

The Effect of Income Risk, Asset Risk and Policy Risk on Household Behaviour

Ben Etheridge

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

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

London

July 2012

For the journey is done and the summit attained,
And the barriers fall,
Though a battle's to fight ere the guerdon be gained,
The reward of it all.

Robert Browning, *Prospice*

Declaration

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

Ben Etheridge

Abstract

This thesis quantitatively examines the types of risk that households face, how they prepare for these risks, and the effect of these risks on inequality.

The first substantive chapter reviews the evolution of inequality over 1978 to 2005 in the UK along several dimensions and serves as an introduction to subsequent chapters. Following the inequality surge in the 1980s, inequality generally rose more slowly in the 1990s on most measures.

The second chapter seeks to explain a puzzling episode in the evolution of inequality in the late 1990s: consumption inequality rose while income inequality fell. I explain this episode by accounting for two features of the UK economy over the period: a house price boom and a sequence of redistributive reforms by the new Labour government. I conclude that asset price movements and government policies can have a noticeable effect on ‘permanent’ (consumption) inequality and that the redistributive effect of the reforms was largely undone by the coincident house price boom.

The third chapter uses panel data over 1991 to 2006 to estimate the transmission of income shocks through to consumption. Only around 50% of ‘permanent’ income shocks are transmitted. This estimate reconciles two views of risk over the period: long-lasting income fluctuations, measured by panel data on incomes alone, were high, while consumption risk, measured by the growth in consumption inequality, was much lower. The results further indicate that such income ‘shocks’ are either not fully permanent or are often foreseen by younger households.

The fourth chapter theoretically examines the precautionary savings motive for consecutive income risks. In most cases (and particularly when facing permanent shocks) households can combine saving for near-term risks with saving for long-term risks. I term this saving ‘complementary’. However, in some interesting cases, the interaction of future risks amplifies the precautionary motive.

Contents

1	Introduction	16
2	Consumption, Income and Earnings Inequality in the UK	20
2.1	Introduction	20
2.2	Macroeconomic Conditions and Data Overview	23
2.2.1	Employment, Growth and Macroeconomic Conditions	23
2.2.2	Data Sources and Definitions	27
2.2.3	Comparisons with UK National Income Product Accounts (NIPA)	31
2.3	Hours, Wages and Earnings Inequality	35
2.3.1	Wages	35
2.3.2	Wage Premia	36
2.3.3	Wage Inequality, Earnings and Labour Supply	37
2.4	From Wages to Disposable Income	39
2.5	Consumption Inequality	43
2.5.1	The Inequality Boom and After	43
2.5.2	From Income to Consumption Inequality	45
2.5.3	The Life-Cycle Dimension	47
2.6	Distributional Dynamics	49
2.7	Top Incomes: The New Inequality	53
2.8	Interpretations and Conclusions	56
A2	Appendix to chapter 2	57
3	Increasing Inequality and Improving Insurance: House Price Booms and the Welfare State in the UK	59

3.1	Introduction	59
3.2	Model	64
3.2.1	Overview	64
3.2.2	The Household's Life-Cycle Programme	65
3.2.3	Modelling the UK Benefit System	68
3.2.4	Modelling Wealth Formation	71
3.2.5	Modelling the House Price Process	73
3.2.6	Solution	74
3.3	Data	74
3.3.1	FES	75
3.3.2	BHPS	75
3.3.3	ODPM House Price Data	76
3.4	Results	77
3.4.1	Overview of Baseline Results	77
3.4.2	The Effect of the Benefit Reforms	80
3.4.3	The Effect of House Price Shocks on Inequality - Results and Intuition	85
3.5	Estimation Procedure	91
3.5.1	Estimating the Income Process	91
3.5.2	Estimating the House Price Process	94
3.5.3	Other Pre-Estimated and Imposed Parameters	94
3.5.4	Estimation by Method of Simulated Moments	96
3.5.5	Choice of Moments	97
3.5.6	Parameter Estimates	97
3.6	Alternative Explanations for the Observed Phenomena	100
3.7	Conclusions	101
A3	Appendix to chapter 3	103
A3.1	Computing Asset Moments Using the BHPS	103
A3.2	Approximating Changes to the Covariance Structure of Income and Consumption in the Presence of Asset Price Shocks	104
A3.3	Deriving the Approximate Income and Consumption Moments in the Presence of Benefit Reform	112

4	The Transmission of Permanent Income Shocks: Evidence from the UK	116
4.1	Introduction	116
4.2	Data from the BHPS	121
4.3	The Model and Estimation	124
4.3.1	The Model of Income and Consumption Dynamics	124
4.3.2	Using Food Expenditures to Infer Consumption Choices	126
4.3.3	Identification and Estimation of Risk and Transmission Parameters	129
4.4	Results	131
4.4.1	The Covariance Structure of Food and Income Changes	131
4.4.2	Results From the Stationary Model	132
4.4.3	The Age Profile of Income Risk	135
4.4.4	Interpreting the Results Using Other Models of Income Risk	137
4.4.5	Comparison with Chapter 3	139
4.5	Taking Midpoints of Food Consumption	139
4.5.1	Analysis of the BHPS data	140
4.5.2	Validation from the PSID	142
4.5.3	Final Remarks on the Data Imputation	144
4.6	Conclusions	145
A4	Appendix to chapter 4	147
A4.1	Using FES Expenditure Data	147
5	For Many a Rainy Day: Precautionary Saving for Consecutive Life-Cycle	
	Risks	151
5.1	Introduction	151
5.2	The Model and Definition of Complementarity	156
5.2.1	The Budget Constraint and Income Process in the Basic 3-Period Model	156
5.2.2	The Consumption Problem	158
5.2.3	Definition of Complementarity	160
5.3	Quantifying Complementarity of Precautionary Saving	161
5.4	Analytical Characterization of Complementarity	166

5.4.1	Complementarity for Permanent Shocks	166
5.4.2	Excessiveness for Small Transitory Risks	167
5.4.3	Further Intuition	170
5.4.4	Classifying Utility Functions	171
5.4.5	Final Remarks on the Characterization of Complementarity	173
5.5	Conclusion	173
A5	Appendix to chapter 5	175
A5.1	Reweighting The Permanent Shock Process	175
A5.2	Derivation of the Approximations for $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$	176
6	Conclusion: Thoughts on Future Research	182
	Bibliography	185

List of Tables

2.1	The Autocovariance Structure of Wage Growth for Male Head	51
3.1	Announcements on the Minimum Wage and Tax Credits	70
3.2	Welfare Effect of All Benefit Reforms, Across the Income Distribution . . .	85
3.3	Wealth-Share Parameters for the 1950s Cohort: All Types and High Edu- cation Alone	88
3.4	Parameter Estimates	98
3.5	Model Fit	99
4.1	Theoretical Consumption and Income Covariance Moments	126
4.2	Covariances of Residual Food and Income Changes	132
4.3	Key Estimates from the Pooled Model	133
4.4	Key Parameter Estimates: Breakdown by Sample	134
5	Comparison of Means, BHPS and FES	148
6	Food Demand in the UK	150
5.1	Simulated Household Wealth at Age 35 as % of Initial Income	165

List of Figures

2.1	The Pattern of Overall Inequality in the UK since 1978 – Gini of Equivalized Disposable Income	21
2.2	Overall Employment and Growth Rates	24
2.3	Employment by education and gender, by year	25
2.4	Proportion of Households with a Working Adult	26
2.5	Employment over the Life-Cycle: Men	27
2.6	Employment over the Life-Cycle: Women	27
2.7	Income Per Capita: FES vs NIPA	32
2.8	Income and Consumption Coverage - FES Totals as a Percentage of NIPA Totals	33
2.9	Consumption Coverage – Selected Categories	33
2.10	Labour Force Participation: FES vs NIPA	35
2.11	Inequality in Wages (Prime-Age Sample)	36
2.12	Wage Premia	37
2.13	Inequality in Labour Supply	38
2.14	Correlation across the life-cycle: Correlation between Male Wages and Hours, by Age (Cubic Polynomial Fit)	39
2.15	Inequality in Equivalized Household Earnings	39
2.16	From Wages to Disposable Income – Heads in Employment	40
2.17	From Earnings to Disposable Income – Heads with Earnings from Employment and Self-Employment	41
2.18	Disposable Income Inequality: Decomposition by Sample	41
2.19	From Earnings to Disposable Income (At least one worker)	42

2.20	Quantiles of the Household Gross Income Distribution	43
2.21	Quantiles of the Household Net Income Distribution	43
2.22	Consumption Inequality: Decomposition	44
2.23	Consumption Inequality: Decomposition by Sample	45
2.24	From Disposable Income to Consumption	46
2.25	Covariance of Disposable Income and Consumption	46
2.26	Life-Cycle Dispersion, Controlling for Cohort Effects	47
2.27	Life-Cycle Dispersion, Controlling for Year Effects	49
2.28	Life-Cycle Household Consumption Dispersion, by Cohort	49
2.29	Variance of Permanent and Transitory Shocks: Labour Earnings Sample . .	52
2.30	Variance of Permanent and Transitory Shocks: Household Disposable Income	53
2.31	Real income growth for the richest 10% and 1% using the SPI, 1996-97 to 2004-05 (GB)	54
2.32	Change in Top Net Income Shares	55
2.33	Income Components for the Top 1%	55
3.1	Variance of Log Income and Consumption in the UK	60
3.2	Effect of Tax Credits on Income: Cohort Born in the 1950s with Low Edu- cation	69
3.3	Empirical and Simulated Inequality	78
3.4	The Evolution of Consumption Inequality	79
3.5	House Prices and Difference Between $\text{Var}(\ln C)$ and $\text{Cov}(\ln C, \ln Y)$ in the Data	79
3.6	House prices and Difference between $\text{Var}(\ln C)$ and $\text{Cov}(\ln C, \ln Y)$ in the Simulations.	80
3.7	Breakdown of the Effect of House Price Shocks and Benefits on the Excess Growth of $\text{Var}(\ln C)$	81
3.8	Covariance Structure of Income Residuals: $\text{Cov}(y_{t+s}, y_t)$	93
3.9	Life-Cycle Consumption Needs By Cohort and Education	95
3.10	Income Profiles by Cohort and Education	95
3.11	Mean Equivalized Consumption Growth	100
4.1	Comparison of food expenditures in the FES and BHPS	124

4.2	Age Profile of Permanent and Transitory Shocks	136
4.3	Age Profile of Transmission Coefficients	138
4.4	Estimating the Joint Distribution of Food and Income Changes	141
4.5	The Distribution of Food Consumption from the PSID	143
4.6	Estimating 2nd Moments of the Food Distribution in the PSID	143
4.7	Estimating the Joint Distribution of Food and Income Changes in the PSID	144
5.1	Precautionary Saving is Complementary for Overall Life-Cycle Risks	165
5.2	Precautionary Saving is Complementary for Permanent Risks Alone	168
5.3	Precautionary Saving is Excessive for Transitory Risks Alone	169
5.4	CARA Preferences: Saving is Neutral for Transitory Risks	171
5.5	...While Saving is Complementary for Permanent Risks	172

Acknowledgements

Funding Arrangements

I am grateful to the Economic and Social Research Council (ESRC) for generous funding through the Institute for Fiscal Studies PhD scholarship programme.

Data Provision

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

Conjoint Work

Two of the chapters of this thesis are joint work:

1. Chapter 2 was written jointly with Richard Blundell
2. Chapter 5 forms part of a joint project with Richard Blundell and Tom Stoker

I thank these co-authors for allowing me to draw on these works in this thesis. The remaining chapters are my own work, as, of course, is responsibility for any errors.

Personal Thank Yous

From far, from eve and morning
 And yon twelve winded sky,
 The stuff of life to knit me
 Blew hither: here am I.

A.E. Housman, UCL Professor of Latin Poetry (1892-1911)

A Shropshire Lad: XXXII

There are numerous people to whom I owe thanks, and without whom this thesis would not have been possible. My greatest thanks go to my primary supervisor Richard Blundell. To learn first-hand from an exceptional scholar is a true privilege. Not only has he guided me through all the intellectual contributions in this thesis, he has maintained faith in me throughout the programme.

Special thanks must also go to Hamish Low, who has acted as both an informal supervisor and mentor over the last couple of years. I feel that I have developed immensely as a researcher under his thoughtful, insightful and expert guidance.

Many thanks also go to my second supervisor, James Banks. It was his inspirational masters course in the economics of ageing that originally motivated me to pursue a PhD on household behaviour using micro data. He also provided lots of ideas and advice at various periods through my research.

I would also like to thank my office mates over the course of the PhD. Thibaut Lamadon, Matthias Parey, Andrew Shephard and Elaine Kelly, all outstanding people, provided me with exciting conversation, excellent company and often much-needed technical help.

And to my family. First I wish to thank my parents and my wider family. During this PhD they have offered amazing support with both housing and childcare. On a more fundamental level, of course, I owe my parents everything. I am extremely grateful.

I give genuine thanks to my beloved sons Max and Rufus, whose arrival precipitated far greater maturity on my part and who have provided me with such happiness away from work, and the motivation to keep going.

My main thanks, however, go to Lisa for her unfailing support over the (many) years of this doctoral research. I am sure I would not have completed this PhD without a partner who had finished one herself. I intend to make sure in the future that it was worth all her love, sympathy and patience.

Chapter 1

Introduction

This dissertation presents four substantive chapters on the economics of the household. Specifically, I study household welfare and consumption behaviour in the face of different risks: income risk, asset price (particularly house price) risk and shocks to government tax and spending policies. Across the dissertation I try to quantify these risks and to assess both theoretically and empirically how households cope with them. In the first half of the dissertation (chapters 2 and 3) I place particular emphasis on how these risks combine to affect inequality and the distribution of living standards in the UK, particularly since 1990. Chapter 4 empirically measures the ex-post response of consumption to income shocks. The final chapter (chapter 5) is theoretical and concerns the ex-ante effect of risk on saving behaviour. A unifying theoretical structure across all chapters is the standard life-cycle model of Brumberg and Modigliani (1954).

Chapter 2 first presents the main datasets used in the rest of the empirical analysis: the Family Expenditure Survey (FES)¹ and the British Household Panel Survey (BHPS). It then presents an analysis of the trends in inequality across income, earnings and consumption in the UK over 1978-2005. The chapter links macroeconomic and microeconomic analyses of inequality. Overall the period is dominated by the inequality boom in the 1980s, which has been studied widely by, for example, Blundell and Preston (1998) and Gosling et al. (2000). Thereafter, the evolution of inequality is more nuanced and episodic. As far as inequality in household incomes is concerned, one episode stands out: after an

¹This has since been renamed the Expenditure and Food Survey, and more recently the Living Costs and Food Survey. For the rest of the thesis I refer to it by its historical name.

increase over the early 1990s, income inequality made a pronounced drop in the late 1990s before rising again in the early 2000s.

Chapter 3 builds on the preceding chapter by examining two related puzzles connected to this drop in income inequality in the late 1990s. First, while income inequality declined, consumption inequality increased. Second, the rise in consumption inequality became dissociated from shocks to permanent income. In a stochastic, life-cycle model of consumption, two factors are needed to explain these movements. First the house price boom exacerbated wealth inequality and caused growth in consumption inequality separate from income inequality. Second, income inequality was reduced by changes to social insurance introduced by the Labour government after 1997 aimed at raising income at the bottom end of the distribution. This compressed the distribution of income, but had less effect on the distribution of consumption: the greater insurance was not matched by increases in lifetime wealth. Introducing these factors into the model explains around 35% of the excess growth in consumption inequality. The extra insurance after 1997 was particularly important in the evolution of inequality for the low educated, whereas the house price boom was more important in the evolution for the high educated.

Chapter 4 also builds on empirical puzzles documented in chapter 2. In this chapter I estimate the transmission of income shocks through to consumption using UK panel data over 1991-2006. I find that only about 50% of permanent income shocks are transmitted, while transitory income shocks are almost completely smoothed. These estimates are more-or-less constant across education groups and across cohorts. These estimates reconcile two views of risk over the period: permanent income risk, measured by panel data on incomes alone, was high, while consumption risk, measured by the growth in consumption inequality, was much lower. In fact, I find that income shocks accounted for around 80% of consumption risk, of which shocks to wages of the head contributed around a half. I conclude by noting that my estimate of the transmission of permanent shocks is lower than is implied by standard models of self-insurance and particularly so for younger groups. One interpretation is the presence of substantial extra consumption insurance. Other plausible interpretations are either the presence of advance information about future income or the absence of a unit root on 'permanent' income shocks, even though neither of these features can be detected directly.

The final substantive chapter (chapter 5) moves away from empirical work to a more theoretical flavour. It concerns the precautionary motive for saving. In particular, I focus on precautionary saving that is driven by ‘prudence’. Intuitively, prudence reflects the strength of the desire to have a higher level of wealth when facing risks. When households are very prudent they are willing to consume less today so that they can have a higher wealth cushion whenever bad shocks may strike.

The chapter itself is motivated by the following observation: households face a variety of motivations for saving, for example for near-term possible emergencies (near-term risk) and for far-off possible emergencies (long-term risk). It seems intuitively plausible that households need not save for different emergencies separately but can combine saving for all future emergencies. I term this behaviour ‘complementarity’ of saving and provide a formal definition. To quantify this complementarity effect I simulate a realistically parametrized life-cycle model with permanent and transitory income risks. The complementarity effect accounts for 8-16% of precautionary savings, depending on the precise specifications used. This effect is driven by the permanent shocks. In order to examine the key mechanisms at work I then look at a simplified 3-period model and characterize saving behaviour analytically without restricting the shape of the utility function or the (within-period) distribution of shocks. I find that permanent shocks induce complementarity for a general class of preferences, including those with constant relative risk aversion (CRRA). However, most preferences in this class, and especially CRRA, display the opposite effect for transitory shocks. In this case the interaction of risks amplifies the precautionary motive, although the effect is small. These results can be interpreted in terms of the inter-temporal connectedness of risks and the pattern of prudence over the wealth spectrum. For example, for CRRA preferences, relative prudence is constant and so the effects are driven by the structure of risks alone: permanent risks allow for complementarity chiefly because the variance of future innovations to life-time wealth declines with bad shocks, and so households need save less for consecutive bad draws. However, when relative prudence is not constant then it can play a key role. The chapter discusses several examples of these types of preferences.

Chapter 6 concludes by discussing ideas for future research. I consider topics either continuing on from or related to those addressed in this thesis and which could be re-

searched using techniques similar to those used here. In particular I discuss questions around the effect of policy uncertainty, the demand for housing, risks around household formation and dissolution and the interaction of different savings technologies for different savings needs.

Chapter 2

Consumption, Income and Earnings Inequality in the UK

2.1 Introduction

Inequality growth in the UK over the past three decades has been episodic. This is clearly illustrated in figure 2.1 which depicts the evolution of the Gini for family income in the UK. There is a well documented¹ inequality ‘boom’ in the early 1980s followed by a period of stability albeit at a higher level of inequality. Then, in the late 1990s, a further rise in inequality occurred largely concentrated at the top of the income distribution and predominantly on employment income in the financial industry.²

This description of inequality growth in Britain refers exclusively to inequality in income and more specifically to earned income inequality. Economic inequality has many linked dimensions – wages, earnings, income and consumption. So, what of inequality in the components of earnings – wages and hours? What of the differences across gender? What of consumption inequality? And what of after tax income and the role of tax and transfers? The aim of this chapter is to provide a coherent analysis of the trends in these various measures of economic inequality.

During the 1980s ‘inequality boom’ the Gini for income rose by a full ten points from around .23 to .33, a large increase by any comparison. We show that this increase in

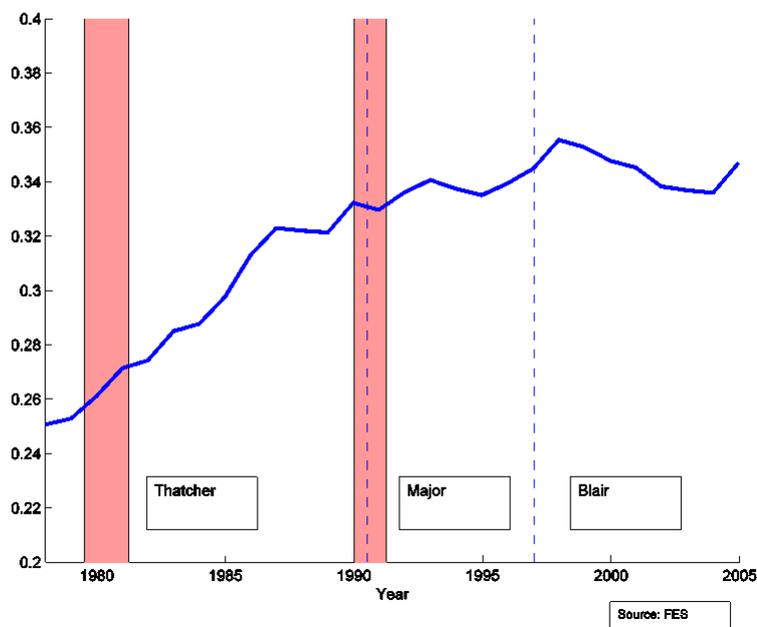
¹Atkinson (1997)

²See Atkinson and Piketty (2007) and Brewer et al. (2007b)

inequality was reflected across the distribution and in the components of income. It is particularly evident in the earnings distribution, reflecting the change in returns to education and skill over this period. Over the inequality boom period, especially in the early 1980s, there was a corresponding sharp rise in consumption inequality, although this tailed off earlier than did the growth in earnings and wage inequality.

To fulfill this task we make use of a number of data sources. However, because we want a consistent series for these underlying variables dating back as far as possible we confine our main analysis to two data sources - the Family Expenditure Survey (FES) and the Labour Force Survey (LFS). The FES has collected data on expenditures, hours, earnings and unearned incomes on a consistent basis for nearly four decades. The LFS, which also has consistent measures of basic labour market variables, is based on a larger sample but has a more limited history of earnings and does not collect data on consumption. We also describe and draw on the British Household Panel Survey (BHPS), which is used more extensively in the rest of the dissertation.

Figure 2.1: The Pattern of Overall Inequality in the UK since 1978 – Gini of Equalized Disposable Income



Notes: The thick bands indicate recessions as defined by a drop of GDP for more than 2 consecutive periods

This study follows a large literature on inequality in the UK across various measures; see Atkinson (1997, 1999). We particularly draw on two previous studies. First, Gosling et al. (2000) who document and analyse changes in the wage structure in the UK over 15 years from the late 1970s using the FES.³ Second, is the Blundell and Preston (1998) study who decompose the income risk faced by different cohorts using FES data on household income and consumption dispersion. Ours is the first study to look closely at the co-evolution over time of wages and hours, through to earnings, to household income and finally to consumption. In addition we present new results on income dynamics for the UK in the 1990s from the BHPS and relate these to our findings from the cross-sectional datasets.

This study is intended to fit into a wider literature studying the relationship between income risk, consumption insurance and inequality. The theoretical backbone to this work originated with the analysis of consumption dispersion in incomplete-market economies by Huggett (1993) and Aiyagari (1994). Around the same time Deaton and Paxson (1994) developed a test of the permanent income hypothesis through the empirical analysis of life-cycle profiles of consumption and income dispersion, using data drawn from a number of economies. Subsequently, a burgeoning literature has attempted to explain the empirical phenomena underlying the observed distributional dynamics and to answer key economic questions: for example, Blundell et al. (2008a), Blundell et al. (2008b), Guvenen (2007), Heathcote et al. (2008, 2010b), Krueger and Perri (2006) and Storesletten et al. (2004). Most of these studies have focused on the US.⁴ The main purpose of our study is to provide ‘key facts’ for the UK over the last three decades, which can feed in to the macroeconomic analysis of distributional dynamics.

We set the scene in the next section by documenting the broad macroeconomic and labour market background for the UK economy over the period since the late 1980s. We then present some details of the data sources used and their ability to match basic aggregate trends. Our attention then turns to the analysis of underlying earnings inequality. We note that the pattern of inequality over the 1980s inequality boom, as in the US, can

³Blundell et al. (2007) show these inequality trends to be largely robust to changes in employment levels and potential for self-selection biases documented in Blundell et al. (2003).

⁴Exceptions are Attanasio et al. (2002), Blundell and Preston (1998) and Blundell et al. (2007), which feature in the discussion below.

be explained by changes in the labour market, in particular to changes in the level and durability of shocks to earnings and changes in female labour supply. We further consider the components of income and earnings and the covariance structure between hours and wages for both men and women. We document a recent strengthening in the relationship between male wages and male hours.

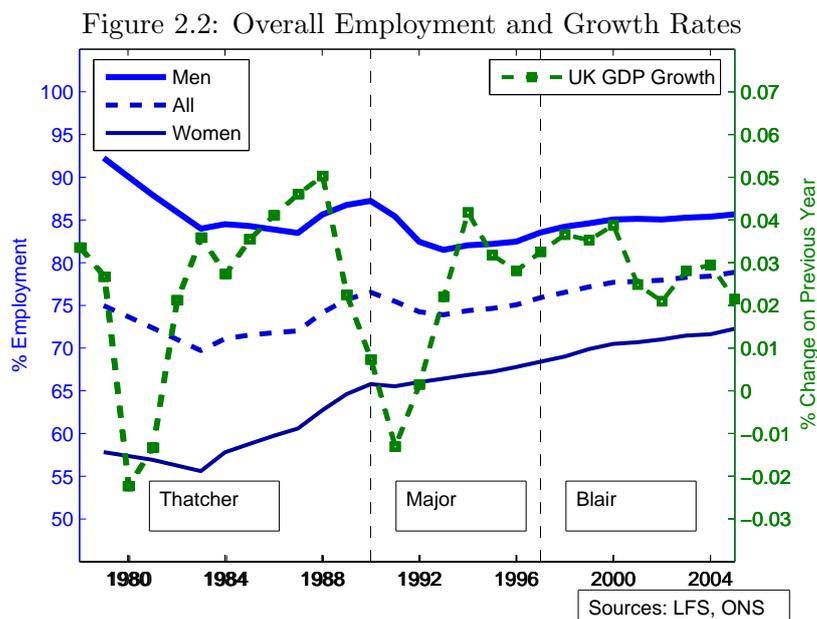
Our analysis continues with an examination of income and consumption inequality over the past three decades. We note the divergence, especially in the late 1980s, between income and consumption inequality. This was originally documented in Blundell and Preston (1998) for the UK and is similar to the findings for the US reported in Cutler and Katz (1992). Blundell et al. (2008b) follow up this study for the US and find that the divergence can be explained by initial growth in the variance of permanent shocks which was then replaced by a continued growth in the variance of transitory income shocks in the late 1980s. Indeed, using consumption and income inequality data for the UK, Blundell et al. (2008a) provide strong evidence of a spike in the variance of permanent shocks to income in the early 1980s. Unfortunately, we do not have panel data on income for the 1980s in the UK and are not able to examine the durability of income and earnings shocks during the inequality boom. However, we are able to examine the dynamics of the various definitions of income and earnings since the early 1990s using the British Household Panel Survey.

Before concluding we finish with a brief discussion of the ‘new inequality’ and the rapid rise in top incomes during the late 1990s.

2.2 Macroeconomic Conditions and Data Overview

2.2.1 Employment, Growth and Macroeconomic Conditions

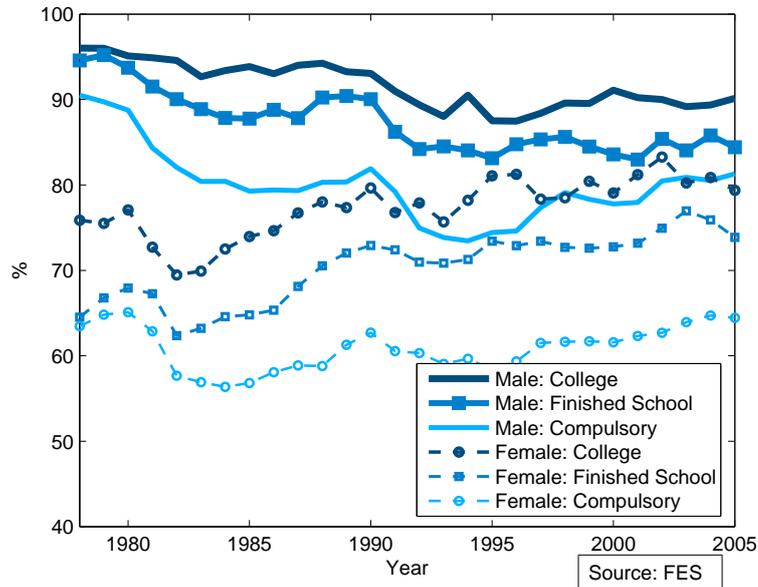
The sharp recession in the very early 1980s in the UK is clearly evident in figure 2.2 by the strong negative real GDP growth rate in 1980 and 1981. This was followed by a severe drop in employment rates for both women and men. Male employment rates never returned to their pre-1980 level in the period up to the financial crisis in 2008, although female employment rates show a strong secular trend upward over the period covered.



The second recession in this period followed soon after the peak growth rates at the end of the 1980s. From late 1993 onwards the economy moved into a period of stable and moderate growth, accompanied by a consistent rise in employment, interrupted only by the recent downturn. This overall growth in employment over this period was offset to some extent by the continued fall in labour market attachment among low skilled workers that extended throughout the first half of the 1990s. This reflected a fall in demand for low skilled workers over this period. This in turn engendered a change in welfare and tax policy that heralded a strong expansion in earned income tax credits and welfare to work policies in the late 1990s under the Blair government.⁵

⁵Blundell (2002)

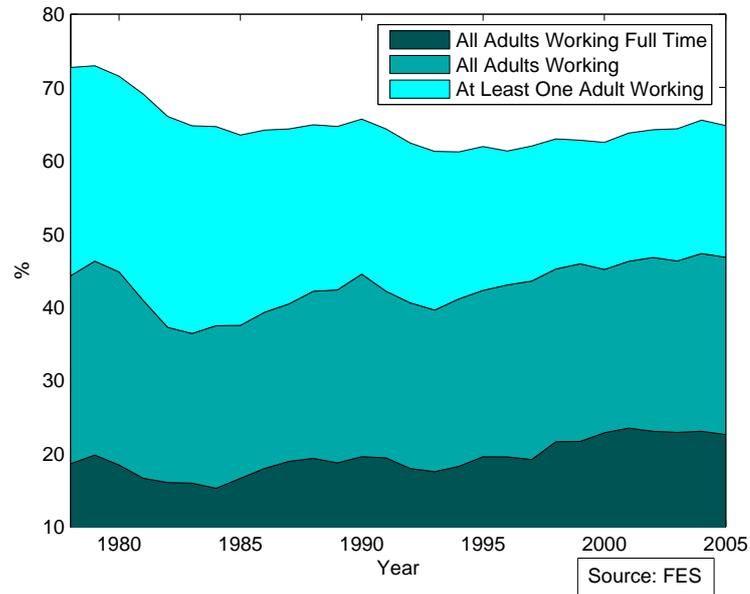
Figure 2.3: Employment by education and gender, by year



The detailed picture of labour market attachment over this period can be seen in figure 2.3. This highlights the impact of the early 1980s recession on the employment of low skilled men and women. Employment rates for lower educated women only very recently returned to the rates of the late 1970s, while for low educated men, employment rates remain below those of three decades ago.⁶ Figure 2.4 shows that this drop in employment among the low educated shows up in a lower level of households with at least one adult working, although the growth in female labour supply continues strongly throughout the period.

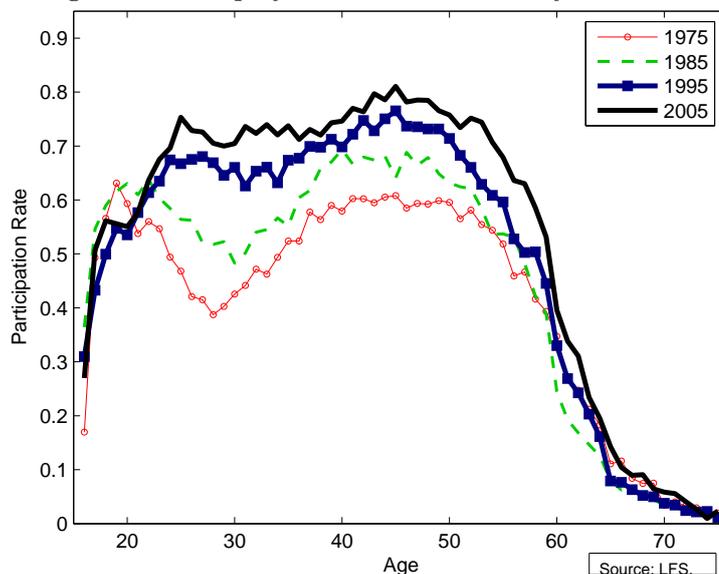
⁶Although not shown here, employment rates for single mothers, also continued to be lower, see Blundell and Hoynes (2004)

Figure 2.4: Proportion of Households with a Working Adult



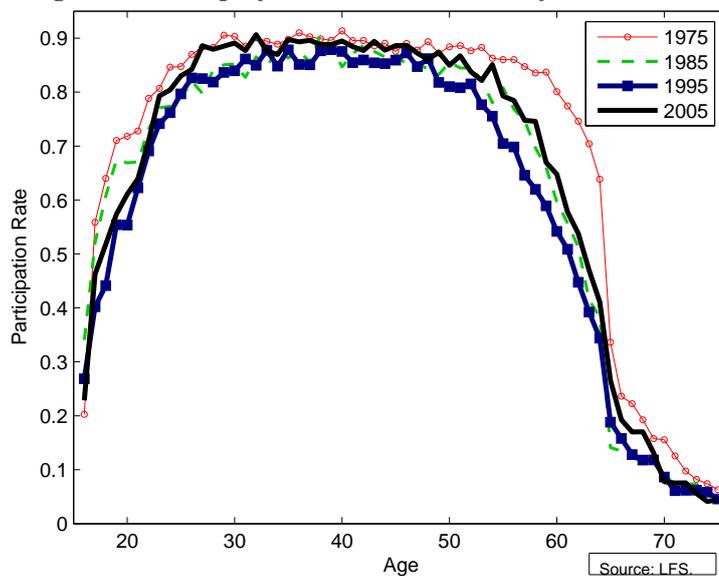
In the analysis that follows we will see that life-cycle changes matter too. The overall changes in working behaviour for men and women by age over this period are perhaps most dramatically documented in figures 2.5 and 2.6. These show that the impact of the 1980s recession on male employment was felt most among the relatively young and old, while the increase in female labour supply has happened most at child bearing years. These are key considerations for understanding changes in inequality across time, across age and across gender.

Figure 2.5: Employment over the Life-Cycle: Men



Source: Blundell, Bozio and Laroque (2008), LFS data.

Figure 2.6: Employment over the Life-Cycle: Women



Source: Blundell, Bozio and Laroque (2008), LFS data.

2.2.2 Data Sources and Definitions

As already noted, there are a number of key data sources used in the analysis reported here; we draw primarily from the consistent repeated cross-section household survey, the Family Expenditure Survey. For our analysis of income dynamics we draw on panel data from the BHPS, although this is only available from 1991 onwards. We analyse the recent

evolution of the top of the income distribution using data from the Survey of Personal Incomes. We also use data on participation from the Labour Force Survey over the entire survey period. In the remainder of this section we briefly describe these data sources and draw some comparison with the national income accounts.

The Family Expenditure Survey - FES

The principal dataset used in this chapter is the UK Family Expenditure Survey (FES). The FES is an annual survey conducted chiefly for determining the basket of goods used to construct the retail price index. The FES has been running since 1957, although it has only collected data in its present form on a consistent basis since the 1970s. In 2001, this dataset merged with the UK National Food Survey to create the Expenditure and Food Survey (EFS), but we shall make reference to the FES for the remainder of the dissertation. In a typical year the FES contains information on around 6500 households. Over the first few decades of the survey, the response rate was consistently over 70%. However, this has declined since the 1980s and fell to 58% in 2000. In general the households form a representative sample, but excluded are those not living in private houses, such as residents of residential homes or students.

For households participating in the FES, each member over 16 is asked to complete a diary detailing all their spending, both home and abroad, over a two week period. In addition to this diary, household members perform an interview in which they are asked questions about their demographic background, and asked to recall expenditures on large infrequently-purchased items (such as cars).

Because data on income have been collected consistently only since 1978, our sample period is 1978-2005. This gives a baseline sample for the analysis in this chapter of 197,190 households (369,599 adults, 496,067 individuals). To each household we allocate a head, in accordance with the guidelines for this project (usually the male in a household consisting of a married couple with children). For the majority of statistics quoted in this study, we use as population all households with heads aged 25-60. The sample is formed as follows: we drop 71,041 households for which the head is outside our age range; we then drop observations where food consumption or disposable income is negative (515 observations), leaving 125,614 households representing 370,343 individuals. For robustness of the results

we trim the top and bottom 0.25% of observations of each distribution. For consistency with the other variables, we follow this same procedure for wages, rather than selecting on the minimum wage or the wage of a typical low-skilled job. It is worth noting, however, that the minimum wage was introduced in the UK in 1999 at £3.60 for over-21s: our trimming point for this year is around 40% of this, at £1.41.

The British Household Panel Survey - BHPS

In order to study wage and income dynamics we use data from the British Household Panel Survey (BHPS). The BHPS is a comprehensive longitudinal study for the UK for general use in the social sciences, running since 1991. Like the US PSID it tracks individuals across household changes and tries to match the population age distribution by taking a refresher sample of new adults in each wave. In the first wave, it achieved a sample size of around 5000 households (10,000 adult interviews), a 65% response rate. The sample size has fallen somewhat since 1991, both because of sample attrition and because of a net outflow of households. In 2000 it achieved around 4200 complete interviews, a 75% response rate.

The BHPS has detailed information on earnings, hours worked and other income, and information on housing and durables, but little information on non-durable expenditure. An auxiliary dataset, documented in Bardasi et al. (1999) contains derived data on net household disposable income, which we use in this study.

We follow similar sample selection procedures for the BHPS as followed for the FES. The baseline sample is 68,027 households, comprising 166,144 individuals. We remove 24,414 households for whom the head is outside our age range. We then trim the bottom 0.5% of the distribution of disposable income and remove observations for which the head's education status is missing (346), leaving 43,017 households, comprising 122,269 individuals. Unlike the FES, where each questionnaire is completed in entirety, the BHPS contains many incomplete observations, so the quoted statistics are computed using fewer observations. For example, the total sample size of observed changes in household income is 24,363.

The Survey of Personal Incomes - SPI

The Survey of Personal Incomes (SPI) is an annual survey conducted by Her Majesty Revenue and Customs (HMRC, the UK equivalent of the US IRS) based on data collected on individuals who could be liable for income tax. We use these data to provide information on top incomes. The dataset is constructed as follows: stratified samples are drawn from three separate HMRC databases (those subject to pay-as-you-earn income taxation, self-assessment and neither of these). Variables that were used to stratify the sample include sex, pay, tax liability, main source of income and occupational pensions in previous years. Individuals with high incomes or rare allowances tend to be over-sampled. In 2004–05, this procedure produced a valid sample of 523,621 cases.

Around 15% of individuals within the SPI are not taxpayers, since their taxable income does not exceed the personal allowance (£4,745 in 2004–05). However, the SPI does not cover all non-taxpayers, since some individuals do not have any interaction with HMRC in a particular year, e.g. individuals without children on non-taxable state benefits.

The SPI contains data pertaining to before-tax income, sources of before-tax income, tax reliefs and some data on individual characteristics, e.g. sex, age group, industry and their marginal rate of income tax. However, the measure of total before-tax income (and some of its components) is incomplete because income that is not subject to tax is not provided to HMRC. Moreover, certain items have to be imputed by HMRC, e.g. investment income where tax has been deducted at source and personal pension contributions.

Certain steps also have to be conducted in order to ensure anonymity. All sources of income, deductions and reliefs are rounded to three significant figures, with tax amounts imputed based on these rounded figures. Unusual combinations of allowances must be examined to ensure no-one can be identified. Some variables are combined to further ensure anonymity. HMRC also ensures that no group has a sampling weight less than 1 in 60 or represents a population of less than 10,000. Finally, individuals with incomes greater than £600,000 are combined to create ‘composite records’ in order to ensure anonymity. This is done by combining cases with similar characteristics (e.g. same stratum and sex) and taking averages for each variable on the file.

The Labour Force Survey – LFS

The Labour Force Survey is a continuous household survey which provides the most detailed data on labour market characteristics such as participation, earnings, training and qualifications. The LFS has been running since 1973 and provides national accounts employment data. It was first collected every two years, then over 1983-1992 it was collected yearly, and since 1992 it has been collected quarterly, as a revolving panel lasting 5 quarters. The sample size in each wave is around 60,000 households covering 140,000 individuals. The survey has complete response to questions on participation; in a typical year, we collect round 100,000 responses for adults between 25 and 60. We do not use the data on earnings and wages, because these data have only been collected since 1992.

2.2.3 Comparisons with UK National Income Product Accounts (NIPA)

Here we present a comparison of per-capita disposable income, expenditure and employment from the UK national accounts and the FES. Owing to definitional and methodological differences, it would be unsurprising to find a difference in levels between the national accounts and FES. Moreover, both datasets are subject to measurement error of different kinds: the FES may include (possibly systematic) mis-reporting by households, while, for example, many national account expenditure items are formed as a residual from income, value-added and trade items in national accounting identities. Of particular interest is the size of any discrepancy, whether any such differences can be accounted for, and whether the two measures have the same time series properties. We give a brief overview of apparent differences between the two datasets: the issues are discussed in further detail in Tanner (1998) and Attanasio et al. (2006).

Figure 2.7 shows per-capita disposable net income in FES and national accounts, deflated by the RPI. The coverage of the FES has been consistently high over the sample period, rarely dropping below 90% of the national accounts level. For most of the period, the FES also matches the dynamics in the national accounts, matching the recession in the 1980s and slowdown in the early 2000s. The FES data departs significantly from the NIPA statistic only in 1992.

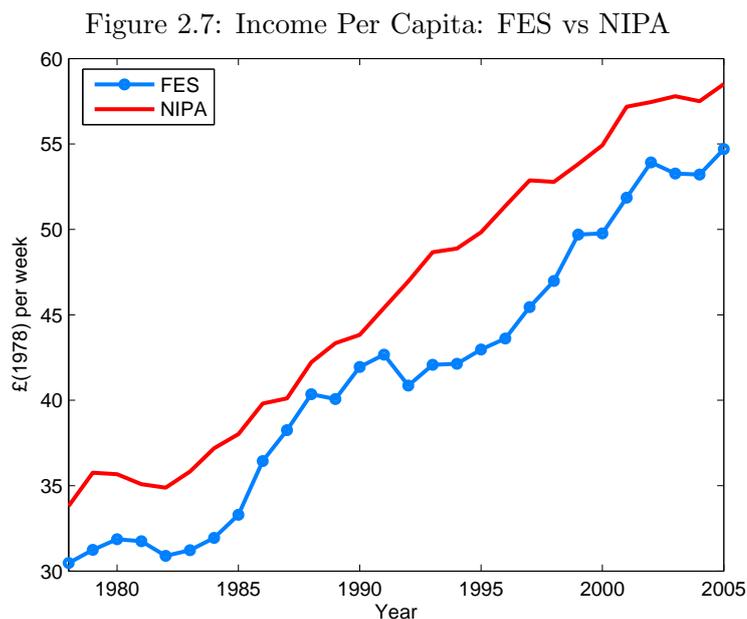


Figure 2.8 shows estimates of per-capita income and total expenditure from the FES as a proportion of national accounts data. The measure of expenditure used here is broader than that used in the rest of this study as we include durable and semi-durable goods, excluding housing and some other small items which are incompatible between the two data sets. The largest departure from national accounts for both income and expenditure occurs in the early 1990s. However, whereas income coverage suffers a pronounced dip in 1992, then recovers later in the decade; the coverage of consumption first begins to decline in 1993, but then to continues to decline.

In order to try to understand what may lie behind the declining performance of the expenditure data, it is worth looking at some of the components behind the total. Figure 2.9 shows the percentage coverage of certain items included in our consumption basket. Expenditure on food, clothing and catering matched the national accounts extremely well, both in levels and in dynamics until the late 1980s (and before the beginning of our sample period). Coverage for these items rarely fell below 90%. On the other hand, alcohol and tobacco have always had low coverage, but this is common for items that carry a social stigma. 1988 saw a sudden collapse in the coverage of catering, which suggests that there was a sudden change in measurement for this category in one of the datasets. However, for all other categories there has been no sudden shift, but a gradual decline in coverage, approximately since 1993. Therefore the decline in coverage of the aggregate since 1993

is not confined to certain items but seems to be a broad trend across many expenditure categories. The case of food expenditure is puzzling since the national accounts data for this item are formed mainly from the FES data. It may therefore be sensible to use the FES food coverage as a basis for comparison.

Figure 2.8: Income and Consumption Coverage - FES Totals as a Percentage of NIPA Totals

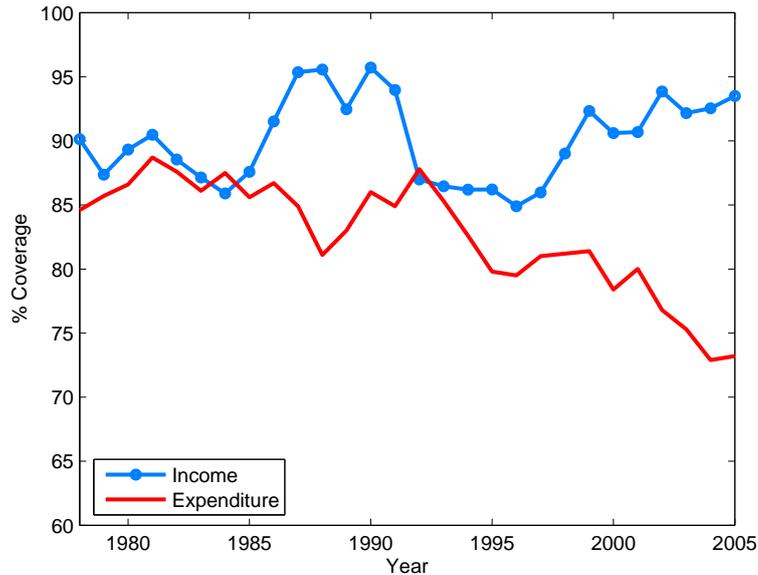
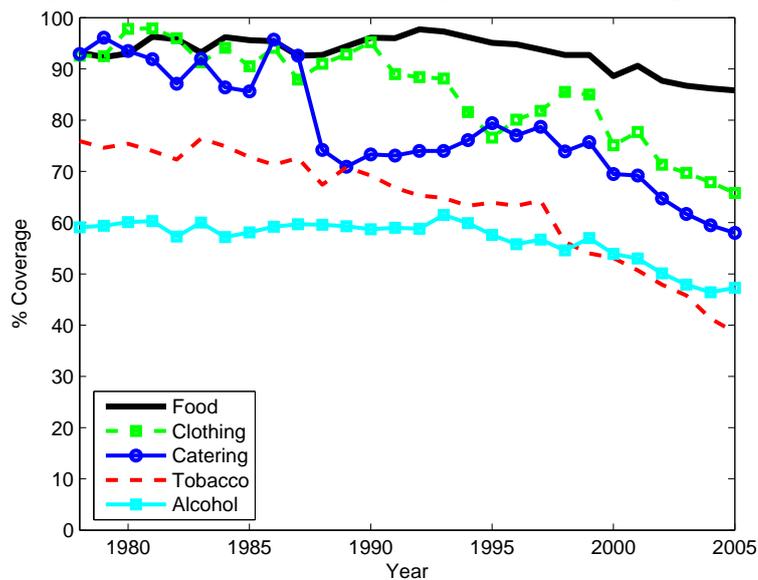


Figure 2.9: Consumption Coverage – Selected Categories



There are several possible explanations for the declining performance of the expenditure data. First, there may be a worsening sampling problem. As mentioned above,

the response rate to the FES has declined from over 70% to under 60% over the past 30 years. It is possible that the survey is systematically selecting out high spenders for some reason. However, the FES continues to cover income well, so the discrepancy would have to be caused by selecting out groups who spend more of their income relative to the rest of the population. We know that FES excludes students and people in residential housing, among others, but it seems unlikely that these two groups can explain all the difference. Second, the departure could be caused by changes in the way people spend money. The 1990s saw the introduction of internet purchasing and a rise in spending on credit cards. Additionally, children's expenditure has become more important: although their expenditure is accounted for, children are not given a diary, so their spending may be under-recorded. Third, spending abroad and spending by NPISH (non-profit institutions serving households e.g. local sports clubs) is not included in the FES. These items are separable from domestic and household spending in the national accounts, though not at the level of individual categories, and there is likely to be high measurement error in recording, for example, foreign spending by UK households. Finally, the decline coincides with the shift from sampling the FES over the calendar year to sampling over the financial year (e.g. from April 1993 to March 1994). However, it is hard to think why this would cause a departure in trend between the datasets, rather than maybe a shift in the coverage. Whatever the cause of this discrepancy, it is interesting to note that the US CEX also displays a more quickly deteriorating coverage for consumption than for income: the comparison of data collection methodology in the FES, the CEX and other consumer surveys seems a promising approach for uncovering the cause of the problem.

