - assumed to be the same for all groups.

In the survey we observe σ_t and N_t for each individual. Survival probabilities are taken from the Italian life tables, by age and gender, for the years 1990 and 2000, so

⁹We define the head of the household as the partner with higher earnings.

¹⁰In the regressions, the reference group is private employees with less than 13 years of education and living in Northern Italy.

¹¹ $t=1989, 1991, 2000, 2002$, the survey years in which the expected replacement rate is elicited.

that the change in life expectancy over time, and in particular before and after the reform, is accounted for.¹² The growth rate of earnings (g_u) is estimated from our data at 0.015 for individuals with university degree and at 0.008 for individuals with less than university degree.¹³ We assume that after retirement pensions are constant in real terms ($g_N = 0$) and that the real interest rate is equal to 2 percent.

Armed with this information, we can compute the expected ratio of pension wealth for each individual in the sample. In households with more than one member, we define the household expected pension wealth-to-income ratio at retirement as the weighted sum of both partners' expected pension wealth-to-income ratio, where each partner is given her relative weight in terms of her income in relation to the income of the couple.

The individual expected pension wealth-to-income ratio is adjusted by the factor suggested by Gale (1998). This factor allows to adjust expected pension wealth for the number of years people have contributed to their pension plan as well as for when in their life cycle they have experienced some shock that should have made them revise their consumption and savings plans (the reforms, in our case). Omitting to adjust for this factor would produce an underestimate of the offset between pension wealth and private wealth, i.e. the estimates for the pension wealth coefficient would be biased towards zero. The adjustment depends on the utility function that is chosen for the underlying life-cycle model and on the values for the discount rate, the interest rate and the time preference rate (see Appendix, section 4.7.2, for further details). We use the adjustment developed in Gale (1998) for a CRRA utility function and set the discount rate and the interest rate equal to 2 percent. Sensitivity analysis is then performed on these values.

4.5.1 Regression estimates

Table 4.5 presents the results obtained from, respectively, OLS (col. 1-2) and median regressions (col. 3-4). To further characterize the distribution of the wealth-income

¹²Data source: *Italian Statistical Annex* (Rome: ISTAT, 1990 and 2000).

¹³The growth rates were obtained from a median regression of log-earnings on sex and employment dummies and full interaction of age with a college dummy. Data source: SHIW, years 1989-2002, individuals with age 20-60.

ratio, we rely on estimates based on least absolute deviations, which are consistent and asymptotically normal in the presence of thick tailed error distributions.¹⁴

Table 4.5: The offset between private wealth and pension wealth: OLS and LAD estimates

	OLS		LAD	
SSW/Disposable Income	-0.403 (0.043)**	-0.167 (0.044)**	-0.474 (0.035)**	-0.254 (0.032)**
Year 1991	0.641 (0.187)**	0.754 (0.182)**	0.596 (0.150)**	0.580 (0.132)**
Year 2000	1.384 (0.183)**	1.464 (0.180)**	0.804 (0.147)**	0.851 (0.131)**
Year 2002	1.700 (0.190)**	1.784 (0.187)**	0.752 (0.152)**	0.909 (0.135)**
Age	0.134 (0.011)**	0.101 (0.011)**	0.161 (0.009)**	0.114 (0.008)**
High School Degree		1.018 (0.133)**		0.643 (0.097)**
University Degree		0.855 (0.206)**		0.454 (0.150)**
Disposable Income		0.012 (0.003)**		0.034 (0.002)**
Self-employed		3.574 (0.159)**		1.974 (0.115)**
Public employee		0.279 (0.154)		0.345 (0.112)**
Central Italy		0.092 (0.160)		0.319 (0.116)**
Southern Italy		-0.279 (0.141)*		-0.065 (0.102)
Constant	-0.990 (0.410)*	-1.869 (0.409)**	-2.543 (0.328)**	-3.202 (0.297)**
Observations	9462	9462	9462	9462
R-squared	0.04	0.10		

Notes: Standard errors robust to unknown form of heteroskedasticity are reported in parentheses. Two stars indicate statistical significance at the 1% confidence level, one star at the 5% level.

¹⁴We also perform trimmed least squares, discarding the top and bottom 1 percent of the private wealth-income ratio distribution. The results are qualitatively unchanged.

In columns (1) and (3) we report the results of a specification that includes only the pension wealth to income ratio, time dummies and age, whereas in columns (2) and (4) we report the results of the full specification as in equation (4.3). The offset between private wealth and pension wealth is, respectively, 40 percent and 47 percent, and statistically different from zero at the 1 percent level, for the OLS and LAD baseline specifications. The full specification gives an offset of, respectively, 17 percent and 25 percent, again significant at the 1 percent level.¹⁵

The estimates indicate that the wealth-income ratio increases with age during the working span (recall that individuals over 50 are excluded), in agreement with the life-cycle model. The extended specifications further signal that private wealth increases with labor income. The latter should not affect the wealth-income ratio if preferences are homothetic. The regression coefficient, on the other hand, can hardly be interpreted as evidence in favor or against homothetic preferences since other variables (education or residence in the South) may also proxy for lifetime earnings. Residence in the South reduces wealth accumulation; education has an opposite effect. These variables are obviously related to household resources. But they may also capture other effects. For instance, there is evidence that the better educated are more likely to report financial assets (Brandolini and Cannari, 1994); households with higher education may have easier access to capital markets and to better investment opportunities; thrift may be correlated with schooling.

The results in Table 4.5, however, understate the offset between pension wealth and private wealth if pension wealth and private wealth are positively correlated. This might be the case if thrift and hard work are correlated tastes, and people with these traits choose to save more and to retire with higher pension wealth. Since the pension reforms provide us with an exogenous source of variation for pension wealth, we can perform instrumental variable estimation and remove this source of bias from our estimated offset.