Figure 2.10: Labour Force Participation: FES vs NIPA

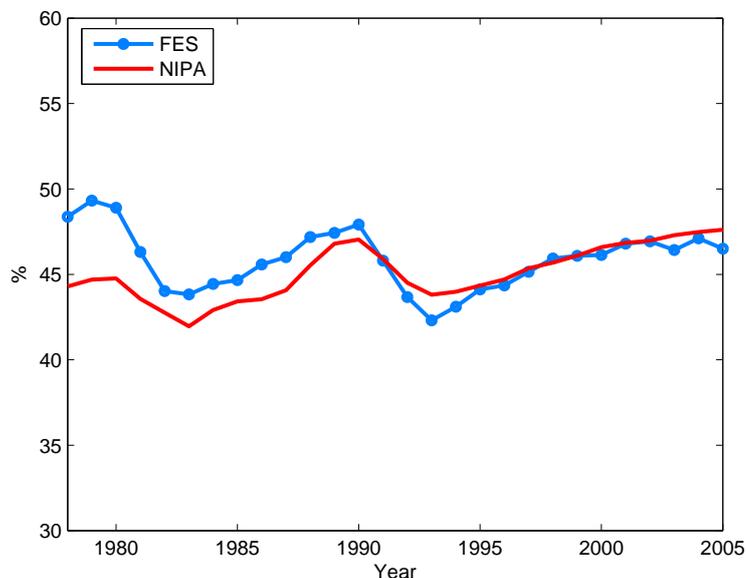


Figure 2.10 shows the employment rate for over-16s in the FES and NIPA data (which derive from the LFS). In contrast to income and expenditure, the match for participation between the FES and NIPA data has improved in the last decade. This is because the demographic weights are now calculated yearly for the FES, while prior to 2001, sampling weights are an interpolation from 10-yearly censuses.

To summarize, the FES seems strong in matching national account income, employment data and to an extent consumption data. However, the departure for expenditure is of growing importance. This raises some puzzles since it occurs for items (food) for which national accounts data uses FES. This is the subject of on-going research as there seems no easy explanation. The discrepancy has increased gradually since the early 1990s, for nearly all items, and it does not seem to have been caused by selecting out high-income households.

2.3 Hours, Wages and Earnings Inequality

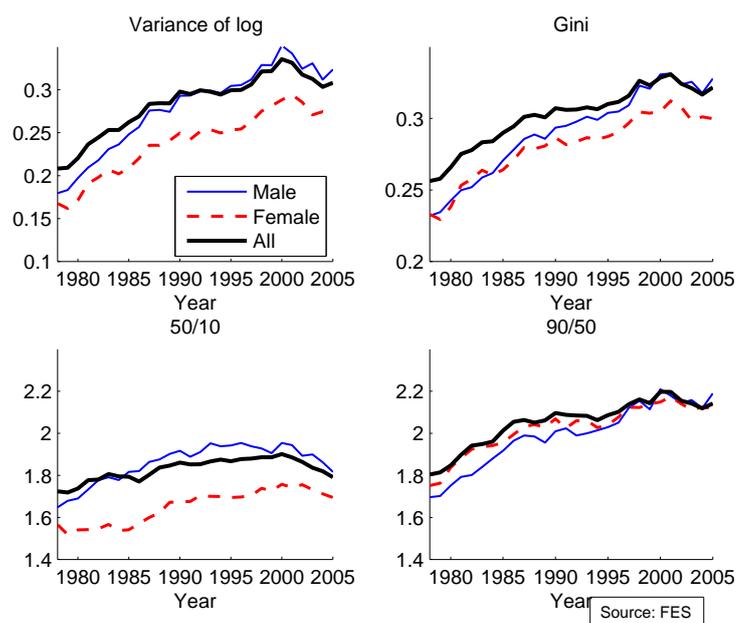
2.3.1 Wages

Our discussion of inequality turns first to the dispersion of wages and labour earnings. Figure 2.11 provides the key measures of inequality in overall hourly wages in the UK over the period 1978 to 2005. The strong growth during the 1980s is clearly visible. As is the

moderation in the early 1990s and the subsequent growth in the late 1990s. This figure is for the prime-age sample (aged 30-59), but the pattern is replicated for the entire sample (not shown).

This general picture of growth in wage inequality especially in the 1980s is reflected in both the variance of the log measure and the Gini measure. The quantile comparisons also show strong growth in inequality across the distribution in the early 1980s. However, the moderation in the early 1990s and subsequent increase in inequality are more marked in the upper-decile comparison (90-50) than in the lower decile comparison (50-10) and inter-quartile range (not shown). Many of the distinguishing features of the evolution of broad wage inequality since the 1980s have occurred primarily in the top quarter of the distribution.

Figure 2.11: Inequality in Wages (Prime-Age Sample)

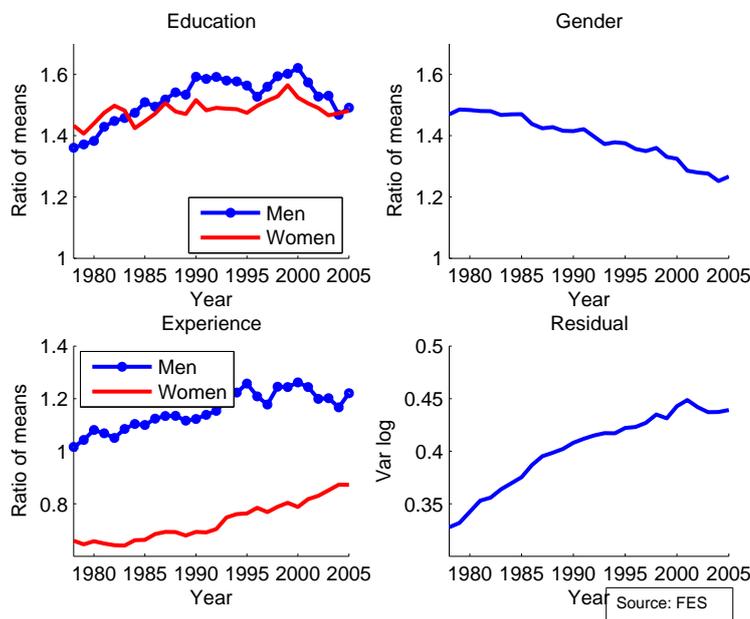


2.3.2 Wage Premia

Education differentials in the UK rose rapidly during the early 1980s and have been reasonably stable thereafter. This is clear from the first panel in Figure 2.12. The experience differential, which here simply measures the time since leaving education, also rose and continued to do so through until the mid-1990s. On the other hand the raw gender differential has fallen secularly over the whole period. The residual term shows that other

factors remain important in explaining the overall growth over this period.

Figure 2.12: Wage Premia



2.3.3 Wage Inequality, Earnings and Labour Supply

The growth of observed wage inequality over this period has been strongest for men, despite the fall in labour market attachment of the low skilled. In contrast, growth in wage inequality for women has been moderated by the fact that growth in the labour supply of women has been strongest for those with medium education levels (see section 2.2.1 above). Figure 2.13 also shows the systematic differences in the variation of hours worked between men and women over this period. This again largely reflects the relative increase in the labour supply of women. Generally male wages are weakly or even negatively correlated with hours of work, although this correlation has been becoming more positive over this period.

This correlation is further investigated in Figure 2.14 which shows that the correlation for men is mostly positive, and increasingly so, at either end of the life-cycle. This is where we expect labour supply elasticities for men to show most responsiveness.⁷⁸ For women

⁷See Blundell and MaCurdy (1999)

⁸For the US the correlation of wages and hours over the life cycle is documented in Kaplan (2010) and in Heathcote et al. (2008). Using PSID data, Kaplan estimates the profile to slope downwards from around -0.1 to -0.2 over the first 25 years of working life before flattening out. He fits a monotonically downward sloping profile with his parameter estimates. Using the same data, Heathcote et al. (2008) estimate the profile to be roughly flat at -0.1 and fit an upward sloping profile.

figure 2.13 shows a strong correlation between wages and hours.

The general picture of inequality growth in wages follows through into household earnings, as can be seen from Figure 2.15 which presents the inequality measures for equivalised household earnings. As with most other variables, the variance-of-log measure responds more to the lower end of the distribution, as reflected in the 50/10 ratio, whereas the Gini is closer to the 90/50 ratio. This feature is observed in other countries (see for example Heathcote et al.'s US study). While the path of inequality at the top end here closely follows the path for the upper half of the wage distribution in figure 2.11, the decrease in dispersion in the lower half is much greater than the corresponding drop in wage dispersion. It is likely this substantial decline is caused by the increase in labor-force attachment among low-skilled workers, as shown in figure 2.3.

Figure 2.13: Inequality in Labour Supply

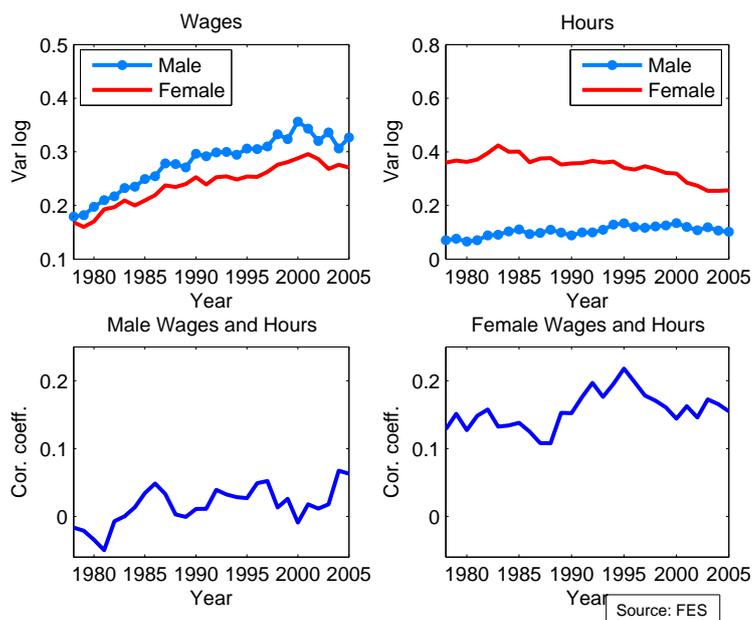


Figure 2.14: Correlation across the life-cycle: Correlation between Male Wages and Hours, by Age (Cubic Polynomial Fit)

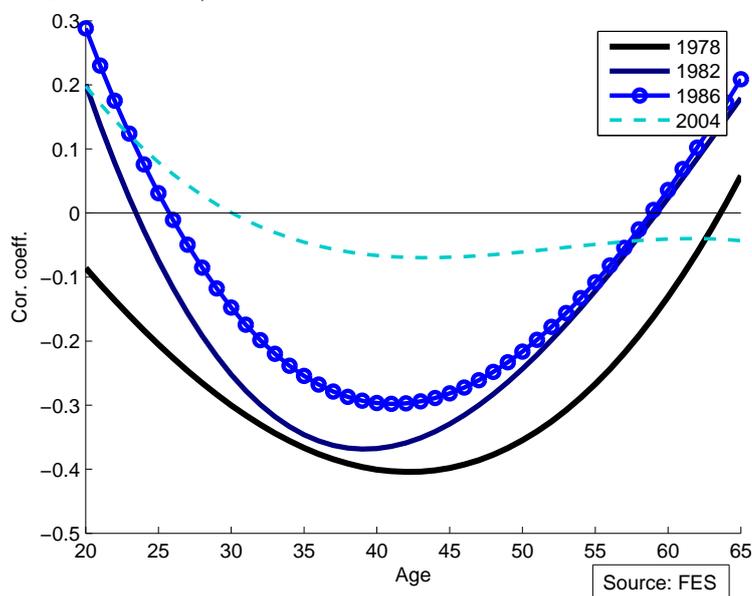
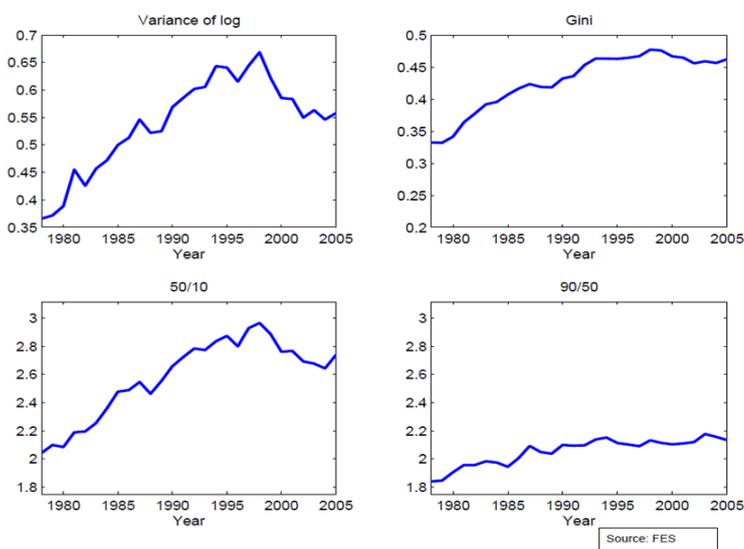


Figure 2.15: Inequality in Equivalized Household Earnings

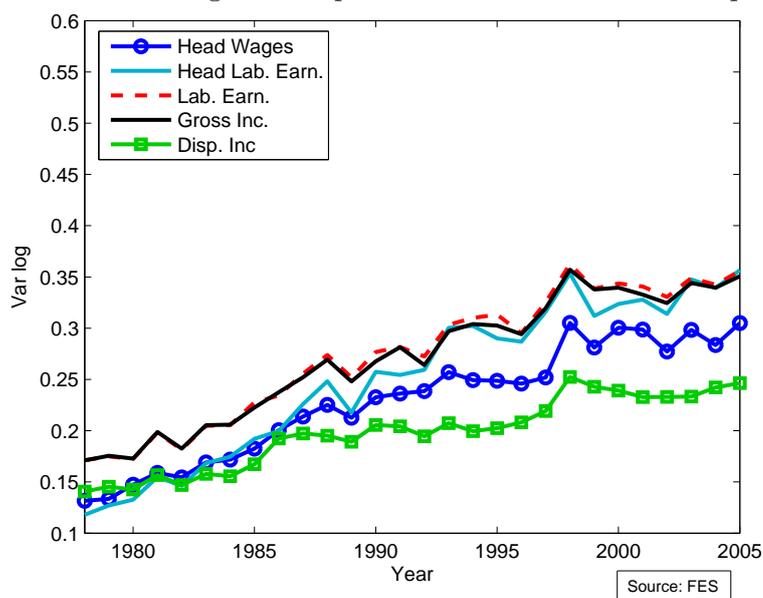


2.4 From Wages to Disposable Income

The linkages between individual hourly wages and family disposable income can be described as a set of ‘insurance’ mechanisms. These are actions that individuals, families and society take in reaction to changes in hourly wages. These insurance mechanisms

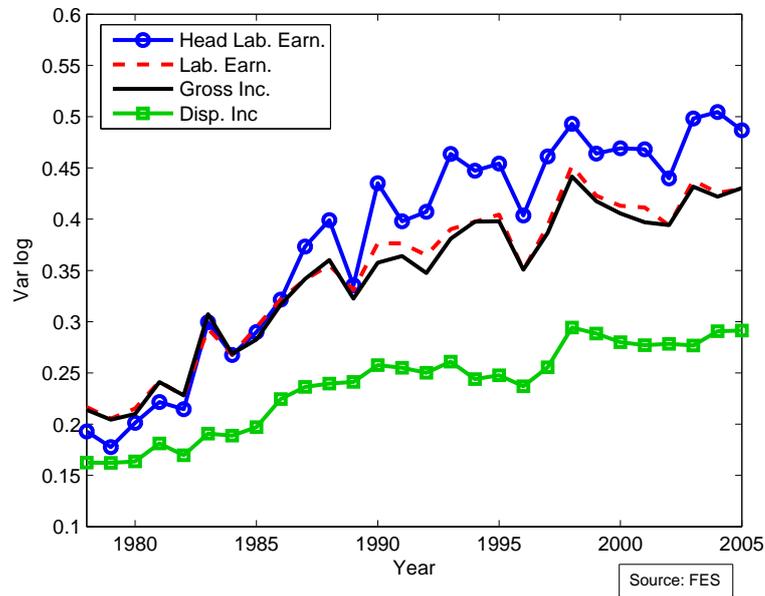
include regular savings and borrowing to smooth out shocks to income. They also include individual and family labour supply responses. They include the workings of the tax and welfare system. These mechanisms place a wedge between the distribution of individual hourly wages and the final distribution of disposable income. To bring these together we have to understand the relationship between income sources and consumption.

Figure 2.16: From Wages to Disposable Income – Heads in Employment



In figure 2.16 we show the overall pattern of the variance of log measure of inequality for the sample of households in which the head is in employment. The sharp rise in inequality for wages through to disposable income in the early 1980s is clearly evident. From 1990 onwards the growth in inequality of household earnings tends to separate from that of the head's wage, pointing to the importance of positive labour supply effects. Inequality in household earnings has grown more slowly than for head earnings, in part because the growth in female labour supply has been strongest amongst those with medium education levels. The slower growth in disposable income inequality highlights the role of taxes and transfers. Figure 2.17 shows the impact of including the self-employed. Here the divergence with disposable income is particularly strong.

Figure 2.17: From Earnings to Disposable Income – Heads with Earnings from Employment and Self-Employment



Inequality is generally much higher and grows more rapidly once we consider the entire sample of households. The impact of including households with no labour income is clear from figure 2.18.

Figure 2.18: Disposable Income Inequality: Decomposition by Sample

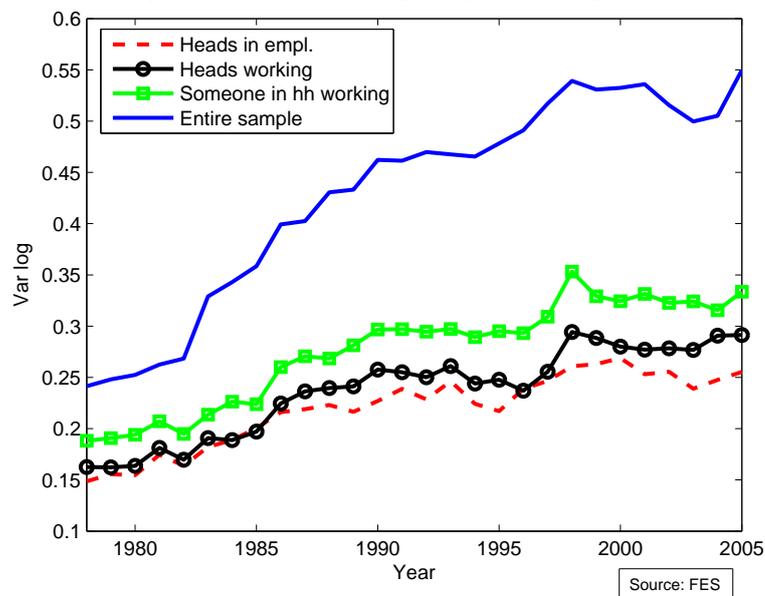
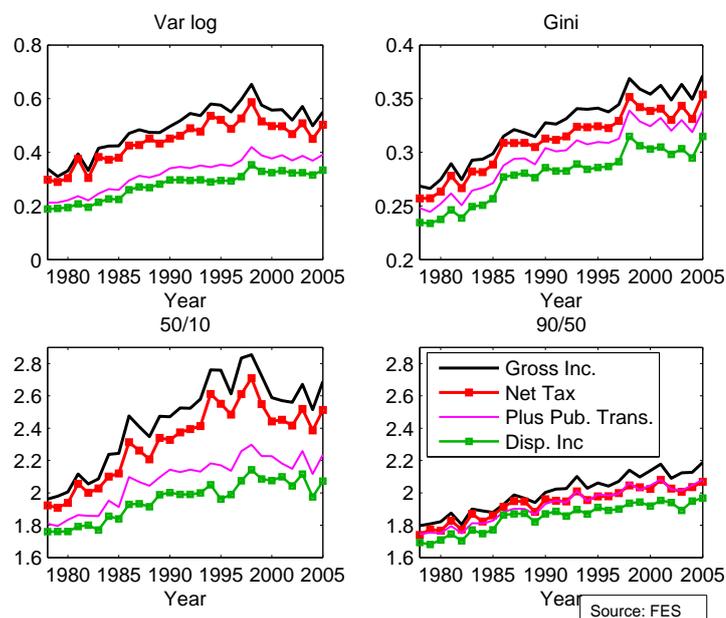


Figure 2.19: From Earnings to Disposable Income (At least one worker)



Not surprisingly perhaps the impact of taxes and transfers is greatest among the lower deciles. Figure 2.19 shows the key differences in the series for the 50-10 ratio.

The key importance of the relationship between the business cycle and inequality is documented in figures 2.20 and 2.21. In the years following each of the two significant recessions in the early 1980s and the early 1990s, inequality in gross income expands, driven largely by deep falls in the lower quantiles of the income distribution. The picture for the distribution of net income is very different. The tax and transfer system plays a key role in off-setting the impact of recessions on the lower quantiles of the income distribution.

Figure 2.20: Quantiles of the Household Gross Income Distribution

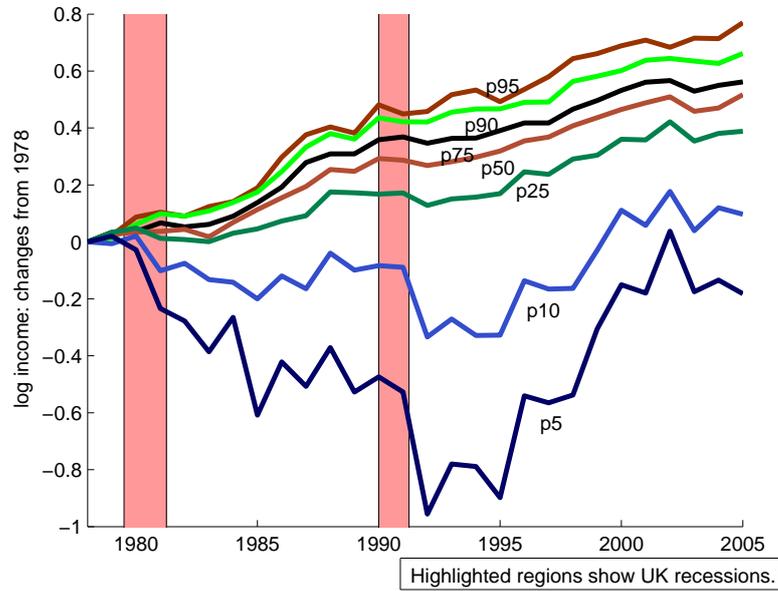
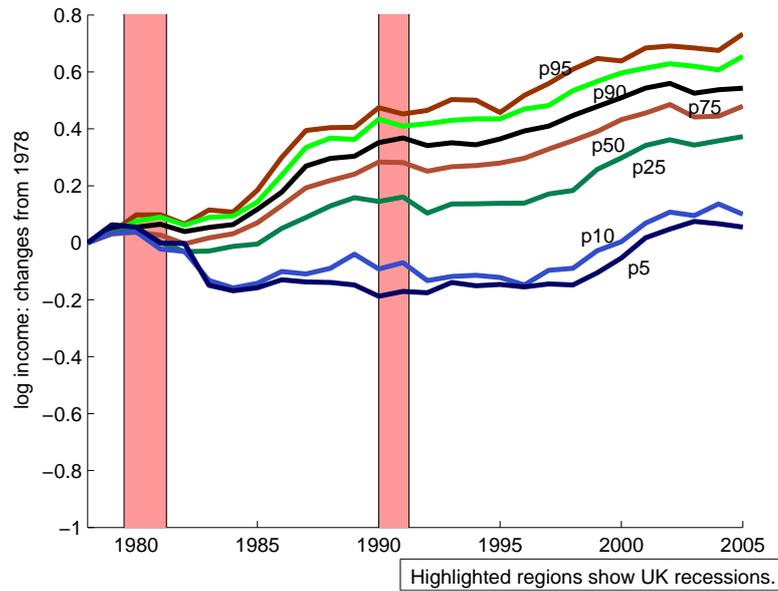


Figure 2.21: Quantiles of the Household Net Income Distribution



2.5 Consumption Inequality

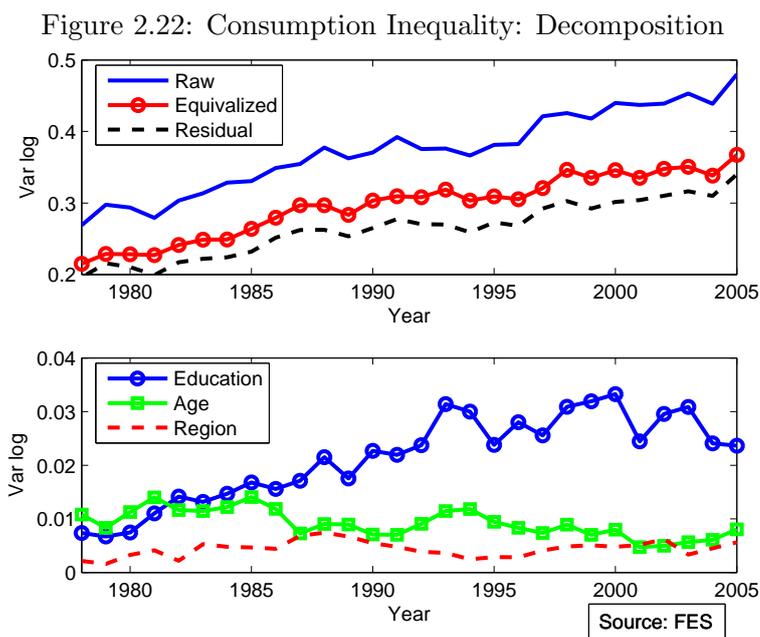
2.5.1 The Inequality Boom and After

Consumption inequality rose strongly in the UK in the early 1980s. This has been documented elsewhere, see Blundell and Preston (1998), but figure 2.22 also points to the episodic nature of consumption inequality growth since the late 1970s. Here we use the

variance of log measure as it decomposes easily.⁹ The systematic growth in consumption inequality gives way to a period of almost no inequality growth in the early 1990s and then an uptake of inequality growth in the late 1990s.

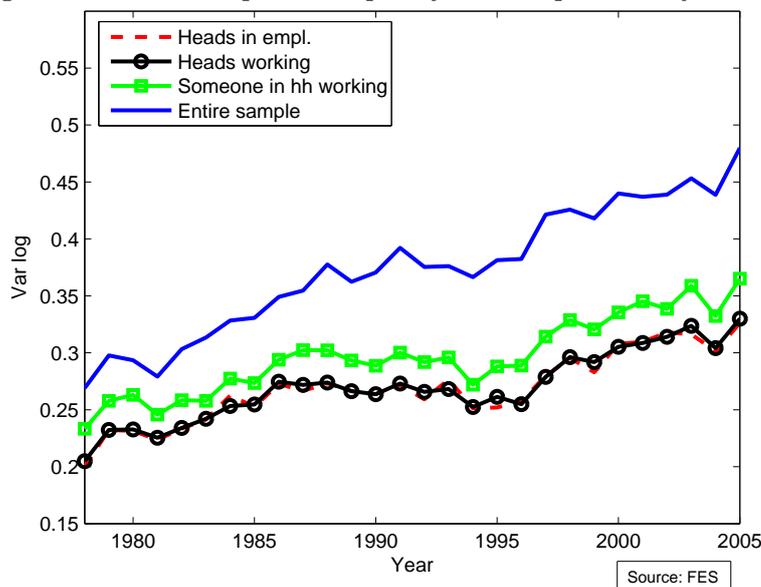
The bottom panel of figure 2.22 shows that the two episodes of inequality growth – the mid-1980s and late 1990s – show distinct patterns with regard to education. Specifically, the 1980s inequality boom followed the education pattern fairly closely but the growth in the late 1990s found no significant counterpart in the education component.

This underlying difference in the nature of the two inequality growth periods in the UK is further revealed in figure 2.23 which considers alternative samples. In the late 1980s and early 1990s there is stronger growth for the entire sample in comparison to the sample with heads working. For the more recent growth in consumption inequality there is very little difference across samples.



⁹It should be noted that log consumption is close to normally distributed, see Battistin et al. (2009).

Figure 2.23: Consumption Inequality: Decomposition by Sample



2.5.2 From Income to Consumption Inequality

The transmission from wages and income through to consumption is of considerable interest in understanding the workings of the economy at both the macro and micro levels. There is a growing literature which seeks to understand these transmission mechanisms, see for example Attanasio and Davis (1996), Blundell et al. (2008b), Guvenen (2007), Heathcote et al. (2008, 2010b), Krueger and Perri (2006). The disjuncture between consumption and income inequality in the UK, documented by Blundell and Preston (1998), is very clear from figure 2.24. At the beginning of the 1980s consumption inequality rose strongly and largely kept pace with the growth in income inequality. By the late 1980s the two series break apart. The two series grow furthest apart in the late 1980s and early 1990s. Income inequality, for all measures, rose strongly in the 1980s, with some further rise in the late 1990s. Consumption inequality, for all measures, rose quite strongly in the early 1980s and then again, although at a slower rate, in the 1990s. Figure 2.25 displays the full variance-covariance structure. This is used in Blundell et al. (2008a) to recover permanent and transitory variances over the 1978-2005 period in the UK for each of the 10 year birth cohorts. They find strong growth in permanent variances in early 1980s and some growth in early 1990s. Transitory variances increase strongly throughout the 1980s and into the 1990s. Birth cohorts show important life-cycle inequality growth, however

these are dominated by the strong growth in permanent shocks in early 1980s with some growth in 1990s, and the strong growth in transitory shocks in late 1980s with milder growth in 1990s. This lines up closely with the results for the US documented in Blundell et al. (2008b).

Figure 2.24: From Disposable Income to Consumption

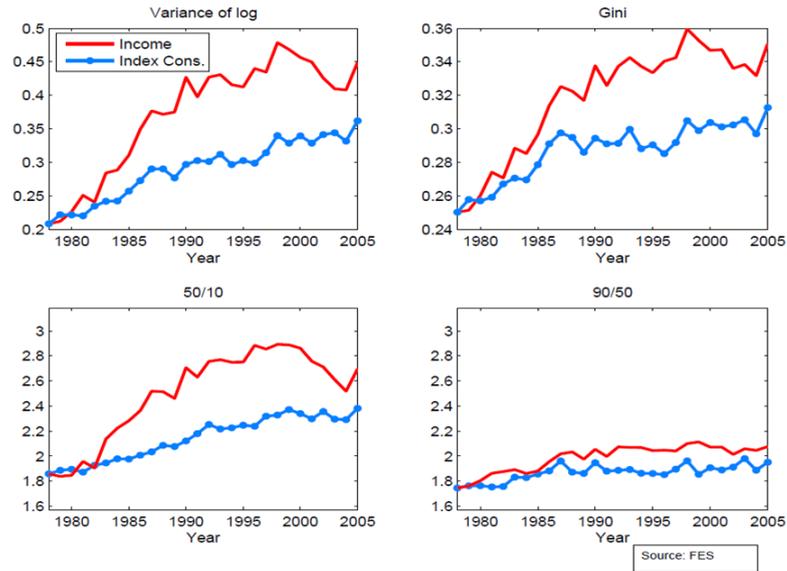
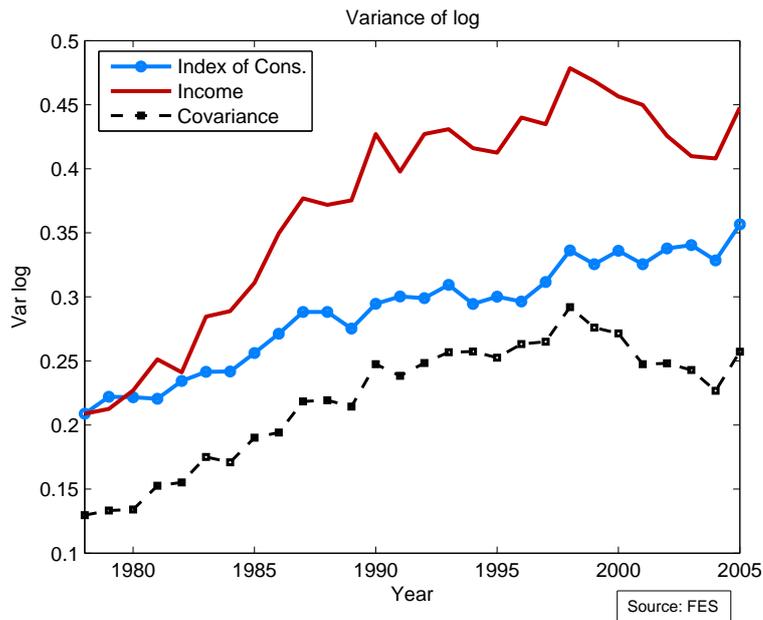


Figure 2.25: Covariance of Disposable Income and Consumption



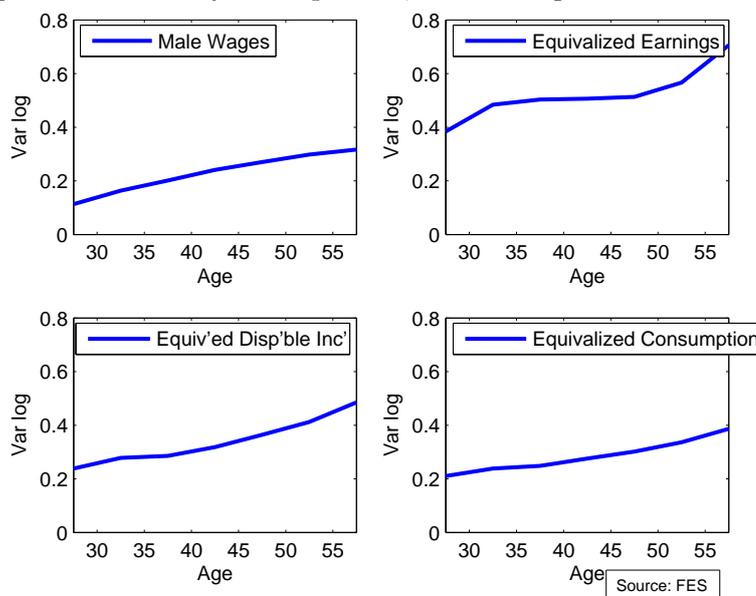
An interesting feature of figure 2.25 is the path of the covariance between income and

consumption. This moves in line with consumption until the mid-1990s. The covariance then begins to fall, suggesting the link between consumption and income is diminishing, but in a way that is consistent with a relative rise in consumption inequality. The strong growth in asset prices especially in the value of real estate which continued to the end of this sample period is one possible explanation. This would drive up expected life-time wealth relative to income and consequently drive up consumption among home owners. Given that home-ownership rates are around 70% in the UK, the inequality this would generate would lie in the 50-10 region, something confirmed in figure 2.24. This hypothesis is explored further in chapter 3.

2.5.3 The Life-Cycle Dimension

We might expect inequality in variables that are subject to permanent shocks to show increasing variance over time. As the analysis in Deaton and Paxson (1994) suggests this is particularly the case for inequality measures over the life-cycle. Figure 2.26 presents these measures over the lifetime, conditioning on cohort effects, for male wages, raw earnings, equalised earnings and equalised consumption.

Figure 2.26: Life-Cycle Dispersion, Controlling for Cohort Effects



One interesting feature of these profiles is that the variance of earnings increases

strongly after 45, while the life-cycle profiles of the variance of wages and consumption are roughly linear over the life cycle. Figure 2.14 above shows that the covariance of wages and hours increases strongly in late working life, implying that labour supply and possibly selection effects are important in explaining the strong increase in variance of earnings up to retirement. Consumption inequality rises consistently with age but at a slower rate than for disposable income. Differences in the rate of growth appear particularly strong at middle and later working ages. Suggesting that uncertainty about longer-run permanent differences in wages becomes less important for individuals in their 40s and early 50s. All profiles are consistent with a wage process driven by idiosyncratic permanent shocks that are at best partially insured and shorter-run fluctuations that are effectively smoothed out.

Figure 2.27 shows the life-cycle profiles conditioning on year effects. Other than male wages, these profiles all show a decreasing profile in mid working life. This highlights the difficulty in identifying time from age effects. To illustrate further, Figure 2.28 shows the variance of log equivalized consumption for four 10-year birth cohorts, first by year, then by age. Clearly in this time period, as each cohort enters working age, consumption dispersion roughly matches that for the previous cohort that is now in its mid 30s. When entering year dummies in a regression, therefore, the secular growth in consumption dispersion is interpreted largely as a time effect. However, we could equally interpret these profiles as steadily-increasing cohort growth in dispersion and a monotonic increasing age effect.

Figure 2.27: Life-Cycle Dispersion, Controlling for Year Effects

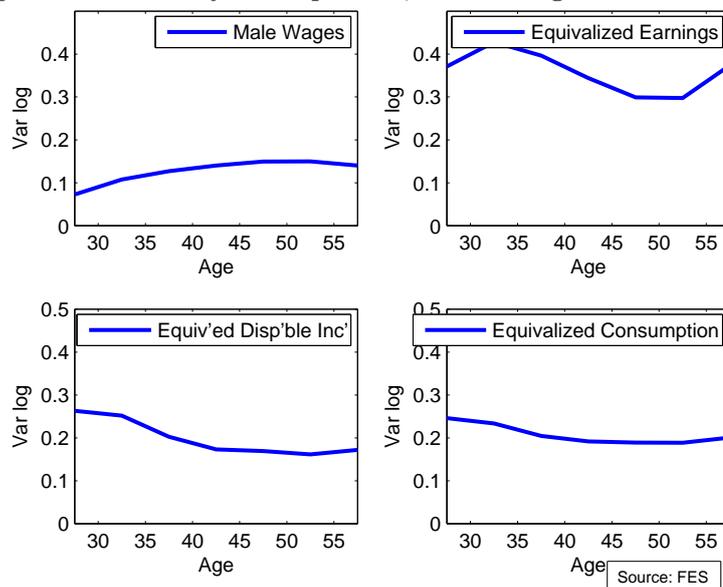
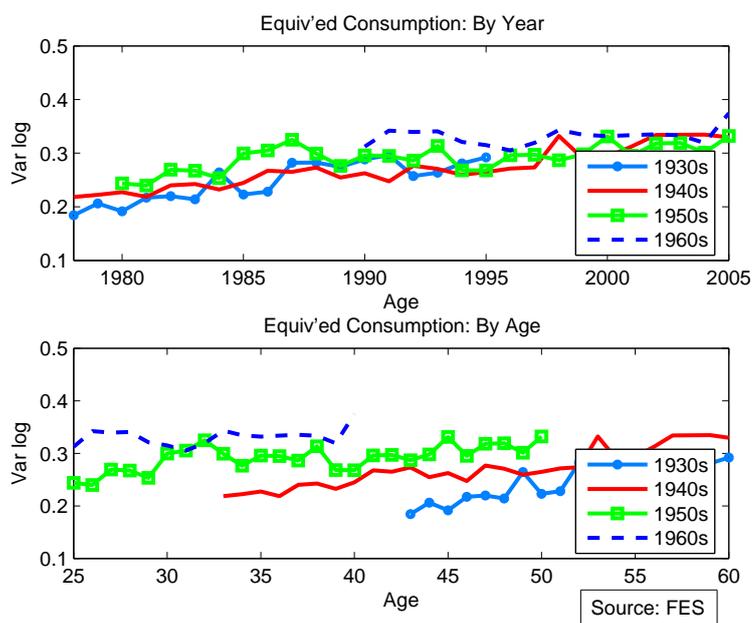


Figure 2.28: Life-Cycle Household Consumption Dispersion, by Cohort



2.6 Distributional Dynamics

In this section we further investigate the dynamics of the distribution of income. First we use panel data on income dynamics from the British Household Panel Data to decompose income into two factors – a persistent and a transitory component. We show that this simple decomposition works well to describe income dynamics in the UK provided

the variances of each component are allowed to be non-stationary and allowed to evolve nonparametrically over time. We then document the path of the variances of the transitory and permanent components over time. Turning first to the panel data dynamics we consider a model of the form:

$$\ln Y_{i,a,t} = Z'_{i,a,t} \lambda + \mu_i + y_{i,a,t}^P + y_{i,a,t}^T$$

The y^P term is the permanent component which follows a martingale process

$$y_{i,a,t}^P = y_{i,a-1,t-1}^P + \zeta_{i,a,t}$$

and y^T is a transitory or mean-reverting component

$$\nu_{i,a,t} = \sum_{j=0}^q \theta_j \epsilon_{i,a-j,t-j} \text{ and } \theta_0 = 1$$

This model implies a simple structure for the autocovariance structure of $\Delta y \equiv \ln Y - Z'\lambda$. In particular that higher order autocovariances in the growth of income should be zero, see Meghir and Pistaferri (2004) for example. This determines the order of the MA component for ν . We argue this model structure provides a good approximation to the UK income data. Alternative models with less persistence or with idiosyncratic trends as in Baker (1997) and Baker and Solon (2003), for example, imply higher-order non-zero autocovariances. The specification of income risk is investigated in more detail in chapter 4.

Unfortunately, the BHPS data has only been collected since 1991 and therefore misses the ‘inequality boom’ of the 1980s. In these results the sample definition is as close as possible to any similar FES statistics: all households (headed by couples or otherwise, but with heads between 25 and 60) are included. ‘Labour earnings sample’ refers to those households where we observe positive household gross labour income.

The results from estimating this model on BHPS data on the growth male hourly wages are provided in Tables 6.1. In this autocovariance analysis we have removed demographic, age and education effects. The autocovariance structure shows significant own and first-

Table 2.1: The Autocovariance Structure of Wage Growth for Male Head

Year	$\text{var}(\Delta y_t)$	$\text{cov}(\Delta y_t, \Delta y_{t+1})$	$\text{cov}(\Delta y_t, \Delta y_{t+2})$	$\text{cov}(\Delta y_t, \Delta y_{t+3})$
1992	0.0636 (0.0053)	-0.0150 (0.0020)	-0.0053 (0.0021)	-0.0037 (0.0022)
1993	0.0529 (0.0028)	-0.0135 (0.0021)	-0.0033 (0.0017)	-0.0011 (0.0015)
1994	0.0599 (0.0046)	-0.0121 (0.0019)	-0.0025 (0.0018)	-0.0016 (0.0016)
1995	0.0653 (0.0061)	-0.0120 (0.0022)	-0.0005 (0.0018)	0.0017 (0.0018)
1996	0.0511 (0.0032)	-0.0125 (0.0016)	0.0000 (0.0016)	-0.0003 (0.0014)
1997	0.0493 (0.0025)	-0.0101 (0.0016)	-0.0015 (0.0015)	0.0015 (0.0016)
1998	0.0515 (0.0024)	-0.0111 (0.0017)	-0.0002 (0.0017)	0.0029 (0.0018)
1999	0.0484 (0.0028)	-0.0107 (0.0020)	-0.0014 (0.0016)	-0.0004 (0.0016)
2000	0.0529 (0.0029)	-0.0185 (0.0021)	0.0005 (0.0015)	0.0002 (0.0017)
2001	0.0555 (0.0029)	-0.0139 (0.0017)	-0.0013 (0.0017)	0.0009 (0.0017)
2002	0.0511 (0.0027)	-0.0137 (0.0017)	0.0001 (0.0018)	- -
2003	0.0506 (0.0034)	-0.0147 (0.0018)	- -	- -
2004	0.0497 (0.0030)	- -	- -	- -

order terms which underlie the simple permanent-transitory model. The second-order terms suggest the possibility of the first-order MA for the transitory component but there is little evidence that further terms are required.

In figures 2.29 and 2.30 we plot the implied estimates of the permanent and transitory variances for household earnings and household disposable income. These show important permanent shocks which show some evidence of falling back in the late 1990s and then tailing off towards the end of the period.

Figure 2.29: Variance of Permanent and Transitory Shocks: Labour Earnings Sample

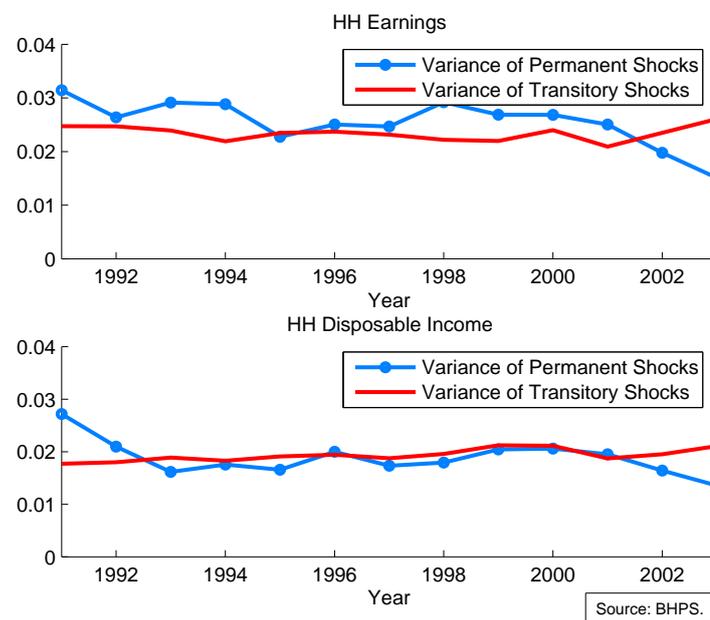
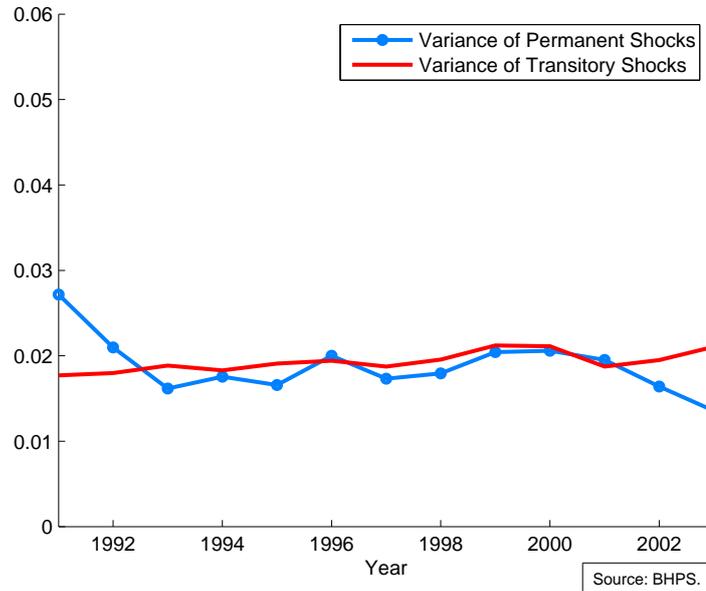


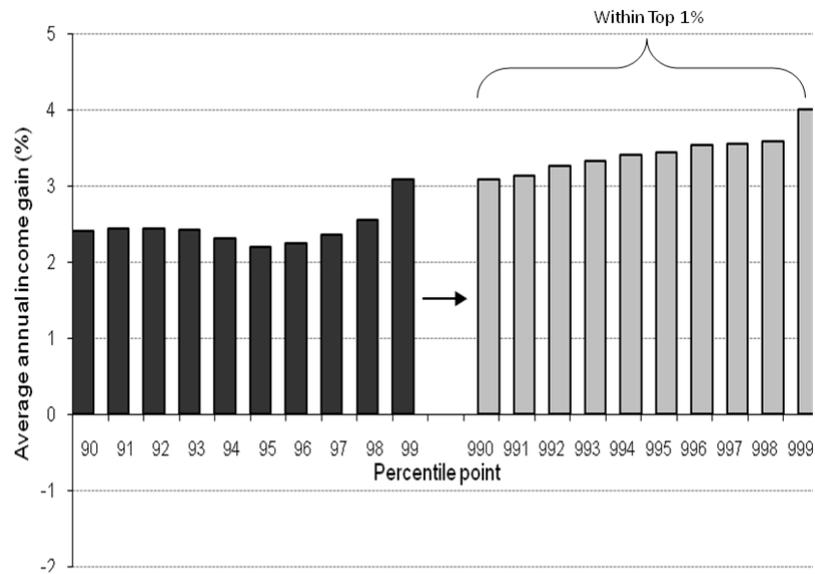
Figure 2.30: Variance of Permanent and Transitory Shocks: Household Disposable Income



2.7 Top Incomes: The New Inequality

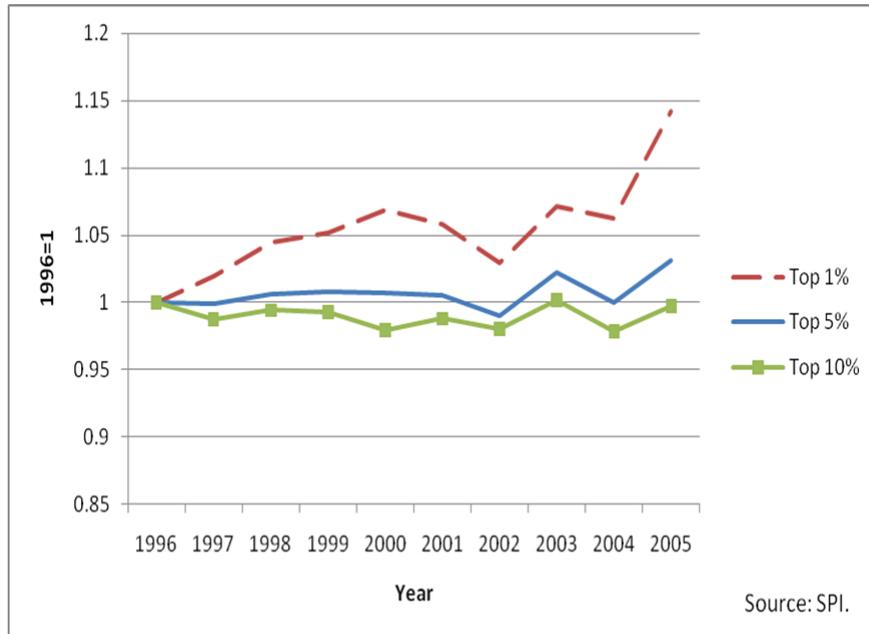
The late 1990s saw highest income growth at the very top of the distribution, and the emergence of a ‘new inequality’ dominated by a growth in employment related incomes, as employment income replaced investment income in the top 1%. This growth in inequality for top incomes is clearly illustrated in Figure 2.31 which uses tax return data to analyse the growth in the top 10 percentiles. The late 1990s sees a strong growth in the top percentiles. Breaking up the top percentile further we see the strongest growth in incomes at the very top of the distribution.

Figure 2.31: Real income growth for the richest 10% and 1% using the SPI, 1996-97 to 2004-05 (GB)



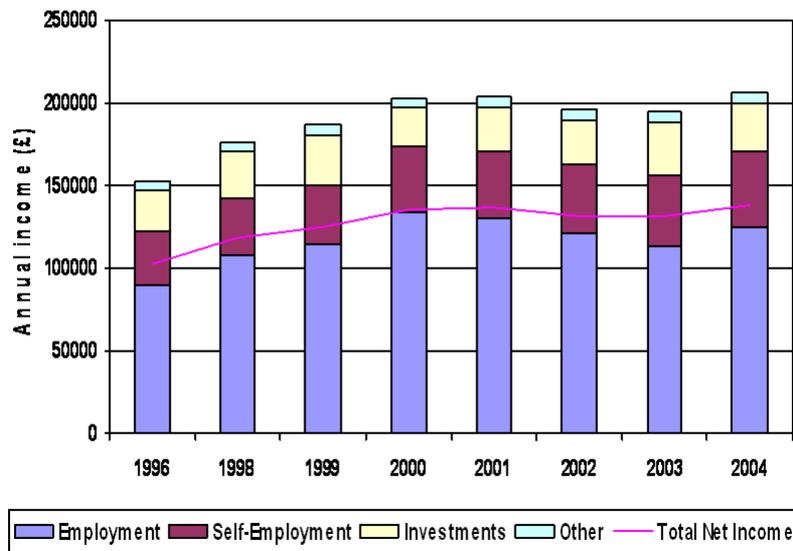
Notes: Incomes are net of income tax but do not include the deduction of council tax or national insurance. Incomes have not been equalised. Percentile incomes are measured as the income of the person on the border of the two percentiles. Source: Brewer et al. (2007b).

Figure 2.32: Change in Top Net Income Shares



Notes: The shares of net income out of total income in 1996 for these groups were: 38.0% for the top 10%: 25.0% for the top 5%, and 10.2% for the top 1%. Source: Brewer et al. (2007a)

Figure 2.33: Income Components for the Top 1%



Notes: Net incomes do not include the deduction of council tax or national insurance. Incomes have not been equalised. Source: Brewer et al. (2007b).

Figure 2.32 shows that the strength of the growth in the top percentile and the strong cyclical nature of these changes. Looking at income components (figure 2.33) we see the importance and cyclical nature of employment remuneration in the top 1% of incomes. The proportion of employment earnings in total gross income for this group grew from 52% in 1985 to a peak of 66% in 2000. It then declined to 58% in 2003 before rising again in 2004.

2.8 Interpretations and Conclusions

The UK has seen significant variation in inequality growth over the last three decades. Income inequality, for all measures, rose strongly in the 1980s, with some further rise in the late 1990s. Consumption inequality, for all measures, rose quite strongly in the early 1980s and then again, although at a slower rate, in the 1990s. The analysis of consumption and income inequality suggests strong growth in the variance of permanent shocks in the early 1980s and some further growth the 1990s. It also points to strong growth in transitory shocks in late 1980s and mild growth in 1990s. Birth cohorts have also shown important life-cycle inequality growth.

We have shown the inequality boom of the 1980s in the UK to be characterised by strong growth in permanent shocks to labour income followed by an increase in transitory volatility leading to a period of moderation. In the late 1990s inequality was dominated by a growth in employment related earnings at the top as employment income replaces investment income in the top 1%. Taxes and transfers have done much to offset losses at the lower end of the earned income distribution.

In this study we have made use of extensive micro-data sources in the UK on consumption, income, earnings, labour market participation, hours of work to study the evolution of the inequality in these series and the relationship between them. On a note of caution we point out that the time series patterns in the household level consumption data have become increasingly different to that documented in national accounts. A further analysis of these differences is warranted.

A2 Appendix to chapter 2

This appendix describes in detail the definitions of the data used from each source and how these data were transformed.

FES Income Data

Wages

The wage variable used is usual labour earnings plus any bonuses, divided by hours worked (see below). We keep only those in employment, omitting the self-employed.

Hours

Our hours variable is usual hours worked plus usual overtime. Again we omit the self-employed.

Earnings and Income

‘Labour earnings’ cover both the employed and self-employed. ‘Labour earnings plus private transfers’ includes regular allowances from outside the immediate family, allowances from a spouse, payment for odd jobs, child income and income from private annuity or trust. ‘Asset income’ is all income from investments minus income from real estate, which is then included in ‘asset income plus residential income’. ‘Gross income’ is the sum of these items. ‘Net disposable income’ consists of ‘gross income’ plus public transfers (social security benefits, state pension, luncheon vouchers, education grants and student top-ups) minus labour and payroll taxes.

BHPS

Income

Data Definitions in the BHPS are almost identical to those for the FES.

Education.

Qualifications are not given in the FES, so we define ‘compulsory education’ as those who left at compulsory leaving age (this has risen from 14 to 16 since WW2), ‘intermediate

education’ as those who attended school up to 18, and ‘high education’ as those who left school after 18. BHPS includes information on educational attainment. We therefore form the following categories: ‘high education’ includes those with an honours degree or equivalent; ‘intermediate education’ includes those with A-levels or equivalent (the equivalent of a US high school diploma), and ‘low education’ is the remainder.

Consumption

Consumption is expenditure on the following items: food, catering, alcohol, tobacco, fuel, household services, clothing, personal goods and services (toiletaries etc.) motoring expenses excluding vehicle purchases, travel expenses, leisure goods (books, music recordings) excluding audiovisual equipment, entertainment and holiday expenses. The main omissions are housing costs, furniture, furnishings and electrical appliances, motor vehicles and garden and audiovisual equipment. In short, our measure of consumption includes non-durable goods and services and excludes durable and semi-durable goods. ‘Consumption plus housing’ includes rent, mortgage interest payments and housing taxes. This is a user-cost measure of housing. The FES does not easily permit a calculation of imputed rents for homeowners as it does not include house prices.

Income and consumption in figures 2.7, 2.8 and 2.9 – comparison with national accounts.

Both income and expenditure data used for these figures differ from those used in the rest of the study. Income is total disposable income minus imputed owner-occupier rental income. Private pension contributions are included but employer pension contributions are excluded.

Expenditure is total household expenditure excluding public transport and housing. These two categories are omitted in order to provide the best fit between FES and national accounts definitions.

Chapter 3

Increasing Inequality and Improving Insurance: House Price Booms and the Welfare State in the UK

3.1 Introduction

As in much of the rest of the world, particularly in the USA, inequality grew strongly in the UK through the 1980s.¹ Figure 3.1 (top two lines) shows this evolution of cross-sectional inequality over 1978-2008, both for household incomes and for consumption.² ³ While a vast literature has looked at factors behind the increase in inequality on the income side,⁴ a parallel literature has focused on consumption dispersion as a measure of welfare (Cutler and Katz, 1992). Since Deaton and Paxson (1994), this literature has sought to explain the link between the two measures. For example, the profiles over the 1980s presented here have been interpreted by Blundell and Preston (1998). In their model, households cannot insure shocks to permanent income so must adjust their consumption,

¹see Krueger, Perri, Pistaferri, and Violante (2010) for a discussion of the cross country evidence and chapter 2 for greater detail on the UK, including documentation of wage, hours and earnings dispersion.

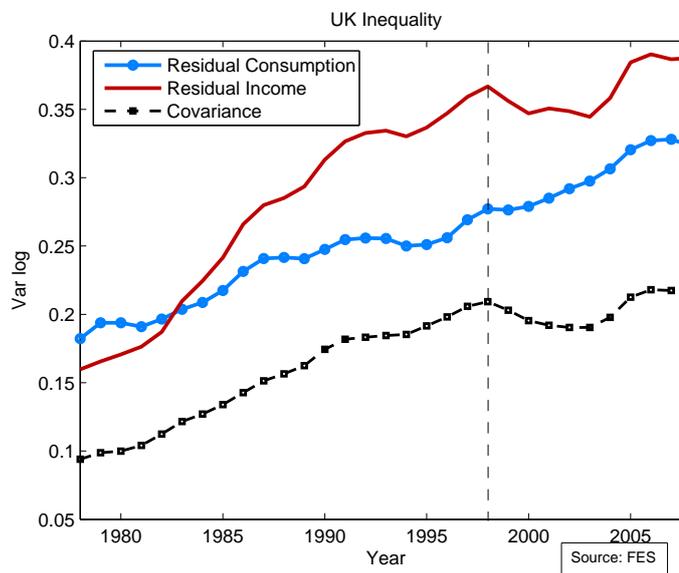
²This figure is a smoothed version of figure 2.25

³Unless otherwise stated, 'inequality' in this chapter refers to the variance of log measure. See chapter 2, for evidence of inequality changes on other measures.

⁴most notably, the effect of skill-biased technical change, see Acemoglu, 2002.

while households can insure transitory shocks through borrowing and saving.⁵ Blundell and Preston therefore argue that consumption inequality is only affected by permanent shocks. Over the 1980s, therefore, the growth in consumption inequality identifies a high variance of permanent shocks. Meanwhile the extra growth in income inequality identifies growth in the variance of transitory incomes. This interpretation of high permanent shocks and an increasing transitory component is corroborated by evidence from panel data on earnings by, for example, Moffitt and Gottschalk (2002) for the USA and Dickens (2000) for the UK.

Figure 3.1: Variance of Log Income and Consumption in the UK



Notes: Residuals are calculated after yearly regressions on: education, household size, region and a quartic in age.

On this analysis the profiles over 1990-1997 can be characterized simply: both permanent and transitory differences between households held steady until around 1995, then the growth in consumption inequality indicates a brief surge in permanent inequality over 1995-1997. However, the experience after 1997 presents two related puzzles: first, while consumption inequality increased, income inequality decreased. An obvious interpretation of this movement would be a continued increase in permanent inequality but a large decline in transitory shocks. This explanation can be discounted by looking at the evolution of the cross-section covariance of consumption and income: under a standard stochastic, life-cycle model of consumption, the covariance will only increase if permanent shocks are

⁵see also Deaton, 1992 for simulations.

present. Figure 3.1 shows that the covariance is declining over time, suggesting an absence of permanent shocks, in contradiction of the increasing cross-section variance of consumption. This is the first puzzle. The second puzzle is the very fact that the covariance is declining: the literature, going back to Deaton and Paxson (1994) has emphasized the role of idiosyncratic income risk, with shocks uncorrelated with the household's position in the income distribution. Under this model, the covariance of income and consumption necessarily increases monotonically over time.⁶⁷ This decline therefore constitutes the second leg of a puzzle.

I explain this episode by accounting for two important features of the UK economy over the period and introducing these into an otherwise standard consumption and savings model. First, the new Labour government, elected in 1997, increased the generosity of welfare benefits in a sequence of measures over 1998-2003 which compressed the distribution of income. Introducing stochastic changes to the benefit regime explains the simultaneous decline in the variance of income with the covariance of income and consumption and induces a smaller decline in the variance of consumption. Second, the UK experienced a strong boom in house prices over 1996-2007 in which real prices grew by 130% nationwide.⁸ Introducing house price shocks into the model induces a growth in the variance of consumption separate from the other moments and further explains its decoupling from the covariance with income. I term this decoupling the 'excess' growth in consumption inequality.

Using cross-sectional data on consumption and income from the Family Expenditure Survey (FES) and panel data on incomes, food consumption and assets from the British Household Panel Survey (BHPS), I study two groups: the high and low education groups for the cohort born in the 1950s. The observed excess growth in the variance of consumption is around 0.04 log points for both groups over 1996-2004. The model explains around 35% of the observed excess growth for both groups. For the high group, the house price boom was the more important factor, explaining around 30% of the observed increase,

⁶Strictly speaking, Deaton and Paxson (1994)'s analysis concerns fixed groups of households, while figure 3.1 inequality profiles for the revolving set of working-aged households. The analysis in the rest of the chapter proceeds with fixed cohorts. The same patterns are evident for these groups.

⁷In some models, e.g. that in Blundell, Low, and Preston (2008a) the covariance need not grow monotonically. Nevertheless, no model in the literature can satisfactorily explain the magnitude of the drop concurrent with the rise in consumption inequality.

⁸House price data from the Office for the Deputy Prime Minister, deflated by the RPI.

while for the low education group, the benefit reforms were relatively more important: around half of the estimated contribution comes from the reforms, half from the house price boom.

A brief intuition for why the house price boom increased consumption inequality is as follows: the elasticity of consumption with respect to housing wealth is approximately given by the share of housing wealth in discounted life-time wealth (including human capital wealth). This wealth share covaries positively with the consumption distribution because those who receive good transitory income shocks accumulate both higher wealth and have higher consumption. Therefore, positive house price shocks exacerbate consumption inequality. A brief intuition for why the benefit reforms compressed consumption inequality less than the covariance of income and consumption is as follows: shocks to permanent income transmit less than one-to-one into consumption changes due to partial insurance (Blundell, Pistaferri, and Preston, 2008b). A compression of permanent income inequality reduces the covariance with consumption proportionally to this transmission factor and reduces the consumption variance proportionally to the square of the transmission factor. Therefore consumption inequality declines by less.

The estimated effects also imply that around 54% of the population benefitted overall from the insurance provided by the reforms, even though only the bottom 6.5% directly benefitted from income subsidies and even though I assume the rest of the population had to pay for the subsidies through a proportional increase in taxation.⁹ In a model without labour supply, the greatest welfare gain is obtained by redistributing income completely. I estimate that the given reforms provided around 3.5% of the welfare gains from complete redistribution.

The analysis proceeds using an extended numerically-solved consumption and saving model, estimated simultaneously on these two groups. I fit this model to the variance of log income, and to mean consumption growth and mean housing and other wealth holdings. I then use the fitted model to predict the moments of interest: the evolution of the variance of log consumption and the covariance of log income and consumption.

⁹Labour supply is exogenous in this model, so the welfare estimates do not include changes to the deadweight loss from labour market distortions. The welfare estimates presented here represent purely the effect of shifts in income and income risk, and should be treated as an upper bound on the true welfare gains.

The model captures the overall profiles of the moments well in addition to explaining the key features of interest. The model abstracts from a formal treatment of home-ownership because the computational burden would preclude a treatment of benefit reforms and the structural estimation, but I argue in section 3.4 that the same mechanisms which drive the distribution of housing wealth in a typical model (such as that presented in Campbell and Cocco (2007)) are present here. In fact, the model captures the distribution of (housing) wealth in key dimensions remarkably well.

I augment this analysis by using approximations to simplified versions of the consumption and savings model. These allow for a more intuitive and closer inspection of the mechanisms driving inequality. In the case of the house price boom, the approximations also allow me to derive sufficient statistics for the effect of house price shocks on consumption inequality in terms of the distributions of income, consumption and assets. These statistics capture the effects of heterogeneity in home-ownership as well as heterogeneity in housing wealth leverage. These estimates imply an effect on consumption inequality of around 0.025 log points over 1997-2004, around 60% of the observed excess growth.

This chapter relates to several literatures. First, after the seminal articles by Deaton and Paxson (1994) and Blundell and Preston (1998) a literature has developed on the effect of risk on various measures of economic inequality, the temperance of this risk by insurance channels and the estimation of unobservable risk and insurance by the evolution of inequality. Recent papers include, for example, Storesletten, Telmer, and Yaron (2004), who examine the insurance value of social security in the US; Krueger and Perri (2006), who examine the role of risk-sharing with limited commitment; Guvenen (2007), who questions the standard permanent-transitory model of the income process driving consumption and Heathcote, Storesletten, and Violante (2010b), who examine the welfare effects in the US of the observed changes in income risk, skill-biased technical change and the decline in the gender wage premium. None of these papers accounts for changes to the tax and benefit regime, nor the effect of asset price shocks. And none of these models can account for the observed profiles.

Second, this chapter adds to those on the effects of (changes to) the tax and benefit system on household consumption and on the income distribution. Krueger, Perri, Pistaferri, and Violante (2010) discuss the evolution of pre-tax and post-tax income across time

and across a range of countries. Johnson and Webb (1993) look in detail at the role of tax and benefit changes to UK income inequality in the 1980s. Many papers look at the consumption smoothing benefits of government insurance programmes, such as Gruber (1997) or Low, Meghir, and Pistaferri (2010). A related literature concerns the effect of changes to government insurance programmes on savings rates. For example Gruber and Yelowitz (1999) examine the effect of increased coverage of medicaid in the US. Sefton, Van De Ven, and Weale (2008) examine the effect on savings rates of changes in state pension provision in the UK.

This chapter also relates to the literature on the wealth effect of house price shocks on consumption. Li and Yao (2007) calibrate a structural model of home-ownership over the life-cycle and emphasize different wealth effects from the house price boom across the life-cycle. I focus instead on differential wealth effects within cohorts. There is a large and continuing empirical literature directly estimating the wealth effect which contextualizes my results, for example Attanasio, Blow, Hamilton, and Leicester (2009) and Campbell and Cocco (2007) who both also use the FES over a similar period.

This chapter proceeds as follows: section 3.2 outlines and discusses the model; section 3.3 discusses treatment of the data, drawn principally from the FES and BHPS; section 3.4 presents the results and provides detailed intuition for what in the model drives these results; section 3.5 discusses the estimation procedure and the choice of moments; section 3.6 briefly evaluates and dismisses alternative hypotheses for the observed phenomena; section 3.7 concludes.

3.2 Model

3.2.1 Overview

To model the effect of house price changes and tax changes on inequality I use a standard consumption and saving model with an exogenous income process and include two important features: asset price risk to mimic the effect of house prices on wealth; and a tax and benefit regime state, changing with government policy. Thus there are three main sources of risk in the model: income risk, divided, as usual, into permanent and transitory components; asset price risk, and ‘benefit regime’ risk. Besides government benefits, the

model is one of self-insurance in that there are no contingent claims markets, in line with the results in Attanasio and Davis (1996). The model is non-stationary and is partial equilibrium: wages and asset returns are taken as given.

3.2.2 The Household's Life-Cycle Programme

Consider the problem of a household which faces both uncertain labour income and asset returns, and chooses a sequence of consumption plans to maximize expected lifetime utility subject to constraints. The household belongs to a cohort, c , indicating the year of birth, and to a group e , denoting either high or low educational attainment. The household dies with certainty in the year $c + T$. The value to the household i at date t , of age $a = t - c$, with assets A_{it} , productivity P_{it} , which faces current government policy S_t is given by:

$$V_{c,e,t}(A_{it}, P_{it}, S_t) = \max_{\{C_{ik}(A_{ik}, P_{ik}, S_k)\}_{k=t}^{c+T}} \mathbb{E}_t \left(\sum_{k=t}^{c+T} \beta^{k-t} v(Z_{c,e,k}) \ln(C_{ik}) \right) \quad (3.1)$$

subject to the evolution of assets:

$$A_{it+1} = \begin{cases} R_{t+1}^*(A_{it} - C_{it}) + (1 - \tau_{c,t+1}[S_{t+1}]) \cdot Y[Z_{c,e,t+1}, S_{t+1}, \tilde{Y}_{it+1}] & \text{if } a < T_w \\ R_{t+1}^*(A_{it} - C_{it}) + 0.4P_{i,c+T_w} & \text{if } a \geq T_w \end{cases} \quad (3.2)$$

$$A_{it} \geq 0 \quad (3.3)$$

$$R_t^* = s_{c,e}^H R_t^H + (1 - s_{c,e}^H) R_t^O \quad (3.4)$$

$$\ln \left(\begin{bmatrix} R_t^H \\ R_t^O \end{bmatrix} \right) \sim N \left(\begin{bmatrix} \mu_H \\ \mu_O \end{bmatrix}, \begin{bmatrix} \sigma_H^2 & \rho_{HO} \\ \rho_{HO} & \sigma_O^2 \end{bmatrix} \right) \quad (3.5)$$

a permanent-transitory process for 'latent' income:

$$\ln \tilde{Y}_{it} = g_{c,e,Z,t} + \ln P_{it} + \epsilon_{it}$$

$$\ln P_{it+1} = \ln P_{it} + \eta_{it+1} \quad (3.6)$$

$$\eta_{it} \sim N(0, \sigma_{\eta,c,e}^2) \quad , \quad \ln P_{i,c} \sim N(0, \sigma_{\alpha,c,e}^2) \quad , \quad \epsilon_{it} \sim N(0, \sigma_{\epsilon,c,e}^2)$$

a specification for government benefits and process for their reform:

$$Y \left[Z_{c,e,t}, S_t, \tilde{Y}_{it} \right] = \max \left(\underline{Y} [S_t] \cdot Z_{c,e,t}, Y_{it} \right) \quad (3.7)$$

$$\Pr(S_{t+1} = x_k | S_t = x_j) = \Pi_{jk} \quad (3.8)$$

and budget balancing of the benefit reforms within a cohort and a time period (suppressing some of the functional dependencies for ease of expression). For each c and t , $\tau_{c,t}$ solves:

$$\underbrace{\tau_{c,t} \sum_{e=1}^2 w_{c,e} \int Y_{c,e,t} dF_{Y_{c,e,t}} [Y_{c,e,t}]}_{\text{Revenue}} = \underbrace{\sum_{e=1}^2 w_{c,e} \int_0^{\underline{Y}[S_t] \cdot Z_{c,e,t}} (\underline{Y} [S_t] \cdot Z_{c,e,t} - \tilde{Y}_{c,e,t}) dF_{\tilde{Y}_{c,e,t}} [\tilde{Y}_{c,e,t}]}_{\text{Gross Income Subsidies}} \quad (3.9)$$

Going into equation 3.1 in more detail: β is the discount factor; E_t the expectations operator conditional on information available in period t (a period being a year); $Z_{c,e,t}$ is a demographic taste shifting parameter, common across individuals of the same cohort, but conditional on education, and assumed to evolve deterministically. $v(\cdot)$ is the modified-OECD equivalence scale, implying that households optimize by equating the discounted marginal utility of equalized consumption across periods. I emphasize the treatment of demographics because these are important to the effect of government redistribution. Individuals live for T periods, working T_w years (from age T_{init} to 65), and face an exogenous mandatory spell of retirement of $T_R = 17$ years at the end of life. I solve the household's problem back to their 26th 'birthday' in 1981, giving 40 years of working life. The date of death is known with certainty and there is no bequest motive.

Moving to the intertemporal budget constraint given by equation 3.2: R_{t+1}^* is a stochastic rate of return on the portfolio; S_t is an indicator of the current state of the tax regime $\{0, 1, 2, 3, 4\}$ corresponding to an initial state in which there is no minimum income support ($S_t = 0$) and 4 states of increasing generosity of income support. During working life, households also receive after-tax labour income: \tilde{Y}_{it} is income before transfers and $Y[\cdot]$ is income after (gross) subsidies. $\tau_{c,t}[S_t]$ is the tax used to pay for the income subsidies and is common across education groups. After retirement, households also get access to a simple pay-as-you-go state pension system. This pays 40% of the household's

final working productivity level each year in retirement.

As given by equation 3.4, the rate of return on the portfolio R_t^* is composed of R_t^H and R_t^O , the interest rates on the two assets (housing and other). $s_{c,e}^H$ is the share of wealth in housing and is common across households of the same cohort and education. This modelling choice is discussed further in subsection 3.2.4. Asset returns, R_{t+1}^H and R_{t+1}^O , are joint i.i.d log normal, so that log asset prices follow a random walk with trend.¹⁰

I set a borrowing constraint that $A_{it} \geq 0$. This has the effect that agents cannot borrow against pensions and cannot borrow against possible minimum income subsidies if they have low productivity. This also implies that agents can have neither negative housing wealth nor negative total current wealth, although other financial wealth can be negative as households finance their home ‘ownership’ through mortgage debt.¹¹

Latent income evolves according to a standard permanent-transitory process as in equations 3.6, such that $g_{c,e,Z,t}$ is the deterministic, forecastable component of income, common across households of the same cohort, age, education, and household size. $\ln P_{it}$ is the stochastic permanent process and ϵ_{it} is the transitory process. I have the usual interpretation that permanent shocks represent long-term productivity changes such as promotions or change in health status within the household, transitory shocks represent bonuses, temporary lay-offs or other short-term changes in hours of work.

The benefit system is described in more detail in subsection 3.2.3. In brief, I model this as a minimum floor to equivalized income ($\underline{Y}[S_t]$ in equation 3.7). The income floor levels, $\underline{Y}[S]$, $S \in \{0, 1, 2, 3, 4\}$ are given exogenously. This income floor evolves according to the Markov process given in equation 3.8.

$w_{c,e}$ in equation 3.9 is the share of the cohort in each education group. A final word on budget balancing: the pension system is not funded by explicit taxation in this model. The concept of latent income in this model is labour income after background government taxes and transfers, such as existing and stable income taxes. I assume the pension system is funded out of this background taxation.

¹⁰To ensure that household wealth doesn’t become negative because of, say, a negative shock to housing wealth I assume limited liability: households must end each period with positive wealth, but can write off negative wealth if necessary at the beginning of the next period. In practice this possibility has no effect on the computation of the model solution, because average leverage and the variance of asset shocks are too small for this ever to occur.

¹¹In the empirical application, the share of wealth in housing, s^H , > 1 , so mortgage debt is higher than other savings and so $A^O \equiv s^O A$ is always negative.

3.2.3 Modelling the UK Benefit System

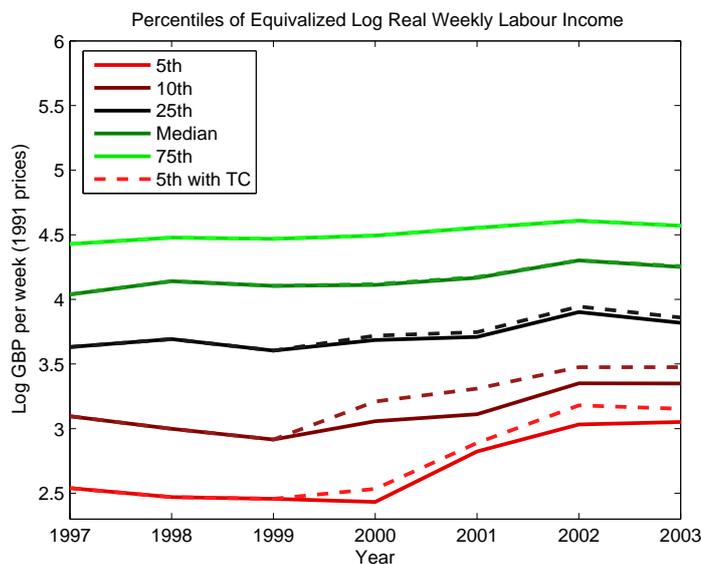
I model the government benefit system as a floor to equalized household income. This modelling decision is specifically designed to capture the effect of changes to the welfare and benefit system introduced by the Labour government after 1997. In its first parliament over 1997-2001 and shortly afterwards, the government introduced a raft of new measures aimed at supporting incomes at the bottom end of the earnings distribution. These can be roughly divided into: active labour market policies, such as the New Deal for Young People and the New Deal for Lone Parents¹²; in-work credits, such as the working family tax credit and child tax credit, and the minimum wage. Henceforth I refer to the combined reforms as ‘benefit’ reforms or the ‘benefit system’. Brewer (2007) provides a comprehensive survey of the details and efficacy of these measures until the mid-2000. He emphasizes that tax credits in particular were focused on families with children. Therefore it seems sensible to model the reforms as applicable to incomes after equalization.¹³ In reality, of course, these policies have a wider impact up the income distribution and receipts are contingent on many more variables than income and household size. However a minimum income floor is in the spirit of the reforms¹⁴ and allowing for greater heterogeneity in effects involves the use of more state variables for little gain. Figure 3.2 shows the direct impact of tax credits on incomes at the bottom end of the distribution. The FES data allow a separation of income before and after the receipt of tax credits. We can see that the effect of tax credits is almost negligible at the 25th centile and above. Moreover the boost to incomes from other measures, which we cannot directly observe, seems large at the 5th centile and also negligible above the 25th centile.

¹²See e.g. DeGiorgi for an analysis of the effect of these policies

¹³In contrast, at the time of writing, the new Conservative government proposes to cap benefits irrespective of total numbers of children. This reform is targetted at absolute rather than equalized income.

¹⁴Dickens and Manning (2004) look at the minimum wage and conclude that there were no effects on unemployment, nor spillovers up the distribution, and although the numbers affected were small, the earnings gains were large for those affected.

Figure 3.2: Effect of Tax Credits on Income: Cohort Born in the 1950s with Low Education



Note: Data do not let us separate tax credits in income after 2003.

I model changes to the minimum income floor as a first-order markov chain with 5 states. The government moves through each level of generosity as a new shock. In the baseline estimation households place high subjective probability on the status quo, i.e. the transition matrix is close to the identity matrix, although all entries are non-zero. Table 3.1 gives some details on announcements and introductions of the minimum wage and tax credits. Over the period 1999 to 2004 the increases to the minimum wage were larger than average earnings growth. Over 1999-2003, tax credits increased in generosity incrementally. I therefore argue that each year brought a new revelation to the generosity of benefit package and that my modelling choice is appropriate. However, I examine the possibility that agents foresaw these changes by extending the estimation to allow different transition probabilities. I discuss the results in section 3.5.¹⁵

The income subsidies are funded by a proportional tax on all income to give a balanced budget within a cohort and within a period. The tax required is small, reaching around

¹⁵A related issue is that the reforms were announced generally a year to 18 months before they were introduced. Blundell, Francesconi, and Van Der Klaauw (2010) look at the anticipation effect of the introduction of the Working Family Tax Credit on participation and hours, and find that hours of work increased before the introduction, confirming forward looking behaviour, and indicating the presence of adjustment costs in labour supply decisions. This would act to decrease income inequality before the formal introduction of the measures. Similarly consumption should have increased at the bottom end on announcement, decreasing consumption inequality before introduction. I ignore such anticipation effects in this analysis, but note that anticipations effects on both income and consumption inequality are in the direction of the reform (i.e. generous reforms decrease inequality), so affect the analysis little.

Table 3.1: Announcements on the Minimum Wage and Tax Credits

Minimum Wage			
Date	Minimum Wage Rate	Growth Rate	Date Increase Announced /Recommended
Apr-99	£3.60		Summer 1998
Oct-00	£3.70	2.8%	"
Oct-01	£4.10	10.8%	Spring 2001
Oct-02	£4.20	2.4%	"
Oct-03	£4.50	7.1%	Spring 2003
Oct-04	£4.85	7.8%	"
Oct-05	£5.05	4.1%	Spring 2005
Oct-06	£5.35	5.9%	"

Tax Credits		
Date	Headline Change	Forecast Revenue Change in Fiscal Year + 2 Years (£m)
Jul-97	No changes announced	0
Mar-98	Working Families Tax Credit (WFTC, from 1999)	2570
Mar-99	Children's Tax Credit (CTC, from 2001)	2955
Mar-00	Increase in WFTC (staggered over 2 years)	1425
Mar-01	Combination of small changes	1140
Mar-02	Increase in CTC and Working Tax Credit (from 2003)	2300
Mar-03	No changes announced	0

2% at its maximum.

The change in benefit regime causes a shift in both the first and second moments of each household's income expectations. At the bottom end, both the increase in income and the reduction in risk have the effect of increasing consumption. Further up the distribution, agents get the benefits of greater insurance, but suffer the withdrawal of expected mean income. The welfare effects at the top end of the distribution are therefore ambiguous.

I model the benefit floor as applying to current incomes, i.e. they act to insure transitory as well as permanent income fluctuations. The measures brought in provided a floor to transitory downwards shocks to wages and the active labour market policies and tax credits were aimed at lowering unemployment amongst low-income families so insured against some lay-off related transitory shocks. However we can imagine that other transitory shocks were not removed by the benefit reforms. They did not remove the risk of short-term lay-offs for example. This modelling decision affects the income and consumption dispersion moments in this way: if benefits truncate current incomes (including the transitory component) as opposed to the permanent component, then raising the benefit

floor lowers income dispersion more than the variance of consumption or the covariance because a part of the effect is to truncate “frothy” transitory shocks, which affect consumption less. If the income floor truncates permanent income alone, and transitory shocks can force current income below the income floor, then benefit changes move all the dispersion statistics more-or-less in tandem.

3.2.4 Modelling Wealth Formation

One aim of this chapter is to investigate how heterogeneity in wealth holdings affects consumption inequality. I divide assets into (gross) housing and other assets for two main reasons. First, housing is the largest source of wealth for the groups I study.¹⁶ Second, the housing market experienced a sustained boom over 1995-2008. In so doing, I am focusing on house price risk. Of course, households face other sources of asset risk, for example pension wealth risk, associated especially with stock market movements.¹⁷ However, it is more difficult to discern how these movements affect inequality because only defined contribution (DC) schemes co-move with the stock market, while defined benefit (DB) schemes shelter the recipient from this risk. Furthermore the BHPS dataset does not distinguish DC from DB pension wealth.¹⁸

I model each household’s non-pension portfolio as being invested as a fixed share in the two assets across the life-cycle. This fixed share is estimated against the portfolio allocations in the data. An alternative approach would be to allow for endogenous portfolio choice in each period and to estimate the coefficient of risk aversion to match portfolio allocations in the data.

When taking the model to data, the consumption concept is matched to non-durable expenditure. Underlying this modelling choice is the assumption that housing is homothetic and separable from non-durable consumption in the utility function. Davis and Ortalo-Magné (2010) provide evidence that housing is a constant budget share across US households. By excluding a treatment of the changing cost of housing services, this model

¹⁶See Banks, Smith, and Wakefield (2002) for a discussion of household portfolios in the UK within an international context.

¹⁷Juster, Lupton, Smith, and Stafford (2006) focus on the direct wealth effect from the stock market boom in the US in the late 1990s and conclude that it dominates that from housing wealth. However, direct stock ownership is smaller in the UK.

¹⁸Even if the dataset did differentiate DC from DB holdings, I would need to make assumptions about the portfolio allocation in the DC schemes.

therefore abstracts from the effect of a house price shock on the mean of non-durable consumption and instead focuses on the effect on consumption dispersion.

I do not explicitly model home-ownership itself. Strictly speaking, the model presented is representative of a continuum of individuals who rent housing and invest in liquid housing securities.¹⁹ To include home-ownership per se would require using a model along the lines of Wakefield (2009), Campbell and Cocco (2007) or Li and Yao (2007). In their models, agents can choose between renting or buying a home. Homeownership essentially provides rental services for free, but incurs a transaction cost, the implicit cost involved in saving for a downpayment and a per-period risk of forced sale²⁰, in which case these costs must be borne again. I do not use these models because the computational burden would preclude a modelling of the benefit regime and would preclude a formal estimation of the parameters. Instead, I argue that the main drivers of the effect of housing wealth on inequality which are present in this model would also be present in one with a formal home-ownership decision. In my model the portfolio share of ‘housing’ wealth becomes positively correlated with consumption over the life-cycle, so that a positive housing wealth shock exacerbates consumption inequality. Similarly, in the model of e.g. Li and Yao (2007) the households who do not own a home and hence who do relatively worse from house price growth are those who have had negative (transitory) income shocks.²¹

3.2.4.1 Capturing Wealth Holdings

Given portfolio shares as an input into the model I require some measure of wealth holdings as an output from the model to match to the data. Ignoring pension wealth I divide total lifetime resources into: human capital wealth; housing wealth, and other financial wealth including savings and mortgage debt. This decomposition ignores smaller categories such as durable holdings, but captures the majority of the household portfolio. I choose, as statistics for the household’s wealth holdings, the shares of housing and other wealth in to-

¹⁹The model therefore departs from e.g. Gourinchas and Parker who treat housing as a consumption commitment and instead use income after housing costs as their income concept and treat wealth as financial wealth outside of housing. Cagetti (2003) in contrast fits his model to the wealth distribution and finds a more plausible fit for preference parameters if he includes housing in his definition of wealth, implying that this is an important store for precautionary savings.

²⁰for example because of job re-location

²¹A model of home-ownership would not allow for multiple homes. 12% of my sample own another property. This may have an important role linking house prices and inequality at the top end of the distribution.

tal lifetime wealth, (formally $\frac{\text{housing wealth}}{\text{housing+other+HC wealth}}$ and $\frac{\text{other wealth}}{\text{housing+other+HC wealth}}$). These are denoted ψ_t^H and ψ_t^O . These statistics therefore serve two purposes. First they serve as a way of fitting the model to the data. In this case, almost any appropriate function of assets and income would work. Second, these statistics have behavioural interpretations. As discussed in section 3.4, ψ_t^H gives the elasticity of consumption with respect to housing wealth. Defining π_t as $\frac{\text{HC wealth}}{\text{housing+other+HC wealth}}$ as in Blundell, Low, and Preston (2008a) then $\pi_t = 1 - \psi_t^H - \psi_t^O$. π_t captures the transmission of a permanent income shock into consumption (the elasticity of consumption with respect to permanent income changes). Appendix A3.1 describes how these statistics are computed from the data.

3.2.5 Modelling the House Price Process

I treat the house price process as common across households and exogenous. Some authors (e.g. Attanasio, Blow, Hamilton, and Leicester (2009)) have emphasized heterogeneity of prices at the regional level. To the extent that I regress on region when looking at consumption residuals in the data, the effect of regional house price movements on mean regional consumption is irrelevant.²² I have also examined the effect of idiosyncratic house price shocks. The level of heterogeneity in house price movements is small though non-negligible. However, true idiosyncratic house price shocks are difficult to identify because they are likely caused by individual investments in the house, such as extensions, refurbishment or alternatively dereliction. This would not be new information to the household, and need not represent a net change in the household's life-time wealth.

A recent literature has looked at the effect of income inequality on house prices. Specifically relating to this chapter Määttänen and Terviö (2010) find that the increase in wage dispersion in Helsinki over 1998-2004 caused a decrease in average prices. However the effect is small and I ignore such a consideration here. Moretti (2008) and Van Nieuwerburgh and Weill (2010) look at the tempering effect of growth in (endogenous) house price dispersion on economic inequality following growth in wage inequality. They argue that increases in income inequality drive an increase in house price inequality as those at the

²²In my model, regions with higher house price growth would exhibit higher growth in consumption inequality. By using the national house price index and national (residual) consumption dispersion I am essentially averaging over these regional changes in consumption inequality.

top bid-up house prices. Such a mechanism is unlikely to be important here as the period I look at features increasing house prices and consumption inequality concurrent with stagnant or declining income inequality.

I model the house price process as a random walk. This assumption is shared by, for example, Campbell and Cocco (2007). A sizeable literature has looked at the precise nature of house price dynamics, with several authors documenting overshooting.²³ Nevertheless, the random walk model is a suitable benchmark.

3.2.6 Solution

There is no analytical solution for the model. Instead, the model must be solved numerically, beginning with the terminal condition on assets, and iterating backwards, solving at each age for the value functions conditional on the state of the benefit regime. I use standard methods for the solution: the income distribution and the distribution of portfolio shares are discretized, so assets are the only continuous state variable. I use the method of endogenous grid points described by Carroll (2006) to form the policy functions.

3.3 Data

I use the Family Expenditure Survey (FES) over 1978-2008 for cross-sectional data on household income and consumption. I obtain data on wealth from the British Household Panel Survey (BHPS) for 1995, 2000 and 2005 and yearly data on income and food consumption over 1991-2006. I stratify both datasets by high and low education. Education in the BHPS is given by qualification level, whereas the FES only has data on age leaving education. I define low education in the BHPS as those with no qualification higher than an O-level. The low education group in the FES comprises those leaving school at the compulsory schooling age (15 for those born before 1957 and 16 thereafter). The high education groups comprise those with higher qualifications or later school-leaving age. The two measures seem broadly comparable: across the sample of heads in the FES after 1991, 50% have low education, (42% when restricted to those born in the 1950s); in the BHPS, 51% have low education (45%).

²³See for example Ortalo-Magné and Rady (2006), and references therein

Finally, I obtain data on national house price movements from the Office for the Deputy Prime Minister (ODPM) over 1969-2008.

3.3.1 FES

The FES is described in detail in chapter 2. Here I just describe features and the treatment of the data relevant to this chapter.

Because data on incomes have been collected consistently only since 1978, I use data over 1978-2008. The main sample period is 1991-2006 to match the available data from the BHPS, but I use data over the longer period for some of the analysis: for example the age-profiles for mean household income and household size. The baseline sample over 1978-2008 contains 109,090 households. Each household is one data point. To each household I allocate a head, the male in a household consisting of a cohabiting couple with children. I use as population all households with heads aged 25-60. The sample is formed as follows: I drop households for which the head is outside the age range, or where food consumption or disposable income is negative, leaving 65,742 households. For robustness of the results I trim the top and bottom 1% of observations of each distribution. There are 64,682 household consumption observations.

The measure of income is total current income: labour earnings net of taxes, plus benefits and private transfers, plus asset returns excluding the drawing down of capital or capital gains. For the consumption dispersion profiles my measure of consumption is total non-durable expenditure. In order to get the right profiles for wealth formation, I include expenditure on all items when constructing the profiles for mean consumption growth. Housing expenditure data in the FES includes rent, mortgage interest payments and maintenance costs. This permits a “user-cost” measure of housing only. Clearly it would be desirable to conduct the analysis using a “real-cost” measure, but this is not possible as the FES contains no measure of housing wealth.

3.3.2 BHPS

The BHPS is described in detail in chapter 2. Again, here I just describe features and the treatment of the data relevant to this chapter.

The BHPS has detailed information on earnings, hours worked and other income, and

information on housing and durables, but little information on non-durable expenditure. An auxiliary dataset contains derived data on net household disposable income (see e.g. Bardasi, Jenkins, and Rigg (1999)), which I use in this study. I follow similar sample selection procedures for the BHPS as followed for the FES. The baseline sample contains 72,069 households. I remove households for whom the head is outside the age range. I then trim the bottom 1% of the distribution of disposable income and remove observations for which the head's education status is missing, leaving 45,798 households. Unlike the FES, where each questionnaire is completed in entirety, the BHPS contains many incomplete observations, so the quoted statistics are computed using fewer observations. For example, the total sample size of observed changes in household income is 32,379.

The BHPS has comprehensive information on housing wealth for most years. Another auxiliary dataset contains estimates of pension wealth for the BHPS sample over 1991-2001 (see Disney, Emmerson, and Tetlow (2009)). The BHPS has comprehensive information on household financial wealth for 1995, 2000 and 2005 only.²⁴ While the value of the first house and the value of all mortgages are reported exactly, the value of second homes and other financial wealth are reported in bands only. Again I use imputed data on the value of each type of asset (see e.g. Banks, Smith, and Wakefield (2002)).

Food consumption is categorized into twelve intervals for all years except the 1991. The top interval is unbounded above and the bottom is bounded by 0, so that the log of food consumption is unbounded below. For all intermediate intervals I assign the midpoint as food expenditure. For the top interval (over £160 per week) I assign £180 spending, for the bottom (less than £10 per week) I assign £5 spending. The results are robust to other sensible imputations. Chapter 4 discusses the use of food consumption data in more detail.

3.3.3 ODPM House Price Data

ODPM (previously the departments: DoE, DETR and DTLR) have published a quarterly house price index since 1968 based on data from the Survey of Mortgage Lenders (SML). For most of its history, the survey has involved a variety of mortgage lenders supplying a

²⁴The data for 1995 do not account for student loans and credit card debts. I ignore this consideration and treat the data as comparable across waves.

five per cent sample of their completions from the preceding month. The advantages of these data over, say, the Land Registry data are chiefly that the survey includes extensive information on the house's characteristics, so the price indices can be weighted correctly to represent the 'typical' house. Furthermore, the data cover the whole of the UK, rather than just Great Britain. The main disadvantage is that these data exclude cash purchases, around 25% of all deals. I use the annual time-series over 1969-2007 and deflate by the UK retail price index.

3.4 Results

I first present the main results of the model. Discussion of the estimation procedure, including details of the results and the parameter estimates is given in section 3.5.