In particular, in equation (4.4) there are two potential sources of endogeneity: the

¹⁵Without the adjustment factor for pension wealth, the offset would be, respectively, 13 percent and 17 percent for OLS and LAD regressions.

subjective replacement rate and the subjective retirement age. In a first IV regression we use as instrument the “statutory pension wealth”, computed by replacing only the expected replacement rate with the statutory rate (derived from legislation before and after the reform as explained in the Appendix, section 4.7.1). In a second IV regression statutory pension wealth is computed in relation to the statutory replacement rate and to the sample median expected retirement age, by gender, before and after the reform.¹⁶

The validity of these instruments rests on the fact that the rules for computing pension wealth change exogenously for the middle-aged after the reform, depending on employment group membership. It also depends on the assumption that the middle-aged did not switch jobs after the reform to offset the impact of the pension reform on their retirement wealth. Under this reasonable assumption, which is corroborated by the evidence on employment transition matrices discussed in Section 4.4 and reported in the Appendix (section 4.7.4), the instruments are also exogenous with respect to private wealth accumulation decisions.¹⁷

Results reported in Table 4.6 indicate that the offset between private wealth and pension wealth is considerably higher than the one resulting from OLS regressions, confirming the idea that the previous estimates are biased toward zero. The estimation with all the controls gives an offset of 31 percent and 52 percent, respectively, depending on the definition of statutory pension wealth used.¹⁸ We check the sensitivity of the results against the inclusion of family size and number of income recipients or a dummy for whether the number of family workers is greater than one among the vector of observable characteristics as in equation (4.3). We also introduce a quadratic term in age and show that the results are not sensitive to any of these changes in specification.

Another set of sensitivity checks relates the interest rates, the discount factor and the coefficient of relative risk-aversion that the computation of the Gale adjustment

¹⁶Median retirement age is set at 60 for males and at 55 for females before the reform and at 65 and 60, respectively, after the reform.

¹⁷Attanasio and Brugiavini (2003) use similar employment-group instruments for pension wealth in their study of the impact of the 1992 Italian pension reform on the household saving rate.

¹⁸We also run a regression of private wealth on statutory pension wealth (the instrument). The estimated offset in the baseline specification is -0.669 (with a standard error of 0.051) and -0.459 (with a standard error of 0.052).

Table 4.6: The offset between private wealth and pension wealth: instrumental variable estimation

	IV1		IV2	
SSW/Disposable Income	-0.542 (0.049)**	-0.309 (0.050)**	-0.761 (0.058)**	-0.525 (0.059)**
Year 1991	0.630 (0.187)**	0.740 (0.182)**	0.614 (0.188)**	0.718 (0.183)**
Year 2000	1.216 (0.186)**	1.288 (0.183)**	0.953 (0.190)**	1.020 (0.188)**
Year 2002	1.512 (0.193)**	1.587 (0.190)**	1.218 (0.197)**	1.288 (0.195)**
Age	0.153 (0.012)**	0.121 (0.012)**	0.182 (0.012)**	0.152 (0.013)**
High School Degree		1.016 (0.133)**		1.013 (0.134)**
University Degree		0.826 (0.206)**		0.782 (0.207)**
Disposable Income		0.010 (0.003)**		0.007 (0.003)*
Self-employed		3.489 (0.160)**		3.360 (0.161)**
Public employee		0.292 (0.154)		0.312 (0.155)*
Central Italy		0.060 (0.161)		0.011 (0.161)
Southern Italy		-0.322 (0.141)*		-0.388 (0.142)**
Constant	-1.328 (0.414)**	-2.156 (0.413)**	-1.856 (0.422)**	-2.594 (0.419)**
Observations	9462	9462	9462	9462
R-squared	0.04	0.10	0.03	0.09

Notes: In the first two regressions the statutory pension wealth-to-income ratio is constructed using the statutory replacement rate and the expected retirement age. In the other regressions the statutory pension wealth-to-income ratio is constructed using the statutory replacement rate and median retirement age before and after the reform (60 for males and 55 for females before the reform and 65 and 60, respectively, after the reform). Standard errors robust to unknown form of heteroskedasticity are reported in parentheses. Two stars indicate statistical significance at the 1% confidence level, one star at the 5% level.

factor involves. In particular, we check the sensitivity of the results to changes of $x = [(r - \delta)/\rho] - r$, where r , δ and ρ are, respectively, the interest rates, the discount factor

and the coefficient of relative risk-aversion. We have assumed earlier that $x = 0.02$, which is consistent, for instance, with $r = \delta = 0.02$. Setting $x = -0.04$ or $x = -0.06$ gives smaller offsets, as expected. For instance, the offset in the full specification is, respectively, 13 percent and 12 percent for the OLS estimates; the first set of IV estimates gives, respectively, an offset of 25 percent and 23 percent for the two values of x , whereas the second set of IV estimates gives an offset of 42 percent and 35 percent.

Previous literature provides some evidence on the effect of the 1992 Italian pension reform on household saving. Using SHIW data for the years 1989-95, Attanasio and Brugiavini (2003) exploit the changes in pension wealth across cohorts and employment groups due to the 1992 reform to estimate the crowding out effect of pension wealth on the household saving rate. They find that a reduction in pension wealth of 1 euro prompts an increase in private saving of between 30 to 40 cents. Although these point estimates are not far from ours, we must bear in mind that there are at least three crucial differences between these two studies: we look at the combined impact of three pension reforms (1992, 1995 and 1997), focus on private wealth, rather than on saving, and rely on a different estimate of pension wealth, based on the expected retirement age and expected replacement rate, rather than computed from legislation.

4.5.2 Implications

The estimated coefficient for the offset between pension wealth and private wealth means that, on average, there is far from full crowding out of private accumulation in our sample of Italian households. Combining this information with our findings that people have revised their expectations on replacement rate and retirement age less than they should have following the reforms, makes the question of whether people are saving enough for their retirement of central interest.

To provide a sense of the magnitudes involved, let us look at the following case: a female aged 40 in the year 2000 (and belonging to the so-called middle-aged group with respect to the reform), single, public employee, with 13 years of education (i.e. non-university graduate), after the reform expects to retire at age 60 with a replacement rate of 75 percent. Using equation (4.4), and applying the survival probabilities for a

female of her age in 2000, her pension wealth-to-income ratio at retirement would be 10.351. Setting her statutory retirement age at 60 and statutory replacement rate at 68 percent,¹⁹ would give a statutory pension wealth-to-income ratio of 9.386, which would mean that her expectations overstate her pension wealth by 26,050 euro (at the annual median disposable income of 27,000 euro). Using the most optimistic estimates of -0.525 for the offset between pension wealth and private wealth, this would mean that this single household would adjust private wealth by 13,677 euro upon an unexpected change in pension wealth of 26,050 euro.

The results have also interesting implications for evaluations of how pension reforms affect private and national saving. On this front, one may be puzzled by the observation that over the past decade the Italian national saving rate has remained roughly constant, despite deep pension reforms reducing considerably pension wealth.²⁰

The aggregate effect of the reform on national saving depends on the reaction of the old, the middle-aged, the young and those already retired, weighted by the respective shares of these groups in total population. The comparison between the statutory and the expected pension wealth reveals that the middle-aged have perceived only two thirds of the reduction in statutory pension wealth; furthermore, from the IV estimates in Table 4.6 only 50 percent^{4.5} of the reduction in perceived pension wealth has been offset by an increase in private wealth. The effect for the old is small, because for this group the change in actual and perceived pension wealth has been far more limited.²¹ On the other hand, our regressions cannot predict the effect of the reform on those that retired before the reform, nor on those that entered the labor market after 1995

¹⁹This would still be 2.6 percent points higher than the statutory replacement rate for a female in her employment and cohort group, as indicated in Table 3.

²⁰In 1981-1990 the average gross national saving rate was 22.3 percent, while in 1991-2000 it was 21 percent (20.4 percent in 2000-01). Clearly pension reform is just one of the main determinants of national saving, so the figures do not rule out that the reduction in pension wealth has increased national saving, with offsetting effects from other sources. But the figures are suggestive of a limited aggregate saving impact of the reform.

²¹Recall from Section 4.2 that pension reforms reduced pension of the elderly, particularly of public employees, through changes in the pension award formula and indexation of benefits to cost of living rather than wages.

(which are not included in the estimation). We speculate that for the former group there should be little effect, because the reform has mainly affected working cohorts and future generations, rather than the retired. As for the young, the effect might be more substantial, but as of today their weight in the total working population is still limited, and therefore an impact on the aggregate wealth and saving might not be visible. Overall, this might explain why the Italian saving rate has not increased despite the decade of pension reforms.

In summary, we find robust evidence that expected pension wealth is a substitute for private wealth. However, we also find that the pension reforms of the last decade did not have a large impact on the household private wealth and, consequently, on national saving. Two factors account for this result. First, the substitution between the two forms of wealth is only imperfect, with offset rates in the order of 0.5. Second, some households do not yet seem to have fully internalized the implications of the reform into their expectations of social security pensions or their retirement plans.

4.6 Conclusions

The Survey of Household Income and Wealth, a large representative sample of the Italian population, elicits expectations of replacement rates from workers interviewed in the years between 1989 and 2002, a period of intense pension reforms. The reforms reduced the replacement rate and increased retirement age, and had different impact on different cohorts and employment groups, providing exogenous variations in replacement rates to study the effect of pension reforms on expectations.

We find that pension reforms indeed affected expectations of retirement benefits. However, the revision in expectations is limited, and many individuals have not yet updated completely their expectations. For instance, while the perceived replacement rate of the self-employed falls by about 10 percentage points between 1989-91 and 2000-02, in reality the rate falls by about 20 points. Moreover, we find that the offset between pension wealth and private wealth is only partial, in the order of 50 percent. This suggests that the effect of pension reform on individual behavior depends critically

on the extent of the knowledge and information that individuals have about the social security system and changes to it, and has three important policy implications.

First, the descriptive and econometric analysis implies that current workers lack crucial information to fully understand the implications of the new pension regime, thus making a clear case for investing public resources in the dissemination of information about pension rights, especially during periods of intense reform. Campaigns to increase financial literacy and the understanding of pension rules, and to provide individuals with regular statements of their expected retirement income, are important steps in this direction. Second, this chapter suggests that if one wants to use observations of past pension reforms to make predictions about likely responses to new reforms, then one needs to estimate how responses in the past were limited by inaccurate updating of expectations, and how the new reform will affect expectations. Finally, given the dramatic reduction in replacement rates implied by the pension reform, combined with an incomplete offset between pension wealth and private wealth, it is likely that some individuals, especially the younger cohorts most affected by the reform, might not be saving enough for their old age. This might have a long-term impact on the well being of future retirees in the coming decades, when the generations affected by the pension reform will start to retire.

4.7 Appendix to chapter 4

4.7.1 The pension award formula before and after the reform

In the *pre-reform* regime social security benefits were computed according to an earnings-based formula:

$$\rho N \bar{w}_R$$

where ρ , N and \bar{w}_R are, respectively, the accrual rate, the years of contributions and the average of the last R years of salary. The accrual rate is 2 percent for private employees and self-employed, and ranges from 2.2 to 2.5 percent for public employees, depending on the years of contribution; R is 5 for private employees, 1 for public employees, and 10 for the self-employed.

In the *post-reform* regime pensions are computed distinguishing between three cases: *earnings model for the old* (more than 18 years of contributions in 1995), *contribution model for the young* (started working after 1995), and *pro-rata model* for the middle-aged (less than 18 years of contributions as of 1995). In each case, different rules apply to public employees, private employees and self-employed.

For the *old*, benefits are the sum of two components. The first component is $\rho \alpha_{92} \bar{w}_R$, where α_{92} is the number of years of contributions at the end of 1992. The second component reflects a gradual increase of R to 10 for private and public employees and to 15 for the self-employed. Namely, for years of contributions between 1992 and 1995, R is increased by 1; for years of contributions between 1995 and the year of retirement, R is increased by the minimum of 5 and $2/3$ of the years of contributions between 1995 and the year of retirement. For instance, for those retiring in 2000 R is increased by 3; for those retiring in 2005 it is increased by 5. The second component is therefore:

$$\rho(\alpha_{95} - \alpha_{92})\bar{w}_{R'} + \rho(N - \alpha_{95})\bar{w}_{R''}$$

where α_{95} is years of contribution at end of 1995, $R' = R + 1$ and

$$R'' = R + \min(5, \text{int}(\frac{2}{3}(N - \alpha_{95}))).$$

Therefore, the pension for the old is:

$$\rho \left[\alpha_{92} \left(1 - \frac{R}{R'} \right) + \alpha_{95} \left(\frac{R}{R'} - \frac{R}{R''} \right) N \frac{R}{R''} \right] \bar{w}_R$$

In practice, for realistic earnings growth rates, the second component has a small impact on the final pension with respect to the pre-reform regime.