3.4.1 Overview of Baseline Results

Figure 3.3 shows the empirical evolution of the variance of log income, the variance of log consumption and their covariance for the high and low education groups, together with their simulated counterparts. The variance of log income forms a subset of the moments used for fitting the model, but the variance of log consumption and the covariance are simulated freely from the model and the estimated parameters. Their proximity to the empirical moments emphasizes the validity of the model.²⁵ After 2000, we see a large dip in the variance of income and the covariance for the low education group, both in the simulations and in the data. The dip in all the moments for the high education is less pronounced. This is because this group has higher (equivalized) income, so the reforms did not compress their income distribution so much.

²⁵I post-estimate measurement error on consumption to match the average simulated variance of log consumption with the empirical moment. This implies classical measurement error of around 0.03 for both groups (in variance of logs). This reflects the usual mis-reporting and also the fact that expenditure is measured in a 2 week diary and contains week-to-week variation.

Figure 3.3: Empirical and Simulated Inequality

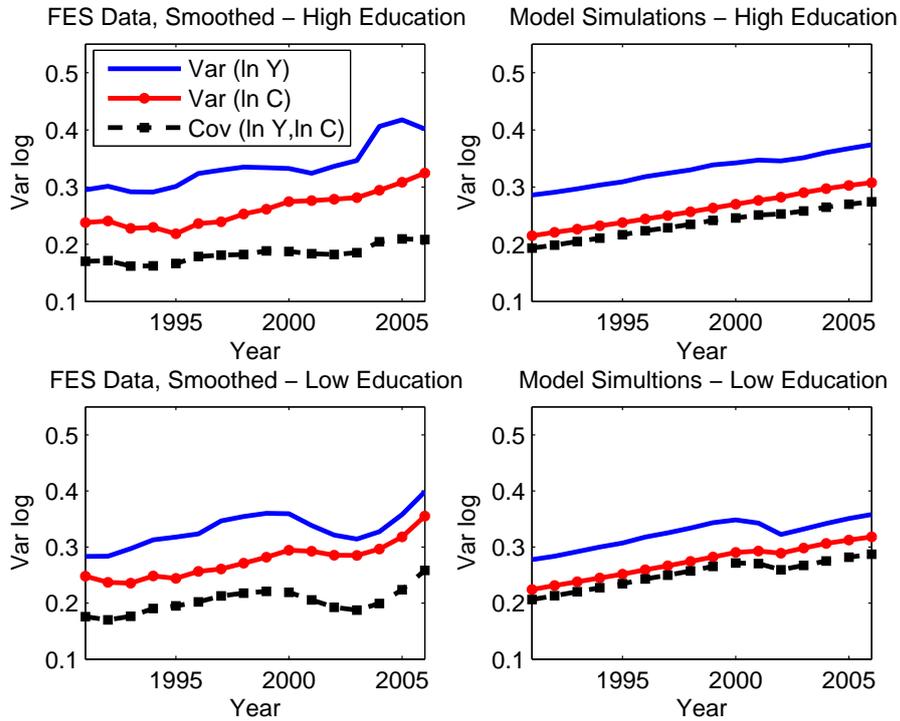


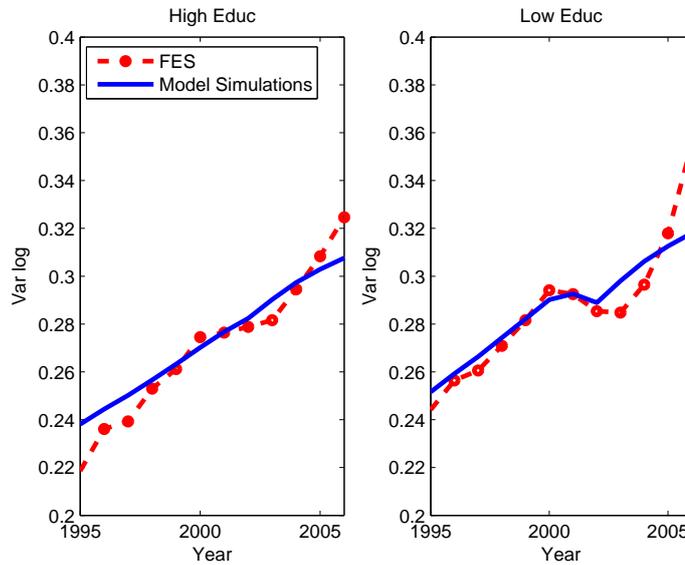
Figure 3.4 shows a key result from the previous figure in closer detail: the variance of log consumption over the main years of house price growth and benefit reform. The simulated variance qualitatively matches its empirical counterpart: the variance of consumption declines absolutely for the low education but not for the high, in both the data and the results.²⁶

As emphasized in the introduction, the covariance of income and consumption is also of primary importance. To a first approximation, this also identifies the permanent differences in economic resources across households.²⁷ A key and puzzling feature of the data is the divergence between the variance of log consumption and the covariance (what I term the ‘excess’ growth in the variance of consumption) over this period. It is to this that I now turn attention. Figure 3.5 shows this excess growth for the two groups plotted alongside the log real house price over 1991-2006, together with vertical lines for the main years of

²⁶The moments in the data all have larger high frequency movements than in the simulations. This partly reflects fluctuations in the variances of transitory and permanent shocks. I model these variances as constant over time. Including time variation in shocks for forward looking households is a delicate feature to model because these variations are presumably unforeseen. The common approach is to allow for a stochastic process to the variance of shocks. See, for example, Bloom (2009). Including this feature in the present model adds needless complexity and does not add to the central point.

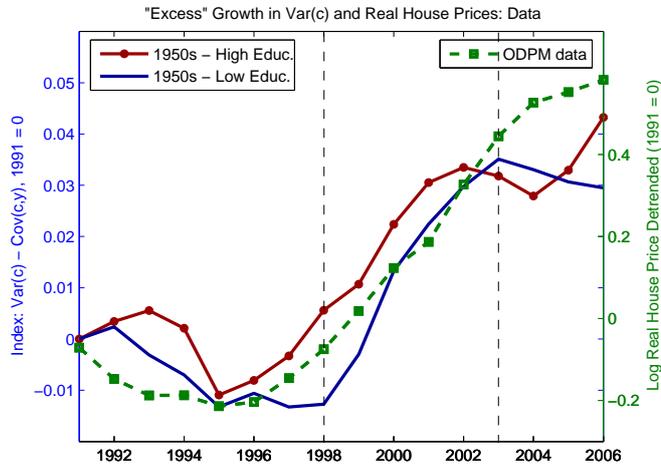
²⁷see Blundell and Preston (1998)

Figure 3.4: The Evolution of Consumption Inequality



benefit reforms 1998-2003. We see a notable correlation between the excess growth and the house price boom, also coinciding with the period of benefit reform.

Figure 3.5: House Prices and Difference Between $\text{Var}(\ln C)$ and $\text{Cov}(\ln C, \ln Y)$ in the Data



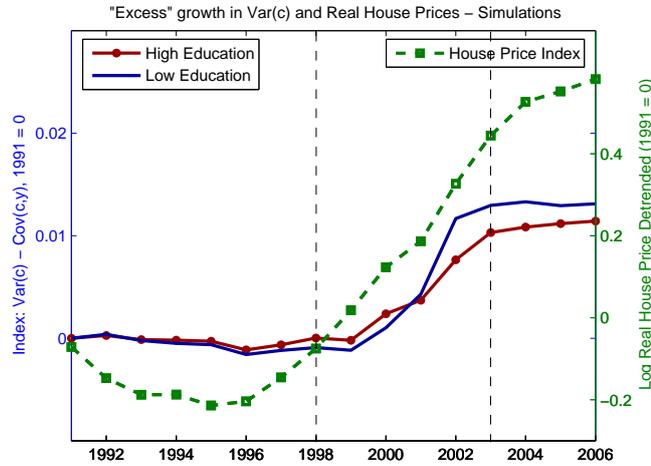
Note: Vertical dashed lines show the beginning and end of the main years of benefit reforms.

Figure 3.6 shows the simulated excess growth in the variance of consumption, also with the key external phenomena indicated. It is clear the simulated growth qualitatively matches that observed in the data. Quantitatively, the simulated peak-to-trough growth is around 35% of that observed for both groups.²⁸

In order to see the forces driving this result in the model, figure 3.7 shows the break-

²⁸The percentages use the growth from 1996 to 2004 from both the simulations and the observed data.

Figure 3.6: House prices and Difference between $\text{Var}(\ln C)$ and $\text{Cov}(\ln C, \ln Y)$ in the Simulations.



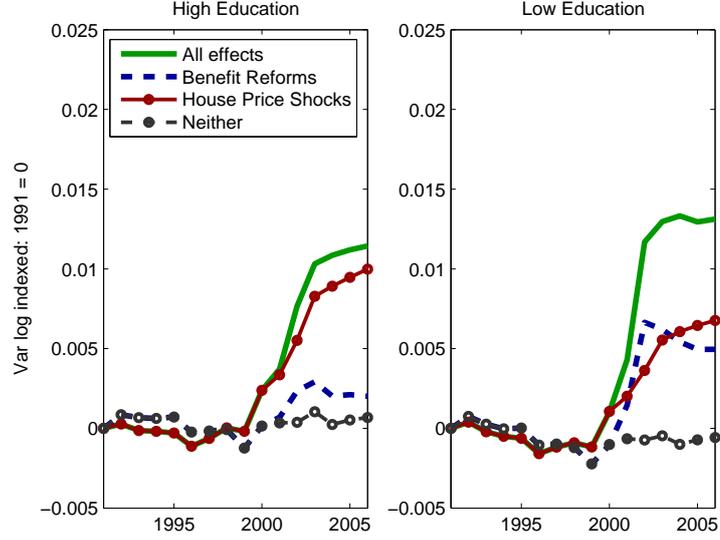
Note: Vertical dashed lines show the beginning and end of the main years of benefit reforms.

down of these results by their cause. In these plots, first I show the excess growth with both benefit reforms and observed house price growth imposed; second I switch off the benefit reforms, leaving just the house price shocks; third, I switch off the house price growth, leaving just the benefit reforms, and finally I run the simulations in the counterfactual world with neither reforms nor house price growth. For the high education group, heterogeneity in housing wealth drives the majority of excess growth in inequality (81% of the growth over 1996-2004); and, as is intuitively plausible, for the low education group, benefit reform is relatively more important (contributing 45% of total growth over 1996-2004).

The reason for these effects is, at first sight, intuitively simple: both house price growth and benefit reform exacerbated the importance of wealth (other than human capital wealth) in economic inequality. House price growth expanded wealth inequality directly. Benefit reforms compressed income inequality leaving wealth inequality constant (at least contemporaneously). To provide more rigour and further insight into this intuition, I now look at each process in turn.

3.4.2 The Effect of the Benefit Reforms

In this subsection I further analyze the effect of the benefit reforms on inequality. It is useful here to employ an analytic approximation to the consumption and saving model, adapting the approach used in, for example Blundell, Low, and Preston (2008a). I then

Figure 3.7: Breakdown of the Effect of House Price Shocks and Benefits on the Excess Growth of $\text{Var}(\ln C)$ 

look at the welfare effects of the reforms across the income distribution.

3.4.2.1 Intuition on the Effect of the Benefit Reforms on Inequality

I now adapt the consumption and saving model described in section 3.2, by abstracting from idiosyncratic shocks and house price shocks and focusing on exogenous shifts in the income distribution. In this section I simplify notation by using lower case letters for logarithms: $c_t \equiv \ln C_t$ and $y_t \equiv \ln Y_t$. Suppose agents have latent (residual, log) income \tilde{y}_t , which for simplicity is constant over periods $t-1$ and t , due to an absence of idiosyncratic risk. Realized income is given by $y_t = \theta_t \tilde{y}_t$, where θ_t is a load factor on residual incomes and represents stretching or compression of the distribution due to, for example, skill-biased technical change, or, as in this example, changes to the tax and benefit system.²⁹ Furthermore suppose $E_{t-1}(\theta_t) = \theta_{t-1}$, because changes to the distribution are unexpected.³⁰ In this case:

$$\begin{aligned} \Delta y_t &= \Delta \theta_t \tilde{y}_t \\ &= \tilde{y}_{t-1} \Delta \theta_t \end{aligned}$$

²⁹The main model stylizes benefit regime as a change in the income floor. Here I deviate by modelling it as a compression of the whole income distribution. Nevertheless, the same intuition should apply.

³⁰Furthermore, for precision, suppose that $E_{t-1}((\Delta \theta_t)^2) \approx 0$, because such shifts are rare. This assumption is required if households at the middle of the distribution are to face the same income risk as those on the periphery.

In appendix A3.3, I show that:

$$\Delta c_t \approx \Gamma_t + \pi_t \tilde{y}_{t-1} \Delta \theta_t \quad (3.10)$$

such that Γ_t is a constant reflecting saving due to the discount rate, interest rates and the precautionary motive (because of possible future income risk), π_t is the share of labour income in life-time wealth, and c_t is log consumption. The household's change in income and consumption relative to the mean is dependent on its position in the distribution: if the distribution of income is compressed, so that $\theta_t < \theta_{t-1}$, then households below mean income see their income grow, because $\tilde{y}_{t-1} < 0$ and $\Delta \theta_t < 0$, while those above mean income see their income decline, because $\tilde{y}_{t-1} > 0$. Consumption has the usual gradient, while the permanent shock to income ($\tilde{y}_{t-1} \Delta \theta_t$) transmits into consumption according to a self-insurance parameter π_t . The intuition for this transmission parameter is the following: if income is dwarfed by wealth (financial wealth and other assets), then a 1% change in income induces a less-than 1% change in consumption. (See Kaplan and Violante (2010) for an analysis of this partial insurance in a simulated Bewley economy).

We focus on $\Delta \text{Var}(c_t) - \Delta \text{Cov}(c_t, y_t)$, the excess growth in the variance of consumption. In appendix A3.3 I further show that this is approximately given by:

$$\begin{aligned} \Delta \text{Var}(c_t) - \Delta \text{Cov}(c_t, y_t) &\approx \Delta \theta_t ((\bar{\pi}_t - 1) \text{Cov}(c_{t-1}, \tilde{y}_{t-1}) \\ &\quad + \bar{\pi}_t (\text{Cov}(c_{t-1}, \tilde{y}_{t-1}) - \text{Var}(\tilde{y}_{t-1}))) \end{aligned} \quad (3.11)$$

If we make the further assumption that $\hat{c}_t = \pi_t \tilde{y}_{t-1}$ ³¹, where $\hat{c}_{it} = c_{it} - E_i(c_{it})$, then we further derive the expression:

$$\Delta \text{Var}(c_t) - \Delta \text{Cov}(c_t, y_t) \approx \Delta \theta_t (\text{Var}(\tilde{y}_{t-1}) \bar{\pi}_t (\bar{\pi}_t - 1)) \quad (3.12)$$

where $\bar{\pi}_t$ is the population mean of π_t . The expression as a whole is negative for $\bar{\pi}_t \in (0, 1)$, the case for positive average asset holdings and positive income flow). In the case where the income distribution is compressed, $\Delta \theta_t < 0$ and the excess growth in consumption

³¹This is not guaranteed from the approximations which concern the *changes* in consumption and income. However consider a household at mean consumption and income, in which case $\hat{c}_{t-1} = \tilde{y}_{t-1} = 0$. After receiving a permanent shock, then $\hat{c}_{t-1} = \pi_t \tilde{y}_{t-1}$.

inequality is pushed up.

Because $\bar{\pi}_t$ captures the transmission of income into consumption it is crucial for understanding equation 3.12. To understand its role we can consider two extremes: first, if there is no wealth, then income changes map one-to-one into consumption changes and income shifts induce the same change in the variance of consumption as in the covariance of consumption and income. This is the case for $\pi = 1$. Second, if wealth completely dominates income then income changes play no role in either the variance of consumption or the covariance. Such is the case when $\bar{\pi}_t = 0$. Only in the case of ‘partial insurance’, when background wealth plays some but not all of the role in financing consumption does the compression of income compress the covariance more than the variance of consumption.

3.4.2.2 Welfare Effects from the Benefit Reforms

I turn now to a brief discussion of the welfare effects of the reforms. There are two main effects. First, all households are affected directly by the reforms and experienced a shift in mean income. Recipients received an unexpected increase in income, while the rest of the population had to balance the state budget through higher taxes, proportional to their income. The minimum income support provides a net transfer from the high education group to the low education group.³² Second, all households experienced a change in the distribution of future income. Those at the bottom of the distribution experienced a reduction in risk,³³ and while those at the top receive little benefit from the income floor.³⁴ Of course, an important part of the welfare effects of such reforms is the dead-weight loss from increased labour market distortions. I do not model labour supply effects here, so the welfare gains presented here should be taken as an upper bound on the true welfare changes.

In the spirit of Lucas (1987), for an agent with preferences over Z_t and C_t defined by

³²I look at just the cohort born in the 1950s. I implicitly assume that benefit reforms were revenue neutral within cohorts. Of course there were possibly net transfers across cohorts. There may also have been net transfers across time because the reforms were not explicitly funded out of current taxation.

³³It is worth noting that a minimum income floor induces asymmetric risk, because the left-hand tail to income shocks is truncated. This change to the third moment of income shocks is likely just as important as the change to the second moment.

³⁴I measure risk as the variance of changes in log income. Therefore the proportional tax increase itself induces no change in risk profile, only the minimum income floor

$v(Z_t)\ln(C_t)$, we define expected utility for a household at time t :

$$E_t U_k = E_t \sum_{s=t}^T \beta^{s-t} v(Z_s) \ln(C_{k,s})$$

where k indexes a consumption stream for a particular scenario. Here different scenarios reflect different ex-post out-turns for the benefit regimes. Ex-ante, households place exactly the same probability distribution on benefit reform.³⁵ We now define $E_t U_{k,\phi}$ to be the utility for scenario k , where consumption is multiplied by a scaling parameter, ϕ :

$$E_t U_{k,\phi} = E_t \sum_{s=t}^T \beta^{s-t} v(Z_s) \ln(\phi C_{k,s})$$

In this notation $E_t U_k \equiv E_t U_{k,\phi=1}$. Letting $k = 2$ represent the scenario without the benefit reforms, and $k = 1$ the scenario with reforms I implicitly define ϕ^* as follows:

$$E_t U_{2,\phi^*} = E_t \sum_{s=t}^T \beta^{s-t} v(Z_s) \ln(\phi^* C_{2,s}) = E_t U_1$$

where ϕ^* is the proportion of consumption in environment 2 needed to give the same utility as scenario 1. Solving for ϕ^* :

$$\phi^* = \exp\left(\frac{E_t U_1 - E_t U_2}{\sum_{s=t}^T \beta^{s-t} v(Z_s)}\right)$$

Table 3.2 shows the welfare effects of all the benefit reforms for the 15 levels of permanent income, expressed as percentages. We derive the welfare measures for households according to their position in the income distribution in 1999. Overall 56% of households benefitted overall from the reforms, even though only 5.5% of households directly received a subsidy. 17% of the population received a benefit bigger than 1% of consumption. There is a strong role for transfers from high to low educated: only 20% of the high education group benefitted from the reforms, compared to 87% of the low education group.

Another way of looking at the welfare effects of the reforms is by comparing their effect with that of a complete redistribution. I run the further scenario: in 1999, the government

³⁵When analysing the welfare effects of benefit reform I abstract from the house price boom. Asset risk is present in each scenario presented here, with no out-turn aggregate shock to asset prices.

Table 3.2: Welfare Effect of All Benefit Reforms, Across the Income Distribution

	High Educ	Low Educ
Highest Income	-0.15	-0.12
	-0.15	-0.07
	-0.14	0.00
	-0.14	0.08

	0.04	1.46
	0.11	1.83
	0.17	2.34
Lowest Income	0.27	3.00

Notes: 1) This table compares the consumption streams from 1999 (age 44) onwards for environments with and without benefit reforms. 2) There are 15 nodes on the (permanent) productivity grid. Each point therefore represents 6.6 % of the population. The middle 7 nodes are omitted from the table.

completely and totally unexpectedly redistributes all income for the rest of working life, and so removes all inequality and income risk. We can think of a putative household about to be assigned a life in 1999 from behind the veil of ignorance. The household may enter one of three worlds: no reform, the actual reforms that were enacted after 1999, or a world with completely equal income. I then compute the compensation required to enter the more redistributive worlds against the laissez-faire world.³⁶ I find that compensation required to deprive the household of the benefit reforms is 3.2% that of the compensation required to deprive the household of complete redistribution.³⁷

3.4.3 The Effect of House Price Shocks on Inequality - Results and Intuition

In this subsection I analyze the effect of the house boom on consumption inequality. To provide intuition I again employ an analytic approximation to the consumption and saving model.

³⁶I abstract from the issue of asset equality by assuming the government doesn't, or can't redistribute assets.

³⁷The putative household needs 0.5% of their consumption good to forego benefit reforms and 17% to forego the world of complete redistribution.

3.4.3.1 Approximation to the Consumption and Saving Model

I now adapt the consumption and saving model described in section 3.2, by abstracting from changes to the benefit system. Households now face just permanent and transitory changes income risk and house price risk. Again, I simplify notation by using lower case letters for logarithms: $c_t \equiv \ln C_t$ and $y_t \equiv \ln Y_t$.

In appendix A3.2 I show that an approximate solution for the growth of log consumption is given by:

$$\Delta c_{it} \approx \Gamma_t + \pi_{it}(\eta_{it} + \alpha_t \epsilon_{it}) + \psi_{it}^H \zeta_t^H \quad (3.13)$$

where Γ_t is a constant reflecting saving due to the discount rate, interest rates and the precautionary motive, ζ_t^H is the realized common shock to housing wealth, η_{it} is the permanent shock to income, ϵ_{it} the transitory shock to income, π_t is the share of labour income in life-time wealth, ψ_{it}^H is the share of gross housing wealth in life-time wealth, defined previously, α_t is an annuitization factor giving the contribution of a transitory shock to life-time wealth, and c_{it} is log consumption.

The intuition for equation 3.13 is simple: a permanent shock to wealth ζ_t^H causes consumption to grow proportionately to how much of the asset the household has. In the case where the household's consumption is financed purely by a housing asset, with no labour income and no financial wealth, then a 1% increase in housing wealth raises consumption by 1%. In the presence of labour income and financial wealth, the elasticity of consumption with respect to housing wealth is its share in total expected life-time wealth. Similarly the elasticity of consumption with respect to a permanent shock to income is π_{it} , the share of permanent income in life-time wealth.³⁸

Appendix A3.2 further shows that this process implies the following moments for in-

³⁸The elasticity of consumption with respect to transitory shocks, u_t , is $\pi_t \alpha_t$: the product of the share of permanent income in lifetime wealth (π_t) with the annuity value of a transitory shock (α_t).

come and consumption:

$$\Delta \text{Var}(y_t) = \text{Var}(\eta_t) + \Delta \text{Var}(\epsilon_t) \quad (3.14)$$

$$\Delta \text{Cov}(c_t, y_t) = \underbrace{\bar{\pi}_t \text{Var}(\eta_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)]}_{\text{Contribution from income}} + \underbrace{\zeta_t^H \text{Cov}(\psi_t^H, y_{t-1}^P)}_{\text{Contribution from housing}} \quad (3.15)$$

$$\Delta \text{Var}(c_t) = \underbrace{(\bar{\pi}_t^2 + \text{Var}(\pi_t))(\text{Var}(\eta_t) + \alpha_t^2 \text{Var}(\epsilon_t))}_{\text{Contribution from income}} + \underbrace{(\zeta_t^H)^2 \text{Var}(\psi_t^H) + 2\zeta_t^H \text{Cov}(\psi_t^H, c_{t-1})}_{\text{Contribution from housing}} \quad (3.16)$$

where $\bar{\pi}_t$ is shorthand for $E_i(\pi_{it})$, and $y_t^P \equiv \ln P_t$, is log permanent income.

The economic intuition for the contribution from housing to growth in the variance of log consumption is as follows: a positive house price shock causes a change in the variance because of variation in the elasticity of consumption, given by ψ_t^H . To provide further intuition we can imagine two distributions of this elasticity. First, if housing wealth is spread uniformly across the distribution and roughly in proportion to households' life-time wealth, then the elasticity is uncorrelated with the consumption distribution and the house price shock induces just an orthogonal shock to the consumption distribution (in addition to the increase in mean consumption) of size $(\zeta_t^H)^2 \text{Var}(\psi_t^H)$. As a second, more realistic case, households who receive good transitory shocks accumulate sufficient funds to put a downpayment on a home and also can afford higher consumption. Therefore, home-ownership should be correlated with consumption. For this reason at least, the housing wealth share (and hence the elasticity) should covary positively with the consumption distribution. In this case the positive shock increases inequality by an additional factor $2\zeta_t^H \text{Cov}(\psi_t^H, c_{t-1})$.³⁹

Turning to the effect on the covariance of income and consumption, the shock induces a change only if the elasticity is correlated with the distribution of permanent incomes. In the standard consumption and saving model, all behaviour is invariant to the level of permanent income (the asset/permanent income ratio suffices as the state variable).

³⁹In contrast, if the elasticity covaries negatively with the consumption distribution, then the positive shock can reduce inequality if $2\zeta_t^H \text{Cov}(\psi_t^H, c_{t-1}) > (\zeta_t^H)^2 \text{Var}(\psi_t^H)$. Such is the case with social security (state pension) wealth. Because of the redistributive nature of the social security system the share of life-time in pensions varies negatively with life-time wealth. A positive shock to such wealth, because of, say, an unexpected increase in generosity would reduce consumption inequality.

Using this as a benchmark, we may think that the housing wealth share should be roughly uncorrelated with the distribution of permanent incomes.⁴⁰

3.4.3.2 The Approximate Effect of House Price Shocks

In order to put empirical flesh on the bones of equations 3.14-3.16, table 3.3 shows relevant moments from the simulations and estimated from the BHPS asset, income and food consumption data for the cohort born in the 1950s. I first show the empirical moments pooled over high and low education, to give better precision. I then show the moments for the high education group and compare them against the simulations. Appendix A3.1 gives details of how the moments in the data were computed.

Table 3.3: Wealth-Share Parameters for the 1950s Cohort: All Types and High Education Alone

	year	Mean ψ^H	Mean ψ^O	Var ψ^H	Var ψ^O	Cov($\psi_t^H, \ln C_{t-1}^{eq}$)	Cov($\psi_t^H, \ln Y_{t-1}^{eq}$)	Var($\ln C_t^{f,eq}$)
All Education Groups								
BHPS data	1995	0.134 (0.003)	-0.047 (0.003)	0.007 (0.000)	0.007 (0.001)	0.017 (0.005)	0.009 (0.002)	0.117 (0.011)
	2000	0.178 (0.006)	-0.054 (0.006)	0.015 (0.001)	0.023 (0.021)	0.024 (0.007)	0.005 (0.006)	0.129 (0.007)
	2005	0.317 (0.007)	-0.041 (0.007)	0.028 (0.002)	0.019 (0.008)	0.038 (0.01)	0.003 (0.009)	0.123 (0.008)
	1995	0.144 (0.005)	-0.048 (0.005)	0.006 (0.001)	0.007 (0.001)	0.011 (0.004)	0.003 (0.004)	0.105 (0.011)
	2000	0.197 (0.005)	-0.059 (0.008)	0.014 (0.002)	0.032 (0.031)	0.018 (0.008)	-0.006 (0.006)	0.134 (0.007)
	2005	0.33 (0.009)	-0.043 (0.01)	0.025 (0.003)	0.023 (0.011)	0.024 (0.011)	-0.005 (0.007)	0.117 (0.008)
Simulations	1995	0.098	-0.025	0.003	0	0.004	-0.001	
	2000	0.202	-0.052	0.01	0.001	0.007	-0.009	
	2005	0.404	-0.104	0.021	0.001	0.011	-0.024	
							Cov($\psi_t^H, \ln Y_{t-1}^P$)	

Notes: ψ^H is the share of lifetime wealth in housing, ψ^O the share in other financial assets (including mortgages). C^{eq} is household equalized consumption, Y^{eq} is household equalized income, $C^{f,eq}$ is equalized food consumption. Cov($\psi_t^H, \ln C_{t-1}^{eq}$) is computed as Cov($\psi_t^H, \ln C_{t-1}^{f,eq}$)/ ξ_f where ξ_f is the elasticity of food consumption with respect to total consumption, estimated to be 0.4. See appendix A3.1 for more details.

The model is fitted to mean ψ^H and mean ψ^O . One striking feature of the data is that other wealth is negative in all years and for both education groups. Mortgage debt exceeds other financial wealth on average for all these cells.⁴¹ The model struggles to match the (absolute) size of both housing wealth, and other wealth/mortgage debt in 1995, when the

⁴⁰In my model wealth, and hence the housing wealth share, is lower among low productivity households because of the minimum income floor. Other reasons why home-ownership may be positively correlated with permanent income level are progressivity in pension provision and absolute (non-proportional) costs of home purchase.

⁴¹I acknowledge that other wealth does not include durables wealth, so the estimate of ‘other’ wealth share is biased downwards.

cohort is aged around 40, but matches well wealth later in the life cycle. It seems that households invest in housing earlier than can be generated by this simple model. Similarly, the model matches well the variation in housing wealth shares later in the life cycle, but understates it earlier. The model always understates the variation in the shares of other wealth.

The penultimate two columns give probably the most important numbers. Equations 3.14-3.16 imply that the covariances of ψ^H with consumption and income are crucial in determining the effect of house prices on inequality. The simulated covariances of ψ^H with consumption are about half the size of the empirical moments for the high education group, indicating that the model is understating the effect on consumption inequality.

It is difficult to identify the covariance with permanent income in the data. Current income does not give a good proxy, because transitory incomes are mostly then stored in wealth, so the covariance of wealth shares with current income should be higher than the covariance with permanent income. Here I show the covariance with current income in the data as a rough upper bound to the covariance with permanent income, and display the covariance with permanent income from the simulations. The covariance of wealth share with lagged current income is small and generally insignificant.

I now use the empirical statistics to derive a first-order approximation of the effect of house price increases on consumption inequality. For example, in 2000, national house price growth was 7.1% above trend. The contribution to growth in the variance of consumption for the whole cohort from $(\zeta_t^H)^2 \text{Var}(\psi_t^H)$ was very small, at $0.071^2 \times 0.015 < 0.0001$. The contribution from $2\zeta_t^H \text{Cov}(\psi_t^H, c_{t-1})$ was $2 \times 0.071 \times 0.024 = 0.0034$. Adding these yearly contributions up we get a point estimate of the effect of house price growth of around 0.025 log points over 1997-2004, around 60% of the observed excess growth in consumption inequality.

3.4.3.3 Comparison with Other Studies on Elasticities and the Marginal Propensity to Consume out of Wealth

As emphasized above, ψ_t^H gives the elasticity of consumption with respect to house price changes in this (approximated) model. Data from the BHPS implies that this number

averages around 0.1-0.15 for households at age 40 and over 0.3 for households at aged 50.⁴² As an empirical exercise, my computation of ψ^H is clearly not the best way of identifying this elasticity. As previously discussed, its absolute magnitude is less important to this study than the way it varies across the population. Nevertheless it is interesting to relate the number obtained to better-identified estimates from the literature.

Campbell and Cocco (2007) find an elasticity of over 1 for older households, a smaller but positive elasticity for younger households and smaller but still positive elasticities for both old and young renters. On my approach, an elasticity of over 1 from a pure wealth effect is only possible if households are implausibly over-leveraged, such that housing wealth is larger than human capital wealth net of mortgage debt. The only explanation for such a high elasticity in my model would be the presence of binding liquidity constraints: then the marginal propensity to consume (MPC) out of housing wealth could be as high as 1, and the elasticity well over 1. It seems implausibly high given the estimates I present on the scale of housing wealth in life-time wealth and the position of the cohort studied in the life cycle.

Attanasio, Blow, Hamilton, and Leicester (2009), in contrast, find an elasticity that declines with age, from the around 0.2 for young households to around 0.13 for middle-aged cohorts, the group I study. Leaving aside the puzzling disparity between these estimates and those in Campbell and Cocco (2007), I note that the absolute size of those Attanasio, Blow, Hamilton, and Leicester (2009) conforms more with the logic of my approach.

Other papers in the literature quote the MPC out of housing wealth. The MPC is the elasticity of consumption multiplied by the consumption/asset ratio. The consumption/asset ratio cannot be computed using a single dataset. Model simulations imply the ratio is around 0.2 for 50-year-old households. This would imply an MPC of around 0.06.

The MPC out of wealth should be the same as that out of transitory income. In reality, transactions costs or behavioural phenomena such as inattention may create differences in MPCs across different types of assets and income. Furthermore there is likely to be a

⁴²Because I am addressing the effect of house price changes on consumption dispersion rather than mean consumption, I have not explicitly modelled the change in the cost of housing services. The true effect on non-durable consumption for renters is likely negative because housing becomes more expensive. The given numbers should therefore be interpreted as compensated elasticities, where the household is compensated for the change in price of the consumption bundle. These therefore ignore income effects of price changes and give pure wealth effects.

large bequest motive for housing in particular. If households plan to bequeath, say, half their housing wealth by downsizing to a smaller property before death whereupon leaving their house to their children, then this would halve the theoretical MPC.

Case, Quigley, and Shiller (2005) look at the MPC out of housing and stock-market wealth and find an MPC out of housing wealth gains of between 0.03 to 0.15. Paiella (2007) finds an MPC out of financial wealth of around 0.08 and an MPC of around 0.025 out of housing wealth for the population of Italian households with heads aged between 25 and 75. Carroll (2006) explicitly distinguishes short from long-run elasticities and finds an MPC of 0.02 over the first year of a house price gain owing perhaps to inattention and short-run adjustment costs, rising to 0.09 over the long-run.

3.5 Estimation Procedure

Estimation proceeds in two main stages. First I pre-estimate several inputs into the household's dynamic programme: parameters of the income process, average income and average household size over the life-cycle, each by cohort and educational achievement. I also estimate the ex-post returns to housing. Then I estimate the full model by solving the household's decision making problem and performing method of simulated moments. I first give details of the initial estimates, then give details of the estimation of the full model.

3.5.1 Estimating the Income Process

I estimate parameters of the household income process using longitudinal data from the BHPS. There is a long literature closely examining the statistical process for (male) earnings, for example, MaCurdy (1982), Abowd and Card (1989), Meghir and Pistaferri (2004), Guvenen (2009), but there are very few studies that empirically test features of the household income process. Blundell, Pistaferri, and Preston (2008b) is an exception.

In line with the process defined in the model (equation 3.6) and with the literatures mentioned, I impose a permanent-transitory decomposition of household disposable in-

come:

$$\begin{aligned} y_{it} &= m_t^a + \beta_t X_{it} + y_{it}^P + y_{it}^T \\ y_{it}^P &= y_{it-1}^P + \eta_{it} \\ y_{it}^T &= g(L)\epsilon_{it} \end{aligned}$$

where: $y_{it} \equiv \ln Y_{it}$ is log current income, m_t^a is an aggregate shock at time t ; X_{it} is a set of household characteristics, which in this application constitute household size, a quartic polynomial in age, education and region; as before, $y_{it}^P \equiv \ln P_{it}$ is log permanent income, η_{it} is the shock to permanent income and ϵ_{it} is the shock to transitory income at time t ,⁴³ and where $g(L)$ is an arbitrary invertible polynomial function of the lags and nests all stationary ARMA processes.

Similarly to the approach in Meghir and Pistaferri (2004) I identify the variance of permanent shocks by the following moment condition:^{44 45}

$$\sigma_\eta^2 = \lim_{\tau \rightarrow \infty} \text{E}(\Delta y_t \sum_{s=-\tau}^{\tau} (\Delta y_{t+s})) \quad (3.17)$$

I identify the variance of transitory incomes under the following moment condition:

$$\sigma_{y^T}^2 = \lim_{\tau \rightarrow \infty} (\text{Var}(y_t) - \text{Cov}(y_{t+\tau}, y_t)) \quad (3.18)$$

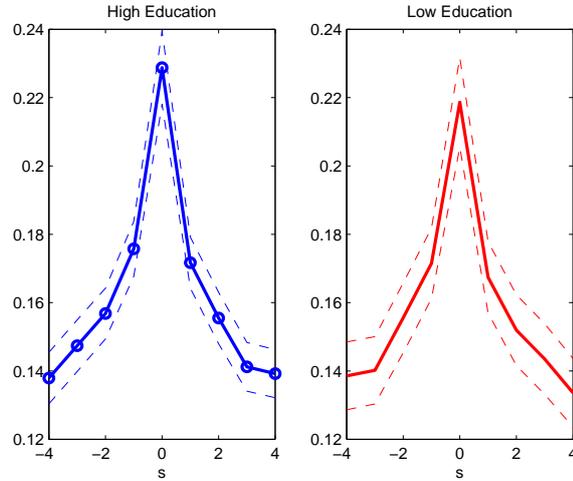
Figure 3.8 shows the variance-covariance at 4 lags for high and low education, for the 1950s cohort for income pooled over 1991-2006. If the long-term component represents the permanent differences between households, we see that the short-term differences between households have a reasonably long-lasting tail, perhaps from some kind of auto-regressive, or high-order moving average process.

To take equations 3.18 and 3.17 to the data I choose $\tau = 3$ which is the correct choice

⁴³For the purposes of inferring risk from this model, I am, of course, supposing the usual assumptions on the household's information set, that it has no advanced knowledge of idiosyncratic changes, that it knows the mean income trajectories for its observable type and that changes to characteristics are known and planned.

⁴⁴Of course, the income process is not infinite, but the meaning should be clear to the reader.

⁴⁵The moment conditions in 3.17 and 3.18 are slightly biased by the effect of the policy reforms, because I am modelling the process to pre-benefit income, and only observe post-benefit income. However, simulations show that the bias is small overall.

Figure 3.8: Covariance Structure of Income Residuals: $\text{Cov}(y_{t+s}, y_t)$ 

Note: Dashed lines give 2 standard-error bands.

if income follows an MA(2) process. i.e.:

$$\begin{aligned} \sigma_{yT}^2 &= \text{Var}(y_t) - \text{Cov}(y_{t+3}, y_t) & (3.19) \\ \sigma_{\eta}^2 &= \text{E}(\Delta y_t \sum_{s=-3}^3 (\Delta y_{t+\tau})) \end{aligned}$$

The results of the overall estimation change little for other sensible choices of τ .

My sample is restricted to households where the head is present for at least 7 years. After regressing on the vector of characteristics and a constant for each year, I pool the sample over all the years (1991-2006) and compute the empirical counterparts of the moments in 3.19.

Estimates are shown along with other inputs into the main estimation in table 3.5, with asymptotic standard errors. The sample size for the high education group is 1908 household observations, for the low education group it is 930.

It is difficult to identify the variance of transitory shocks separately from measurement error in this model without imposing structure on the form of measurement error.⁴⁶ I assume there is no measurement error on income and assign all period-by-period variation to transitory shocks.

⁴⁶see Meghir and Pistaferri (2004)

3.5.2 Estimating the House Price Process

In accordance with the discussion in section 3.2 I assume real house prices follow a random walk with drift. I estimate an average real return on housing of 0.034 with a standard deviation of shocks of 0.089. These are estimated over 1969-2008 from the ODPM data. As a simple investigation of the house price time series I run an OLS regression on:

$$\Delta lHP_t = \mu + \beta t + \gamma lHP_{t-1} + \zeta_t$$

where lHP_{t-1} is the log real house price and u_t is an innovation. I estimate $\hat{\gamma} = -0.135$ with a t-value of -1.55 (39 observations). A Dickey-Fuller test fails to reject the presence of a unit root at any reasonable level of significance.

When simulating the model, I impose log real house price changes as the ex-post return to housing wealth.

3.5.3 Other Pre-Estimated and Imposed Parameters

I impose a return of 0.018 for the other ‘safe’ asset with a standard deviation of 0.033. These statistics are derived from Barro (2006), from data on real bond returns in the UK over 1954-2004.⁴⁷

I use a utility function that is separable in Z_{it} and C_{it} and use logarithmic preferences over consumption: $u(Z_{it}, C_{it}) = v(Z_{it})\ln(C_{it})$. The consumption felicity function implies an elasticity of intertemporal substitution (EIS) in line with the micro literature and higher than in the macro literature.⁴⁸

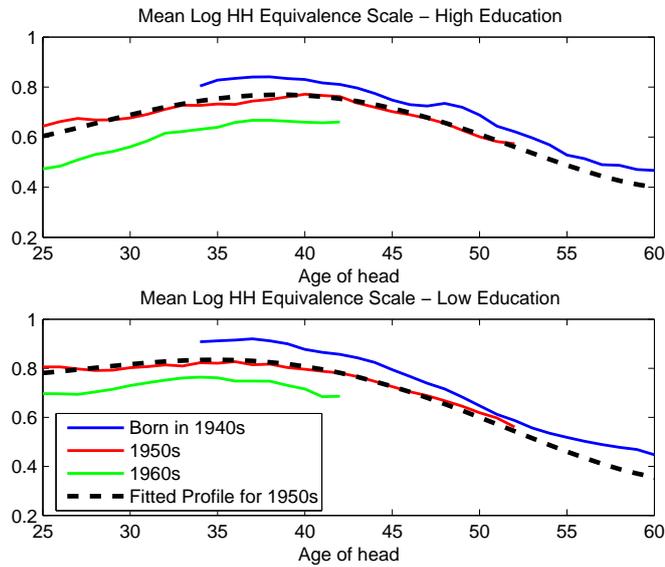
For the function $v(Z_t)$ I use the modified-OECD equivalence scale. With log preferences this implies that households equate the expected marginal utilities of equivalized consumption. Figure 3.9 shows the equivalence scales for three cohorts for the low educated group.

The evolution of expected income is very important to the results I gather. Life cycle wealth formation affects the extent to which house-price shocks affect households. Figure 3.10 shows the mean raw (un-equivalized) real income profiles by education and cohort.

⁴⁷Barro (2006), Table IV.

⁴⁸See, for example, Attanasio and Weber (1995) on the micro side and Barro (2006) on the macro.

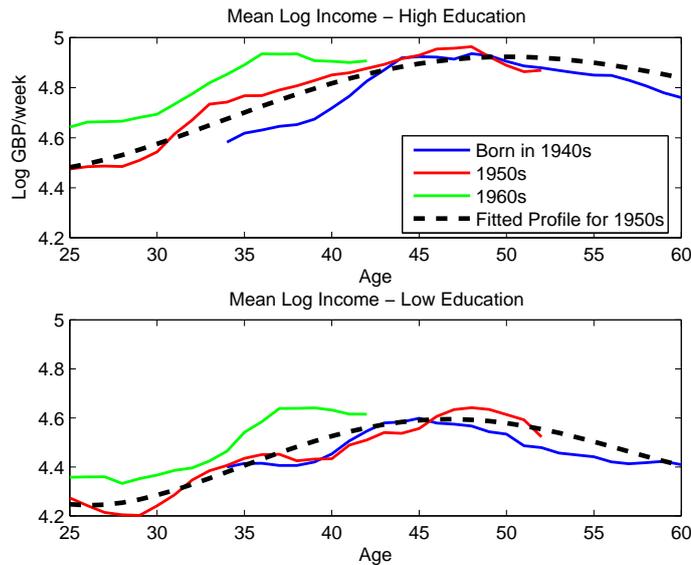
Figure 3.9: Life-Cycle Consumption Needs By Cohort and Education



Note: Equivalence scale is 1 for first adult, 0.7 for additional adults, 0.5 for children

I fit a stylized income profile for each education group by regressing the data on cohort dummies and a quartic polynomial in age. According to this profile, real income grows 56% for the high education group from age 25 until the peak at age 50 before declining by 14% up to retirement. For the low education group income grows by 41% until the peak at age 47 before declining by 21% up to retirement.

Figure 3.10: Income Profiles by Cohort and Education



I pre-estimate the initial wealth endowment to the median asset/income ratio for those aged 25-30 in 1995. This is 0.23 for the high education group and 0.17 for the low education

group. This initial endowment of wealth affects the results little.

3.5.4 Estimation by Method of Simulated Moments

In the baseline model I estimate 12 main parameters: the minimum income levels in the first and fourth states (\underline{Y}_1 and \underline{Y}_4); the variance of initial permanent income for high and low education; the variance of permanent shocks; variance of transitory shocks; the rate of time preference, and the average leverage, also all for high and low education.⁴⁹ I assume that \underline{Y}_2 and \underline{Y}_3 are linear interpolants of \underline{Y}_1 and \underline{Y}_4 .

I estimate using method of simulated moments. Because consumption is the only choice variable and is continuous, the objective function is concave. However, because the income grid is discrete and I use a finite number of simulations to generate the distribution of incomes the approximation of the objective function is locally non-concave. Nevertheless, I proceed with gradient methods and overcome the local non-smoothness by performing numerical differencing with larger step size compared to when solving a completely smooth problem.

I estimate using the following criterion function:

$$\hat{\phi} = \arg \min_{\phi} (\hat{\alpha}^D - \hat{\alpha}^S(\phi))\Omega(\hat{\alpha}^D - \hat{\alpha}^S(\phi))$$

where $\hat{\alpha}^D$ are the moments in the data and $\hat{\alpha}^S(\phi)$ are the corresponding simulated moments for given parameter values ϕ . The simulated moments, $\hat{\alpha}^S(\phi)$, are computed from 30000 draws. The optimal weighting matrix under the null is the inverse of the covariance matrix from the data, $\text{var}\hat{\alpha}^D$. I use the diagonal of this matrix to reduce well-known bias.⁵⁰ Standard errors can be computed using the formula in Smith Jr (1993):

$$\text{var}(\hat{\phi}) = (J'\Omega J)^{-1}J'\Omega V\Omega J'(J'\Omega J)^{-1}$$

where $J = \frac{\partial \hat{\alpha}^S(\phi)}{\partial \phi}$ and $V = \text{var}(\hat{\alpha}^D - \hat{\alpha}^S(\hat{\phi}))$. Lee and Ingram (1991) show that V reduces to $(1 + \frac{1}{K})\text{var}(\hat{\alpha}^D)$ at the null where K is the ratio of the number of simulated draws to

⁴⁹I also estimate measurement error in consumption. However this is separable in the criterion function from the other parameters so I estimate it after the main procedure by minimizing the squared differences of the simulated and empirical moments.

⁵⁰see Altonji and Segal (1996) for an analysis of small-sample bias for the optimal minimum-distance estimator, a close cousin of the estimator used here.

the number of data points. The moments I use are derived from different sample sizes, but they are dwarfed by the number of simulation draws. When computing the variances directly, I find the contribution of the error from the simulations to the overall variance is very small. At its largest, simulation error contributes 0.5% to the standard errors on the estimates of the variance of permanent shocks.

3.5.5 Choice of Moments

To fit the model I use the following moments: the variance of log income over 1991-2006 for both the high and low education groups (2x16 moments); mean equivalized consumption growth over 1981-2006 for both groups (2 moments); estimates of the variance of permanent shocks and variance of transitory shocks from the BHPS (2x2 moments) and mean ψ_t^H and mean ψ_t^O for 1995, 2000 and 2005 (2x6 moments).

A discussion of the periods chosen for these moments is warranted. The 1950s cohort enters adult life in 1981. I use consumption growth from the beginning of adult life in order to capture life-cycle wealth formation. However, I do not use the dispersion statistics over the 1980s because: first, this was a period of higher latent idiosyncratic income risk, and second, it is possible that tax and benefit reforms also affected the income and consumption distributions similarly to the way I am examining over 1999-2003. To model the 1980s properly one would need to have distinct episodes of permanent income risk and to have a full treatment of the tax reforms at the time. 1991-2006 is also the period of BHPS data, which are important inputs into the estimation, especially for the identification of income risk and asset dispersion.

3.5.6 Parameter Estimates

Table 3.4 shows the parameter estimates for the baseline model. One striking feature is how similar are the estimates for both groups. This reflects that growth in the variance of consumption, growth in mean consumption and wealth holdings are roughly similar across the groups.