For the *young*, benefits are computed according to a contribution model:

$$\gamma \tau \sum_{t=0}^{N-1} w_t (1+g)^{N-1-t}$$

where τ is the contribution rate (0.33 for private and public employees and 0.20 for self-employed) and g a 5-year moving average of the GDP growth rate. Contributions are therefore proportional to earnings, capitalized on the basis of a 5-year moving average, and then transformed in flow benefits using a coefficient (γ), set by legislators, that depends on retirement age and life expectancy. Currently, γ increases from 4.720 percent for somebody retiring at 57 to 6.136 percent for somebody retiring at 65.

For the *middle-aged*, benefits are computed using the earnings model for years of contributions before 1995, and the contribution model for years of contributions after 1995. The pension is then equal to:

$$\rho \alpha_{92} \bar{w}_R + \rho (\alpha_{95} - \alpha_{92}) \bar{w}_N + \gamma \tau \sum_{t=\alpha_{95}}^{N-1} w_t (1+g)^{N-1-t}$$

4.7.2 The adjustment factor for pension wealth

As in Gale (1998), we adjust pension wealth multiplying each individual's expected pension wealth by a factor that takes into account people's position in the life cycle and years of service in the pension as well as the position in people's life cycle when a change in pension benefits takes place (the reforms, in our case). The underlying idea for the simplest theoretical model is that people plan their consumption at the beginning of their working career, and consumption is a function of total lifetime resources, that is earnings and pension benefits. Since decisions are based on total lifetime resources, the true offset between pension wealth and private wealth is 100 percent (coefficient of -1). However, as pointed out by Gale (1998, pp. 708-710) an estimate of the coefficient of

pension wealth in a regression of private wealth on earnings to date, lifetime earnings and pension benefits, would not produce the true offset. The pension wealth coefficient would instead be between -1 and 0, and a function of the years of service in the pension and of the expected life horizon. In particular, in the case of a CRRA utility function, the coefficient would be as follows:

$$Q = \frac{\exp(xS - 1)}{\exp(xT - 1)} = \frac{\exp(xS - 1)}{\exp(x(le + S) - 1)}$$

where $x = \frac{r-\delta}{\rho} - r$, and r =interest rate, δ =time preference rate, ρ =coefficient of relative risk aversion, S =years of service in the pension, T =life span, and le =life expectancy.

Therefore, one would need to adjust pension wealth by this factor in order to recover the true offset in the regression. Intuitively, this factor adjusts pension wealth to account for the fact that a change in pension wealth that takes place at the beginning of one's career translates into a change in the consumption plan (and therefore in non-pension wealth) over the life span. At time S , the reduction in non-pension wealth is captured by Q , and Q increases with S to reflect the fact that the later in life we observe individual's decisions, the more of the initial plan has already taken place.

A further aspect to be taken into account is given by the time at which the change in pension benefit is realized. For a generic time t^* , Gale's adjustment factor is:

$$Q^* = \frac{\exp(x(S - t^*) - 1)}{\exp(x(le + S - t^*) - 1)}$$

This accounts for the fact that individuals had to revise their plans at time t^* and the remaining horizon over which they can realize their plans is shorter.

In our setting, we assume that $r = \delta = 0.02$ and apply different adjustment factors according to which group the individual belongs to. In particular, the so-called "Old" group is not affected by the reform, and therefore we apply a version of Q , corrected for the fact that individuals start contributing to the pension system at different ages

(we observe this in the data), i.e.:

$$G = \frac{\exp(-r(\text{age} - \text{agew}) - 1)}{\exp(-r(\text{le} + \text{age} - \text{agew}) - 1)}$$

where age = age at which observed and agew = age at which started working. The adjustment factor for the group affected by the reform (“Middle-aged”) instead needs to take into account of the year in which the reform took place and is therefore a version of Q^* . We assume that the year of the reform is 1995 and adjust pension wealth of individuals belonging to this group, and observed after the reform, by the following factor:

$$G^* = \frac{\exp(-r(\text{age} - \text{ageref}) - 1)}{\exp(-r(\text{le} + \text{age} - \text{ageref}) - 1)}$$

where ageref = age at which the individual faced the reform.

4.7.3 The Survey of Household Income and Wealth

The SHIW is a representative sample of the Italian resident population. Sampling is in two stages, first municipalities and then households. Municipalities are divided into strata defined by regions and classes of population size (less than 20,000, 20,000 to 40,000, more than 40,000). Households are then randomly selected from registry office records, see Biancotti et al. (2004) for details. From 1987 onward the survey is conducted every other year and covers about 24,000 individuals and 8,000 households, defined as groups of individuals related by blood, marriage or adoption and sharing the same dwelling. Interviews are conducted by a specialized agency with professional interviewers, and are preceded by extensive training and meetings with Bank of Italy representatives. Interviews take place in person, by visiting the residence of the household. Because of its sample design and its collection of detailed wealth statistics, the SHIW is similar to the Survey of Consumer Finances (SCF), which is representative of the U.S. population. The English version of the questionnaire, data and survey documentation can be downloaded from the Bank of Italy web site: <http://www.bancaditalia.it>

Table 4.7 compares the population and sample means of selected demographic variables (age, gender and region) in 1989, 1991, 2000 and 2002. The population means

are obtained from the *Italian Statistical Annex* (Rome: ISTAT, 1989, 1991, 2000 and 2002). Sample statistics are computed using the SHIW population weights, defined as the inverse of the probability of inclusion in the sample. Overall, the comparison indicates that the sample reflects fairly well the demographic characteristics of the Italian population and is stable over time.

Table 4.7: Population and sample means of selected variables

	1989		1991		2000		2002	
	Popul.	Sample	Popul.	Sample	Popul.	Sample	Popul.	Sample
Males	48.6	48.7	48.6	48.9	48.5	48.6	48.6	48.5
Females	51.4	51.3	51.4	51.1	51.5	51.4	51.4	51.5
Age <24	32.7	32.9	32.0	31.6	26.2	26.3	25.8	25.7
Age 25-44	28.6	28.8	29.2	28.7	30.8	30.7	30.9	29.1
Age 45-64	24.1	26.1	24.0	25.9	25.0	25.0	25.1	27.1
Age >65	14.5	12.3	14.8	13.8	18.0	18.0	18.2	18.1
North	44.3	46.5	44.8	44.3	44.6	44.6	44.7	44.7
Center	19.1	18.9	19.1	19.0	19.2	19.2	19.3	19.3
South	36.6	34.7	36.1	36.5	36.2	36.2	36.0	36.1

Notes: Sample statistics are computed using SHIW population weights, defined as the inverse of the probability of ample inclusion. Sources: *Survey of Households Income and Wealth, Bank of Italy, and Italian Statistical Annex* (Rome: ISTAT, 1989, 1991, 2000 and 2002).

4.7.4 Employment transition matrix

The rotating panel component of the SHIW allows us to compute sector transition matrices in 1989-91, 1991-93, 1993-95, 1995-98, 1998-2000, and 2000-02. Table 4.8 reports two such transition matrices and shows that mobility across sector is stable over time. All elements on the main diagonal are close to 90 percent. The Shorrocks mobility index is 12.5 percent in 1989-91 and 13 percent in 2000-02. The hypothesis that the mobility index did not change after the reform cannot be rejected at the 1 percent level (the statistic associated to the difference test is 0.12 and is distributed

as a standard normal). Comparison with mobility matrices for other years confirms this result. We interpret the stability of the employment transition matrix as indirect evidence that the pension reform did not change the pattern of mobility across the three sectors.

Table 4.8: Transition matrix among employment groups

		1991		
1989		Private sector	Public sector	Self-employed
Private sector		0.89	0.09	0.02
Public sector		0.07	0.92	0.01
Self employed		0.05	0.02	0.93

		2002		
2000		Private sector	Public sector	Self-employed
Private sector		0.88	0.08	0.03
Public sector		0.08	0.91	0.01
Self employed		0.00	0.03	0.97

Notes: The upper left cell in each panel reports the probability of working in the private sector in years t and $t + 1$. The other cells have similar interpretation. The generic cell i, j gives the probability of working in sector i in year t and in sector j in year $t + 1$.

Chapter 5

Conclusion

This thesis has focussed on some open issues within the life cycle framework. Chapters 2 and 3 dealt with home-ownership decisions and their relation to labour supply, the former in an empirical analysis, and the latter in a structural lifecycle model. Chapter 4 took up the theme of how expectations should be incorporated in an empirical analysis, and provided some evidence that relates to the so called “retirement-consumption puzzle”.

More specifically, chapter 2 presented empirical evidence for the UK of a positive relationship between female participation in the labour market and mortgage commitments. This relationship is seen to persist even after controlling for family characteristics and unobserved heterogeneity, although the negative effect on female participation of having a young child can outweigh the positive effect due to mortgage commitments. The evidence of a positive relationship between female participation and mortgage commitments is in line with previous international research. A test of endogeneity of the mortgage choice with respect to labour supply, rejects the null hypothesis.

Chapter 3 studies the implications of a structural life cycle model of household consumption, labour supply and saving together with home-ownership decisions. The empirical motivation is provided by chapter 2 but theoretical implications, particularly for labour supply, appear to be smaller than in the data. One potential reason why the scale of effects in the model, particularly for labour supply, appear to be smaller than

in the data, is that education is the only source of observed heterogeneity in the model. Natural extensions of the model would allow for a richer consideration of demographics over the life cycle. In particular, introducing a second earner in the household could prove fruitful for explaining variability in the labour supply of the household; the two earners would face separate wage rates each one with an idiosyncratic component, rather than the single wage rate that is assumed in the current version of the model. A second observable demographic characteristic that it would possibly be important to introduce is the presence and arrival of children, particularly given that fertility, home-ownership and labour supply decisions are thought of as being interrelated if not simultaneous. As well as including other dimensions of observable heterogeneity, it would be important to include a mortgage repayment schedule for home-owners with a mortgage. This would track widespread UK “repayment mortgage” scheme that seem to be a more common institutional feature found in the data.

Chapter 4 analyses expectations over retirement outcomes to try and answer two questions: first, to assess how accurate individuals’ expectations are, by examining changes across various changes in pension legislation; second, to estimate the offset between expected private wealth accumulation and (state) pension wealth. Individual’s expectations are found to overestimate the expected ratio of pension benefits to pre-retirement income. Expectations following changes in pension legislation are found to respond to the effects implied by the reforms but the adjustment is not complete. Further years of data will clarify whether expectations are adjusting slowly and so will adjust more fully with time. Our findings on the offset between private wealth and expected pension wealth suggest that this is far from complete.

Bibliography

- [1] AIYAGARI, S.R. (1994), "Uninsured idiosyncratic risk and aggregate saving", *Quarterly Journal of Economics*, 109(3):659-684.
- [2] AIYAGARI, S.R. AND E.R. McGRATTAN (1998), "The optimum quantity of debt", *Journal of Monetary Economics*, 42: 447-469.
- [3] ALDERSHOF, T., R. ALESSIE AND A. KAPTEYN (1997), "Female labour supply and the demand for housing", Progress Report 46, CentER, Tilburg University.
- [4] ALESSIE, R., M. P. DEVEREUX AND G. WEBER (1997), "Intertemporal consumption, durables and liquidity constraints: a cohort analysis", *European Economic Review*, 41: 37-59
- [5] ALESSIE, R., B. MELENBERG AND G. WEBER (1988), "Consumption, leisure and earnings-related liquidity constraints. A note." *Economics Letters*, 27: 101-104.
- [6] ARELLANO, M. (2000), "Discrete choices with panel data", mimeo, CEMFI, Madrid.
- [7] ARELLANO, M. AND B. HONORÉ (2000), "Panel data models: some recent developments", Working Paper No. 0016, CEMFI, Madrid, now in Heckman J. J. and E. Leamer (eds): *Handbook of Econometrics* (5), Ch. 53, Elsevier Science.
- [8] ATTANASIO, O. (2000), "Consumption", in J. Taylor and M. Woodford (eds): *Handbook of Macroeconomics*, ch.11, North Holland.
- [9] ATTANASIO, O., J. BANKS, C. MEGHIR AND G. WEBER (1999), "Humps and bumps in lifetime consumption", *Journal of Business and Economic Statistics*, 17(1): 22-35.