The minimum income guarantee is quite imprecisely measured in the first year (1999). The subsidy doesn't actually apply in this year, it is mainly identified through its effect on the minimum income in later years (2000 and 2001).

Table 3.4: Parameter Estimates

Variable	Fixed λ		Description
	High Educ	Low Educ	
Min Inc 1999	5.8798 (11.7117)	-	
Min Inc 2002	20.4302 (1.7796)	-	
λ	0.0010 (Imposed)	-	Probability of regime change
β	0.9782 (0.0010)	0.9783 (0.0012)	Time preference
Var η	0.0064 (0.0008)	0.0081 (0.0012)	Variance of perm. shocks
Var α	0.1392 (0.0143)	0.1356 (0.0178)	Initial variance
Var ϵ	0.0882 (0.0048)	0.0669 (0.0097)	Variance of trans. shocks
Leverage	1.3449 (0.0488)	1.2645 (0.0540)	

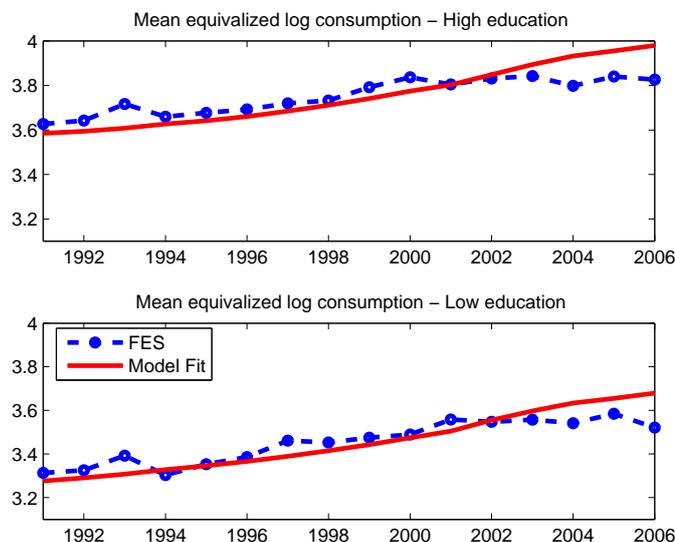
Table 3.5 shows the empirical and simulated moments, and the contribution of the distance to the criterion function. Figure 3.11 shows the fit of mean equivalized consumption growth over 1991-2006. Growth in mean equivalized consumption and wealth holdings have the largest influence on the criterion function. There is a tension between the two, met by estimation of β . The simulated growth in mean consumption is too high, implying β should be lower (less patience). On the other hand, simulated wealth holdings when the cohort is aged 40, in 1995, are too small, implying that β should be higher (more patience). Simulated wealth then overshoots the empirical wealth holdings in 2005 implying β should be lower, but this has less effect on the criterion function. A formal treatment of home-ownership would likely fit these data better. For a given β , households would save early to try to buy a house (simulated wealth holdings in mid-life would be higher). Once a house is purchased, wealth accumulation would slow down, so later wealth holdings would be little affected.

Table 3.5: Model Fit

Variable	Empirical Moment	Simulated Moment	Weighted Distance
High Education			
Var(ln Y_{1991})	0.295	0.286	0.432
...
Var(ln Y_{2006})	0.401	0.374	1.422
$E_t(\Delta \ln C_t)$	0.017	0.025	49.778
Var η	0.014	0.006	1.881
Var ϵ	0.088	0.088	0.006
$E_i \psi_{1995}^O$	-0.048	-0.025	11.067
$E_i \psi_{2000}^O$	-0.059	-0.052	1.136
$E_i \psi_{2005}^O$	-0.043	-0.104	4.502
$E_i \psi_{1995}^H$	0.144	0.098	37.706
$E_i \psi_{2000}^H$	0.197	0.202	0.143
$E_i \psi_{2005}^H$	0.33	0.404	8.403
Low Education			
Var(ln Y_{1991})	0.283	0.278	0.159
...
Var(ln Y_{2006})	0.398	0.358	2.508
$E_t(\Delta \ln C_t)$	0.015	0.027	76.281
Var η	0.017	0.008	2.463
Var ϵ	0.075	0.067	0.906
$E_i \psi_{1995}^O$	-0.045	-0.014	20.026
$E_i \psi_{2000}^O$	-0.043	-0.034	1.649
$E_i \psi_{2005}^O$	-0.037	-0.073	1.559
$E_i \psi_{1995}^H$	0.117	0.068	43.771
$E_i \psi_{2000}^H$	0.138	0.162	2.531
$E_i \psi_{2005}^H$	0.289	0.349	5.569

Notes: Var η is the variance of permanent shocks to income, Var ϵ is the variance of transitory shocks to income. ψ^H is the share of lifetime wealth in housing, ψ^O the share in other financial assets (including mortgages). See appendix A3.1 for more details.

Figure 3.11: Mean Equivalized Consumption Growth



3.6 Alternative Explanations for the Observed Phenomena

One could try to explain the observed phenomena in other ways. Here I propose a few candidates and briefly discuss why they fail to explain the profiles.

First I claim that no model from the heterogeneous-agent macro literature can explain these time series. The closest fit comes perhaps from the unitary model of life-cycle consumption and labour supply in Kaplan (2010). In fact, in his model the covariance of earnings (which I take as the equivalent income concept) and consumption declines towards the end of the life-cycle while the variance of consumption continues to grow. These trends are due to increasing wealth effects over the life-cycle. Nevertheless, this feature of his model arises from a life-cycle trend, while I document the divergence both in cross-section and for a particular cohort, and as an acute episode.

Such a particular and acute episode could only be caused by a brief change in the size of income shocks, according to such a model. The decrease in the variance of income suggests a decline in the variance of transitory shocks. Assuming that households are generally well-insured against transitory shocks, this decline in the variance of shocks would only impact the variance of consumption and the covariance if consumption and leisure are non-separable. Even with non-separabilities the variance of consumption would necessarily decline with the variance of income.⁵¹ This contradicts the cross-sectional

⁵¹This is because transitory shocks add orthogonal variation to wage and earnings dispersion. Whether consumption and leisure are complements or substitutes, such orthogonal variation induces orthogonal

evidence.

Turning to hypotheses related to the growth in house prices, it is worth briefly considering two other purported links between house price growth and consumption. First, it is argued that house prices drive consumption growth not through a wealth effect but through the alleviation of credit constraints (Muellbauer and Murphy (1997)). If this were relevant here, we would expect the variance of consumption to decline when house prices grow, due to a reduction in numbers of households at a binding constraint.

Alternatively it is suggested that house price growth is caused by increased income expectations, which also drive consumption growth. One can imagine a link to growth in the variance of consumption through, for example, a model such as Guvenen (2009). Suppose, as in Guvenen, that agents have heterogeneous income trends. If those at the top of the distribution receive a boost to their income expectations, this would drive an increase in the variance of consumption coincident with equilibrium growth in house prices. However, income inequality was flat or declining over this period.

3.7 Conclusions

I document an empirical puzzle, that for the population as a whole, the covariance between log income and log consumption declined over the late 1990s and early 2000s in the UK, while the variance of log consumption increased. This implies contradictory profiles for the evolution of differences in permanent income. When stratifying the sample by education and for a particular cohort (those born in the 1950s) I find that both the variance of consumption and its covariance with income declined over the relevant period for the low education group but remained relatively stable for the high education group. Nevertheless for both groups there remained a puzzling divergence between the variance of consumption and the covariance with income.

I explain this episode by accounting for two important features of the UK economy over the period and introducing these into an otherwise standard consumption and savings model. First, the new Labour government, elected in 1997, increased the generosity of benefits in a sequence of measures over 1998-2003 which compressed the distribution of variation in the consumption distribution. A decline in transitory variances therefore induces a decline in the variance of consumption.

income. Introducing stochastic changes to the benefit regime explains the simultaneous decline in the variance of income with the covariance of income and consumption and induces a smaller decline in the variance of consumption. Second, the UK experienced a strong boom in house prices over 1996-2007 in which real prices grew by 130% nationwide. Introducing house price shocks into the model induces a growth in the variance of consumption separate from the other moments and further explains its decoupling from the covariance with income.

I introduce these features into an otherwise-standard consumption and savings model, and estimate against data from the FES and BHPS using the method of simulated moments. I find that the model explains the features of interest well: the benefit reforms affected the low education group particularly strongly, while the effect of house price growth was comparatively modest and affected both groups roughly equally.

It will be interesting to see how the house price declines after the sample period affected consumption inequality. Figure 3.1 shows that consumption inequality dipped in 2008, and, in fact, converged with the covariance with income. At the time of writing, house price growth has been flat after the decline in 2008 and brief bounce back in 2009. If house prices move significantly in the near future it would provide an interesting test of the hypotheses presented. It would also be interesting to examine the effect of the house price boom in other countries where the appropriate data are available.

A3 Appendix to chapter 3

A3.1 Computing Asset Moments Using the BHPS

This appendix documents how I compute π_t , ψ_t^H , ψ_t^O for $t \in \{1995, 2000, 2005\}$. π_{it} is defined as $\frac{\text{Discounted Labour Income}}{\text{Wealth+Discounted Labour Income}}$ for a household indexed by i . ψ_{it}^H is defined as $\frac{\text{Household Wealth}}{\text{Wealth+Discounted Labour Income}}$. ψ_{it}^O is defined as $\frac{\text{Other Wealth}}{\text{Wealth+Discounted Labour Income}}$. We compute total household wealth as the sum of net housing wealth and financial wealth of the head and spouse. We ignore pension wealth in our baseline estimates because this is very illiquid and it is unlikely households can borrow against this in the case of an adverse shock. Also we do not have data on pensions for 2005. We ignore financial wealth of other tax units because this is unlikely to be used to insure head and spouse shocks. We then compute expected future income by the following procedure. First we restrict the sample to households headed by a couple, in order to eliminate multi-tax unit households. We then estimate permanent income by averaging income at time t , $t+1$ and $t+2$, to smooth measurement error and transitory shocks. We take the twice-forward income rather than $t-1$ income, because later we calculate the covariance with time $t-1$ income. We assume constant future net income until the head mandatorily retires at 65, then no labour income thereafter. We discount this income stream at the rate of the expected return on housing, 3.4%pa. We perform robustness checks against all these assumptions; they change the results little.

Table 3.3 displays estimates of $\text{Cov}(\psi_t^H, \ln y_{t-1}^P)$, which is an input into equation 3.15. Of course we do not observe y_{t-1}^P , only y_{t-1} . Theory suggests that wealth stores should be more positively correlated with lagged actual income than with permanent income, because transitory fluctuations will be mainly stored as wealth. If $\text{Cov}(\psi_t^H, \ln y_{t-1})$ deviates from $\text{Cov}(\psi_t^H, \ln y_{t-1}^P)$ for this reason, then estimating the latter by averaging $\ln y_{t-1}$ and $\ln y_{t-2}$ or by instrumenting $\ln y_{t-1}$ with $\ln y_{t-2}$ will not help, because transitory income in $t-2$ will cause similar biases. Simulations from simple consumption and saving dynamic programmes suggest that while the covariance of income and wealth is similar in magnitude to the covariance of consumption and wealth, the covariance of wealth with permanent income is zero. Therefore, for the approximate estimates we impose $\text{Cov}(\psi_t^H, \ln y_{t-1}^P) = 0$, but check robustness by estimating with $\text{Cov}(\psi_t^H, \ln y_{t-1}^P) = \text{Cov}(\psi_t^H, \ln y_{t-1})$.

Finally, we require an income elasticity of food consumption, because the BHPS does not contain data on total expenditure but only food consumed within the house. We therefore also estimate a food demand equation using the FES data. We pool the data over the sample period (1990-2007) and regress food expenditure on: the relative price of food, household size, head's age, year dummies and total expenditure instrumented by asset income. We estimate an elasticity of 0.35. When estimating the main model we ignore sample correlation between estimates of the food demand elasticity and the inequality moments.

A3.2 Approximating Changes to the Covariance Structure of Income and Consumption in the Presence of Asset Price Shocks

I derive an expression for changes to the covariance structure of consumption and income in the presence of income and asset risk. The proof follows that in Blundell, Low, and Preston (2008a) (henceforth referred to as BLP). My derivation is conceptually very similar and requires only minor technical changes. I give the derivation here in reasonable detail for completeness. I follow the following plan: first I sketch the key ideas; second I present a stripped down version of the model displayed in section 3.2, and finally I show that the mechanics of the derivation work in the same way to BLP while emphasizing the parts which differ.

Sketch Proof

The proof revolves around equating the consumption account and the income account of the distribution of (the log of) future life-time resources. To derive a relationship between the shocks to consumption and income I then take the following steps:

1. I take a Taylor-type expansion of the distribution of future resources around expected resources and period-by-period innovations.
2. By taking the difference between expectations at time t and $t-1$, I generate expressions for innovations to future resources first in terms of (percentage) consumption innovations, then in terms of (percentage) income innovations. To first order, the equality between the two takes a simple and attractive form.

3. Finally, I bound the size of the higher-order terms to show that the first-order terms can indeed be approximately equated.

The Model

I now specify a stripped-down version of the model used in section 3.2. In this subsection I suppress i subscripts to make clear that asset returns could be idiosyncratic, or common across groups of households.

Households are born at time $t = 0$, work until $t = T_w$ and die at time $t = T$. The household maximises lifetime utility:

$$V_t(A_t, P_t) = \max_{\{C_k(A_k, P_k)\}_{k=t}^T} \mathbb{E}_t \left(\sum_{k=t}^T \beta^{k-t} \ln(C_k) \right)$$

where β is a subjective discount factor, assumed to be common across households. I ignore deterministic changes to consumption needs here for simplicity. These could be re-introduced and would affect the (common) gradient on consumption growth. The value function is homothetic, so the state space could be rewritten as one variable: $\frac{A_{it}}{P_{it}}$.

We have the law of motion for assets and terminal condition:

$$A_{t+1} = \begin{cases} R_{t+1}^* (A_t - C_t) + Y_{t+1} & \text{if } t < T_w \\ R_{t+1}^* (A_t - C_t) & \text{if } t \geq T_w \end{cases}$$

$$A_{T+1} \geq 0$$

with the following process for asset returns:

$$R_t^* = s^H R_t^H + (1 - s^H) R_t^O$$

$$\begin{bmatrix} R_t^H \\ R_t^O \end{bmatrix} \sim \log\text{-}N \left(\begin{bmatrix} \mu_H \\ \mu_O \end{bmatrix}, \begin{bmatrix} \sigma_H^2 & \rho_{HO} \\ \rho_{HO} & \sigma_O^2 \end{bmatrix} \right)$$

where s^H is the share of the portfolio invested in housing. For clarity, we also distinguish between beginning-of-period assets A_t and end-of-period assets $\mathbb{M}_t \equiv A_t - C_t$, so that the

law of motion before retirement can be written:

$$A_t = \mathbb{M}_{t-1} R_t^* + Y_t$$

The life-time budget constraint at time t can be written:

$$\sum_{s=0}^{T-t} \frac{C_{t+s}}{\prod_{k=1}^s (R_{t+k}^*)} = \sum_{s=0}^{T_w-t} \frac{Y_{t+s}}{\prod_{k=1}^s (R_{t+k}^*)} + \mathbb{M}_{t-1} R_t^*$$

Income evolves as in the standard permanent-transitory model:

$$\begin{aligned} \ln Y_t &= g_t + \ln P_t + \epsilon_t \\ \ln P_{t+1} &= \ln P_t + \eta_{t+1} \\ \eta_t &\sim N(0, \sigma_\eta^2) \quad , \quad \ln P_0 \sim N(0, \sigma_\alpha^2) \quad , \quad \epsilon_t \sim N(0, \sigma_\epsilon^2) \end{aligned}$$

such that g_t is the deterministic component of income, (later assumed common across households with the same observable characteristics).

An Approximate Consumption Growth Equation

With log preferences, the standard arguments of log-linearization apply. I now write $c_{it} \equiv \ln C_{it}$. Re-instating i subscripts we have that the change to log consumption is approximately a martingale with drift:

$$\Delta c_{it} = v_{it}^C + \Gamma_t + \mathcal{O}(\mathbb{E}_{t-1} |v_{it}^C|^2) \quad (20)$$

v_{it}^C is the innovation to consumption. For log preferences, Γ_t is constant across consumption levels and hence across consumers with the same preferences.

Approximating Lifetime Resources

As in BLP I define a function $F : \mathbb{R}^{N+1} \rightarrow \mathbb{R}$ by $F(\xi) = \ln \sum_{j=0}^N \exp(\xi_j)$. By exact Taylor expansion around an arbitrary point $\xi^0 \in \mathbb{R}^{N+1}$

$$\begin{aligned} F(\xi) &= \mathcal{K} + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0} (\xi_j - \xi_j^0) \\ &\quad + \frac{1}{2} \sum_{j=0}^N \sum_{k=0}^N \frac{\partial^2 F(\bar{\xi})}{\partial \xi_j \partial \xi_k} (\xi_j - \xi_j^0) (\xi_k - \xi_k^0) \end{aligned} \quad (21)$$

where $\mathcal{K} = \ln \sum_{j=0}^N \exp \xi_j^0$ is constant.

Approximating the Consumption Account of Lifetime Resources We now expand the consumption account of lifetime resources around $\mathcal{K}_c = \ln \sum_{j=0}^{T-t} \mathbb{E}_{t-1} \frac{C_{it+j}}{\prod_{k=1}^j R_{it+k}^*}$, the logarithm of expected discounted expenditures. Again I write $c_{it} \equiv \ln C_{it}$. I define:

$$\begin{aligned} \xi_j &= c_{it+j} - \sum_{k=1}^j \ln R_{it+k}^* \\ \xi_j^0 &= \mathbb{E}_{t-1} c_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^* \end{aligned}$$

Applying the approximation formula in 21, and taking expectations with respect to information set \mathcal{I} :

$$\begin{aligned} \mathbb{E}_{\mathcal{I}} \ln \sum_{j=0}^{T-t} \frac{C_{it+j}}{\prod_{k=1}^j R_{it+k}^*} &= \mathcal{K}_c \\ &+ \sum_{j=0}^{T-1} \theta_{it+j} [(\mathbb{E}_{\mathcal{I}} c_{it+j} - \mathbb{E}_{\mathcal{I}} \sum_{k=1}^j \ln R_{it+k}^*) - (\mathbb{E}_{t-1} c_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*)] \\ &\quad + \mathcal{O}(\mathbb{E}_{\mathcal{I}} \|v_{it}^T\|^2) \end{aligned}$$

such that:

$$\theta_{it+j} = \frac{\exp[\mathbb{E}_{t-1} c_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*]}{\sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} c_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*]}$$

are the shares of discounted consumption in total lifetime consumption and $\sum_{j=0}^{c-t+T} \theta_{it+j} = 1$ and v_{it}^T is the vector of future innovations to consumption. These formulae differ from those in BLP only in having a stochastic interest rate.

Approximating the Income Account of Lifetime Resources Similarly to above, we now expand the income account around $\mathcal{K}_y = \ln \sum_{j=0}^{T_w-t} \mathbb{E}_{t-1} \left[\frac{Y_{it+j}}{\prod_{k=1}^j R_{it+k}^*} + \mathbb{M}_{it-1} R_{it}^* \right]$, the logarithm of expected discounted incomes. I write $y_{it} \equiv \ln Y_{it}$. Letting $N = T_w - t + 1$, I define:

$$\begin{aligned}\xi_j &= y_{it+j} - \sum_{k=1}^j \ln R_{ik}^* \\ \xi_j^0 &= \mathbb{E}_{t-1} y_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{ik}^* \\ \xi_N &= \ln(\mathbb{M}_{it-1}) + \ln R_{it}^* \\ \xi_N^0 &= \mathbb{E}_{t-1} (\ln(\mathbb{M}_{it-1}) + \ln R_{it}^*)\end{aligned}$$

Applying the approximation formula in 21, and taking expectations with respect to information set \mathcal{I} :

$$\begin{aligned}\mathbb{E}_I \ln \left(\sum_{j=0}^{T_w-t-1} \frac{Y_{it+j}}{\prod_{k=1}^j R_{it+k}^*} + \mathbb{M}_{t-1} R_t^* \right) \\ &= \mathcal{K}_y \\ &\quad + \pi_{it} \sum_{j=0}^{T_w-t-1} \alpha_{t+j} [(\mathbb{E}_I y_{it+j} - \mathbb{E}_I \sum_{k=1}^j \ln R_{it+k}^*) - (\mathbb{E}_{t-1} y_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*)] \\ &\quad + (1 - \pi_{it}) [\mathbb{E}_I (\ln(\mathbb{M}_{it-1}) + \ln R_{it}^*) - \mathbb{E}_{t-1} (\ln(\mathbb{M}_{it-1}) + \ln R_{it}^*)] \\ &\quad + \mathcal{O}(\mathbb{E}_I \|\nu_{it}^R\|^2)\end{aligned}$$

where ν_{it}^R is the vector of future innovations to income and:

$$\begin{aligned}\alpha_{t+j} &= \frac{\exp[\mathbb{E}_{t-1} \ln y_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*]}{\sum_{m=0}^{T_w-t-1} \exp[\mathbb{E}_{t-1} \ln y_{it+m} - \mathbb{E}_{t-1} \sum_{k=1}^m \ln R_{it+k}^*]} \\ \pi_{it} &= \frac{\sum_{j=0}^{T_w-t-1} \exp[\mathbb{E}_{t-1} \ln y_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*]}{\Lambda_{it}} \\ \Lambda_{it} &= \sum_{j=0}^{T_w-t-1} \exp[\mathbb{E}_{t-1} \ln y_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*] + \exp \mathbb{E}_{t-1} \ln (\mathbb{M}_{it-1} R_{it}^*)\end{aligned}$$

Intuitively, α_{t+j} is an annuitization factor for income for which $\sum_{j=0}^{T-t} \alpha_{t+j} = 1$, π_{it} is the share of human capital wealth in lifetime wealth, and Λ_{it} is total lifetime wealth.

Equating Innovations to the Consumption and Income Accounts

Due to the lifetime budget constraint, the distributions of income and consumption accounts can be equated with respect to any information set, \mathcal{I} . Applying the operator $\mathbb{E}_t - \mathbb{E}_{t-1}$ to the consumption account:

$$\begin{aligned}
(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \ln \sum_{j=0}^{T-t} \frac{C_{it+j}}{\prod_{k=1}^j R_{it+k}^*} &= \sum_{j=0}^{T-t} \theta_{it+j} [(\mathbb{E}_t c_{it+j} - \mathbb{E}_t \sum_{k=1}^j \ln R_{it+k}^*) - (\mathbb{E}_{t-1} c_{it+j} - \mathbb{E}_{t-1} \sum_{k=1}^j \ln R_{it+k}^*)] + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2) \\
&= \sum_{j=0}^{T-t} \theta_{it+j} [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ c_{it+j}] + \sum_{j=0}^{T-1} \theta_{it+j} \left[(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \sum_{k=1}^j \ln R_{it+k}^* \right] + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2) \\
&= \sum_{j=0}^{T-t} \theta_{it+j} v_{it}^C + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2) \\
&= v_{it}^C + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2)
\end{aligned} \tag{22}$$

Applying the operator $\mathbb{E}_t - \mathbb{E}_{t-1}$ to the income account and rearranging:

$$\begin{aligned}
(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \ln \sum_{j=0}^{T-t} \frac{Y_{it+j}}{\prod_{k=1}^j R_{it+k}^*} &= \pi_{it} \sum_{j=0}^{T-t-1} \alpha_{t+j} [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ y_{it+j}] + (1 - \pi_{it}) [(\mathbb{E}_t - \mathbb{E}_{t-1}) \circ \ln R_{it}^*] + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2) \\
&= \pi_{it} (\eta_{it} + \alpha_t \epsilon_{it}) + \psi_{it}^H \zeta_t^H + \psi_{it}^O \zeta_t^O + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2)
\end{aligned} \tag{23}$$

where:

$$\begin{aligned}
\psi_{it} &\equiv (1 - \pi_{it}) = \frac{\exp \mathbb{E}_{t-1} \ln M_{it-1} R_{it}^*}{\Lambda_{it}} \\
\psi_{it}^H &= s^H \psi_{it} \\
\psi_{it}^O &= s^O \psi_{it}
\end{aligned}$$

Intuitively, ψ_{it}^H and ψ_{it}^O are the shares of housing and other wealth in discounted total resources.

Putting together equations 22 and 23, and inserting into equation 20 gives:

$$\Delta c_{it} = \Gamma_t + \pi_{it}(\eta_{it} + \alpha_t \epsilon_{it}) + \psi_{it}^H \zeta_t^H + \psi_{it}^O \zeta_t^O + \mathcal{O}(\mathbb{E}_I \|v_{it}^T\|^2)$$

or dropping the shock to other assets and approximating to first order:

$$\Delta c_{it} \approx \Gamma_t + \pi_{it}(\eta_{it} + \alpha_t \epsilon_{it}) + \psi_{it}^H \zeta_t^H$$

as in equation 3.13.

Deriving the Approximate Cross-Sectional Covariance Structure of Income and Consumption

Here I derive the formulae for changes to the variance of consumption and covariance of income and consumption. Now we assume that the house price shock is common to all households.

As before, $v_{i,t}^C$ is the period- t idiosyncratic shock to consumption; $v_{i,t}^{inc}$ the change to income; ζ_t^H the shock to rates of return on housing; η_{it} the permanent shock to incomes, and ϵ_{it} the transitory shock to incomes. Dropping i subscripts on household-level variables:

$$v_{i,t}^C = \pi_{it}(\eta_{it} + \alpha_t \epsilon_{it}) + \psi_{it}^H \zeta_t^H$$

$$v_{i,t}^{inc} = \eta_t + \Delta \epsilon_t$$

We make frequent use of the following formula from Goodman (1960), that for independent variables:

$$\text{Var}(xy) = \mathbb{E}^2(x)\text{Var}(y) + \mathbb{E}^2(y)\text{Var}(x) + \text{Var}(x)\text{Var}(y)$$

We also frequently use the formula in Bohrnstedt and Goldberger (1969), that for any three variables:

$$\text{Cov}(xy, v) = \mathbb{E}(x)\text{Cov}(y, v) + \mathbb{E}(y)\text{Cov}(x, v) + \mathbb{E}(\tilde{x}\tilde{y}\tilde{v})$$

Where $\tilde{z} = z - \mathbb{E}z$. Specifically, if one of v, x or y has zero mean and is independent of the other two, then $\text{Cov}(xy, v) = 0$.⁵²

⁵²Suppose, for example, x is independent of the other two and has zero mean. Then $\mathbb{E}(x) = \text{Cov}(x, v) = 0$

Deriving an Expression for $\Delta\text{Var}(c_t)$ By simple re-arrangement:

$$\begin{aligned}\Delta\text{Var}(c_t) &= \text{Var}(c_{t-1} + \Delta c_t) - \text{Var}(c_{t-1}) \\ &= \text{Var}(\Delta c_t) + 2\text{Cov}(c_{t-1}, \Delta c_t)\end{aligned}$$

Then:

$$\begin{aligned}\text{Var}(\Delta c_t) &= \text{Var}(\pi_t(\eta_t + \alpha_t\epsilon_t) + \psi_t^H \zeta_t^H) \\ &= (\bar{\pi}_t^2 + \text{Var}(\pi_t))(\text{Var}(\eta_t) + \alpha_t^2 \text{Var}(\epsilon_t)) + (\zeta_t^H)^2 \text{Var}(\psi_t^H)\end{aligned}$$

where the last line follows from the first by application of Goodman's formula.

$$\begin{aligned}2\text{Cov}(c_{t-1}, \Delta c_t) &= 2\text{Cov}(c_{t-1}, \pi_t(\eta_t + \alpha_t\epsilon_t) + \psi_t^H \zeta_t^H) \\ &= 2\text{Cov}(c_{t-1}, \pi_t(\eta_t + \alpha_t\epsilon_t)) + 2\text{Cov}(c_{t-1}, \psi_t^H \zeta_t^H) \\ &= 2\zeta_t^H \text{Cov}(c_{t-1}, \psi_t^H)\end{aligned}$$

by application of Bohrstedt's formula. Putting the terms together we have:

$$\begin{aligned}\Delta\text{Var}(c_t) &= (\bar{\pi}_t^2 + \text{Var}(\pi_t))(\text{Var}(\eta_t) + \alpha_t^2 \text{Var}(\epsilon_t)) + (\zeta_t^H)^2 \text{Var}(\psi_t^H) \\ &\quad + 2\zeta_t^H \text{Cov}(\psi_t^H, c_{t-1})\end{aligned}\tag{24}$$

as in equation 3.16.

Deriving an Expression for $\Delta\text{Cov}(c_t, y_t)$: We have:

$$\begin{aligned}\Delta\text{Cov}(c_t, y_t) &= \text{Cov}(c_{t-1} + \Delta c_t, y_{t-1} + \Delta y_t) - \text{Cov}(c_{t-1}, y_{t-1}) \\ &= \text{Cov}(c_{t-1}, \Delta y_t) + \text{Cov}(\Delta c_t, y_{t-1}) + \text{Cov}(\Delta c_t, \Delta y_t)\end{aligned}$$

and $E(\tilde{x}\tilde{y}\tilde{v}) = E(x\tilde{y}\tilde{v}) = E(E(x\tilde{y}\tilde{v}|\tilde{y}\tilde{v})) = E(\tilde{y}\tilde{v}E(x|\tilde{y}\tilde{v})) = 0$. The argument follows similarly if either y or v is independent of the other two and has zero mean.

Looking at each term in sequence:

$$\begin{aligned}\text{Cov}(c_{t-1}, \Delta y_t) &= \text{Cov}(c_{t-1}, \eta_t + \Delta \epsilon_t) \\ &= \text{Cov}(c_{t-1}, -\epsilon_{t-1}) \\ &= -\bar{\pi}_{t-1} \alpha_{t-1} \text{Var}(\epsilon_{t-1})\end{aligned}$$

And for the 2nd term:

$$\begin{aligned}\text{Cov}(\Delta c_{it}, y_{t-1}) &= \text{Cov}(\pi_t(\eta_t + \alpha_t \epsilon_t) + \psi_t^H \zeta_t^H, y_{t-1}) \\ &= \text{Cov}(\pi_t(\eta_t + \alpha_t \epsilon_t), y_{t-1}) + \zeta_t^H \text{Cov}(\psi_t^H, y_{t-1})\end{aligned}$$

Finally:

$$\begin{aligned}\text{Cov}(\Delta c_{it}, \Delta y_{it}) &= \text{Cov}(\pi_t(\eta_t + \alpha_t \epsilon_t) + \psi_t^H \zeta_t^H, \eta_t + \Delta \epsilon_t) \\ &= \text{Cov}(\pi_t \eta_t, \eta_t) + \text{Cov}(\pi_t \alpha_t \epsilon_t, \Delta \epsilon_t) + \text{Cov}(\psi_t^H \zeta_t^H, \Delta \epsilon_t) \\ &= \bar{\pi}_t \text{Var}(\eta_t) + \bar{\pi}_t \alpha_t \text{Var}(\epsilon_t) - \zeta_t^H \text{Cov}(\psi_t^H, \epsilon_{t-1})\end{aligned}$$

Putting this together we get:

$$\begin{aligned}\Delta \text{Cov}(c_t, y_t) &= \bar{\pi}_t \text{Var}(\eta_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)] + \zeta_t^H \text{Cov}(\psi_t^H, \ln y_{t-1}) - \zeta_t^H \text{Cov}(\psi_t^H, \epsilon_{t-1}) \\ &= \bar{\pi}_t \text{Var}(\eta_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(\epsilon_t)] + \zeta_t^H \text{Cov}(\psi_t^H, \ln y_{t-1}^P)\end{aligned}\quad (25)$$

where $\ln y_t^P = \ln y_t - \epsilon_t$ is log permanent income, as in equation 3.15.

A3.3 Deriving the Approximate Income and Consumption Moments in the Presence of Benefit Reform

I derive an expression for changes to the covariance structure of consumption and income when the income distribution is compressed or expanded exogenously. The proof is similar to that in appendix section A3.2 and to the derivation in Blundell, Low, and Preston (2008a), so I omit most of the details. I specify the model in full, then sketch the path to obtain the expressions in the main body of the text.

The Model

For completeness, I re-specify the model. It is very similar to that specified in section 3.2. However I abstract from asset risk and abstract from transitory shocks to income in order to simplify the analysis.

The household maximises discounted lifetime utility by choosing its consumption stream:

$$V_t(A_{it}, P_{it}) = \max_{\{C_{ik}(A_{ik}, P_{ik})\}_{k=t}^{c+T}} \mathbb{E}_t \left(\sum_{k=t}^T \beta^{k-t} \ln(C_{ik}) \right)$$

where β is a subjective discount factor, assumed to be common across households. Assets have the following the law of motion and terminal condition:

$$\begin{aligned} A_{it} &= (A_{it-1} - C_{it-1})R_t + Y_{it} \\ A_{i,T+1} &\geq 0 \end{aligned}$$

where R_t is a risk-free interest rate.

Income evolves according to permanent shocks, a deterministic trend and a load-factor on the stochastic component representing stretching or compression due, for example, to tax changes:

$$\begin{aligned} \ln \tilde{Y}_{it} &= \ln \tilde{Y}_{it-1} + \eta_{it} \\ \ln Y_{it} &= g_t + \theta_t \ln \tilde{Y}_{it} \\ \eta_{it} &\sim N(0, \sigma_\eta^2) \quad , \quad \ln P_{i,o} \sim N(0, \sigma_\alpha^2) \end{aligned}$$

We partially define the following stochastic process for the load factor on income:

$$\begin{aligned} \mathbb{E}_{t-1} \theta_t &= \theta_{t-1} \\ \mathbb{E}_{t-1} (\Delta \theta_t)^2 &\approx 0 \end{aligned}$$

i.e. the load factor follows a martingale with negligible second moment. This second condition is required so that households at the centre of the distribution attach the same probability distribution to income changes as those at the periphery. Henceforth I use

$$c_{it} = \ln C_{it}, y_{it} = \ln Y_{it}, \tilde{y}_{it} = \ln \tilde{Y}_{it}.$$

The Approximate Consumption Growth Equation

The arguments used in appendix section A3.2 apply identically here. Consumption growth is approximately a martingale with drift. We can then approximate the innovation to the consumption account and to the income account of lifetime resources. The discounted lifetime percentage innovation to the consumption account is approximately the time- t percentage innovation to consumption. The discounted lifetime percentage innovation to the income account is $\pi_{it}(\tilde{y}_{t-1}\Delta\theta_t + \theta_{t-1}\eta_{it})$. Ignoring the permanent shocks to latent income we derive the following approximate consumption growth equation:

$$\Delta c_{it} = \Gamma_t + \pi_{it}\tilde{y}_{it-1}\Delta\theta_t + \mathcal{O}(\|v_{it}^{inc,T_w}\|^2)$$

where Γ_t is a gradient reflecting, for example, discounting and intertemporal substitution, and v_{it}^{inc,T_w} is the vector of future innovations to income. To first-order this can be expressed:

$$\Delta c_{it} \approx \Gamma_t + \pi_{it}\tilde{y}_{it-1}\Delta\theta_t \quad (26)$$

as in equation 3.10.

Deriving the Approximate Cross-Sectional Covariance Structure of Income and Consumption

Examining the change to consumption inequality

$$\begin{aligned} \Delta \text{Var}(c_t) &= \text{Var}(\Delta c_t) + 2\text{Cov}(c_{t-1}, \Delta c_t) \\ &\approx \mathcal{O}((\Delta\theta_t)^2) + 2\text{Cov}(c_{t-1}, \pi_t\tilde{y}_{t-1}\Delta\theta_t) \\ &= 2\Delta\theta_t [\text{Cov}(c_{t-1}, \pi_t\tilde{y}_{t-1}) + \text{E}_i(\hat{c}_{t-1}\hat{\pi}_t\tilde{y}_{t-1})] + \mathcal{O}((\Delta\theta_t)^2) \end{aligned}$$

where $\hat{c}_{t-1} = c_{t-1} - E_i(c_{t-1})$ and $\hat{\pi}_t \equiv \pi_t - E_i(\pi_t)$. Now examining the change to the covariance:

$$\begin{aligned}
\Delta \text{Cov}(c_t, y_t) &= \text{Cov}(\Delta c_t, \Delta y_t) + \text{Cov}(c_{t-1}, \Delta y_t) + \text{Cov}(\Delta c_t, y_{t-1}) \\
&\approx \mathcal{O}((\Delta \theta_t)^2) + \text{Cov}(c_{t-1}, \tilde{y}_{t-1} \Delta \theta_t) + \text{Cov}(\pi_t \tilde{y}_{t-1} \Delta \theta_t, \theta_{t-1} \tilde{y}_{t-1}) \\
&= \mathcal{O}((\Delta \theta_t)^2) + \Delta \theta_t \text{Cov}(c_{t-1}, \tilde{y}_{t-1}) + \theta_{t-1} \Delta \theta_t \text{Cov}(\pi_t \tilde{y}_{t-1}, \tilde{y}_{t-1}) \\
&= \mathcal{O}((\Delta \theta_t)^2) + \Delta \theta_t \text{Cov}(c_{t-1}, \tilde{y}_{t-1}) + \theta_{t-1} \Delta \theta_t [\bar{\pi}_t \text{Var}(\tilde{y}_{t-1}) + E(\hat{\pi}_t \tilde{y}_{t-1})]
\end{aligned}$$

Ignoring the terms of order $(\Delta \theta_t)^2$ we have:

$$\begin{aligned}
\Delta \text{Var}(c_t) - \Delta \text{Cov}(c_t, y_t) &\approx \Delta \theta_t (2\bar{\pi}_t \text{Cov}(c_{t-1}, \tilde{y}_{t-1}) + 2E(c_t, \hat{\pi}_t \tilde{y}_{t-1}) - \theta_{t-1} \text{Cov}(c_{t-1}, \tilde{y}_{t-1}) \\
&\quad - \bar{\pi}_t \theta_{t-1} \text{Var}(\tilde{y}_{t-1}) - E(\hat{\pi}_t \tilde{y}_{t-1}))
\end{aligned}$$

We can normalize θ_{t-1} to 1. Furthermore $E(c_t, \hat{\pi}_t \tilde{y}_{t-1}) \approx 0$ and $E(\hat{\pi}_t \tilde{y}_{t-1}) \approx 0$ because consumption and saving decisions are homothetic in this model, so π_t covaries little with permanent income. Simplifying accordingly we get:

$$\Delta \text{Var}(c_t) - \Delta \text{Cov}(c_t, y_t) \approx \Delta \theta_t ((\bar{\pi}_t - 1) \text{Cov}(c_{t-1}, \tilde{y}_{t-1}) + \bar{\pi}_t (\text{Cov}(c_{t-1}, \tilde{y}_{t-1}) - \text{Var}(\tilde{y}_{t-1}))) \tag{27}$$

as in equation 3.11.

Chapter 4

The Transmission of Permanent Income Shocks: Evidence from the UK

4.1 Introduction

For a young family embarking on life together, the future is riddled with uncertainties. There is uncertainty about the success of marriage and future family. Uncertainty about future health. Uncertainty, even, about where the family will end up living. Among these uncertainties, probably the greatest concerns the families' disposable resources. Who in the family will work? How successful will they be? Do they have a secure job for life or will they be forced to move firms or even sectors? Will they endure spells of unemployment? The evolution of income will be critical in determining lifetime well-being.

As critical in determining lifetime well-being will be how much these income fluctuations transmit into consumption. If there is little consumption response because, for example, households can insure themselves, then the fluctuations will be benign. Different consumption responses are implied by different models of intertemporal allocation. The extreme models of autarkic consumption at one end and full insurance at the other have long been rejected.¹ Researchers have more recently focused on the spectrum of consump-

¹Jappelli and Pistaferri (2006) for example reject the myopic or rule-of-thumb model in which households consume income in each period. They also reject full insurance. See also Attanasio and Davis, 1996.

tion models in between. Attanasio and Pavoni (2010), for example, examine consumption insurance with private information. In their set-up, households can trade the full range of Arrow securities, but they are denied full insurance because hidden information about productivity and savings induces partial market failure.²

Blundell et al. (2008b) (henceforth BPP) provide an authoritative empirical assessment of the consumption response. They study the inequality boom of the 1980s in the US and estimate that 65% of permanent shocks transmit through to consumption. This represents slightly lower transmission than is generated by simple consumption and savings models using plausible parameters (Kaplan and Violante, 2010 and Carroll, 2009). In keeping with Attanasio and Pavoni and the related literature, BPP interpret this finding as evidence of *partial* insurance: households can insure themselves against shocks more than is achievable using purely their own wealth. BPP further use this estimate to account for the co-evolution of income and of consumption inequality in the US over their sample period.

In this chapter, I study income risk and its transmission in the UK over 1991-2006. Over and above providing independent estimates of key parameters, I have two specific motivations. First, as documented in chapter 2 and by Heathcote et al. (2010a), the 1980s was atypical, with a substantial rise in inequality in both the UK and the US, because of widespread structural change in both countries. These authors further report that inequality grew more slowly in both nations in the 1990s. It therefore seems important to examine income risk and its transmission in this later period. BPP only study the period until 1992 and the PSID offers yearly consumption data only until 1996. Therefore, the BHPS dataset in the UK provides an opportunity for this study. Second, and more pertinently, chapter 2 documents that permanent income risk in the UK in the 1990s and early 2000s was substantial, when measured by panel data on incomes. Meanwhile the slow growth in cross-sectional consumption inequality implies that consumption risk was much lower.^{3 4} The cross-sectional evidence either contradicts the evidence from the

²Related papers in this literature include Krueger and Perri (2006), who examine insurance when households can walk away from mutually-agreed contracts. Golosov and Tsyvinski (2007) use a set up similar to Attanasio and Pavoni to study optimal public insurance.

³Attanasio and Jappelli (2001) use this moment to identify the variance of shocks to the marginal utility of wealth. Blundell and Preston (1998) use this moment to identify the variance of permanent income shocks. In effect, they assume that all shocks to marginal utility come from income and that permanent income shocks transmit fully into consumption.

⁴Chapter 3 argues that other risks (the house price boom and the large increase in government redistribution) are important in driving the movements in inequality. I can ignore such considerations here: that

income panel or indicates that income shocks had little effect on consumption.

I estimate the transmission of permanent income shocks through to consumption to be 0.49. This is lower than BPP's estimate of 0.64. Permanent income risk is comparable to that in BPP: I estimate the variance of permanent shocks for my sample to be 0.019 (14% standard deviation of yearly shocks) compared to 0.020 in BPP. From the repeated cross-section, I also estimate the variance of consumption shocks from the repeated cross section to be 0.0055 (7.5% standard deviation of yearly shocks). These results have three immediate and important implications. First, the implied contribution from income risk is around 80% of total consumption risk. Income shocks therefore provide the bulk of shocks to consumption.⁵ Second, the risk to consumption induced by income shocks is about half that in the BPP sample. Insofar as we can compare risk across countries, this provides fresh evidence that consumption risk was lower in the 1990s than in the 1980s. Third, the results imply that focusing on the variance of permanent income shocks as a measure of the cost of risk is misleading. If the cost of risk is proportional to the variance of shocks (as in Lucas, 1987), then ignoring consumption smoothing overstates the cost measure by a factor of 4.

Among other results, I find little difference in the transmission of permanent shocks across the subgroups I study, although the estimates are too imprecise to make any firm conclusions. Transitory shocks are almost completely smoothed through borrowing and saving. I find that the level of permanent income risk seems to follow a U shape over the life-cycle, indicating that the precautionary saving motive is particularly high early in working life. I estimate that head earnings shocks contribute about half of the consumption risk induced by total household income. This estimate implies a sizable contribution from other channels: labour supply of other household members, asset income and changes to taxes and benefits. On the other hand, an examination of head wage shocks implies that head labour supply neither amplifies nor dampens permanent wage shocks.

As mentioned, the central estimate of 0.49 on the transmission of permanent shocks

chapter notes that these other factors do not substantially change the estimates of income risk because they largely cancel each other out. Subsection 4.4.5 further compares the approach and results in the current chapter to those in chapter 3.

⁵Of course, consumption can change for many reasons other than revisions to life-time resources: there are also taste shifts and credit constraints. But most interesting is consumption changes owing to shocks to the marginal utility of wealth. This is my definition of 'consumption risk' for the remainder of the chapter.

is lower than is predicted by models of self-insurance. Following BPP, one interpretation is that households achieve substantial extra insurance through other channels. However given that my definition of income includes all transfers and gifts, it is difficult to think of extra mechanisms which could bridge the gap between self-insurance and the observed transmission. Moreover I find that transmission seems particularly low early in the life-cycle. These findings can be plausibly explained by two other hypotheses: first, if young households have advance information about career choice and career success (as advocated by Cunha et al. (2005)), but later income fluctuations reflect more unforeseen news, then the transmission parameter should be lower in early working life. Second, if income shocks are not permanent but only very persistent (as advocated, for example, by Guvenen (2007)), then again, the transmission of income shocks would be lower in early life. I find no direct evidence in favour of either alternative hypothesis, although Guvenen (2009) argues that the tests employed here on typical household surveys lack sufficient power to isolate the correct income process.

It is important to stress that, even though I estimate a model without advance information and with a unit root (a permanent shock), my estimates of transmission are robust to misspecification in either direction. One corollary is that the estimate of the consumption risk induced by income shocks is more robust than the estimates of income risk itself. Unfortunately, the interpretation of the results in terms of insurance is not robust to misspecification in either direction.⁶ We need to identify advance information and the degree of persistence on shocks in order to infer insurance from transmission. Given the difficulties discussed, therefore, in the rest of the chapter I interpret the results in terms of transmission rather than insurance.

The questions posed in this chapter can only be answered using panel data on both income and consumption. However, the quality of consumption data in panel data sets is generally poor. BPP (documented further in Blundell et al. (2004)) have made advances in the use of the PSID for the US and argue persuasively that the consumption data under their treatment give reliable results. I follow the spirit of their techniques here: I use repeated cross-sectional data on both food and other non-durable consumption from the FES and panel data on food expenditure and incomes from the BHPS. I combine these to

⁶See Meghir and Pistaferri (2010) for a further discussion of this identification problem.

infer the dynamic relationship between income and total non-durable consumption.

My final contribution, therefore, is to show that the food data in the BHPS can be used for solid economic research on consumption behaviour. Until now no related research has used these BHPS data, while numerous studies have used the data from the PSID.⁷ The BHPS data are thought to be of lower quality because they are collected in bands, and because they cover a smaller subset of consumption items.⁸ To deal with the first concern, I simply take midpoints of the banded data. I show, through a validation study using PSID data, that using midpoints is as good as a technique as any, and that the efficiency loss from banded observations over exact observations is small. This being said, the transmission of income through to consumption is estimated less precisely than in BPP's study. A sensible conclusion is that the drawback with the BHPS data is the definition of food, which has a much lower income elasticity than does the PSID definition, rather than the fact that they are in bands.

Besides the papers already mentioned, this study fits into a long literature examining consumption and income dynamics using microdata, going back to Hall and Mishkin (1982). In a recent and related study of households in Russia, Gorodnichenko et al. (2010) find a transmission of permanent shocks similar that in BPP. In other related research, Gorbachev (2010) argues that consumption volatility increased steadily since the late 1970s in the US. She extends the sample in BPP and uses the biennial data after 1996 to construct volatility measures until 2004. She uses a different approach to mine, concentrating on the volatility in consumption itself whereas I look at the risk induced by income. My approach can be thought of as an instrumental variables estimator which removes extraneous and less important factors such as measurement error and temporary fluctuations in expenditure.⁹ These extraneous factors seem to account for the vast majority of consumption changes.

This chapter proceeds as follows. Section 4.2 discusses the key features of the food panel data in the BHPS and describes treatment of the income data. Section 4.3 describes the model of income and consumption dynamics and the procedure for taking it to the

⁷Recent examples include Gorbachev (2010) and Guvenen and Smith Jr (2010) who use BPP's imputed data for total consumption

⁸The subset in the BHPS is food consumed in the home, whereas the PSID also includes all food consumed outside.

⁹This interpretation owes to Kaplan and Violante (2010).

data. Section 4.4 gives the key results and discusses the relationship between this chapter and chapter 3. Section 4.5 discusses further the technical details of dealing with the BHPS data and provides the validation against the PSID. Section 4.6 concludes.

4.2 Data from the BHPS¹⁰

The analysis uses data from the BHPS and FES. Other chapters of this thesis have detail on both these data sets. Here I describe just those features of the BHPS survey and my treatment that are particularly important for this specific analysis. A brief discussion of the FES data is contained in appendix A4.1.

Despite its status as the main UK household panel survey, the BHPS has limited data on consumption. The survey only contains questions about food consumed within the home and about energy use and small durables purchases such as TVs and kitchen appliances. Within this set, only for food does consumption plausibly equal expenditure. I therefore focus on these responses. In comparison, the PSID survey for the US includes food purchased outside the house. Food ‘in’ has a much lower income elasticity than food ‘out’ because high income households substitute towards restaurant meals, so the signal from changes in food consumption to total consumption and living standards is weaker than in the equivalent US analysis. Consequently any hypothesis test will have lower power than those in, say, Blundell et al. (2008b). An advantage of the BHPS data is that it covers a period over which PSID data became weak. To the extent that levels of risk and insurance reflect global changes in capital and labour markets, then these data inform about the global economy in this ‘missing’ period.

The specifics of the BHPS questions about food expenditures are as follows: the first wave of the BHPS asks ‘Thinking about your weekly food bills approximately how much does your household usually spend in total on food and groceries?’ Respondents give exact answers. Respondents include all of food, bread, milk, soft drinks etc. They also include take-aways eaten in the home. Respondents exclude pet food, alcohol, cigarettes and meals out. From wave 2 onwards spending information is collected according to 12 bands. For these waves I impute consumption to be the mid-point of each interval. For

¹⁰Thanks to Zoe Oldfield for sharing the fruits of earlier unpublished research comparing food consumption data in the FES and BHPS

the bottom interval (£0-£10) I assign £5 spending. For the top band (£160+) I assign £180. The alternative is to estimate moments of interest using maximum likelihood and assuming an underlying distribution such as the normal. As a robustness check, figure 4.1 shows a comparison of the cross-sectional mean and variance using both the midpoints and maximum likelihood estimates, together with estimates from the FES. It shows that both treatments of the BHPS data give similar results in these dimensions. Section 4.5 discusses the use of midpoints in more detail, including analysis of the autocovariances and a validation using data from the PSID.

An advantage of the BHPS over the PSID is that the timing of the questions is less problematic for analysis. While some have suggested that income and consumption measures in the PSID refer to different time periods¹¹, all relevant questions in the BHPS ask about current circumstances. Income and consumption observations are therefore likely to be synchronized.

Almost all interviews are carried out between September 1 and December 1 in each survey year (less than 10% carry on into the new year). While the gap between interviews could be a minimum of 9 months and a maximum of 15 months within this main period, I neglect this variation in timing and consider that all first differences indicate yearly changes in variables.

I use a variety of income concepts in the analysis. Wages are defined as usual earnings in the current job divided by usual hours. I remove wages and earnings that have been imputed by the BHPS compilers. These imputations are based on growth rates of the variables from similar individuals and so affect the estimation of dynamics. I obtain the measures for total household labour income and household net disposable income from an auxiliary data set (see Bardasi et al. (1999) for more documentation, and Jenkins (2010) for a discussion). For both these two variables I use current measures (usual monthly income at the time of interview) rather than annual incomes. Net disposable income is defined as the sum of earned income, asset income and transfers (public and private) minus state taxes (income tax and national insurance contributions). Capital gains, or the drawing down of capital, is excluded in this definition. Pension income, which is often derived from the drawing down of capital is included in the definition, but because my sample consists

¹¹see Hall and Mishkin (1982) for a detailed discussion.

of heads of working age, its contribution to income is small.

The sample selection proceeds as follows: I use only the core BHPS sample and ignore the low-income booster sample. Following BPP, I take only households headed by a stable and heterosexual couple in the BHPS¹², but allow for entry and exit of children. Naturally, this makes the discussion relevant to couples only. The dynamics of income and consumption for unstable households are potentially more interesting and important.¹³ I exclude households with heads aged between 25 and 65 and take only those heads born between 1940 and 1969. Finally, I exclude responses from Northern Ireland in the BHPS because they are not represented in the FES. I form an unbalanced panel by selecting households for whom the first difference of income appears at least 5 times over the course of the survey (16 years). Therefore a household appears in the covariance matrices with a minimum of 6 appearances for income, though households could conceivably appear 9 times and still be dropped from the sample. Food expenditure is almost always observed.

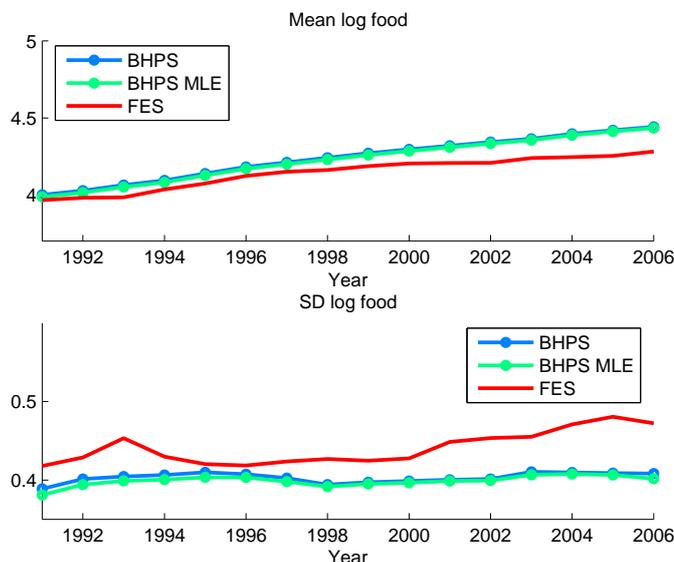
I trim the top and bottom 0.5% of the distribution of all income variables.¹⁴ I do not trim the food consumption distribution: since expenditures are assigned to 12 bands there is not the same chance of reporting implausibly high or low observations through miscoding, or omission of a component. Such trimming of the levels of income does not theoretically make a difference to the central estimates, but improves precision. The results are robust to close alternative procedures. I choose not to trim according to the growth rates of the variables. Even though trimming according to growth rates improves the precision of the estimates, simulations show that even slight trimming noticeably biases down the sample covariance between income and consumption changes. Hence such trimming biases down the estimated transmission of income shocks through to consumption. The initial sample comprises 116,111 household-year observations with 96,787 income observations. The final sample comprises 20,552 observations with 19,157 income observations.

¹²here I select on a dynamic aspect of the data. This may cause differential selection between the BHPS and the FES. Nevertheless, one would think that including only stable couples in the BHPS (and all couples in the FES) would bias estimates for permanent risk in the BHPS downward. Therefore, sample selection does not weaken the motivation behind this paper, which is that permanent risk from the time-series evidence seems higher than from the cross-sectional evidence.

¹³see, for example, Voena (2010) for an analysis of the effect of divorce on consumption, savings and labour supply

¹⁴By taking logs of all the income and expenditure variables I also remove any negative observations. These comprise 0.6% of the initial sample themselves.

Figure 4.1: Comparison of food expenditures in the FES and BHPS



Notes: ‘BHPS’ gives the statistics using the imputation described above. ‘BHPS MLE’ gives maximum likelihood estimates using the observed bands assuming that food is distributed normally.

4.3 The Model and Estimation

4.3.1 The Model of Income and Consumption Dynamics

I use a standard model of income dynamics, exactly as in Blundell et al. (2008a), and commonly referred to as the restricted income profiles model by, for example Guvenen (2009) and Hryshko (2010).¹⁵ Income is assumed to be composed of three parts: a deterministic component reflecting the lifetime shape of the wage profile and life-cycle labour supply; a stochastic permanent component evolving as a random walk, and a stochastic short-lived ‘transitory’ component, evolving as an MA(1) process. This transitory component might include measurement error, which I do not attempt to identify separately. Formally:

$$\ln Y_{it} = g_{c,e,Z,t} + \ln P_{it} + \epsilon_{it} + \theta \epsilon_{it-1}$$

$$\ln P_{it} = \ln P_{it-1} + \zeta_{it}$$

¹⁵The identification of income process here is along the lines of Meghir and Pistaferri (2004) using the test in MaCurdy (1982). Much recent work models income dynamics in more detail. See for example, Altonji et al., 2009 who provide a rich statistical specification which allows for different types of employment transition, Low et al. (2010), who allow for labour force participation frictions in a more structural setting, and Postel-Vinay and Turon (2010), who model productivity shocks to the firm and the renegotiation of employment contracts.

where $g_{c,e,Z,t}$ is the deterministic component, depending on observable characteristics such as cohort, education, demographic variables and time. P_{it} is permanent income for household i at time t , ζ_{it} is the innovation to permanent income. ϵ_{it} is the time- t innovation to transitory income (measurement error) and θ is the moving average parameter governing duration of the transitory shock. I make the usual assumptions that ϵ_{it} and ζ_{it} represent genuine time- t innovations to the household and that households can perfectly distinguish transitory from permanent shocks.

If we define $y_{it} = \ln Y_{it} - g_{c,e,Z,t}$ to be the log of the stochastic component of household income, then changes to this ‘residual’ income are given by:

$$\begin{aligned}\Delta y_{it} &= \zeta_{it} + \Delta \epsilon_{it} + \theta \epsilon_{it-1} \\ &= \zeta_{it} + \epsilon_{it} + (\theta - 1) \epsilon_{it-1} - \theta \epsilon_{it-2}\end{aligned}\tag{4.1}$$

An approximate solution to the standard household’s optimization problem is given in Blundell et al. (2008a). Defining c_{it} to be household log consumption, net of predictable components (depending mainly on demographic variables), then the approximate solution for observed consumption changes is:

$$\Delta c_{it} \approx \Gamma_t + \phi_{it} \zeta_{it} + \psi_{it} \epsilon_{it} + \xi_{it} + \Delta \nu_{it}\tag{4.2}$$

where Γ_t is a constant reflecting saving due to the discount rate, interest rates and the precautionary motive, and is constant across households within the cohort. ϕ_{it} captures the transmission of permanent shocks into consumption. ψ_{it} captures the transmission of transitory shocks into consumption, and ξ_{it} is an idiosyncratic to consumption due to, say, idiosyncratic portfolio returns. ν_{it} is measurement error; here it is modeled as classical, but we could, for example, impose an MA(1) structure. Blundell et al. (2008a) give accurate approximations under a specific specification of the asset market - only a risk free bond is available. In chapter 3, I extend this model to allow for risky assets. Here, as in BPP I allow ϕ_{it} and ψ_{it} possibly to reflect other types of insurance; for example those provided by poorly measured asset markets such as those implicit in extended family networks. Unless otherwise stated, I assume ϕ_{it} and ψ_{it} to be common across households and time

and denoted by ϕ and ψ .

The equations 4.1 and 4.2 together with the assumptions that all shocks are uncorrelated and unforeseen provide all the covariance restrictions for growth moments implied by the model. The covariance restrictions are most accessibly summarized in table 4.1. The vertical axis gives lagged and current consumption and income changes, while future and current consumption and income changes are given on the horizontal axis. The covariance of consumption and non-contemporaneous income changes is zero, while the covariance of consumption changes is non-zero only at one lag/lead; and then only because of measurement error. All covariances of variables at more than two periods' distance are zero.

Table 4.1: Theoretical Consumption and Income Covariance Moments

	Δc	Δc_{+1}	Δy	Δy_{+1}
Δc	$\phi^2\sigma_\zeta^2 + \psi^2\sigma_\epsilon^2 + \sigma_\xi^2 + 2\sigma_\nu^2$	$-\sigma_\nu^2$	$\phi\sigma_\zeta^2 + \psi\sigma_\epsilon^2$	$-(1-\theta)\psi\sigma_\epsilon^2$
Δc_{-1}		$-\sigma_\nu^2$	$-(1-\theta)\psi\sigma_\epsilon^2$	$-\psi\theta\sigma_\epsilon^2$
Δy		$\phi\sigma_\zeta^2 + \psi\sigma_\epsilon^2$	$\sigma_\zeta^2 + g(\theta)\sigma_\epsilon^2$	$-(1-\theta)^2\sigma_\epsilon^2$
Δy_{-1}		0	$-(1-\theta)^2\sigma_\epsilon^2$	$-\theta\sigma_\epsilon^2$

Notes: ϕ captures the transmission of permanent income shocks into consumption
 ψ captures the transmission of transitory income shocks into consumption
 σ_ζ^2 is the variance of permanent shocks, σ_ϵ^2 the variance of transitory shocks
 σ_ξ^2 is the variance of heterogeneous growth on consumption
 σ_ν^2 is the variance of measurement error on consumption
 θ is the MA(1) coefficient
 I define $g(\theta) \equiv 2(1-\theta+\theta^2)$ to save space in the table

4.3.2 Using Food Expenditures to Infer Consumption Choices

As discussed, I do not observe total non-durable consumption, only food consumption. In order to make inference about the response of non-durable consumption to shocks, I form a measure of 'adjusted' food as follows. I begin with a specification for food demand almost identical to that used by BPP:

$$f_{i,t} = \mathbf{W}'_{i,t}\mu + \mathbf{p}'_t\Theta + \beta(q_{i,t})c_{i,t} + e_{i,t} \quad (4.3)$$

where \mathbf{W}_i is vector of household fixed effects, \mathbf{p}_t is a vector of prices, μ and Θ are vectors of coefficients. $\beta_{q_{i,t}}$ is the income elasticity of demand for food, for group q , to which household i belongs. $e_{i,t}$ is an error term uncorrelated with total consumption and reflecting, for example, taste shocks. Appendix A4.1 discusses estimation of this equation and gives specification tests. The income elasticity is estimated to be around 0.4 for all relevant groups, principally those separated by cohort and education. Using 4.3 we can define ‘adjusted’ food as:

$$\begin{aligned}\tilde{f}_{i,t} &= f_{it} - \mathbf{W}'_{i,t}\mu + \mathbf{p}'_t\Theta \\ &= \beta(q_{i,t})c_{i,t} + e_{i,t}\end{aligned}$$

If we assume that the income elasticity does not vary much between consecutive years, then for a group with the same value of $q_{i,t}$ and hence the same income elasticity of demand:

$$\Delta c_{i,t} \approx \frac{1}{\beta_{q,t}} \left(\Delta \tilde{f}_{i,t} - \Delta e_{i,t} \right)$$

I use these equations to translate the moments in table 4.1 into moments of food changes. I now absorb variation in taste for food ($e_{i,t}$) into measurement error ($\nu_{i,t}$). The non-zero moments on the left-hand side of the table, for example, then become:

$$\begin{aligned}\text{Var}(\Delta \tilde{f}_{qt}) &= \phi^2 \beta_{q,t}^2 \sigma_\zeta^2 + \psi^2 \beta_{q,t}^2 \sigma_\epsilon^2 + \sigma_{\xi_{q,t}}^2 + 2\sigma_{\nu_{q,t}}^2 \\ \text{Cov}(\Delta \tilde{f}_{qt}, \Delta \tilde{f}_{qt+1}) &= -\sigma_{\nu_{q,t}}^2 \\ \text{Cov}(\Delta \tilde{f}_{qt}, \Delta y_t) &= \phi \beta_q \sigma_\zeta^2 + \psi \beta_q \sigma_\epsilon^2\end{aligned}$$

for group indexed by q .¹⁶

This method is styled on and closely relates to that used by BPP. It contrasts with other methods of imputing total consumption such as Skinner (1987), who regresses consumption on observable features (such as food and durables) that are present in both the panel and the cross-section, and Ziliak (1998), who uses income and changes in wealth to calculate consumption as a residual. To give further explanation for my treatment of the data it

¹⁶I can also vary the other parameters (such as ϕ , ψ etc) by group but I suppress these subscripts in the present discussion.

is useful to compare it in detail to BPP's treatment. BPP translate food demands into non-durable consumption by fully inverting equation 4.3. Blundell et al. (2004) show that this procedure preserves the mean of non-durable consumption and replicates the time-series of the variance up to an intercept shift. I do not replicate this procedure because my definition of food has a far lower income elasticity and so the denominator in the inversion is much closer to zero. When I invert fully, the error in food demands (e_{it}) is magnified much more. The variance of changes in this imputed 'non-durable' consumption is implausibly large (around 0.4) and dwarfs that from income (around 0.1).¹⁷ However, my procedure is little different to BPP's. The only substantive difference is that I cannot pool observations of households with different cross-sectional income elasticities of food demand. But I can still deploy different elasticities across time when estimating a non-stationary model, and when I estimate on different groups (such as by cohort or education), I deploy different elasticities with each group. However, appendix A4.1 shows that the elasticity does not vary significantly across groups or over time.¹⁸ A drawback of my method is that I cannot assess external validity of the procedure by comparing the distribution of imputed consumption in the BHPS with that from the FES.

In practice, when I remove the predictable components of consumption changes, as discussed in 4.3.1, I regress on a very similar vector of controls as in the demand equation.¹⁹ Therefore I do not need to impute adjusted food as an intermediate step, but instead perform one regression on observed food demands. Nevertheless methodologically, my analysis is based around a demand specification. And to emphasize, I estimate a demand equation for food in appendix A4.1 in order to derive income elasticities.

¹⁷A related point is that taste variation for food consumed in the home may be larger than for all food and non-durable consumption.

¹⁸Time-variation in the elasticity is crucial to BPP's argument. That argue that assuming a constant elasticity implies an increase in insurance over time whereas, in fact, insurance stayed constant, while the elasticity varied over time. The evolution of the elasticity over time does not appear so important to my analysis.

¹⁹I do not regress on price in this vector, but this is common across all households so has no effect on idiosyncratic variation.

4.3.3 Identification and Estimation of Risk and Transmission Parameters

Following Blundell et al. (2008b) and Kaplan and Violante (2010) I pursue the following simple and intuitive identification strategy to estimate the risk and transmission parameters. As they discuss, identification of the transmission coefficient on permanent shocks to income is best considered as a regression of Δc_{it} on Δy_{it} , with Δy_{it} instrumented by $\sum_{k=-2}^2 \Delta y_{it+k}$. The strategy works because the instrument contains only the time-t permanent shock and other shocks that do not affect time-t consumption growth. Specifically, the instrument holds time-t transitory shocks constant. Formally:

$$\begin{aligned} \text{Cov} \left(\Delta y_{it}, \sum_{k=-2}^2 \Delta y_{it+k} \right) &= \sigma_{\zeta}^2 \\ \text{Cov} \left(\Delta c_{it}, \sum_{k=-2}^2 \Delta y_{it+k} \right) &= \phi \sigma_{\zeta}^2 \\ \implies \phi &= \frac{\text{Cov} \left(\Delta c_{it}, \sum_{k=-2}^2 \Delta y_{it+k} \right)}{\text{Cov} \left(\Delta y_{it}, \sum_{k=-2}^2 \Delta y_{it+k} \right)} \end{aligned} \quad (4.4)$$

This estimator has other attractive properties: for example, as Kaplan and Violante (2010) discuss, the estimator on ϕ is robust to advance information of one period. I make three adjustments to the estimator in practice: First, I drop Δy_{it-2} and Δy_{it-1} from equation 4.4 and exploit that $\text{Cov}(\Delta c_{it}, \Delta y_{it} + \Delta y_{it+1} + \Delta y_{it+2}) = \phi \sigma_{\zeta}^2$. This works because the covariance of Δc_{it} with lagged income changes is zero under the PIH. I make this adjustment because this choice of moments is more efficient. Furthermore, using $\text{Cov}(\Delta c_{it}, \Delta y_{it} + \Delta y_{it+1} + \Delta y_{it+2})$ is robust to habit formation.²⁰ In fact, the choice makes no substantive difference because table 4.2 in the results section shows that $\text{Cov}(\Delta c_{it}, \Delta y_{it-1})$ and $\text{Cov}(\Delta c_{it}, \Delta y_{it-2})$ are insignificant and of opposite sign. Second, I adapt the estimator for the use of the unbalanced panel. When estimating the model I do not require that 6 years of consecutive observations be present. Therefore I identify $\phi \sigma_{\zeta}^2$ as $\sum_{k=0}^2 \text{Cov}(\Delta c_{it}, \Delta y_{it+k})$ and σ_{ζ}^2 as $\sum_{k=-2}^2 \text{Cov}(\Delta y_{it}, \Delta y_{it+k})$ (ie. I take the summation

²⁰ $\text{Cov}(\Delta c_{it}, \Delta y_{it-1})$ is positive if it takes more than one period for consumption to respond fully to permanent income shocks. Therefore including this moment in the estimation of ϕ will cause bias if there are habits.

outside the covariance operator). Third and finally, in my main estimation, I pool observations over all time periods. This yields reliable results because the income elasticity of food demands is almost constant over the period. In summary, my estimators for σ_ζ^2 and ϕ in terms of moments of adjusted food and income are:

$$\hat{\sigma}_\zeta^2 = \sum_{k=-2}^2 \hat{\text{Cov}}(\Delta y, \Delta y_{+k}) \quad (4.5)$$

$$\hat{\phi} = \frac{1}{\hat{\beta}} \frac{\sum_{k=0}^2 \hat{\text{Cov}}(\Delta \tilde{f}, \Delta y_{+k})}{\hat{\sigma}_\zeta^2} \quad (4.6)$$

where the sample covariances are taken across individuals and time and $\hat{\beta}$ is the average income elasticity across time for the relevant group.

Likewise I identify the transmission of transitory shocks through the regression of Δc_{it} on Δy_{it} , instrumented by Δy_{it+1} . Mirroring the case for permanent shocks, variation in Δy_{it+1} induces change in time-t transitory income and holds fixed the time-t permanent income shock. Interestingly, this identification condition works identically no matter the structure of serial correlation on transitory/short-lived income. The estimator is:

$$\hat{\psi} = \frac{1}{\hat{\beta}} \frac{\hat{\text{Cov}}(\Delta \tilde{f}, \Delta y_{+1})}{\hat{\text{Cov}}(\Delta y, \Delta y_{+1})} \quad (4.7)$$

Identification of the other parameters given in table 4.1 is less straightforward and requires minimum distance techniques. Of these, the variance of transitory shocks and the MA(1) coefficient can be identified through the income moments alone. The variance of other idiosyncratic shocks to consumption and measurement error on consumption, however, requires fitting the variance of consumption growth.

Estimation proceeds in the following distinct stages. First I estimate the food demand equation using the FES. Second, I regress food and income in the BHPS on vectors of controls to form residuals. These controls are: demographic characteristics of the household (the logs of number of adults, children under 4, children age between 5 and 11, and children aged between 12 and 18); educational attainment of the head interacted with year, regional dummies and a quartic in the head's age. Finally I estimate the parameters of interest using the covariance restrictions described in equations 4.5, 4.6 and 4.7.

4.4 Results

4.4.1 The Covariance Structure of Food and Income Changes

Table 4.2 presents the key panel data moments in two columns. The left hand column shows the covariances of changes to residual food expenditure with changes to residual disposable income. I now examine these to assess the basic consumption model and compare it to some simple alternatives. The 3rd row shows the contemporaneous covariance between consumption and income changes. This covariance is significantly positive indicating that income changes do indeed have traction on food expenditure. The 1st and 2nd rows show the covariances of consumption changes with lags of income changes. The theoretical counterparts are zero under the permanent income hypothesis. The empirical moment $\text{Cov}(\Delta \tilde{f}, \Delta y_{-1})$ can therefore be used to test two main alternative models. Under the alternative hypothesis of excess sensitivity due to, say, liquidity constraints, this moment should be negative.²¹ Under the alternative hypothesis of habit formation, this moment should be positive, because consumption takes more than one period to adjust to a permanent income shock. The empirical covariance is insignificant, indicating that neither effect is present and dominant. Rows 4 and 5 show the covariances of consumption changes with leads of income changes. The theoretical covariances corresponding to these rows are slightly negative in the current model, because transitory shocks to income induce small changes to consumption. These covariances can be positive, however, under the alternative hypothesis that households receive sufficient advance information of income shocks. The empirical moments $\text{Cov}(\Delta f_t, \Delta y_{t+1})$ and $\text{Cov}(\Delta f_t, \Delta y_{t+2})$ are again insignificant, indicating that neither the effect of transitory shocks nor advance information is present and dominant.²² BPP also fail to distinguish these covariances from zero. All the empirical covariances of food and income changes therefore support the basic model of consumption.

The income moments in the right hand column of table 4.2 display the classic features

²¹See Flavin (1981)

²²Of course both factors may be present but cancel each other out. We can test this possibility by computing the theoretical effect of transitory shocks on consumption. This is done by estimating the size of transitory shocks from income data alone and calculating their annuity value. We can then net out this effect to estimate the effect of advance information. Similarly, by computing the annuity value of a persistent shock we can identify the persistence of any AR(1) component by studying the size of $\text{Cov}(\Delta f, \Delta y_{+k})$ for $k > 0$.

of the permanent-transitory model specified. All autocovariances are significantly different from zero, except for the third lag: the key signature of a unit-root-permanent and MA(1)-transitory process.

Table 4.2: Covariances of Residual Food and Income Changes

	$\Delta \tilde{f}$	Δy
Δy_{-2}	-0.0005 (0.0009)	
Δy_{-1}	0.0007 (0.0009)	
Δy	0.0032*** (0.0009)	0.1166*** (0.0031)
Δy_{+1}	-0.0003 (0.0008)	-0.0418*** (0.0020)
Δy_{+2}	0.0007 (0.0009)	-0.0072*** (0.0016)
Δy_{+3}		0.0013 (0.0015)

Notes: Asymptotic standard errors in brackets.

4.4.2 Results From the Stationary Model

Table 4.3 shows the main estimates of permanent income risk and its transmission through to consumption. These estimates use an average income elasticity for food of around 0.4.²³ The variance of permanent shocks, at 0.019, is economically substantial. It sums to a variance of 0.56 (a standard deviation of 0.75) over a 30-year career. According to these estimates, a household has a 17% chance of more than doubling its permanent income (and the same chance of more than halving its income) over a 30-year period, relative to its expected income growth. The transmission coefficient at 0.49 is lower than BPP's central estimate of 0.64, though the standard error is too high to distinguish the two. My estimate is also lower than is suggested by a basic consumption and saving model. For example, Kaplan and Violante (2010) quantify consumption smoothing in a Bewley model with a risk free bond and suggest a theoretical transmission coefficient of around 0.8 when households have access to good credit markets. Similarly, Carroll (2009) generates

²³The estimate of the variance of permanent shocks is insensitive to variation of this elasticity. The transmission coefficient, however, is more sensitive. Although the estimates of these elasticities presented in appendix A4.1 are quite precise, I neglect this variation when computing standard errors for the main parameters. The standard error on ϕ , in particular is therefore slightly biased downwards.

an MPC out of permanent income of between 0.75 and 0.92 using a life-cycle model in partial-equilibrium. This disparity invites the obvious question: how do households smooth consumption more than is possible with just a risk-free bond? Before attempting to answer this question, remember that the income definition used includes payments from all contingent asset markets, in particular all (public and private) transfers and gifts. So it seems unhelpful to appeal to a more complex asset structure. Of course, measured income may be a poor indication of access to resources provided by, for example, extended family networks, especially for the poorest households.²⁴ I discuss further the disparity between the evidence and the self-insurance model later in this section.

Table 4.3: Key Estimates from the Pooled Model

Variable	Description	Estimate	Cons. X-Section	BPP Estimates
ϕ	Transmission of perm. shocks	0.49 (0.15)		0.64
$\text{Var}(\zeta)$	Variance of perm. shocks	0.0187 (0.0047)		0.0202
$\phi^2 \text{Var}(\zeta)$	Contribution to cons. risk from income	0.0044* (0.0027)		0.0083
	Consumption risk		0.0055(0.0013)	

Notes: Asymptotic standard errors in brackets

'Cons. X-Section' gives the average growth in the variance of log consumption from the FES

* This is equal to $0.49^2 \times 0.0187$, i.e. the first row squared times the 2nd row.

Table 4.3 also shows the implied contribution to consumption risk from income shocks, measured by $\phi^2 \text{Var}(\zeta)$. Alongside this, I present a measure of total consumption risk, the average growth in the cross-sectional variance of log consumption.²⁵ Of course, consumption risk comes from sources other than income, such as risk to health status and demographic needs. I quantify the contribution from asset (house price) risk and government policy reforms using the same data in chapter 3.²⁶ According to my estimates, the implied contribution from income shocks is around 80% of the total.

²⁴See Meyer and Sullivan (2003).

²⁵This is the moment used by Deaton and Paxson (1994) and Blundell and Preston (1998)

²⁶The contributions to consumption risk from other sources is an important and open question. Gorodnichenko et al. (2010) discuss how consumption volatility from income risk is a tiny component of total consumption volatility, as measured by $\text{Var}(\Delta c_t)$. However it is extremely difficult to unpick how much of this volatility represents permanent changes to consumption. We can estimate this permanent variation from the consumption time series alone by specifying a lag structure on measurement error and temporary taste shocks, but the estimates are highly dependent on the specification. A better strategy is to enumerate factors that might change life-time wealth, such as health or children, and to study the effect of each of these in turn on consumption.

The final column of table 4.3 shows the comparable estimates from BPP for the US over the 1980s. Both the variance of permanent shocks and the transmission coefficient are higher than the estimates from my sample. The consumption risk induced by income shocks in the BPP sample is around 90% higher, reflecting that the 1980s was a period of greater structural upheaval, higher risk and higher inequality growth across the developed world.

Table 4.4: Key Parameter Estimates: Breakdown by Sample

	Var(ζ)		ϕ		ψ	
Income - All groups	0.0187	(0.0047)	0.49	(0.15)	0.02	(0.05)
High Educ	0.0153	(0.0075)	0.42	(0.26)	0.05	(0.08)
Low Educ	0.0237	(0.0070)	0.53	(0.20)	0.02	(0.08)
Born in 1940s	0.0240	(0.0081)	0.47	(0.22)	0.01	(0.09)
1950s	0.0148	(0.0065)	0.39	(0.27)	0.01	(0.07)
1960s	0.0176	(0.0104)	0.63	(0.32)	0.05	(0.11)
Early period	0.0199	(0.0060)	0.54	(0.20)	0.00	(0.08)
Late period	0.0175	(0.0069)	0.47	(0.23)	0.04	(0.07)
Head Wage	0.0146	(0.0025)	0.32	(0.14)	0.08	(0.10)
Head Earnings	0.0156	(0.0034)	0.34	(0.14)	-0.09	(0.13)

Notes: Asymptotic standard errors in brackets

Var (ζ) is the variance of permanent shocks

ϕ is the transmission of permanent income shocks into consumption

ψ is the transmission of transitory income shocks into consumption

Table 4.4 presents estimates of the key parameters for different groups and for different income concepts. Broadly speaking the estimates are too imprecise to provide any firm conclusions on differences between groups, but I now discuss them informally.

In the first row I repeat the key estimates from table 4.3. The second and third rows separate the sample by the head's education status.²⁷ The low education group has a substantially higher variance of permanent income risk and higher transmission of shocks than the better educated group. The implied contribution to consumption risk (not shown) is 2.5 times greater for the low education group than the high education, indicating that those with low education might be in greater need of further social insurance. In the final column, I present estimates of the transmission through to consumption of transitory

²⁷I define high education as having A-levels or above. i.e. the head is educated until at least 18 years old. This comprises roughly half the sample.

income shocks. This coefficient is economically close to zero and insignificant for these groups and for all other sample breakdowns.

Rows four to six of table 4.4 show the estimates by cohort. The oldest cohort faces the highest permanent income risk. The oldest and youngest cohorts also have the highest transmission of permanent shocks through to consumption. Although the standard errors are too large to make firm statements, the point estimates conflict with the basic life-cycle model of self-insurance. In the basic model, the transmission of income shocks is governed by the size of asset holdings relative to human capital wealth. Because households accumulate assets until retirement, the transmission coefficient should therefore decline correspondingly over the working life. I discuss these results further in the next subsection.

The next two rows of table 4.4 show results when I split the sample period into two halves. The estimated variance of permanent shocks is slightly higher in the first half, which includes the recession of the early 1990s. The transmission of permanent shocks is again estimated imprecisely, but it seems, as in BPP's analysis, that insurance is stable over the survey period. The final rows show the estimates when replacing household income with head wages and head earnings. The transmission coefficients should now be thought of as estimates from a factor model of consumption changes along the lines of Altonji et al. (2002).²⁸ The variance of permanent shocks to wages and to earnings are lower than to net disposable income for this sample of stable households. The implication, therefore, is that the positive contribution from other components of income, such as asset income and other labour income, is greater than the negative contribution to the variance of disposable income shocks from tax-and-benefit progressivity. Moreover, the transmission coefficients on permanent wage and earnings are lower than on disposable income. Combining the variance of permanent shocks and their transmission implies that head wage risk accounts for only about a half of the contribution to consumption risk from total income risk.

4.4.3 The Age Profile of Income Risk

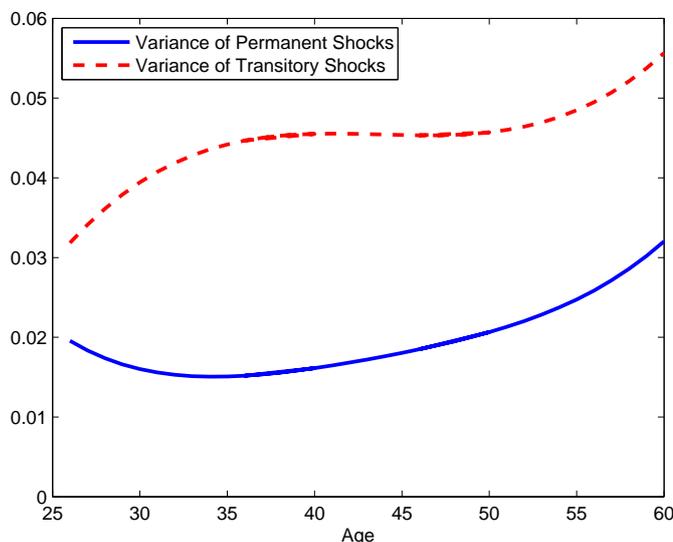
Chapter 2 documents that, apart from during the recession in the early 1990s, the variance of permanent shocks is more-or-less constant over the sample period. Given the absence

²⁸In this case there is no underlying theoretical model of consumption because we are not closing the budget constraint.

of strong time effects, the period therefore seems a useful one to study age effects in the processes of interest: the variance of income shocks, and the transmission coefficients.

Figure 4.2 shows smoothed profiles of the variances of permanent and transitory shocks by age. I form these plots by computing the yearly variances by minimum distance, then fitting a quadratic polynomial. I do this for each cohort and use the income data alone. This procedure uses the non-stationary equivalents of the bottom right hand corner of table 4.1. I estimate age profiles assuming no cohort effects or time effects.²⁹ The profiles of income shock dispersions themselves tell an interesting story. The variance of both permanent and transitory shocks rises towards the end of working life. Earlier in life a higher proportion of shocks is permanent, then the shocks tend to become more transitory in nature. That the level of transitory shocks is much higher than of permanent shocks should itself be disregarded because, of course, we cannot distinguish transitory shocks from measurement error.

Figure 4.2: Age Profile of Permanent and Transitory Shocks



Notes: Plots show a quartic polynomial fit through year- and cohort-specific estimates

Figure 4.3 shows a quadratic polynomial fit of estimates of the transmission of permanent income shocks. In the same graph I plot calculations of the proportion of human

²⁹Because there is so little overlap of observations across cohorts, it is difficult to test whether I am merely picking up differences across cohorts or genuine age effects. As for the exclusion of time effects: an idiosyncratic year component to the variance of permanent shocks should not affect the estimated age effects much. Time trends may affect the profiles, but chapter 2 shows that the average variance of permanent shocks is more-or-less constant over the period.

capital wealth in lifetime wealth. As discussed above, in the simple self-insurance model, these asset moments provide a first-order theoretical approximation of the transmission coefficient.³⁰ I present calculations of this asset moment both including and excluding pension wealth in the definition of financial wealth. I do so for two reasons: first, I have fewer time periods of data on pension wealth, and second, pensions may play less role in consumption smoothing if households cannot borrow against them earlier in the life-cycle.³¹ The figure shows that the estimated transmission parameter is too low at the beginning of the life-cycle. After mid-working age, the estimated transmission coefficient corresponds better to the theoretical prediction, although the standard errors are large, and no sensible hypothesis on the shape of the plot can be rejected. Taking the given plot as a basis for discussion, however, it is worth considering why the transmission of shocks to consumption may be so low, and particularly early in the life-cycle. An explanation for the disparity must lie with either the specification of income dynamics or the consumption process. Several authors (Cunha et al., 2005, Keane and Wolpin, 1997 and Primiceri and Van Rens, 2009) have argued that young people have a lot of information about future outcomes. Advance information about career choice or likely success would explain the low transmission in early working life. Another possibility is that long-lived shocks are persistent but not fully permanent. In this case, the persistent shock has less of an effect on consumption earlier in the life-cycle than later, when the horizon is shorter and both types of shock have similar net present value. It is to this explanation that I now turn.

4.4.4 Interpreting the Results Using Other Models of Income Risk

I argued above that the covariance structure in table 4.2 implies the presence of a unit-root permanent component to income. This argument is the classic Macurdy test against the alternative of a long-lived autoregressive component.³² However, Guvenen (2009), for example, has questioned such evidence and argues that this test lacks power in typical household surveys.³³ The presence or not of a unit root is important in many areas:

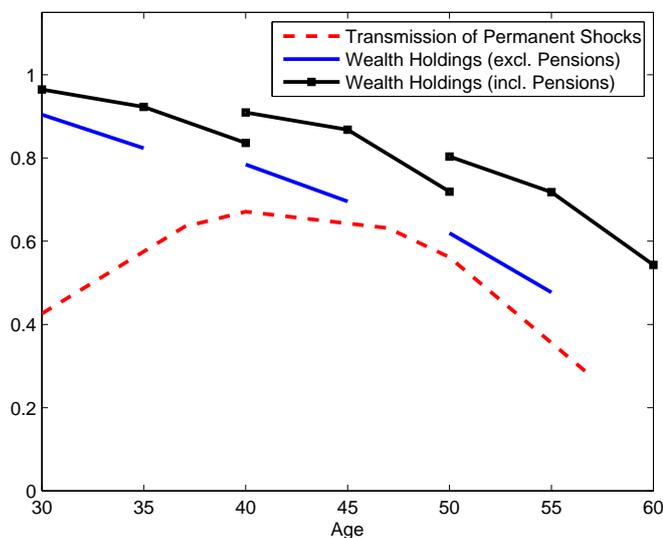
³⁰The approximation also works with risky assets, so long as the household cannot trade unobservable securities which condition payment on the evolution of household income.

³¹The calculations are as in chapter 3.

³²MaCurdy (1982)

³³The statistical properties of income dynamics are a focus of ongoing research. Hryshko (2010), for example, argues that competing models can be distinguished using PSID income data and argues for the unit-root permanent and transitory process. Guvenen and Smith Jr (2010) use a consumption panel and

Figure 4.3: Age Profile of Transmission Coefficients



Notes: ‘Wealth Holdings’ is the mean share of human capital wealth in total life-time wealth. These are calculated both including and excluding pension wealth. ‘The Transmission of Permanent Shocks’ is calculated as a quadratic fit through estimates for each cohort for the first and second halves of the sample period.

precautionary saving behaviour is very different with a unit root compared to a persistence parameter of, say, 0.9. The presence or not of a unit root is particularly important for the present analysis. As Kaplan and Violante (2010) discuss, a persistence parameter even slightly below 1 can generate a substantial reduction in the theoretical transmission of the shock. As a simple example, the net present value of an infinitely-lived AR(1) shock ζ with persistence ρ and rate of return r is $\zeta \frac{1+r}{r+(1-\rho)}$. If $(1-\rho) = r$ (for example with $\rho = 0.97$ and $r = 0.03$) the net present value of the shock is half that of the unit root shock. More realistically, over a 30-year horizon, if $r = 0.03$, then ρ can be as high 0.93 for the AR(1) shock to be worth only half as much as the unit-root shock. Under this parametrization and time horizon, the central estimate of 0.49 on the transmission coefficient corresponds to almost complete transmission of the persistent shock, and hence no insurance.³⁴ The sensitivity of the theoretical transmission coefficient to choice of income model confirms the importance of research into income processes and limits the interpretation of present results in terms of insurance somewhat.

the restrictions imposed by the standard life-cycle model and argues for heterogeneous trends and an AR(1) component.

³⁴Kaplan and Violante (2010) discuss bias in the estimation procedure under the presence of an AR(1) component. They point out that there is little bias to estimation of ϕ , even though the moment conditions are misspecified.

4.4.5 Comparison with Chapter 3

Because the analysis here uses data and model very similar to that in chapter 3 it seems sensible to explain how the approach and results cohere with those in the earlier chapter. One area of overlap is the estimates of permanent shocks to income. In chapter 3, I obtained ‘pre-estimates’ of this variance which fed into the main structural estimation. I estimated these from an MA(2) model with estimates around 0.015.³⁵ In the main estimation I found that an estimate of 0.006 fit the cross-sectional data. This estimation used a numerically-solved life-cycle consumption and saving model, in which the transmission of a permanent shock to consumption would correspond to that in Kaplan and Violante, 2010 and Carroll, 2009. The current chapter investigates the transmission directly by looking at panel consumption data. It finds similarly that the variance of consumption shocks is much lower than the variance of permanent shocks to income implied by the panel income data under the permanent-transitory model.

4.5 Taking Midpoints of Food Consumption

The food data in the BHPS are a potentially valuable resource, but have not been used widely. An important contribution of this chapter, therefore, is to demonstrate that these data do in fact convey useful economic information. In this section I assess the validity of my treatment of food expenditures, as discussed in section 4.2. I compare to alternative treatments and argue that taking midpoints of consumption yields empirically accurate results. The argument I present has 2 strands: first I show that taking the midpoints corresponds empirically well to performing maximum likelihood estimation using the normal distribution and that the normal is the natural choice for this type of analysis. Second, I perform a validation exercise using PSID data, for which we have point observations of household expenditure. I show that banding these data, then using the midpoints, makes little difference to estimates of the relevant variances and covariances.

³⁵In the current chapter we use an MA(1) model. This choice comes from the formal test and rejection of the presence of an MA(2) component.

4.5.1 Analysis of the BHPS data

Taking midpoints of the food points is arbitrary and performed for convenience. An alternative is to specify the underlying distribution of expenditures and to estimate the 2nd moments using maximum likelihood. Here I specify an underlying normal distribution, joint across food and income and across time. This assumption has a theoretic rationale. The normal distribution is the natural choice because the maximum likelihood estimator for a cross section of continuous data is just the sample mean, sample variance and the sample correlation.³⁶ Therefore taking the (non-parametric) covariance matrix of data is akin to estimating the covariance matrix by maximum likelihood under the assumption that the data are normally distributed.

The likelihood function used is

$$LL(\mu, \Sigma) = \sum_{i=1}^n (\Phi(\tilde{x}_i^U, \tilde{y}_i^U, \rho) + \Phi(\tilde{x}_i^L, \tilde{y}_i^L, \rho) - \Phi(\tilde{x}_i^U, \tilde{y}_i^L, \rho) - \Phi(\tilde{x}_i^L, \tilde{y}_i^U, \rho))$$

where μ is the (2x1) vector of means; Σ the covariance matrix; $\Phi()$ is the bivariate standard normal cdf for observations $(\tilde{x}_i, \tilde{y}_i)$ with correlation coefficient ρ ; $\tilde{x}_i = \frac{x_i - \mu_x}{\sigma_x}$, and x_i^U and x_i^L are upper and lower limits of the band containing x_i . As the number of bands increases, in the limit the log likelihood tends towards the standard likelihood function, and the solution for μ and Σ is the sample mean and variance as above. The cdf of the bivariate normal distribution, however, has no analytic expression. I therefore approximate it using the method in Owen (1956).³⁷ The derivatives of the cdf can be expressed analytically, however, so the optimization in the maximum-likelihood estimation relies mostly on precise analytical expressions.

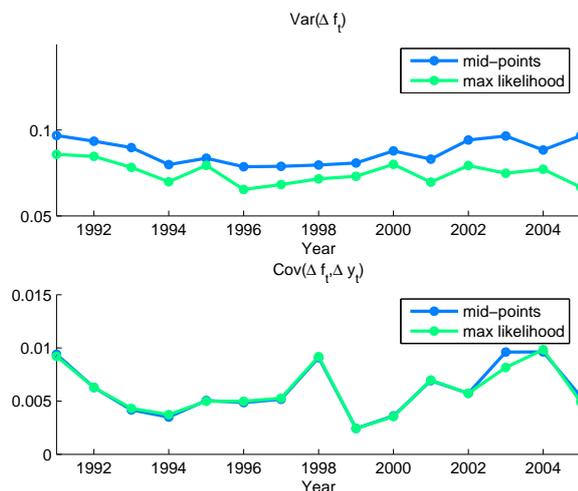
The top panel of figure 4.4 shows the variance of changes in consumption using both midpoints and MLE. The MLE estimates are computed estimating the joint distribution of (c_t, c_{t+1}) , and then using $\text{Var}(\Delta c_t) = \text{Var}(c_t) + \text{Var}(c_{t+1}) - 2\text{Cov}(c_t, c_{t+1})$. The figure shows that the variance of changes using the midpoints is slightly higher than using MLE.

³⁶Similarly the maximum likelihood estimator of the linear regression model with normal disturbances is the OLS estimator.

³⁷This transforms computations of the bivariate normal cdf to a formula of two parameters. I then compute a table (2-dimensional grid) of values of the cdf using numerical integration and then compute all intermediate values using interpolation. It is easy to store enough data in the grid to leave the approximation error of the interpolation negligible.

This is likely because taking midpoints induces extra measurement error. However, both sets of estimates have similar dynamics so it seems this extra measurement error is constant over time.

Figure 4.4: Estimating the Joint Distribution of Food and Income Changes



More importantly, the bottom panel of figure 4.4 shows the covariance of food changes with income changes. This is the key moment used in the main analysis. I estimate these by performing separate bivariate normal estimations for $\text{Cov}(f_{it}, \Delta y_{it})$ and $\text{Cov}(f_{it+1}, \Delta y_{it})$, then subtracting. Here the MLE estimates are almost identical to the midpoint estimates. This is likely because the extra (non-standard) measurement error induced by assigning each band to its mid-point is orthogonal to measurement error in income.

Using the midpoints instead of the maximum likelihood estimates comes at no real loss of efficiency: the standard error on $\text{Cov}(\Delta f_{it}, \Delta y_{it})$ from the MLE estimates, given by bootstrapped estimates, is almost identical to that when using the midpoints.³⁸

As discussed, the normal distribution assumption has a theoretical appeal. However for a finite number of bands, the accuracy of the method depends on the true distribution and it is important to quantify the error under my approach. For this, I perform parallel computations with the most similar data set for which we observe the panel of food consumptions. For this we turn to the PSID data.

³⁸The standard errors on $\text{Var}(\Delta f_{it})$ show more difference. In the first year of the survey, for example, the standard error of $\text{Var}(\Delta f_{it})$ using the MLE (as derived using the inverse of the hessian) is 0.0031. When using the midpoints, the standard error is 0.0048.