- [10] ATTANASIO, O., L. BLOW, R. HAMILTON AND A. LEICESTER (2005), "Consumption, House Prices and Expectations," Bank of England Working Paper no. 271
- [11] ATTANASIO, O. AND M. BROWNING (1995), "Consumption over the life cycle and over the business cycle", *American Economic Review*, 85(5): 1118-1137.
- [12] ATTANASIO, O. AND A. BRUGIAVINI (2003), "Social security and households savings", *Quarterly Journal of Economics*, 118(3): 1075-1120.
- [13] ATTANASIO, O. AND S. J. DAVIS (1996), "Relative wage movement and the distribution of consumption", *Journal of Political Economy*, 104(6): 1227-1262.
- [14] ATTANASIO, O., P. GOLDBERG AND E. KYRIAZIDOU (2000), "Credit constraints in the market for consumer durables: evidence from micro data on car loans", NBER Working Paper no. 7694.
- [15] ATTANASIO, O. AND S. ROHWEDDER (2003), "Pension wealth and household saving: evidence from pension reforms in the United Kingdom", *American Economic Review*, 93(5): 1499-1521.
- [16] ATTANASIO, O. AND G. WEBER (1994), "The UK Consumption boom of the late 1980's: aggregate implications of microeconomic evidence," *The Economic Journal*, 104(427): 1269-1302.
- [17] ATTANASIO, O. AND G. WEBER (1995), "Is consumption growth consistent with intertemporal optimization? Evidence for the Consumer Expenditure Survey", *Journal of Political Economy*, 103, 1121-1157.
- [18] BANKS, J., R. BLUNDELL, J. SMITH, AND Z. SMITH (2004), "House price volatility and housing ownership over the lifecycle," University College London, Department of Economics Discussion Paper 04-09.
- [19] BANKS, J., Z. SMITH AND M. WAKEFIELD (2002), "The distribution of financial wealth in the UK: evidence from 2000 BHPS data", IFS Working Paper no. 02/21.

- [20] BERNHEIM, D. (1990), "The timing of retirement: a comparison of expectations and realizations", in Wise, D. (ed): *The economics of aging*, The University of Chicago Press, Chicago.
- [21] BIANCOTTI, C. G. D'ALESSIO, I. FAIELLA AND A. NERI (2004), "I bilanci delle famiglie italiane nell'anno 2000" (in Italian), *supplementi al Bollettino Statistico*, 12, Bank of Italy, Rome.
- [22] BLANCHARD, O. AND G. MANKIW (1988), "Consumption: beyond certainty equivalence", *American Economic Review, Papers and Proceedings*, 78(2): 173-177.
- [23] BLUNDELL, R., M. BROWNING AND C. MEGHIR (1994), "Consumer demand and the life-cycle allocation of household expenditures", *Review of Economic Studies*, 61(1): 57-80.
- [24] BLUNDELL, R. AND T. MACURDY (1999), "Labour supply: a review of alternative approaches", in Ashenfelter, O. and D. Card (eds): *Handbook of Labour Economics* (3), Ch. 27: 1560-1695, Elsevier Science.
- [25] BLUNDELL, R., L. PISTAFERRI AND I. PRESTON (2004), "Consumption inequality and partial insurance", Working Paper n.04/28, Institute for Fiscal Studies, London.
- [26] BLUNDELL, R. AND J. POWELL (2001), "Endogeneity in semiparametric binary response models", Working Paper n.05/01, Cemmap, London.
- [27] BOERI, T., A. BOERSCH-SUPAN AND G. TABELLINI (2001), "Would you like to shrink the welfare state? A survey of European citizens", *Economic Policy*, 32, 9-50.
- [28] BOTTAZZI, R. (2004), "Labour market participation and mortgage related borrowing constraints", IFS Working Paper no. 04/09, Institute for Fiscal Studies, London.
- [29] BOTTAZZI, R., T. JAPPELLI AND M. PADULA (2005), "Retirement expectations, pension reforms and their effect on private wealth accumulation", CEPR Discussion Paper no. 4882, Centre for Economics Policy Research, London.

- [30] BRANDOLINI, A. AND L. CANNARI (1994), "Methodological appendix: the Bank of Italy's Survey of Household Income and Wealth", in Ando, A., L. Guiso and I. Visco (eds): *Saving and the accumulation of wealth. Essays on Italian household and government saving behaviour*, Cambridge University Press, Cambridge.
- [31] BROWNING, M. (1992), "Children and household economic behaviour", *Journal of Economic Literature*, XXX(3): 1434-1475.
- [32] BROWNING, M. AND T. CROSSLEY (2001), "The life-cycle model of consumption and saving", *Journal of Economic Perspectives*, 15(3): 3-22.
- [33] BROWNING, M., A. DEATON AND M. IRISH (1985), 'A profitable approach to labour supply and commodity demands over the life-cycle"', *Econometrica*, 53(3): 503-544.
- [34] BROWNING, M., L.P. HANSEN AND J. HECKMAN (1999), "Micro data and general equilibrium models", in J. Taylor and M. Woodford (eds): *Handbook of Macroeconomics*, ch.8, North Holland.
- [35] BROWNING, M. AND A.M. LUSARDI (1996), "Household saving: micro facts and micro theories", *Journal of Economic Literature*, 34(4): 1797-1855.
- [36] BROWNING, M. AND C. MEGHIR (1991), "The effects of male and female labour supply on commodity demands", *Econometrica*, 59(4): 925-951.
- [37] BRUGIAVINI, A. (1999), "Social security and retirement in Italy", in Gruber, J. and D. Wise (eds): *Social security and retirement around the world*, The University of Chicago Press, Chicago.
- [38] CARROLL, C. (1994), "How does future income affect current consumption?", *Quarterly Journal of Economics*, 109(1): 111-147.
- [39] CARROLL, C. (1997), "Buffer-stock saving and the life cycle/permanent income hypothesis", *Quarterly Journal of Economics*, 112(1): 1-55.
- [40] CARROLL, C. (2001), "A theory of consumption function, with or without liquidity constraints", *Journal of Economic Perspectives*, 15(3): 23-45.