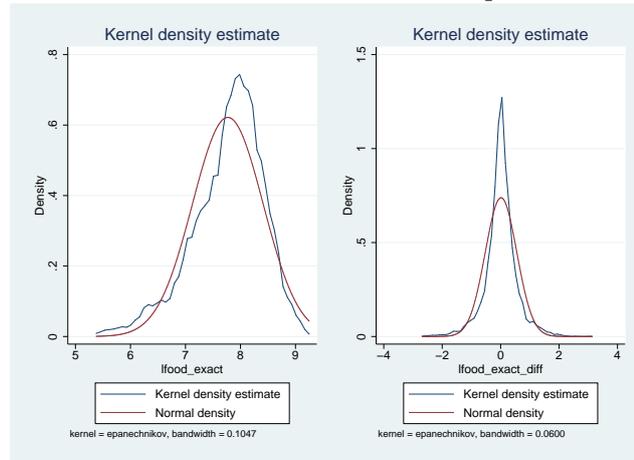
4.5.2 Validation from the PSID

We can test how close this estimator comes to the sample covariance for data distributed as usual by performing a validation exercise with other data sets. Here I pick food data from the Michigan Panel Study of Income Dynamics (PSID). The data were downloaded from the data archive for the BPP paper. The PSID is the standard dataset for studies of the present type. The reader can go to BPP for a description of the dataset.

I perform the following actions on the data. As in BPP I use only households for whom the head is born between 1919 and 1960. My definition of food is food in, to keep comparability with the BHPS data. To remove outliers I first remove households with an annual income less than \$10. I then trim the top and bottom 0.5% of the cross-sectional distributions of food and income. I also remove those observations for which the change in log food consumption is greater than 1.6 or less than -1.6. As in the main analysis, I do not perform this on the income data.³⁹ I assign expenditures to bands in the following way: I set thresholds for the top and bottom band each to capture 0.075% of the distribution, in line with the proportions in the BHPS. I then set the intervals at equal spaces in log space. The induced distribution of expenditures is similar to that in the BHPS; for example the modal band in both datasets captures around 25% of observations. I then assign midpoints as the geometric mean of the interval thresholds. For the top and bottom band I assign each observation so that all the observations are equally spaced. This assignment is, of course, entirely arbitrary, but in line with that from the BHPS. The results that follow are robust to other sensible assignments.

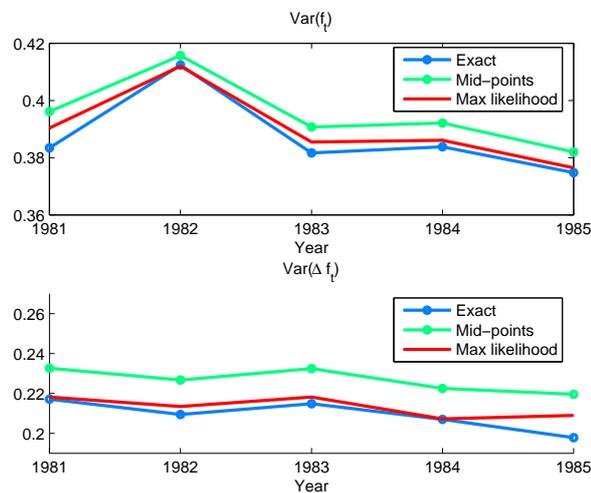
³⁹I do not perform regressions on household characteristics. These change the size of the variances and covariances, but likely do not affect the accuracy of the approximations, which depend on the shape of the joint distribution of income and food consumption. This joint distribution is not affected so much by the first-stage regressions.

Figure 4.5: The Distribution of Food Consumption from the PSID



To assess normality in the underlying data, figure 4.5 shows kernel density estimates of the cross-sectional distribution and the distribution of consumption changes, accompanied by fitted normal distributions. Both distributions clearly deviate from the normal: the cross-sectional distribution is skewed with a long left-hand tail, while the distribution of changes is symmetric but clearly leptokurtic.

Figure 4.6: Estimating 2nd Moments of the Food Distribution in the PSID



The top panel of figure 4.6 shows estimates of the cross sectional variance of food expenditures using the exact data, the imposed midpoints, and maximum likelihood estimates using the imposed bands. I pick 1981-1985 as an example sub-period. Both the midpoints and the maximum likelihood estimates slightly overstate the variance, but they capture the dynamics well. The bottom panel of 4.6 shows estimates of the variance of

changes of food expenditure using the three different methods. As for the BHPS, the estimates using mid-points are higher because of the extra measurement error. The maximum likelihood is closer to the exact variance. Both the approximations (using the midpoint and the maximum likelihood) follow the dynamics of the PSID very closely.

Figure 4.7: Estimating the Joint Distribution of Food and Income Changes in the PSID

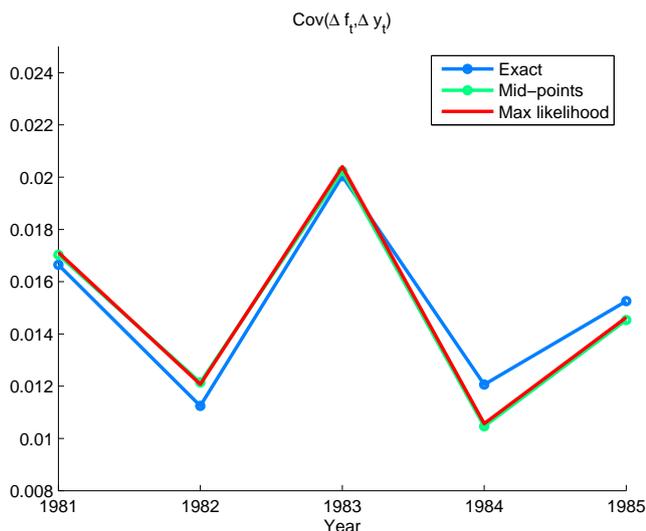


Figure 4.7 shows the covariance of food changes and income changes using the different estimation methods. To repeat, this is the crucial moment for the identification of transmission parameters. Again the mid-points and the bands give almost exactly the same answer. They also capture the level and the dynamics of the precise estimates extremely well. The standard errors on the covariances are almost identical (at 0.0035 in 1981) when using either the precise observations or the midpoints. There is therefore no loss of efficiency when using the midpoints.

4.5.3 Final Remarks on the Data Imputation

I conclude that taking midpoints of the banded food data yields empirically accurate results. As a final word I discuss further econometric alternatives. An obvious alternative when using banded data is to identify bounds on the relevant variances and covariances non-parametrically. The advantage of this approach is that it doesn't require imputing food data at all nor does it require placing parametric assumptions on the underlying distribution. Stoye (2010), for example, discusses identification of spread parameters using

(univariate) banded data. There are, however, several problems with such an approach. First, the top and bottom bands in the data are unlimited. The variance is therefore unbounded without at least some minimal further restrictions on the distribution. Second, even with limits on the top and bottom band, the implied non-parametric bounds on the variance are quite large. They are derived by allocating the observations to the extremities of the observed bands which yield minimal and maximal variance.⁴⁰ We know from all other datasets, however, that food expenditures are smoothly distributed. A simple bounds analysis therefore greatly overstates reasonable ignorance about the exact variance. A more sophisticated approach would allow for including statistical restrictions on the shape of the distribution. However, I know of no econometric theory developed in this area which would be suitable for the present study.

4.6 Conclusions

In this chapter, I study income risk and its transmission through to consumption in the UK over 1991-2006. I am motivated by trying to reconcile two different views of risk over the period. Permanent income risk for my sample was substantial: the average variance of permanent shocks estimated from panel income data was 0.019 (a standard deviation of 14%). Meanwhile, consumption risk appears much lower: the variance of shocks estimated from the growth in consumption inequality was 0.0055 (a standard deviation of 7.5%). I use techniques similar to those in Blundell et al. (2008b) (BPP) for the analysis; I use both cross-sectional data on food and total non-durable expenditures from the FES and panel data on food and income from the BHPS.

I estimate the transmission of permanent shocks to be 0.49. This is lower than BPP's estimate of 0.64. The variance of permanent shocks is also slightly lower than for the BPP sample (0.019 compared to 0.020). These results have three immediate and important implications. First, the implied contribution from income risk is around 80% of total consumption risk, as estimated from the repeated cross-section. Income shocks therefore provide the bulk of shocks to consumption. Second, the risk to consumption induced by income shocks is about half that in the BPP sample. Insofar as we can compare risk across

⁴⁰For a univariate distribution the maximum bound is obtained by placing all observations furthest away from the mean band, and the minimum bound by placing all observations closest to the mean band.

countries, this provides fresh evidence that consumption risk was lower in the 1990s than in the 1980s. Third, the results imply that focusing on the variance of permanent income shocks as a measure of the cost of risk is misleading. If the cost of risk is proportional to the variance of shocks (as in Lucas (1987)), then ignoring consumption smoothing yields a cost measure that is 4 times too high.

A smaller contribution of this chapter is to show that the banded data on food expenditures in the BHPS can be used for solid economic research on consumption behaviour. By performing a validation study using data from the PSID, I show that the banded data are almost as useful as precisely observed data.

Finally I note that the transmission of permanent income shocks is lower than is obtained in standard models of self-insurance (such as those in Kaplan and Violante (2010) and Carroll (2009)). The gap between theory and evidence is particularly high early in the life-cycle. On one interpretation this gap implies substantial extra insurance, and particularly for younger people. However, given that my income definition includes all gifts and transfers it is hard to think of more mechanisms that provide the extra insurance. The gap can be explained by two other hypotheses: first, if households have advance information about career success, but later income fluctuations capture more unforeseen news, then the transmission parameter should be lower early in life. Second, if income shocks are not permanent but only very persistent, then again, the transmission of income shocks would be lower early in life. I find implicit support either for the presence of advance information about future income changes or for the absence of a unit root on income shocks. I find no direct evidence, however, for either of these features.

Both the presence of advance information and the specification of the income process are subject to much current research. The fact that no consensus has been reached on either of these topics shows how difficult a research problem these provide. The results from this chapter suggest other more tractable areas of research, however. First, I note that estimates of consumption risk induced by income risk are more robust than estimates of income risk alone. Future research could follow this path, because quantifying consumption risk remains an important task in its own right.⁴¹ In interesting area of future research

⁴¹For example, the level of consumption risk determines the optimal intertemporal savings distortion. See Farhi and Werning (2009).

would be to assess the contributions from both the components of income and from other sources. Finally, it is worth repeating that this chapter concerns only stable households headed by a couple. Non-stable households are probably more interesting, but are, of course, harder to study. Research on their behaviour and circumstances is needed.

A4 Appendix to chapter 4

A4.1 Using FES Expenditure Data

This appendix gives more details on data from the FES. First I give a brief description of food questions in the dataset, then I give further details of my demand estimation.

Comparing BHPS and FES food consumption data

As mentioned in section 4.2, the food questions in the BHPS are based on recall. In contrast, the FES collects data on expenditure in a diary survey. Each household details all their spending, both home and abroad, over a two week period.⁴² Several papers discuss the relative merits and characteristics of recall versus diary methods, such as Battistin et al. (2003). I include both food and groceries in ‘food’ because this gives a closer match to the BHPS data.

Table 5 shows the characteristics of the final samples in the FES and BHPS datasets, pooled across the first half and then the second half of the sample period. There are some levels differences between the datasets: notably, households in the BHPS appear to have more adults and fewer children than those in the FES. However, the trends from the first half to the second half of the period are similar for all measures across both datasets.

Estimating the Food Demand Equation

My analysis requires a uniform income elasticity of demand across each group I study. A uniform income elasticity is a controversial claim. At a raw theoretical level, it is well known that the implied log-linear demand function fails adding up (Deaton and Muellbauer, 1980). More generally, most demand analyses estimate a concave elasticity

⁴²In addition to this diary, household members perform an interview in which they are asked to recall expenditures on large infrequently-purchased items, such as cars.

Table 5: Comparison of Means, BHPS and FES

	1991-1998 BHPS	1991-1998 FES	1999-2006 BHPS	1999-2006 FES
head's age	40.9416	40.8875	46.7315	47.4973
hh size	3.472	3.411	3.375	3.208
Number of adults	2.416	2.199	2.497	2.228
Number of children	1.040	1.212	0.906	0.980
Compulsory level of education	0.453	0.506	0.346	0.484
Working	0.890	0.934	0.883	0.877
Retired	0.008	0.004	0.040	0.047
Other labour force status	0.101	0.061	0.077	0.076
0 cars	0.072	0.090	0.043	0.068
1 car	0.448	0.466	0.335	0.384
2 cars	0.401	0.369	0.481	0.438
>2 cars	0.078	0.075	0.139	0.109
Homeowner	0.826	0.799	0.869	0.844

Notes: Rows "Compulsory education" and below give proportions.

The means are simple pooled averages, unweighted by the sample sizes in each year.

(for example Browning and Meghir, 1991). Nevertheless, I present evidence that any non-linearity is negligible and does not substantially affect the analysis.

Table 6 gives the results from the estimation of the main food demand equation. I instrument expenditure variables by log income and interactions to remove attenuation bias from measurement error. The main point of this regression is to back out income elasticities. The base elasticity is 0.39 for the low education group, born in the 1940s, in year 1991. I allow the elasticity to vary by education, cohort and allow for a linear effect across time. The coefficients on all these interactions are small and insignificant. These estimates indicate that the income elasticity does not vary much across different parts of the income distribution.

The results don't change when I allow for a full set of interactions between expenditure and year. When I allow for a quadratic term in total expenditure (and keep the other interactions with the linear term), the coefficient on total expenditure squared is -0.0663 , and on total expenditure is 1.098 , with standard errors of 0.018 and 0.20 . The implied elasticity at the 10th centile of the expenditure distribution in 2000 is 0.45 , and at the 90th centile is 0.27 . The food demand equation does therefore display some curvature.

Another important consideration is the effect of participation on food demands. An effect of participation on the intercept has a large impact on the implied transmission of income through to consumption.⁴³ However, dummies for male and female participation have small coefficients. The coefficient on female participation is negative and significant but quantitatively very small. When controlling for participation in both the intercept and the elasticity, the effects are larger, but very imprecisely estimated and not significant.

The analysis so far depends on the exogeneity of total expenditure. I test for endogeneity in the main equation (other than by measurement error) by including asset income and its interactions in the set of instruments. The exclusion of asset income in the determination of food demands is based on the two-stage budgeting framework. A Sargan test of the over-identifying restrictions has a p-value of 4.2%⁴⁴, providing some evidence of misspecification at the 5% level. The estimated income elasticity is slightly lower when instrumenting with asset income alone. The estimated elasticity is 0.385 for the base group in 1999 compared to 0.405 in the main equation. The estimates of the transmission of income shocks are therefore biased slightly downwards. This does not affect the main result much, however. Allowing for joint determination of food and total expenditure, on the other hand, implies larger standard errors on the estimated transmission coefficient.

⁴³This is because a large fraction of income variation comes through participation effects. If hours and food demands are non-separable and correcting for participation induces extra variation in total non-durable consumption, because of, say, a negative coefficient on participation, then implied transmission of income shocks through to total consumption will be higher.

⁴⁴Chi-squared statistic of 11.53 with 5 degrees of freedom

Table 6: Food Demand in the UK

Variable	Estimate	Variable	Estimate	Variable	Estimate
ln c	0.386*** (0.0191)	Age spouse ³	-0.00201 (0.00207)	yr = 1993	-0.0492** (0.0236)
ln c × Born 1950s	0.0213 (0.0198)	Age spouse ⁴	0.000112 (0.000117)	yr = 1994	-0.0469 (0.0324)
ln c × Born 1960s	-0.0384* (0.0216)	Yorkshire	-0.00317 (0.00897)	yr = 1995	-0.0625 (0.0409)
ln c × High education	0.00143 (0.000961)	North West	-0.0201* (0.0103)	yr = 1996	-0.0500 (0.0499)
ln c × (year-1991)	0.00237 (0.00175)	East Midlands	0.0107 (0.0115)	yr = 1997	-0.0959 (0.0600)
ln p_{food}	-0.734*** (0.220)	West Midlands	-0.0152 (0.0108)	yr = 1998	-0.138* (0.0703)
ln Number Adults	0.423*** (0.0139)	East Anglia	-0.0386*** (0.0105)	yr = 1999	-0.145* (0.0804)
ln # kids aged 0-4	0.210*** (0.00753)	London	-0.0394*** (0.00999)	yr = 2000	-0.171* (0.0916)
ln # kids aged 5-10	0.200*** (0.00625)	South East	-0.0562*** (0.0105)	yr = 2001	-0.221** (0.100)
ln # kids aged 11-18	0.248*** (0.00606)	South West	-0.0289** (0.0123)	yr = 2002	-0.249** (0.110)
Age head	0.0421** (0.0167)	Wales	-0.0241* (0.0123)	yr = 2003	-0.251** (0.121)
Age head ²	-0.00840** (0.00385)	Scotland	0.00403 (0.0106)	yr = 2004	-0.284** (0.131)
Age head ³	0.000562* (0.000288)	Born 1950s	-0.105 (0.111)	yr = 2005	-0.314** (0.142)
Age spouse	-0.0189 (0.0367)	Born 1960s	0.237* (0.121)	yr = 2006	-0.320** (0.152)
Age spouse ²	0.0117 (0.0133)	year = 1992	-0.0309* (0.0158)	Constant	0.320 (0.405)
Observations	36,411				
R-squared	0.320				

Standard errors in brackets

Instrumented: ln c and interactions.

Instruments are: ln y and interactions.

Age^2 is divided by 10, Age^3 by 100 and Age^4 by 1000 for readability of coefficients

Chapter 5

For Many a Rainy Day: Precautionary Saving for Consecutive Life-Cycle Risks

5.1 Introduction

Economists have long recognized the variety of motives for saving.¹ An important debate over the last twenty years has concerned the relative importance of the precautionary motive for emergencies versus the life-cycle motive.² ³ It is tempting to consider precautionary saving as driven by short-term possible needs as opposed to life-cycle saving which, by definition, has a long-term focus. However, households face risks of different types and magnitudes over the whole life-cycle. Within the variety of motivations for general saving, therefore, households have competing motivations for precautionary saving, for near-term as well as for far-off emergencies.

An issue intimately related to assessing the importance of each saving motive is how well these motives complement each other. While some forms of wealth, such as illiquid pension wealth, seem targeted for a particular saving motivation (life-cycle saving), liquid

¹see Browning and Lusardi (1996) for an exhaustive list.

²See Kimball (1990) and Caballero (1990) for early theoretical formulations of the precautionary motive in terms of prudence. Deaton (1991) and Carroll (1997) provide classic analyses focusing on explaining patterns of life-cycle consumption and wealth formation.

³Carroll and Samwick (1997) and Gruber and Yelowitz (1999), for example, emphasize the importance of precautionary saving in total wealth, while Dynan (1993) and Hurst et al. (2005), for example, find precautionary saving to be less important.

wealth in particular can meet many motivations simultaneously: a dollar in the bank can be spent in retirement, can be used in a near-term emergency, or can be used eventually in a far-off emergency. It seems intuitively plausible that households need save less when they can combine saving in this way. Understanding whether this is the case is important for understanding household, and hence aggregate, wealth accumulation.

In this chapter we examine this complementary aspect of saving. We focus on precautionary saving for consecutive future income risks. Precautionary saving for consecutive risks seems particularly likely to be complementary because of its contingent nature. Emergencies occur only rarely and precautionary wealth is only rarely needed. Therefore the accumulated wealth stock can presumably be put to other ends, and specifically as rolled-over precautionary wealth against subsequent income risk. On the other hand, the theoretical literature on background risk (discussed, for example, in Gollier (2004)) emphasizes that the presence of multiple risks amplifies risk aversion. On this intuition, the presence of multiple risks might amplify the need for precautionary wealth.

We formalize the intuitive notion of complementarity in the following way: first we quantify precautionary saving when the household faces risk both in mid-life and then late-life consecutively. We then quantify precautionary saving when the household faces risk either in mid-life or late-life in isolation. We label the sum of saving for these two isolated periods of risk the ‘total’ precautionary effect of the risks. We say that saving is complementary, or exhibits complementarity, if initial precautionary saving for consecutive risks is less than the total precautionary effect.

Our results are both quantitative and analytical. On the quantitative side, we simulate the standard life-cycle consumption and saving model with both permanent and transitory income risks. In a range of realistic parametrizations, the total precautionary effect of the risks is 8-16% higher than precautionary saving for the consecutive risks. This complementarity effect is driven almost entirely by permanent shocks: the effect of transitory shocks is negligible.

In order to understand how complementarity arises we then simplify the model and examine saving behaviour analytically. We focus on a 3-period horizon and first admit small permanent risks alone. We find that complementarity depends on the shape of relative

prudence over the wealth spectrum.⁴ Utility functions with constant relative prudence (those in the CRRA class) always permit complementary saving. More generally, so do a large subset of utility functions with harmonic absolute risk aversion (HARA).⁵ This class includes not only CRRA but also constant absolute risk aversion (CARA) and quadratic functions.⁶ On the other hand, saving might not be complementary whenever relative prudence is locally strongly declining. This occurs whenever households have minimum consumption needs and a household is on the ‘breadline’. With these preferences, households are so averse to consecutive bad shocks that the interaction of the risks exacerbates precautionary saving. We term this behaviour excessiveness of saving. Such excessive saving is only local, however. When the household is sufficiently far from the breadline, precautionary saving is complementary again.

To extend intuition we examine the case of small transitory risks alone. For these risks, first-period saving is not complementary but is in fact excessive for standard preferences, such as CRRA. For transitory shocks we can derive a more interpretable condition on the shape of prudence behaviour: saving is excessive if absolute prudence is declining and convex.⁷ In this case, convexity of the prudence function intuitively contributes to excessive saving because it implies that average prudence is greater the more independent risks are added, an application of Jensen’s inequality. Convexity and a negative slope are attractive properties for the prudence function to possess.⁸

We gain the most intuition for these results by considering preferences of the CARA form and transitory risks. Here saving is neutral; it is neither excessive nor complementary. We can understand this result by considering the following informal argument. First, innovations to wealth from transitory risks are independent across time. Second, wealth

⁴Relative prudence is defined as $-\frac{xu''(x)}{u'(x)}$ for utility function $u(x)$ and consumption level x . This can be interpreted as the strength of the desire to have higher wealth when facing a given proportional gamble (such as $\pm 10\%$). See Kimball (1990).

⁵HARA preferences have absolute risk aversion of the form $\frac{1}{ax+b}$, where a and b are constant, x is the consumption level and risk aversion is defined to be $-\frac{u''(x)}{u'(x)}$.

⁶CRRA preferences can be represented by utility function $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ for some parameter γ . These preferences have both constant relative risk aversion and constant relative prudence. CARA preferences can be represented by utility function $u(x) = \text{Exp}(-\gamma x)$ for some parameter γ . These preferences have both constant *absolute* risk aversion and constant *absolute* prudence.

⁷Absolute prudence is defined as $-\frac{u''(x)}{u'(x)}$. This differs from relative prudence by employing an *absolute* gamble in the definition, such as $\pm \$10$.

⁸If the prudence function is declining and convex, then wealthier people are less averse to an gamble of absolute size, but the rate of this decline is reducing with wealth. As stated, all HARA preferences except CARA and quadratic preferences have a convex and negative absolute prudence.

level is irrelevant when saving for a given risk. From the viewpoint of the middle period therefore, the precautionary motive for the final-period risk does not depend on the outcome of the most recent shock. Returning to the first period, the household must save for both middle- and final-period risks. However, because the middle-period outcome will not change behaviour, the size of the middle risk does not affect the precautionary motive for the final risk. Symmetrically, the size of the final risk does not affect the precautionary motive for the middle risk. In short, the risks don't interact in the initial saving decision.⁹

We can understand the more complex environments by considering where this argument breaks down. For example, consider two consecutive permanent risks. In this case the variances of the innovations to life-time wealth are no longer independent. In fact the final-period variance is increasing quadratically in the middle-period outcome: the variance of subsequent wealth innovations is reduced following a bad income shock. This dependence arises even though the raw income shocks are independent. The risk process itself therefore provides a kind of insurance that limits the need for precautionary saving for standard preferences. Permanent risks therefore permit complementarity much more readily than do transitory risks. This is true not only for CARA preferences but more generally for CRRA preferences. It is worth emphasizing that this feature of permanent risks arises not because *mean* income is dependent over time. It arises because of the dependence over time of the *variance*.

After the early works cited above, the theoretical and empirical literature on precautionary savings has developed steadily in, for example, Carroll and Kimball (2001) and Carroll (2004). These papers emphasize that when households discount the future strongly, display prudence and face income risk, they save to meet a target wealth holding, the buffer stock. This buffer stock intuitively conforms to the notion that a fixed level of wealth can meet all subsequent risks. The models in these papers, however, include constant income risk over the planning horizon and so do not explicitly explore the distinction between near-term and far-term risk.

Aside from the literature on precautionary savings, this chapter contributes to a growing literature on the interaction of risks. Classic theoretical works by Kimball (1993), Pratt and Zeckhauser (1987), Gollier and Pratt (1996) characterize conditions on the

⁹See Caballero (1990) for the full mathematical treatment.

utility function for the introduction of background risks to affect the desirability of risk bearing. These papers focus on risk bearing and so characterize results in terms of the coefficient of risk aversion. Our work relates to precautionary saving so characterizes the results in terms of the coefficient of prudence. There is a growing recent empirical literature on background risks, particularly the effect of risk at the end of the life cycle on prior behaviour. For example, Goldman and Maestas (2005) look at the effect of medical expenditure risk on portfolio decisions earlier in retirement. De Nardi et al. (2010) look at the effect of medical expenditure and mortality risk on the precautionary saving motive earlier in retirement. Their model begins in retirement, but presumably the risks discussed also affect savings motives earlier in the life cycle. Guiso et al. (2009) look at the effect of pension risks on portfolio allocation earlier in the life-cycle. All these papers find that the background risks have an important impact on household behaviour.

This chapter further builds upon the work in Blundell and Stoker (1999). This earlier paper is more concerned with how consumption changes track income changes ex-post and obtaining a complete, but approximate, description of the consumption plan in a 3-period environment. We focus on ex-ante saving decisions in the first period, and are more specific about the risk environment the agent faces, restricting it to being the empirically plausible permanent-transitory process.

This chapter proceeds as follows. In section 5.2 we lay out the basic three-period model through which all the results can be understood. We also discuss our definition of complementarity. In section 5.3 we present the first, quantitative results. Here we use a more realistic life-cycle model which includes both transitory and permanent income risk and a longer-term planning horizon. In section 5.4 we present the analytic results. First we discuss savings behaviour for small permanent risks alone and standard classes of preferences. We then build intuition by considering transitory shocks alone. Meanwhile we extend intuition by constructing exotic preferences which induce different saving behaviour. Section 5.5 concludes and discusses other features which may affect complementarity of saving: liquidity constraints and holdings of illiquid assets such as housing and pensions.

5.2 The Model and Definition of Complementarity

In this section we lay out the consumption and saving framework in a three-period environment. This model is standard. The purpose of the exposition is to define notation and to draw attention to those features of the model that are important for the later analysis. The following section (section 5.3) presents quantitative results for a longer planning period, but all the analysis is intelligible in terms of the shorter model.

5.2.1 The Budget Constraint and Income Process in the Basic 3-Period Model

Households must choose consumption in three periods (indexed by $t = 0, 1, 2$) subject to the following budget constraint:

$$c_0 + \frac{c_1}{R} + \frac{c_2}{R^2} = y_0 + \frac{y_1}{R} + \frac{y_2}{R^2} \quad (5.1)$$

where c_t is consumption at time t , y_t is income at time t . Households save in a risk-free bond with interest rate R .

Income follows a standard stochastic multiplicative permanent-transitory process:

$$\begin{aligned} y_t^P &= y_{t-1}^P G_t \psi_t \\ y_t &= y_t^P \xi_t \\ \mathbb{E}(\xi_t) = \mathbb{E}(\psi_t) &= 1 \quad , \quad \text{Var}(\xi_t) = \sigma_{\xi_t}^2, \text{Var}(\psi_t) = \sigma_{\psi_t}^2 \quad t = 1, 2 \end{aligned}$$

where G_t is deterministic growth, y_t^P is latent permanent income, ψ_t represents the permanent shock to income, and ξ_t a transitory shock to income.¹⁰ These shocks are uncorrelated with each other and uncorrelated with other shocks across time. This process nests the standard lognormal process, in which case, for example, $\ln \psi_t \sim \mathcal{N}\left(-\frac{\sigma^2}{2}, \sigma^2\right)$, where $\sigma^2 = \ln\left(\sigma_{\psi_t}^2 + 1\right)$. The notation here differs from that in earlier chapters. Note that in this chapter we express the income process in levels rather than logs, so the specification of shocks differs from that in earlier chapters.

¹⁰In this 3-period model all shocks have a large impact on life-time wealth, so transitory shocks have greater impact on consumption than in a one-year-per-period model. The quantitative importance of transitory risks therefore cannot be determined easily from the 3-period model.

We now show how the variances of the shocks affect the variances of income. Defining $\bar{y}_t \equiv \mathbb{E}_0(y_t) = y_0 \prod_{j=1}^t G_j$:

$$\begin{aligned} y_1 &= \bar{y}_1 \psi_1 \xi_1 \\ y_2 &= \bar{y}_2 \psi_1 \psi_2 \xi_2 \end{aligned}$$

To simplify the exposition, we focus on the cases first with only transitory shocks, then only permanent shocks. We define $\text{Var}_t(y_2)$ to be the variance of period-2 income conditional on the period- t information set ($t \in \{0, 1\}$). First, when there are only transitory shocks:

$$\begin{aligned} \text{Var}(y_1) &= (\bar{y}_1)^2 \sigma_{\xi_1}^2 \\ \text{Var}_0(y_2) = \text{Var}_1(y_2) &= (\bar{y}_2)^2 \sigma_{\xi_2}^2 \end{aligned}$$

in which case the variance of second period income is independent of the first period innovations, but depends on the size of expected income growth. When there are only permanent shocks:

$$\begin{aligned} \text{Var}(y_1) &= (\bar{y}_1)^2 \sigma_{\psi_1}^2 \\ \text{Var}_1(y_2) &= (\bar{y}_2 \psi_1)^2 \sigma_{\psi_2}^2 \\ \text{Var}_0(y_2) &= (\bar{y}_2)^2 (\sigma_{\psi_1}^2 + (1 + \sigma_{\psi_1}^2) \sigma_{\psi_2}^2) \end{aligned}$$

where the last line can be derived using the formula for the variance of products (as in Goodman (1960)).¹¹

Our analysis concerns innovations to life-time wealth. Therefore we now relate income shocks to these innovations. The 2nd-period innovation to life-time wealth when there are only permanent shocks is:

$$\zeta_2^* \equiv \frac{y_2 - \mathbb{E}_1 y_2}{R^2} = \frac{\bar{y}_2 \psi_1 (\psi_2 - 1)}{R^2}$$

¹¹For uncorrelated random variables $\text{Var}(xy) = \text{Var}(x) \mathbb{E}(y)^2 + \text{Var}(y) \mathbb{E}(x)^2 + \text{Var}(x) \text{Var}(y)$.

then

$$\text{Var}_1(\zeta_2^*) = \frac{(\bar{y}_2\psi_1)^2}{R^4}\sigma_{\psi_2}^2 \quad (5.2)$$

$$\text{Var}_0(\zeta_2^*) = \frac{\bar{y}_2^2}{R^4}(1 + \sigma_{\psi_1}^2)\sigma_{\psi_2}^2 \quad (5.3)$$

Here we make two related points. First, as equation 5.2 shows, the variance of the second-period innovations, from the viewpoint of period 1, depends on the realization of the first period shock. This connection of innovations given by equation 5.2 is an important feature of multiplicative risk. Second, as equation 5.3 shows, the variance of these innovations, from the viewpoint of period 0, depends on the variance of period-1 risk. We want the ex-ante (period-0) variance of period-2 innovations not to depend on the period-1 variance. We therefore generally work with a scaled and recentred shock $\tilde{\psi}_2 = \frac{\psi_2 + \sqrt{1 + \sigma_{\psi_1}^2} - 1}{\sqrt{1 + \sigma_{\psi_1}^2}}$. Then $\mathbb{E}_0(\tilde{\psi}_2) = 1$ and $\text{Var}(\tilde{\psi}_2) = \frac{\sigma_{\psi_2}^2}{1 + \sigma_{\psi_1}^2}$. If $y_2 = \bar{y}_2\psi_1\tilde{\psi}_2$, and $\tilde{\zeta}_2^*$ is the wealth innovation for the scaled shock then:

$$\text{Var}_1(\tilde{\zeta}_2^*) = \frac{(\bar{y}_2\psi_1)^2}{R^4(1 + \sigma_{\psi_1}^2)}\sigma_{\psi_2}^2 \quad (5.4)$$

$$\text{Var}_0(\tilde{\zeta}_2^*) = \frac{\bar{y}_2^2}{R^4}\sigma_{\psi_2}^2 \quad (5.5)$$

Equations 5.2 and 5.4 show that the connection between period-1 and period-2 wealth innovations is the same for both $\tilde{\psi}_2$ and ψ_2 . Equation 5.5 shows that the ex-ante variance of scaled wealth innovations no longer depends on period-1 risks. Using $\tilde{\psi}_2$ instead of ψ_2 does not affect the results substantially, but makes the analysis more interpretable. See appendix A5.1 for an extensive discussion.

5.2.2 The Consumption Problem

The agent faces the following value function problem:

$$V_0(W_0) = \max_{\{c_t(W_t):t=0,1,2\}} u(c_0) + \mathbb{E}_0(\beta u(c_1) + \beta^2 u(c_2)) \quad (5.6)$$

subject to the budget constraint given in equation 5.1 and the process for wealth innovations. Here $u(x)$ is the per-period felicity function, \mathbb{E}_t is the expectations operator at time

ρ and β is the rate of time preference, assumed to be common across both future periods. $V_0(W_0)$ is the value of the programme to the household at time 0 at wealth level W_0 . c_0 is period-0 saving. We focus on how initial-period saving responds to changes in future risk. Therefore we could rewrite programme 5.6 as:

$$V_0(W_0, \Gamma) = \max_{\{s_0(W_0, \Gamma)\}} u(W_0 - s_0) + \mathbb{E}_0(V_1(s_0, \Gamma)) \quad (5.7)$$

where $V_1(\cdot)$ represents the value function at period 1, Γ represents the parameters of the problem, including, for example, time preference, interest rates and the distributions of future income risk, and s_0 represents period-0 saving.

Interest features solely on the effect of changes in income risk. It is well known that assigning a one-dimensional measure of riskiness is tricky. The standard approach is to use the notion of mean-preserving spreads, which provides a partial ordering of income distributions in terms of second-order stochastic dominance. However, in general, higher order features of the distribution will also affect saving.¹² In this chapter we will refer to the size of income risk purely in terms of the variance or standard deviation. This is justified by two main considerations. First, in the quantitative analysis in section 5.3 we perform numerical simulations using a log-normal distribution. Log-normal shocks (with unit mean) are characterized completely by their 2nd moment. Second, in the analytical analysis in section 5.4, we study behaviour following the introduction of small, mean zero risks. For small risks, again only the 2nd moment is relevant: the effect of higher-order moments vanishes.

With this in mind we rewrite programme 5.7 further as:

$$V_0(W_0, \sigma_1, \sigma_2 | \Gamma) = \max_{\{s_0(W_0, \sigma_1, \sigma_2 | \Gamma)\}} u(W_0 - s_0) + \mathbb{E}_0(V_1(s_0, \sigma_1, \sigma_2 | \Gamma))$$

for some parameters σ_1 and σ_2 governing the 2nd moment of risks in periods 1 and 2. These could govern permanent risk, transitory risk or some combination of the two. A solution to this programme is given by the function $s_0(W_0, \sigma_1, \sigma_2 | \Gamma)$, or more simply $s_0(\sigma_1, \sigma_2)$.

¹²See Eeckhoudt and Schlesinger (2008) for an analysis in a 2-period environment.

5.2.3 Definition of Complementarity

Let $s_0(\sigma_1, \sigma_2)$ be first-period saving as a function of future risks, with standard deviations given by σ_1 and σ_2 . To repeat, we will assume that, given other features of the environment (such as the shape of the distribution of shocks), we can classify and order risks purely in terms of the standard deviations of income shocks. We further assume that initial-period saving is twice differentiable in these standard deviations. Given a counterfactual risk profile $\langle \sigma_1^*, \sigma_2^* \rangle$, we consider changes in saving from the base case:

$$\Delta s_0(\sigma_1^*, \sigma_2^*) \equiv s_0(\sigma_1^*, \sigma_2^*) - s_0(\sigma_1, \sigma_2)$$

Similarly we denote changes in saving from the base case for each risk in isolation by $\Delta s_0(\sigma_1^*, \sigma_2)$, and $\Delta s_0(\sigma_1, \sigma_2^*)$, (so, for example, $\Delta s_0(\sigma_1^*, \sigma_2) = s_0(\sigma_1^*, \sigma_2) - s_0(\sigma_1, \sigma_2)$). We say that savings exhibit complementarity for this environment if, for $\sigma_1^* > \sigma_1$ and $\sigma_2^* > \sigma_2$:

$$\Delta s_0(\sigma_1^*, \sigma_2^*) < (\Delta s_0(\sigma_1^*, \sigma_2) + \Delta s_0(\sigma_1, \sigma_2^*)) \quad (5.8)$$

In terms of the earlier discussion, the right hand side of 5.8 gives the total precautionary effect of the risks. Taking a Taylor-series expansion of these terms gives:

$$\Delta s_0(\sigma_1^*, \sigma_2^*) - (\Delta s_0(\sigma_1^*, \sigma_2) + \Delta s_0(\sigma_1, \sigma_2^*)) = \frac{1}{2} \Delta \sigma_1 \Delta \sigma_2 \frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} + \mathcal{O} \left((\Delta \sigma_1)^3, (\Delta \sigma_2)^3 \right) \quad (5.9)$$

where $\Delta \sigma_j \equiv (\sigma_j^* - \sigma_j)$. Therefore, for small changes in risk, this complementarity depends on the sign of the cross partial derivative $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$. This gives the following definition:

Definition 1. Savings (s_0) display *complementarity* for utility function $u(\cdot)$, for future risks parametrized by $\langle \sigma_1, \sigma_2 \rangle$ and for environment Γ , if $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} < 0$. If $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} > 0$ then savings display *excessiveness*.

During the rest of the chapter we sometimes refer to $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$ as the complementarity function.

Further discussion of definition 1 is merited. First, the definition is in terms of the standard deviation of future risks rather than the variance. Of course $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} = 4\sigma_1\sigma_2 \frac{\partial^2 s_0}{\partial \sigma_1^2 \partial \sigma_2^2}$, and so $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$ and $\frac{\partial^2 s_0}{\partial \sigma_1^2 \partial \sigma_2^2}$ always have the same sign when σ_1 and σ_2 are positive. Therefore

the distinction is inconsequential.

Finally, it is helpful to compare definition 1 to other possible formalizations of the intuitive notion of complementarity. An obvious alternative is to wonder whether saving is higher or lower if a given amount of future risk is ‘bunched’ in one particular period or spread out more evenly. Using the notation above, and allowing ϵ to denote a small deviation in a risk parameter we might try to study

$$s_0(\sigma + \epsilon, \sigma - \epsilon) - s_0(\sigma, \sigma) \quad (5.10)$$

In this case, we are bunching risk in the middle period. Corresponding to our intuitive notion of complementarity, we might say that saving is complementary if agents save less when risk is spread out, ie expression 5.10 is negative.¹³

Similarly as before, taking a Taylor-series expansion of expression 5.10 gives

$$\epsilon \left(\frac{\partial s_0}{\partial \sigma_1}(\sigma, \sigma) - \frac{\partial s_0}{\partial \sigma_2}(\sigma, \sigma) \right) + \frac{1}{2} \epsilon^2 \left(\frac{\partial^2 s_0}{(\partial \sigma_1)^2}(\sigma, \sigma) - 2 \frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}(\sigma, \sigma) + \frac{\partial^2 s_0}{(\partial \sigma_2)^2}(\sigma, \sigma) \right) + O(\epsilon^3)$$

We immediately see that this definition is, in fact, less interpretable. It combines first-order and second-order derivatives of the savings function: it is not clear whether or not $\left(\frac{\partial s_0}{\partial \sigma_1}(\sigma, \sigma) - \frac{\partial s_0}{\partial \sigma_2}(\sigma, \sigma) \right)$ is zero and, if not, what sign it takes. Its sign might depend on other features of the environment. In contrast, we will see that $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$ can be characterized more cleanly. This latter definition is more useful for related but distinct questions, such as what is the arrangement of risk such that precautionary saving is minimized? This may be an interesting question if the amount of precautionary saving provides a good approximation to the welfare cost of risk.

5.3 Quantifying Complementarity of Precautionary Saving

We begin the analysis by numerically simulating a realistic life-cycle consumption and savings model. In the following section (section 5.4) we return to the three period model to examine the conditions for complementarity in more theoretical detail.

We choose the following basic parametrization: the household is ‘born’ at age 20,

¹³We must take care in this definition that total ex-ante wealth is held constant, and specifically that risk sequence $\langle \sigma + \epsilon, \sigma - \epsilon \rangle$ yields the same risk to lifetime wealth as does $\langle \sigma, \sigma \rangle$.

and receives labour income for 45 years (from 20 to 65). The household then retires for another 15 years, in which time it finances consumption out of savings. We set $R = 1.02$ and $\beta R = 1$. Initial income is 1 unit. Income grows at 1% per year over working age. The household has CRRA preferences with coefficient of relative risk aversion of 2. This is closer to the standard coefficient estimated using microdata such as in Attanasio and Weber (1995) and is lower than estimates used in the macro literature (see Barro (2006)).

In keeping with thinking of this as a 3-period model, we divide the planning horizon into three tranches of fifteen years, corresponding to ages 21-35, 36-50 and 51-65.¹⁴ We assume no income risk in the first tranche and then a constant yearly variance of permanent and transitory shocks in the final two tranches: the yearly variance of permanent shocks is 0.02 (standard deviation of 14%); the yearly variance of transitory shocks is 0.04 (standard deviation of 20%). Income fluctuations are lognormally distributed. We exclude risk in the first 15 periods because we are interested in precautionary saving for risk in the medium and long term. By excluding risk in the first tranche we are able to identify pure life-cycle saving when we switch off risk in these latter periods. Allowing risk in the first tranche does not affect results but obscures their interpretation somewhat.

We then run the following experiments. First we switch off all risk over all periods. Second we keep risk (both permanent and transitory shocks) in the middle tranche alone. Third we allow for risk in the final tranche alone. Finally we allow for risk in both the middle and late tranches as standard. As discussed in section 5.2, when running these experiments we are careful to account for the change in the variance of life-time wealth precisely. Using unadjusted shocks, the variance of life-time wealth is higher when facing both tranches of risk than the total variance from each risk in isolation because the variances of a multiplicative process do not sum precisely. We deal with this problem by reweighting the shocks when the household faces both tranches of risk. Further details are given in appendix A5.1.

Table 5.1 shows the results from these simulations. It shows saving at the end of the 15th period, just before income risk kicks in. The first column shows results for the baseline environment. The first row shows savings when there is no income risk. This

¹⁴The household lives for another 15 years after retiring but faces no risk, so faces a trivial planning problem.

therefore represents pure life-cycle saving caused by the pattern of life-cycle income and the retirement period. The second row shows accumulated saving in the standard model with risk in both middle and final tranches. This accounts for both life-cycle and precautionary saving. The third row subtracts life-cycle saving (row 1) to leave precautionary saving alone. The fourth row shows saving when there is just risk in the middle tranche. The fifth row shows saving when there is risk in the final tranche alone. Note that precautionary saving is much larger in the 4th row than the 5th, because permanent shocks in mid life persist until retirement and so have a greater effect on life-time wealth. The sixth row shows the sum of the precautionary saving in these two scenarios (row 4 + row 5 - 2*row 1). This can be thought of as the total precautionary effect of the two tranches of risk. The seventh row shows the difference between the total precautionary effect and standard precautionary saving (row 6 - row 3). We interpret this as the complementarity effect of precautionary saving for this environment. The household can save the equivalent of 40% of its initial yearly income less because of the sequencing of risks. The final row represents this complementarity effect as a percentage of precautionary saving (in row 3): complementary saving is around 16% of total precautionary saving, a noticeable sum.

The remaining columns of table 5.1 show the equivalent results when we vary the parameters. The second column shows saving and complementarity when we reduce the variance of permanent shocks to 0.01 and the variance of transitory shocks to 0.02. The complementarity effect is reduced to around 15% of initial income and a little over 9% of precautionary wealth. The lower variance of permanent shocks itself reduces precautionary saving by around 40%. The next two columns show results with the same configurations of income risk but an enhanced life-cycle motive. In this scenario the household lives for 20 years after retirement (dies at 85). Moreover we set income to be flat in the final third of working life (income still grows at 1% pa until age 50). With the higher variance of income shocks the complementarity effect is slightly reduced to 14.2%, indicating that the higher life-cycle saving crowds out the complementarity effect from precautionary saving. Comparing the 4th to the 2nd columns we see a similar effect: complementarity is reduced from 9.3% to 8.4%.

The 5th, 6th and 7th columns of table 5.1 all show results when the baseline variance of permanent shocks is held at 0.02 and the variance of transitory shocks at 0.04, but some

other feature of the environment is changed. In the 5th column, we remove transitory shocks. We see that saving behaviour is almost identical showing that the vast majority of the effect comes through permanent shocks. For the 6th column we reduce the coefficient of relative risk aversion to 1. Total precautionary saving is reduced by 25% and the complementarity effect is reduced to 12.5% of precautionary saving. The final column shows results when we don't reweight permanent shocks. The results are very similar here to the base case, indicating that the results are not caused by how we account for the variance of life-time wealth.

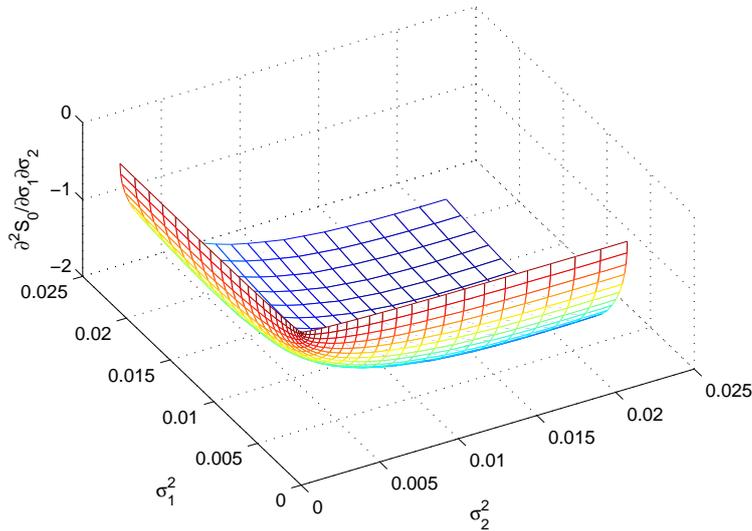
Figure 5.1 links our results to the formal definition of complementarity. This figure shows numerical calculations of the complementarity function $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$ at various levels of risk. The horizontal axes display the variances of permanent shocks. We tie the variances of transitory shocks to be double those of the permanent shocks. The vertical axis shows the amount of complementarity to small changes in risk around these levels. We term this 'local' complementarity. This figure is computed by solving the model numerically on a 30-by-30 grid of income risks. A negative amount here shows local complementarity. A positive amount would show local excessiveness. Here there is local complementarity at all levels of risk except at zero. Of course, the results in table 5.1 could be obtained by integrating this surface over larger changes in risk.