- [41] CHAMBERLAIN, G. (1984), "Panel data", in *Handbook of Econometrics* (2), Ch. 22: 1248-1318, Griliches, Z. and M. D. Intriligator (eds), Elsevier Science.
- [42] DEATON, A. (1985), "Panel data from a time-series of cross-sections", *Journal of Econometrics*, 30: 109-126.
- [43] DEATON, A. (1991), "Saving and Liquidity Constraints", *Econometrica*, 59(5): 1221-1248.
- [44] DEATON, A. AND J. MUELLBAUER (1980), *Economics and consumer behaviour*, Cambridge University Press, Cambridge.
- [45] DEL BOCA, D. AND A. LUSARDI (1997), "Finanziamento della casa e partecipazione al lavoro: un'analisi empirica per l'Italia" (in Italian), in *Ricerche Quantitative per la Politica Economica 1997*, Convegno Banca d'Italia-CIDE, Bank of Italy.
- [46] DEL BOCA, D. AND A. LUSARDI (2001), "Credit market regulations changes and labour market decisions", Working Paper n.22/2001, Child, Turin.
- [47] DIAZ, A. AND LUENGO-PRADO, M. J. (2003), "Precautionary savings and wealth distribution with durable goods", *mimeo*, Northeastern University, Boston, MA.
- [48] DISNEY, R., A. HENLEY AND D. JEVONS (2003) "House price shocks, negative equity and household consumption in the UK", *mimeo*, University of Nottingham.
- [49] DISNEY, R. AND S. TANNER (1999), "What can we learn from retirement expectations data?", IFS Working Paper no. W99/17, Institute for Fiscal Studies, London.
- [50] DOMINITZ, J. (1998), "Earnings expectations, revisions, and realizations", *Review of Economics and Statistics*, 80(3): 374-388.
- [51] DOMINITZ, J. (2001), "Estimation of income expectations models using expectations and realization data", *Journal of Econometrics*, 102: 165-195.
- [52] DOMINITZ, J. AND C. MANSKI (1994), "Using expectations data to study subjective income expectations", *mimeo*.

- [53] DOMINITZ, J., C. MANSKI AND J. HEINZ (2002), "Social security expectations and retirement savings decisions", NBER Working Paper no. 8718.
- [54] ENGELHARDT, G. (1996), "House prices and home owner saving behaviour", *Regional Science and Urban Economics*, 26: 313-336.
- [55] FELDSTEIN, M. (1974), "Social security, induced retirement and aggregate capital accumulation", *Journal of Political Economy*, 82: 357-374.
- [56] FERNANDEZ-VILLAVARDE, J. AND D. KRUEGER (2001) "Consumption and saving over the life cycle: how important are consumer durables?", *mimeo*
- [57] Financial Service Authority, "FSA guide to mortgages", London, July 2001.
- [58] FLAVIN, M. (1981), "The adjustment of consumption to changing expectations about future income", *Journal of Political Economy*, 89(5): 974-1009.
- [59] FLAVIN, M. AND S. NAKAGAWA (2004), "A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence", NBER Working Paper no. 10458.
- [60] FLAVIN, M. AND T. YAMASHITA (2002), "Owner-occupied housing and the composition of the household portfolio", *American Economic Review*, 92(1): 345-362.
- [61] FORTIN, N. (1995), "Allocation inflexibilities, female labour supply, and housing assets accumulation: are women working to pay the mortgage?", *Journal of Labour Economics*, 13 (3): 524-557.
- [62] FRIEDMAN, M. (1957), "A theory of the consumption function", Princeton University Press, Princeton.
- [63] GALE, W. (1998), "The effect of pension wealth on household wealth: a reevaluation of theory and evidence", *Journal of Political Economy*, 106: 706-723.
- [64] GOURINCHAS, P.O. AND J. PARKER (2002), "Consumption over the life cycle", *Econometrica*, 70(1): 47-89.

- [65] GROSSMAN, S. AND G. LAROQUE (1990), "Asset pricing and optimal portfolio choice in the presence of illiquid housing consumption goods", *Econometrica*, 58(1): 25-51.
- [66] GRUBER, J. AND D. WISE (1999), "Social security and retirement around the world", The University of Chicago Press, Chicago.
- [67] GUISO, L., T. JAPPELLI AND D. TERLIZZESE (1992), "Earnings uncertainty and precautionary saving", *Journal of Monetary Economics*, 30: 307-337.
- [68] GUISO, L., T. JAPPELLI AND D. TERLIZZESE (1992), "Income risk, borrowing constraints, and portfolio choice", *American Economic Review*, 86: 158-172.
- [69] GUSTMAN, A. AND T. STEINMEIER (2001), "Imperfect knowledge, retirement and saving", NBER Working Paper no. 8406.
- [70] HAURIN, D. R., P. H. HENDERSHOTT AND S. M. WACHTER (1996), "Borrowing constraints and the tenure choice of young households", NBER Working Paper no. 5630.
- [71] HAURIN, D. R., S. M. WACHTER AND P. H. HENDERSHOTT (1995), "Wealth accumulation and housing choices of young households: an exploratory investigation", NBER Working Paper no. 5070.
- [72] HEATHCOTE, J., K. STORESLETTEN AND G. VIOLANTE (2005), "Insurance and opportunities: the welfare implications of rising wage dispersion", *Journal of Labour Economics*, 13 (3): 524-557.
- [73] HECKMAN, J. (1974), "Life cycle consumption and labour supply: an explanation of the relationship between income and consumption over the life cycle", *American Economic Review*, 64(1): 188-194.
- [74] HECKMAN, J. (1979), "Sample selection bias as a specification error", *Econometrica*, 47(1): 153-61.
- [75] HECKMAN, J. AND T. MACURDY (1980), "A life cycle model of female labour supply", *Review of Economic Studies*, 64(1): 188-194.

- [76] HONORÉ, B. E. (2002), "Nonlinear models with panel data", mimeo, Department of Economics, Princeton University.
- [77] HUBBARD, R.G., J. SKINNER AND S. ZELDES (1994), "Precautionary saving and social insurance", *Journal of Political Economy*, 103(2): 360-399.
- [78] HUGGETT, M. AND G. VENTURA (1999), "On the distributional effects of social security reform", *Review of Economic Dynamics*, 2: 498-531.
- [79] HURD, M. AND K. MCGARRY (1995), "Evaluation of the subjective probabilities of survival in the Health and Retirement Study", *Journal of Human Resources*, 30: S268-292.
- [80] JAPPELLI, T. (1990), "Who is credit constrained in the US economy?", *Quarterly Journal of Economics*, 105: 219-234.
- [81] JAPPELLI, T. AND L. PISTAFERRI (2000), *Risparmio e scelte intertemporali* (in Italian), Il Mulino, Bologna.
- [82] JONES, L.D. (1990), "Current wealth constraints on the housing demand of young owners", *Review of Economics and Statistics*, LXXII: 424-432.
- [83] KIMBALL, M. (1990), "Precautionary savings in the small and in the large", *Econometrica*, 58(1): 53-73.
- [84] KRUEGER, D. AND F. PERRI (2003), "On the welfare consequences of the increase in inequality in the United States", NBER working paper no. 9993.
- [85] LaFAYETTE, W. C., D. R. HAURIN AND P. H. HENDERSHOTT (1995), "Endogenous mortgage choice, borrowing constraints and the tenure decision", NBER Working Paper no. 5074.
- [86] LEECE, D. (2000), "Choice of mortgage instrument, liquidity constraints and the demand for housing debt in the UK", *Applied Economics*, 32: 1121-1132.
- [87] LI, W. AND R. YAO (2004), "The life-cycle effects of house price shocks", *mimeo*.

- [88] LOW, H. W. (2005) "Self-insurance in a life-cycle model of labour supply and saving" *Review of Economic Dynamics* (forthcoming)
- [89] LOW, H.W., MEGHIR, C. AND PISTAFERRI, L. (2005) "Wage risk and employment risk over the life-cycle" *mimeo*.
- [90] LUENGO-PRADO, M. J. (2004), "Durables, nondurables, down payments and consumption excesses' ', *mimeo*, Northeastern University, Boston, MA.
- [91] LUSTIG, H. AND S. VAN NIEUWERBURGH (2002), "Housing collateral, consumption insurance and risk premia", *mimeo*.
- [92] MODIGLIANI, F. (1986), "Life cycle, individual thrift, and the wealth of nations", *American Economic Review*, 76(3): 297-313.
- [93] MEGHIR, C. AND G. WEBER (1996), "Intertemporal nonseparability or borrowing restrictions? A disaggregate analysis using a U.S. consumption panel", *Econometrica*, 64(5): 1151-1181.
- [94] MINIACI, R. AND G. WEBER (2000), "Econometric Issues in the estimation of household portfolio models", in Guiso, L., M. Haliassos and T. Jappelli (eds): *Household portfolios*, MIT Press.
- [95] MODIGLIANI, F. AND R. BRUMBERG (1954), "Utility analysis and the consumption function: an interpretation of cross-section data", in K. Kurihara (ed.): *Post-Keynesian Economics*, Rutgers University Press, New Brunswick.
- [96] MUELLBAUER, J. AND A. MURPHY (1997), "Booms and Busts in the UK Housing Market," *The Economic Journal*, 107(445): 1701-1727.
- [97] NICOLETTI-ALTIMARI, S. AND M. D. THOMSON (1995), "The effect of liquidity constraints on consumption and labour supply: evidence from Italian households", *Temi di Discussione n. 252*, Bank of Italy.
- [98] NORDVIK, V. (2001), "A Housing Career Perspective on Risk," *Journal of Housing Economics*, 10: 456-471.

- [99] ORTALO-MAGNÉ, F. AND S. RADY (2002), "Tenure Choice and the riskiness of non-housing consumption," *Journal of Housing Economics*, 11: 266-279.
- [100] ORTALO-MAGNÉ, F. AND S. RADY (2005), "Housing market dynamics : on the contribution of income shocks and credit constraints", Discussion Paper 2005-01, Department of Economics, University of Munich.
- [101] PISTAFERRI, L. (2003), "Anticipated and unanticipated wage changes, wage risk, and intertemporal labour supply", *Journal of Labour Economics*, 21(3).
- [102] SKINNER, J. (1988), "Risky income, life cycle consumption, and precautionary savings", *Journal of Monetary Economics*, 22(2): 237-255.
- [103] SMITH, R. AND R. BLUNDELL (1986), "An exogeneity test for a simultaneous equation Tobit model with an application to labor supply", *Econometrica*, 54 (3): 679-685.
- [104] TAUCHEN, G. (1986), "Finite state Markov-chain approximations to univariate and vector autoregressions", *Economics Letters*, 20: 177-181.
- [105] WEBER, G. (1993), "Earnings-related borrowing restrictions: empirical evidence from a pseudo panel for the U.K.", *Annales d'économie et de statistique*, 29:157-173
- [106] WOOLDRIDGE, J. M. (2002), *Econometric analysis of cross section and panel data*, Cambridge, MA: MIT Press.
- [107] YOSHIKAWA, H., AND F. OHTAKE (1989), "An analysis of female labour supply, housing demand and the saving rate in Japan", *European Economic Review*, 33: 997-1030.
- [108] ZELDES, S. (1989), "Consumption and liquidity constraints: an empirical investigation", *Journal of Political Economy*, 97(2): 305-346.