Table 5.1: Simulated Household Wealth at Age 35 as % of Initial Income

	Environment Variance of perm. shocks	Baseline		High life-cycle		No trans. shocks	Low risk aversion	No re-weighting
		High	Low	High	Low			
(1)	No risk	68	68	176	176	68	68	68
(2)	Risk throughout career	322	225	388	306	321	261	325
(3)	Precautionary wealth*	254	157	212	131	253	193	257
(4)	Risk in mid-career only	306	207	381	296	305	244	306
(5)	Risk in late-career only	124	100	212	197	123	109	124
(6)	Total prec. effect**	294	172	242	141	292	217	294
(7)	Complementarity effect***	41	15	30	11	40	24	37
(8)	as % of prec. wealth****	16.0	9.3	14.2	8.4	15.7	12.5	14.4

Notes: * = row(2)-row(1)
 ** = (4)-(1) + (5)-(1)
 *** = (6)-(3)
 **** = 100*(7)/(3)
 ‘High’ variance of permanent shocks is 0.02 per year. ‘Low’ is 0.01 per year.
 In ‘baseline’ scenario expected income grows at 1% pa; households are ‘born’ at age 20 and work for 45 years, then live for 15 years in retirement. In ‘high life-cycle’ scenario income growth is flat in the last third of working life and retirement lasts 20 years. In ‘no re-weighting’ scenario, the variances of permanent shocks are not reweighted.
 See text for more details.

Figure 5.1: Precautionary Saving is Complementary for Overall Life-Cycle Risks



Notes: $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$ on the vertical axis is the complementarity function. Horizontal axes show baseline variances of permanent shocks in mid-career (σ_1^2) and in late career (σ_2^2). I display variances on the horizontal axes but emphasize that the differentiation is with respect to the standard deviation.

We set $\beta = 0.98$ and $\beta R = 1$. Initial income is 1 unit and income growth is 1% per year over the whole life time. Income fluctuations are lognormally distributed. The coefficient of relative risk aversion is 2. Households work for 45 years then live for 15 years in retirement. Initial-period saving is used. See text for more details.

5.4 Analytical Characterization of Complementarity

In this section we study the complementarity effect by obtaining an analytical characterization of household saving in a stripped-down version of the model.

5.4.1 Complementarity for Permanent Shocks

We consider the 3-period consumption and saving model given in section 5.2 when the agent faces only permanent risk. For small risks in period 1 and 2 with standard deviations $\sigma_1 = \sigma_2 = \sigma$, and for $\beta = R = 1$ we obtain the following result, derived in appendix A5.2:

$$\begin{aligned} \frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} \Big|_{\sigma_1 = \sigma_2 = \sigma} = & \\ & -\sigma^2 \frac{\mathcal{A} c_0^3 \left(-c_0 u^{(3)}(c_0)^3 - 2c_0 u^{(4)}(c_0) u^{(3)}(c_0) u''(c_0) + 12u^{(4)}(c_0) u''(c_0)^2 + 3c_0 u^{(5)}(c_0) u''(c_0)^2 \right)}{u''(c_0)^3} \\ & + \mathcal{O} \left((\sigma)^3 \right) \end{aligned} \quad (5.11)$$

where \mathcal{A} is some positive constant $u^{(n)}(\cdot)$ denotes the n^{th} derivative of $u(\cdot)$ and $c_0 = W_0 - s_0$ is period-0 consumption. σ_1, σ_2 are the standard deviations of income innovations in each subsequent period, and we look at $\sigma_1 = \sigma_2 = \sigma$.¹⁵ $(\sigma)^3$ denotes the cube of the standard deviation (and not some parameter of the third moment of the income distribution). We re-weight the risks as described in section 5.2 and appendix A5.2.

According to equation 5.11, complementarity is linear in the variance of risk, σ^2 . We can most conveniently reformulate equation 5.11 in terms of the coefficient of relative prudence, $p_r(c)$, defined as $-\frac{cu'''(c)}{u''(c)}$:

$$\frac{\partial^2 s}{\partial \sigma_1 \partial \sigma_2} \Big|_{\sigma_1 = \sigma_2 = \sigma} \approx \mathcal{A}_1 \sigma^2 c \left(c^2 p_r''(c) - \frac{7}{3} c p_r'(c) p_r(c) + 2c p_r'(c) - 2p_r(c) - \frac{5}{3} p_r(c)^2 \right) \quad (5.12)$$

where \mathcal{A}_1 is another positive constant. The right hand side of equation 5.12 is hard to interpret but we can state conditions to put a sign on it. First note that for CRRA preferences of the form $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$, the coefficient of relative prudence is constant at $1 + \gamma$. The first and second derivatives of the relative prudence function are therefore zero and the complementarity function in 5.12 simplifies to $-\mathcal{A}_1 \sigma^2 c \left(2p_r(c) + \frac{5}{3} p_r(c)^2 \right) =$

¹⁵The income innovations here are actually transformed from those in section 5.2. See appendix A5.2 for more details.

$-\mathcal{A}_1 \frac{3}{4} \sigma^2 c (7\gamma^2 + 20\gamma + 13) < 0$. CRRA preferences therefore display complementarity and this complementarity is increasing in the prudence parameter. Figure 5.2 shows $\frac{\partial^2 s}{\partial \sigma_1 \partial \sigma_2}$ for permanent risks and for CRRA preferences, obtained from simulations. This figure shows complementarity for small changes of risk not just from the perfect certainty baseline, but also at higher levels of risk. The figure also shows complementarity when the baseline variances of innovations (σ_1^2 and σ_2^2) are not equal. The negative sign on the surface indicates local complementarity at all levels of risk.

More generally, equation 5.12 states that the complementarity function is negative as long as the relative prudence function is near flat. For example, the class of preferences which are HARA and for which risk tolerance is defined at zero wealth induces complementarity. CRRA functions are included in this class. This subclass of HARA functions has coefficient of relative prudence $p_r = \frac{x}{ax+b}$ for $b \geq 0$ and $0 \leq a < 1$. These preferences therefore have relative prudence which is weakly increasing and weakly concave. For these functions, complementarity is stronger the smaller are a and b .¹⁶

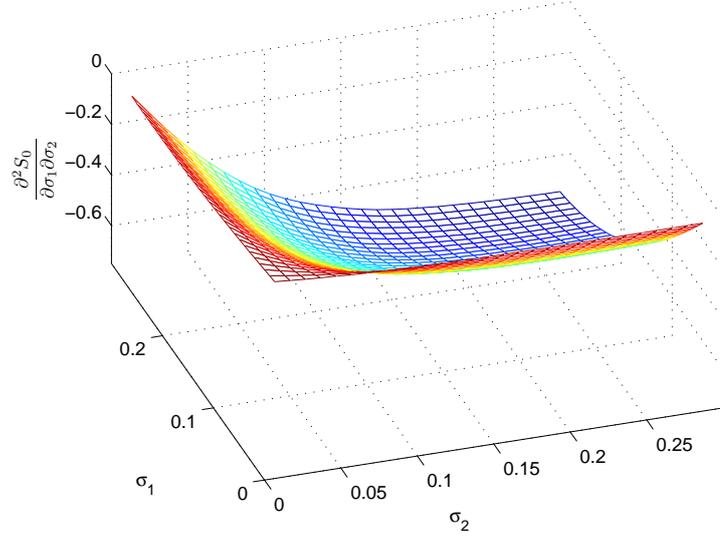
We gain further intuition into these standard cases by considering a counter-example. For Stone-Geary preferences of the form $u(c) = \ln(c - \underline{c})$, where \underline{c} is a minimum consumption need, savings are in fact excessive near \underline{c} . To see this note that relative prudence term is $p_r = \frac{2c}{c-\underline{c}}$ and the complementarity function given by equation 5.12 simplifies to $-\frac{\mathcal{A}c^3(4c-9\underline{c})}{(c-\underline{c})^3}$. This function is positive for any $\underline{c} < c < \frac{9}{4}\underline{c}$. For these preferences, relative prudence rises to infinity as c approaches \underline{c} from above. It seems that near the consumption floor households are particularly averse to a series of consecutive negative shocks, and so the consecutive risks amplify precautionary behaviour.

5.4.2 Excessiveness for Small Transitory Risks

We now turn to the characterization of first-period saving when income shocks in periods 1 and 2 are transitory. For small risks with standard deviations $\sigma_1 = \sigma_2 = \sigma$, and for $\beta = R = 1$ we obtain the following result, also derived in appendix A5.2:

¹⁶The complementarity function simplifies to $-\frac{x^3(12(a+1)b+a(6a+5)x)}{3(ax+b)^3}$

Figure 5.2: Precautionary Saving is Complementary for Permanent Risks Alone



Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$: the complementarity function. On the horizontal axes are standard deviations of income innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Preferences are CRRA with coefficient of relative risk aversion of 2.

$$\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} \Big|_{\sigma_1 = \sigma_2 = \sigma} = - \frac{\mathcal{A} \sigma^2 (u^{(3)}(c_0)^3 + 2u^{(4)}(c_0)u^{(3)}(c)u''(c_0) - 3u^{(5)}(c_0)u''(c_0)^2)}{u''(c_0)^3} + \mathcal{O}((\sigma)^3) \quad (5.13)$$

Given that the denominator, $u''(c_0)^3$, < 0 for a concave utility function, savings exhibit complementarity if and only if:

$$u^{(3)}(c_0)^3 + 2u^{(4)}(c_0)u^{(3)}(c_0)u''(c_0) - 3u^{(5)}(c_0)u''(c_0)^2 < 0 \quad (5.14)$$

We can most conveniently reformulate inequality 5.14 in terms of absolute prudence concepts. Letting $p_a(c)$ denote the coefficient of absolute prudence, given by $-u'''(c)/u''(c)$, and dropping the 0 subscript, then:

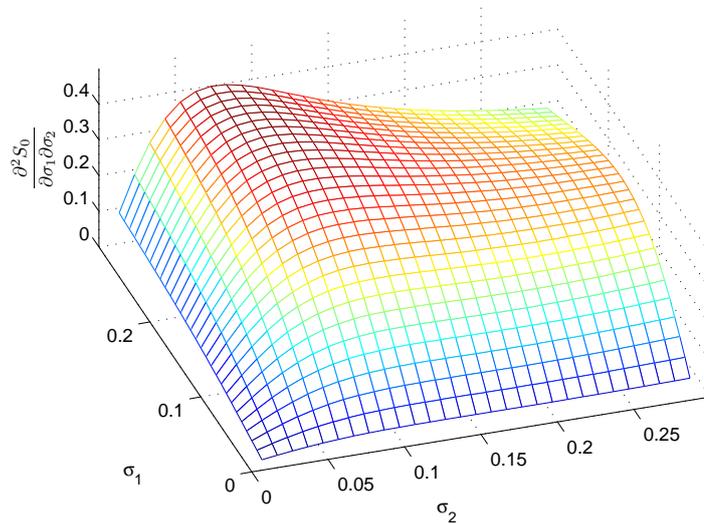
$$\frac{\partial^2 s}{\partial \sigma_1 \partial \sigma_2} \approx \mathcal{A}_1 \sigma^2 \left(p_a''(c) - \frac{7}{3} p_a'(c) \cdot p_a(c) \right) \quad (5.15)$$

for some positive constant \mathcal{A}_1 . A necessary and sufficient condition for complementarity to small risks here is $p_a''(c) \leq \frac{7}{3} p_a'(c) \cdot p_a(c)$.

For CRRA preferences the coefficient of absolute prudence is $\frac{1+\gamma}{c}$. For these preferences

$p_a''(c) = \frac{2(1+\gamma)}{c^3} > 0$ and $p_a'(c) = -\frac{(1+\gamma)}{c^2} < 0$. Therefore, $\frac{\partial^2 s}{\partial \sigma_1 \partial \sigma_2} > 0$ and savings exhibit not complementarity in this case but excessiveness. Figure 5.3 shows $\frac{\partial^2 s}{\partial \sigma_1 \partial \sigma_2}$ for transitory risks and for the CRRA utility function, obtained from simulations. The figure shows that saving is excessive at all levels of baseline risk. It further shows that the absolute height of the function is comparable to that for permanent shocks given in figure 5.2. However, transitory risks in the 3-period model have a far greater effect on life-time wealth than they do in a model with a longer time horizon. As discussed in section 5.2 we emphasize that the effect of transitory shocks is quantitatively small for realistic parametrizations.

Figure 5.3: Precautionary Saving is Excessive for Transitory Risks Alone



Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$: the complementarity function. On the horizontal axes are standard deviations of income innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Preferences are CRRA with coefficient of relative risk aversion of 2.

More generally, saving is excessive if prudence is declining and convex, because then always $p_a'(c) < 0$ and $p_a''(c) > 0$. An intuitive reason for this sufficient condition is as follows: the wealth innovations caused by transitory income shocks are independent. If prudence is declining and convex, average prudence is greater the more independent risks are added (a consequence of Jensen's inequality). Therefore the interaction of risks amplifies the precautionary saving motive. It seems intuitively plausible that prudence be declining and convex because this implies that wealthier households have less need to avoid a gamble (of constant variance) but that the rate of decline of prudence is diminishing with wealth. Indeed this condition on prudence is satisfied by all preferences in the HARA class for which

risk tolerance is not constant, i.e. notably excluding CARA and quadratic preferences.¹⁷

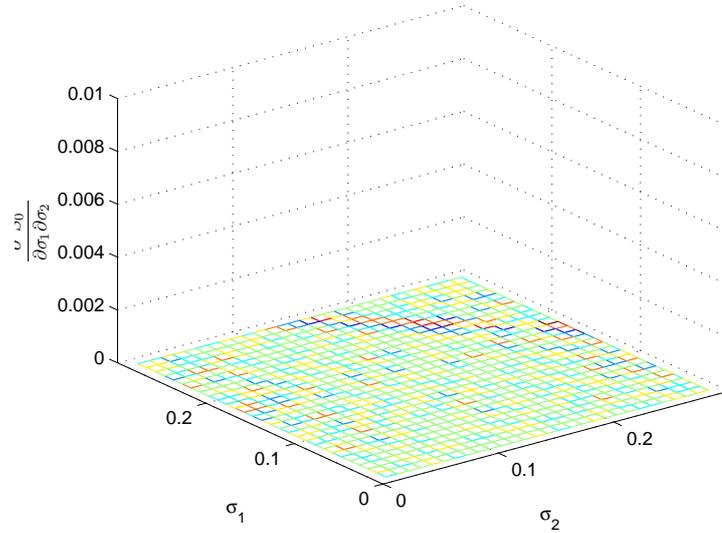
5.4.3 Further Intuition

The obvious question is why CRRA preferences induce excessive saving for transitory shocks but complementary saving for permanent shocks. Paradoxically we gain the best intuition for CRRA preferences by considering saving under CARA preferences. For CARA preferences, the first and second derivatives of absolute prudence are zero. Therefore, expression 5.15 implies that, for transitory risks, saving is neutral: it is neither excessive nor complementary. It is easily checked that saving is also neutral for larger risks (see Caballero (1990)). Figure 5.4 shows this result graphically. This result is obvious when we remember two facts. First, transitory risks are independent across time. Second, the wealth level is irrelevant under CARA preferences when saving for a given risk. From the viewpoint of the middle period therefore, the precautionary motive for the final-period risk does not depend on the outcome of the most recent shock. Returning to the initial period, the household must save for both middle- and final-period risks. However, because the middle-period outcome will not change behaviour, the size of the middle risk does not affect the precautionary motive for the final risk. Symmetrically, the size of the final risk does not affect the precautionary motive for the middle risk. In short, the risks don't interact in the initial saving decision.

In contrast, and just like CRRA preferences, saving is complementary for permanent shocks. Figure 5.5 shows this result. The argument above for why CARA preferences induce neutral savings no longer holds: in this case the variance of final-period shocks now depends on the middle-period outcome. The risks are no longer independent. In fact the variance of middle-period shocks now decreases, the lower is the outcome in the first period. A prudent saver is most concerned about the lowest possible outcomes and places most weight on these outcomes for decision making. Therefore, precautionary savings to cover both risks combined need not to be as high as the sum of savings to match each risk on its own. In short, it seems the risk pattern itself provides a kind of insurance.

¹⁷The result for HARA preferences is easily shown by differentiating the coefficient of risk tolerance $-\frac{u'(c)}{u''(c)}$ with respect to c . Re-arranging we find that $p(c) = k.r(c)$ for some constant k and where $r(c)$ is the coefficient of absolute risk aversion $= -\frac{u''(c)}{u'(c)}$. Note that $r(c)$ is decreasing and convex for HARA functions (except those with constant risk tolerance), therefore so must be $p(c)$.

Figure 5.4: CARA Preferences: Saving is Neutral for Transitory Risks



Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{u'(s_0)}{\partial \sigma_1 \partial \sigma_2}$: the complementarity function. On the horizontal axes are standard deviations of income innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Coefficient of absolute risk aversion is 3.

5.4.4 Classifying Utility Functions

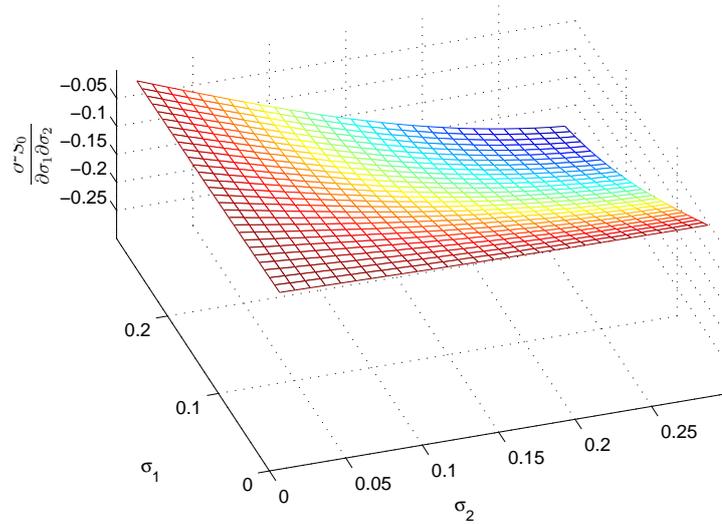
We can classify utility functions according to whether they induce complementary or excessive saving more generally. Comparing equations 5.13 and 5.11, the complementary function for permanent shocks is lower than for transitory shocks by a term in $u^{(4)}(c)u''(c)^2$.

We therefore classify utility functions as follows:

Proposition. *As long as $u^{(4)}(c)$ is negative then if a utility function induces excessive saving for permanent shocks then it also induces excessive saving for transitory shocks.*

Saving is excessive for Stone-Geary preferences and permanent risks when consumption is near the ‘breadline’. Therefore saving is excessive for these preferences and transitory shocks also. On the other hand, all other HARA preferences (i.e. except Stone-Geary) lie in the middle ground: they induce complementarity for permanent shocks but excessiveness for transitory shocks. To complete the classification it is instructive to construct a utility function that is complementary for both permanent and transitory shocks. According to the discussion in section 5.4.2, a utility function induces complementarity if it has

Figure 5.5: ...While Saving is Complementary for Permanent Risks



Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$: the complementarity function. On the horizontal axes are standard deviations of wealth innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Coefficient of absolute risk aversion is 3.

increasing absolute prudence. An example is the following:

$$u(c) = -e^{-\frac{c+1}{2}} - \sqrt{\frac{\pi}{2}}(c+1) \times \operatorname{erf}\left(\frac{c+1}{\sqrt{2}}\right) + 2(c+1)$$

where $\operatorname{erf}(x)$ represents the error function, the anti-derivative of $\frac{2}{\sqrt{\pi}}e^{-x^2}$. This utility function has positive third derivative and negative fourth derivative and coefficient of absolute prudence of c . Such a utility function has many undesirable features.¹⁸ Nevertheless, a household with these preferences has complementary saving even for transitory shocks because it is little affected by a negative shock in the middle period. In fact, the household has less desire to save following a bad shock, because absolute prudence is lower with lower wealth. Returning to the initial period, the household therefore need not save much more for consecutive risks than for just a single risk. It is happier to save little and to have more equal consumption across time in expectation.

¹⁸For example, it does not obey inada conditions because $u(0) = 2$. Moreover at high levels of wealth it displays high prudence but low risk aversion. The high prudence arises because, even though households do not lose much expected utility from risk, their elasticity of substitution is so high that they equally lose little from allocating consumption to the future. The latter affect dominates so they precautionary save and increasingly so at higher levels of wealth.

5.4.5 Final Remarks on the Characterization of Complementarity

To our knowledge, no such similar conditions have been derived before. However the results for transitory shocks relate to the literature on multiple risk bearing in static settings in papers by Pratt and Zeckhauser (1987), Kimball (1993) and Gollier and Pratt (1996), summarized neatly in the last paper. These papers elucidate the related and intuitively attractive notions of standard risk aversion (Kimball), proper risk aversion (Pratt and Zeckhauser) and risk vulnerability (Gollier and Pratt). All are formalizations of the idea that background risks should make agents more averse to new risks; for example an agent should be more averse to investing in equities if exposed to high labour-market risk.¹⁹ While this literature concerns risk aversion and portfolio choice, it seems intuitive that such effects carry over to precautionary saving when risks are independent (i.e. transitory). The results for permanent shocks differ of course because the (independent) shocks induce dependence in the innovations to life-time wealth. Nevertheless it is striking that for these two standard specifications of risk the results should be so contrasting.

5.5 Conclusion

In this chapter we define and examine the concept of complementarity in precautionary saving. Intuitively, precautionary saving should be complementary because of its contingent nature. Emergencies occur only rarely and precautionary savings are only rarely needed. Therefore the accumulated stock can presumably be put to other ends, specifically as rolled-over precautionary savings against subsequent income risk.

On the quantitative side, we simulate a standard life-cycle consumption and saving model with both permanent and transitory income processes. In a range of realistic parametrizations, we calculate complementarity to account for around 8-16% of precautionary savings. This effect is driven almost entirely by permanent shocks: the effect from transitory shocks is negligible.

We then study the complementarity effect in more detail by analyzing a stripped-down version of the model with only 3 periods. We show that permanent shocks admit comple-

¹⁹These restrictions on utility functions can be summed up neatly in notation related to that above. If $r(c) \equiv -u''(c)/u'(c)$ is the coefficient of absolute risk aversion at consumption/wealth level c , then a necessary and sufficient condition for standard risk aversion is that both $r(c)$ and $p(c)$ be decreasing.

mentary savings for a general class of preferences, notably including the standard CRRA form. However, we find two instances which depart from our basic intuition. First, consecutive transitory shocks amplify precautionary savings for standard preferences. Second, even consecutive permanent risks can amplify precautionary savings when households have minimum consumption needs. The effect from transitory shocks is small, however, for empirically-plausible risk sizes. Moreover, the ‘breadline’ effect is only local: saving is complementary for wealthier households. We conclude that complementary saving is the norm.

The present study could be extended in several ways. First, it would be interesting to see how liquidity constraints affect the results. In general, constraints exacerbate the precautionary motive (Carroll and Kimball (2001)). However, constraints will not bind along the expected income path. The intuition above therefore carries through: households facing a constraint should not need to save much more for consecutive risks than for just a single risk.

A related question is how the distribution of shocks affects the results. As a specific example, what if households can receive zero (or very small) income in any period, for example due to unemployment? The analysis presented in section 5.4 cannot answer this question because it applies only to small, local risks. This question could, however, be studied quantitatively.

In this chapter we have focused on permanent and transitory processes. Recent work (such as Guvenen (2009)) has argued that idiosyncratic durable income shocks are not permanent, but only very persistent. The results here apply to the extreme cases of a more general autoregressive process. Moreover, we have interpreted the results purely in terms of the structure of income shocks and the shape of the instantaneous felicity function. Future work could explore the intrinsic dynamics of the problem in more detail.

Finally, this chapter suggests a much broader research agenda. Recalling the introductory comments we may ask the following questions: how do other savings motivations complement each other? How does the presence of different savings technologies affect this complementarity? What are the implications for household saving if pension and housing wealth become more liquid, perhaps because of financial innovation?

A5 Appendix to chapter 5

A5.1 Reweighting The Permanent Shock Process

This appendix gives more detail on the reweighting of permanent income shocks discussed in section 5.3.

We consider permanent shocks ψ_i for $i = 1 \dots T$ with mean 1 and variance $\sigma_{\psi_i}^2$. By Goodman's rule²⁰ it is easy to see that $\text{Var}(y_t) = \text{Var}(\prod_{i=1}^t \psi_i) = \prod_{i=1}^t (\sigma_{\psi_i}^2 + 1) - 1$.²¹ In this discussion we ignore the first, risk-free tranche of the life-cycle and consider two tranches: $1..T_0$ and $T_0..T$, both containing permanent risk.

The problem is the following. When the household faces risk in the first tranche of working life alone then at time $T_0 + 1$, we have $\sigma_{\psi_{T_0+1}}^2 = 0$ and so:

$$\text{Var}(y_{T_0+1}) = \text{Var}\left(\prod_{i=1}^{T_0} \psi_i\right)$$

Similarly when the household faces risk in the second tranche alone then

$$\text{Var}(y_{T_0+1}) = \sigma_{\psi_{T_0+1}}^2$$

When the household faces both risks combined then

$$\begin{aligned} \text{Var}(y_{T_0+1}) &= \text{Var}\left(\prod_{i=1}^{T_0+1} \psi_i\right) \\ &= \prod_{i=1}^{T_0+1} (\sigma_{\psi_i}^2 + 1) - 1 \\ &= \prod_{i=1}^{T_0} (\sigma_{\psi_i}^2 + 1) (\sigma_{\psi_{T_0+1}}^2 + 1) - 1 \\ &= \prod_{i=1}^{T_0} (\sigma_{\psi_i}^2 + 1) - 1 + \prod_{i=1}^{T_0} (\sigma_{\psi_i}^2 + 1) \sigma_{\psi_{T_0+1}}^2 \\ &> \prod_{i=1}^{T_0} (\sigma_{\psi_i}^2 + 1) - 1 + \sigma_{\psi_{T_0+1}}^2 \\ &= \text{Var}\left(\prod_{i=1}^{T_0} \psi_i\right) + \sigma_{\psi_{T_0+1}}^2 \end{aligned}$$

Therefore income risk in period $T_0 + 1$ is greater when the household faces both risks combined. By a simple induction argument we can show that the same holds for risks in time t for any $t > T_0$.

²⁰For uncorrelated random variables $\text{Var}(xy) = \text{Var}(x) E(y)^2 + \text{Var}(y) E(x)^2 + \text{Var}(x) \text{Var}(y)$.

²¹We ignore mean income growth here without loss of generality.

To solve this problem we reweight the risks to make the variance of lifetime wealth the same in all scenarios. To do this we create an alternative sequence of income shocks $\tilde{\psi}_i$ with variances $\sigma_{\tilde{\psi}_i}^2$ for $i > T_0$. We do this iteratively from period T_0 onwards as follows: when facing risk in the first tranche only, the variance of income $t > T_0$ is given as before by $\text{Var}\left(\prod_{i=1}^{T_0}\psi_i\right)$. When facing risk only in the second tranche then income risk at time t is given again as before by $\text{Var}\left(\prod_{i=T_0+1}^t\psi_i\right)$. When facing both risks combined and with the alternative income process the variance of income is

$$\begin{aligned}\text{Var}(\tilde{y}_t) &= \text{Var}\left(\prod_{i=1}^t\tilde{\psi}_i\right) \\ &= \text{Var}\left(\prod_{i=1}^{t-1}\tilde{\psi}_i\right)\sigma_{\tilde{\psi}_t}^2 + \sigma_{\tilde{\psi}_t}^2 + \text{Var}\left(\prod_{i=1}^{t-1}\tilde{\psi}_i\right)\end{aligned}\quad (16)$$

Setting this equal to the sum of the two isolated risks added together gives:

$$\text{Var}\left(\prod_{i=1}^{t-1}\tilde{\psi}_i\right)\sigma_{\tilde{\psi}_t}^2 + \sigma_{\tilde{\psi}_t}^2 + \text{Var}\left(\prod_{i=1}^{t-1}\tilde{\psi}_i\right) = \text{Var}\left(\prod_{i=1}^{T_0}\psi_i\right) + \text{Var}\left(\prod_{i=T_0+1}^t\psi_i\right)$$

solving for $\sigma_{\tilde{\psi}_t}^2$:

$$\sigma_{\tilde{\psi}_t}^2 = \frac{\text{Var}\left(\prod_{i=1}^{T_0}\psi_i\right) + \text{Var}\left(\prod_{i=T_0+1}^t\psi_i\right) - \text{Var}\left(\prod_{i=1}^{t-1}\tilde{\psi}_i\right)}{1 + \text{Var}\left(\prod_{i=1}^{t-1}\tilde{\psi}_i\right)}$$

A5.2 Derivation of the Approximations for $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$

This appendix gives more details of the derivation of the formulae in section 5.4

Translating the Budget Constraint

For the 3 period model with $\beta = R$ and with constant expected income we have household problem

$$V_0(W_0) = \max_{\{c_t(W_t):t=0,1,2\}} u(c_0) + \mathbb{E}_0(u(c_1) + u(c_2))$$

subject to

$$c_0 + c_1 + c_2 = y_0 + y_0\psi_1\xi_1 + y_0\psi_1\psi_2\xi_2$$

It is convenient to rewrite the budget constraint in terms of relative wealth innovations.

In the case of transitory risk the budget constraint is:

$$c_0 + c_1 + c_2 = y_0 + y_0\xi_1 + y_0\xi_2$$

We make the transformation that $a_0 = 3y_0$, $\epsilon_1 = \frac{1}{3}(\xi_1 - 1)$, and $\epsilon_2 = \frac{1}{3}(\xi_2 - 1)$.²² We can then write the budget constraint as

$$c_0 + c_1 + c_2 = a_0 + a_0\epsilon_1 + a_0\epsilon_2$$

such that ϵ_1 and ϵ_2 have mean zero and $\text{Var}(\epsilon_1) = \frac{1}{9}\sigma_{\xi_1}^2$ and $\text{Var}(\epsilon_2) = \frac{1}{9}\sigma_{\xi_2}^2$.

In the case of permanent shocks alone we can write the constraint

$$c_0 + c_1 + c_2 = y_0 + y_0\psi_1 + y_0\psi_1\tilde{\psi}_2$$

We make the transformations $a_0 = 3y_0$, $\epsilon_1 = \frac{2}{3}(\psi_1 - 1)$, and $\epsilon_2 = \frac{1}{3}\frac{(\psi_2 - 1)}{\sqrt{1 + \sigma_{\psi_1}^2}}$.²³ We can then write the budget constraint as:

$$c_0 + c_1 + c_2 = a_0 + a_0\epsilon_1 + a_0\frac{(1 + \frac{3}{2}\epsilon_1)}{\sqrt{1 + \frac{9}{4}\text{Var}(\epsilon_1)}}\epsilon_2$$

such that ϵ_1 and ϵ_2 have zero mean and $\text{Var}(\epsilon_1) = \frac{4}{9}\sigma_{\psi_1}^2$ and $\text{Var}(\epsilon_2) = \frac{\sigma_{\psi_2}^2}{9(1 + \sigma_{\psi_1}^2)}$. We have weighted ϵ_1 and ϵ_2 appropriately so that the variance of life-time wealth is held constant when we sum risks.

And we nest the permanent and transitory cases by writing the budget constraint:

$$c_0 + c_1 + c_2 = a_0 + a_0\epsilon_1 + a_0\frac{(1 + \frac{3}{2}\theta\epsilon_1)}{\sqrt{1 + \frac{9}{4}\text{Var}(\theta\epsilon_1)}}\epsilon_2$$

When $\theta = 0$ this collapses into the standard transitory process. When $\theta = 1$ it collapses to the case for permanent shocks. where e_1 and e_2 are random variables with mean zero and variances σ_1^2 and σ_2^2 . For the remainder of the proof we consider variation in σ_1^2 and

²²More generally, we can allow for non-constant expected income over the life cycle by allowing for $\epsilon_1 = \frac{y_1}{\sum y_i}(\xi_1 - 1)$, $\epsilon_2 = \frac{y_2}{\sum y_i}(\xi_2 - 1)$.

²³Again more generally, we can allow for non-constant expected income over the life cycle by allowing for $\epsilon_1 = \frac{y_1 + y_2}{\sum y_i}(\psi_1 - 1)$, $\epsilon_2 = \frac{y_2}{\sum y_i}\frac{(\psi_2 - 1)}{\sqrt{1 + \sigma_{\psi_1}^2}}$.

σ_2^2 rather than $\sigma_{\xi_1}^2$ and $\sigma_{\xi_2}^2$, and $\sigma_{\psi_1}^2$ and $\sigma_{\psi_2}^2$. The approximation can be adjusted for the prior variances by scaling up by a constant.

For the rest of this derivation I assume that the random variables e_1 and e_2 are symmetric and binary, with outcomes $\pm\sigma_1$ and $\pm\sigma_2$ each with probability $\frac{1}{2}$. This restriction is without loss of generality: for small mean-zero risks, only the second moment matters, other aspects of the distribution of risk are irrelevant.

Obtaining the Approximation

In period 0, the agent therefore faces an optimization condition (Euler equation) of the following form:

$$\begin{aligned} u'(W_0 - s_0) - \mathbb{E}_0 \left(u' \left(s_0 + e_1 - s_1 \left(s_0 + e_1, W_0 \frac{(1 + \frac{3}{2}\theta\epsilon_1)}{\sqrt{1 + \frac{9}{4}\text{Var}(\theta\epsilon_1)}} \sigma_2 \right) \right) \right) \\ = f(s_0(\sigma_1, \sigma_2), \sigma_1, \sigma_2) \\ = 0 \end{aligned}$$

where $s_0(\sigma_1, \sigma_2)$ represents optimal savings as a function of income risk standard deviations, and the savings problem at $t = 1$, $s_1(s_0 + e_1, \sigma_2)$, is solved. Using the implicit function theorem twice, we can derive $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$:

$$\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} = \frac{-\frac{\partial^2 f}{\partial \sigma_1 \partial \sigma_2} - \frac{\partial s_0}{\partial \sigma_2} \frac{\partial^2 f}{\partial s_0 \partial \sigma_1} - \frac{\partial s_0}{\partial \sigma_1} \left(\frac{\partial^2 f}{\partial s_1 \partial \sigma_2} + \frac{\partial s_0}{\partial \sigma_2} \frac{\partial^2 f}{\partial s_0^2} \right)}{\frac{\partial f}{\partial s_0}} \quad (17)$$

We want to derive a Taylor-series approximation of this expression for $\sigma_1 = \sigma_2 = \sigma$ around 0. To break this expression up we should bear in mind the following algebra for such expansions:

If $\text{Tayl}(f(x), n, x_0)$ represents the Taylor expansion of $f(x)$ to order n at x around x_0 , i.e. $\text{Tayl}(f(x), n, x_0) = f(x_0) + \sum_{k=1}^n \frac{(x-x_0)^k}{k!} f^{(k)}(x_0) + \mathcal{O}((x-x_0)^{n+1})$, then:

$$\text{Tayl}(f(x) + g(x), n, x_0) = \text{Tayl}(f(x), n, x_0) + \text{Tayl}(g(x), n, x_0)$$

$$\text{Tayl}(f(x)g(x), n, x_0) = \text{Tayl}(f(x), n, x_0) \text{Tayl}(g(x), n, x_0)$$

and if $f(x_0) \neq 0$, then:

$$\text{Tayl}\left(\frac{1}{f(x)}, n, x_0\right) = \text{Tayl}\left(\frac{1}{\text{Tayl}(f(x), n, x_0)}, n, x_0\right)$$

Therefore, we can break the expression down and perform successive approximations on constituent parts to gain an accurate approximation to the whole.

Without displaying all calculations, we give an illustration of the approximations made, picking the first term in the numerator in equation 17. We now let a_{1h} , a_{1l} , s_{1h} , s_{1l} , c_{1h} , c_{1l} represent assets at the start of period 1, saving and consumption after resolution of high/low period-1 shocks.

We can expand each term in 17 in terms of savings functions and utility. For example:

$$\begin{aligned} \frac{\partial^2 f}{\partial s_0 \partial \sigma_2} &= \frac{(1 + \frac{3e_1\theta}{2})}{2\sqrt{\frac{9e_1^2\theta^2}{4} + 1}} W_0 \left(u^{(3)}(c_{1h}) \frac{\partial s_{1h}}{\partial \sigma_2} \left(1 - \frac{\partial s_{1h}}{\partial a_{1h}} \right) + u''(c_{1h}) \frac{\partial^2 s_{1h}}{\partial a_{1h} \partial \sigma_2} \right) \\ &\quad + \frac{(1 - \frac{3e_1\theta}{2})}{2\sqrt{\frac{9e_1^2\theta^2}{4} + 1}} W_0 \left(u^{(3)}(c_{1l}) \frac{\partial s_{1l}}{\partial \sigma_2} \left(1 - \frac{\partial s_{1l}}{\partial a_{1l}} \right) + u''(c_{1l}) \frac{\partial^2 s_{1l}}{\partial a_{1l} \partial \sigma_2} \right) \end{aligned} \quad (18)$$

We can further derive expressions for each component part in terms of the utility function:

$$\begin{aligned} \frac{\partial s_1}{\partial \sigma_2} &= \frac{\frac{1}{2}(u''(c_{2h}) - u''(c_{2l}))}{-u''(c_1) - E[u''(c_2)]} \\ \frac{\partial s_1}{\partial a_1} &= \frac{u''(c_1)}{-u''(c_1) - E[u''(c_2)]} \\ \frac{\partial^2 s_1}{\partial a_1 \partial \sigma_2} &= \frac{u^{(3)}(c_1)(1 - \frac{\partial s_1}{\partial \sigma_2} \frac{\partial s_1}{\partial a_1}) + u^{(3)}(c_{2h})(1 - \frac{\partial s_1}{\partial \sigma_2} \frac{\partial s_1}{\partial a_1}) - u^{(3)}(c_{2l})(1 - \frac{\partial s_1}{\partial \sigma_2} \frac{\partial s_1}{\partial a_1})}{-u''(c_1) - E[u''(c_2)]} \end{aligned}$$

We now begin the approximations. These are made in several stages. First, note that $c_{2h} = s_1 + \sigma$ and $c_{2l} = s_1 - \sigma$, so, to second order: $f(c_{2i}) \approx f(s_1) \pm \sigma \cdot f'(s_1) + \frac{\sigma^2}{2} \cdot f''(s_1)$.

Using these approximations we get:

$$\frac{\partial s_1}{\partial \sigma_2} = -\sigma \frac{u'''(s_1)}{u''(c_1) + u''(s_1)} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial s_1}{\partial a_1} = \frac{u''(c_1)}{u''(c_1) + u''(s_1)} - \sigma^2 \frac{u''(c_1)u'''(s_1)}{2(u''(c_1) + u''(s_1))^2} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial^2 s_1}{\partial a_1 \partial \sigma_2} = \sigma \frac{(u''(s_1)u'''(c_1)u'''(s_1) + u''(c_1)u'''(s_1)^2 - u''(c_1)^2 u'''(s_1) - u''(c_1)u''(s_1)u'''(s_1))}{(u''(c_1) + u''(s_1))^3} + \mathcal{O}(\sigma^3)$$

Second, we now relate c_1 to s_1 for small σ . As shown by Kimball (1990), the growth in consumption depends on the coefficient of absolute prudence, $s_1 \approx c_1 + \frac{\sigma^2}{2} \frac{u'''(c_1)}{u''(c_1)}$. Inserting this expression into our three formulae gives:²⁴

$$\frac{\partial s_1}{\partial \sigma_2} = -\sigma \frac{u'''(c_1)}{2u''(c_1)} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial s_1}{\partial a_1} = \frac{1}{2} - \sigma^2 \frac{u''''(c_1)}{8u''(c_1)} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial^2 s_1}{\partial a_1 \partial \sigma_2} = \sigma \frac{(u'''(c_1)^2 - u''(c_1)u''''(c_1))}{4u''(c_1)^2} + \mathcal{O}(\sigma^3)$$

Finally, we relate c_{1h} to c_{1l} , and apply to equation 18. Considering first-period sav-
ings/consumption as a function of assets:

$$s_{1i} = s_1(s_0 \pm \sigma) \approx s_1(s_0) \pm \sigma \cdot s_1'(s_0) + \frac{\sigma^2}{2} \cdot s_1''(s_0)$$

and similarly:

$$\begin{aligned} c_{1i} = c_1(s_0 \pm \sigma) &\approx c_1(s_0) \pm \sigma \cdot c_1'(s_0) + \frac{\sigma^2}{2} \cdot c_1''(s_0) \\ &= c_1(s_0) \pm \sigma \cdot (1 - s_1'(s_0)) - \frac{\sigma^2}{2} \cdot s_1''(s_0) \end{aligned}$$

Inserting these approximations into equation 18 and simplifying leads to the final expres-

²⁴These derivatives can all be expressed in terms of risk prudence and tolerance concepts. $\frac{\partial s_1}{\partial e_2} = \frac{\epsilon}{2} \cdot p(c) + \mathcal{O}(e^3)$, $\frac{\partial s_1}{\partial a_1} = \frac{1}{2} - \frac{\epsilon^2}{8} p(c)t(c) + \mathcal{O}(e^3)$, and $\frac{\partial^2 s_1}{\partial a_1 \partial e_2} = \frac{\epsilon}{4} p'(c) + \mathcal{O}(e^3)$, where $p(c)$ is the coefficient of absolute prudence, and $t(c)$ is the coefficient of absolute tolerance, defined to be $-\frac{u''''(c)}{u'''(c)}$.

sion. To illustrate these calculations we show the result for the last expression in equation 18:

$$u''(c_{1t}) \frac{\partial^2 s_{1t}}{\partial a_{1t} \partial \sigma_2} = -\frac{W_0 \sigma}{16u''(c_1)^2} [\sigma u^{(3)}(c_1)^3 + u''(c_1)^2 (\sigma u^{(5)}(c_1) + u^{(4)}(c_1)(3\theta\sigma + 2)) - u^{(3)}(c_1)u''(c_1) (2\sigma u^{(4)}(c_1) + u^{(3)}(c_1)(3\theta\sigma + 2))] + \mathcal{O}(\sigma^3)$$

Simplifying equation 18 we get:

$$\frac{\partial^2 f}{\partial s_0 \partial \sigma_2} = -\frac{1}{4} W_0 \sigma u^{(4)}(c_1) + \mathcal{O}(\sigma^3)$$

And finally:

$$\begin{aligned} \frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} \Big|_{\sigma_1 = \sigma_2 = \sigma} = & \\ & -\sigma^2 \frac{W_0^3 (-W_0 u^{(3)}(c)^3 + 36\theta u^{(4)}(c)u''(c)^2 - 2W_0 u^{(4)}(c)u^{(3)}(c)u''(c) + 3W_0 u^{(5)}(c)u''(c)^2)}{36u''(c)^3} \\ & + \mathcal{O}(\sigma^3) \end{aligned}$$

Letting $\theta = 1$ and noting that $c_0 = \frac{1}{3}W_0 + \mathcal{O}(\sigma^2)$ gives us the result for permanent shocks given in equation 5.12. Letting $\theta = 0$ gives us the result for transitory shocks given in equation 5.13.

Chapter 6

Conclusion: Thoughts on Future Research

This dissertation has presented my research on the risks faced by UK households since the early 1990s and the effect of these risks on the distribution of welfare. I have stated conclusions from the separate pieces of research at the end of each chapter. Therefore, rather than re-stating the conclusions again I use this section to discuss possible future research on these topics. The suggestions here could be researched using techniques and data similar to those used in the rest of the dissertation.

Chapter 3 contains a very stylized model of policy changes. An obvious task for future research is to try to understand better how households form beliefs over future policy changes and how this policy uncertainty affects household behaviour. Chapter 3 concerns the introduction of tax credits. Probably the area where long-term policy uncertainty affects welfare most critically is in pensions arrangements - both the provision of state pensions and treatment of private pensions. Analysing this uncertainty is hard. Compared to idiosyncratic income risks and even aggregate income risks, where data can inform us how much objective risk is present, clearly no-one can assign probabilities to the chances of particular changes to the pension system. Moreover changes to the pension system might happen over a range of dimensions, from the scope of means testing, to the time profile of benefit payments, to the absolute generosity of the system and the strength of the relationship to prior contributions. However, data on subjective expectations over

the pension system (for example in the US Health and Retirement Survey and in the the English Longitudinal Survey of Aging) can inform the mapping from subjective expectations to behaviour. Furthermore we can tackle the mapping from objective uncertainties to subjective beliefs by studying the effect on expectations of news and announcements of the pension system. Recent work on financial literacy¹ indicates that many are over-optimistic about future state pension generosity. It is therefore likely that the process of forming beliefs is complex and varies across the population according to financial ‘ability’.

Chapter 3 also touches on the question of housing choices across the income distribution. This is also an important area of research. In both the US and the UK, the policy environment from the 1980s onwards encouraged universal homeownership.² Given the role of sub-prime mortgages in the recent financial crisis, housing policy will likely move towards the provision of affordable housing of all tenure types. The demand for housing in general is an active area of research but demand at the bottom end of the income distribution seems particularly pertinent if the supply side at the bottom end is such a mutable area of government policy. Important considerations here are that those at the bottom of the distribution face higher employment risk, so mortgage default becomes a greater threat.

Chapter 4 fits into a very established stream of research where the research problems are clearly defined: we need to understand better the availability of consumption insurance, the specification for income risk and to identify the household’s information set. These research problems are clear, well-known, important but difficult. However, related questions have received less attention and may be worthy of more research. I discussed some of these at the end of the chapter, and repeat them briefly here. First, this chapter examined only stable households headed by a couple. Other types of households are becoming increasingly prevalent: their circumstances should be researched more. This chapter argues that for the stable households, income risk is the most important risk. It may be that, for the population as a whole, demographic risk such as divorce risk is as important.

Chapter 5 is very much an initial attempt at a little-understood area and warrants much future research. The obvious route is to research empirically the effect of consecutive

¹for example in Lusardi and Mitchell (2011)

²See Rajan (2010), chapter 1, for a discussion of the US.

risks. Also, in this chapter we have only considered a single liquid asset. There are of course multiple savings vehicles of various types of liquidity: most importantly, housing and pension wealth. It would be interesting to research to what extent these technologies can each meet the various motivations for saving.

Bibliography

- ABOWD, J. AND D. CARD (1989): “On the covariance structure of earnings and hours changes,” *Econometrica: Journal of the Econometric Society*.
- ACEMOGLU, D. (2002): “Technical change, inequality, and the labor market,” *Journal of economic literature*.
- AIYAGARI, S. (1994): “Uninsured idiosyncratic risk and aggregate saving,” *The Quarterly Journal of Economics*, 109, 659.
- ALTONJI, J., A. MARTINS, AND A. SIOW (2002): “Dynamic factor models of consumption, hours and income,” *Research in Economics*, 56, 3–59.
- ALTONJI, J. AND L. SEGAL (1996): “Small-sample bias in GMM estimation of covariance structures,” *Journal of Business & Economic Statistics*.
- ALTONJI, J., A. SMITH, AND I. VIDANGOS (2009): “Modeling earnings dynamics,” *NBER Working Paper*.
- ATKINSON, A. (1997): “Bringing income distribution in from the cold,” *The Economic Journal*, 107, 297–321.
- (1999): “The distribution of income in the UK and OECD countries in the twentieth century,” *Oxford Review of Economic Policy*.
- ATKINSON, A. AND T. PIKETTY (2007): *Top incomes over the twentieth century: a contrast between continental European and English-speaking countries*, Oxford University Press, USA.

- ATTANASIO, O., E. BATTISTIN, AND A. LEICESTER (2006): "Consumption, income and the well-being of families," *mimeo*.
- ATTANASIO, O., G. BERLOFFA, R. BLUNDELL, AND I. PRESTON (2002): "From earnings inequality to consumption inequality," *Economic Journal*.
- ATTANASIO, O., L. BLOW, R. HAMILTON, AND A. LEICESTER (2009): "Booms and Busts: Consumption, house prices and expectations," *Economica*.
- ATTANASIO, O. AND S. DAVIS (1996): "Relative wage movements and the distribution of consumption," *Journal of Political Economy*.
- ATTANASIO, O. AND T. JAPPELLI (2001): "Intertemporal choice and the cross-sectional variance of marginal utility," *Review of Economics and Statistics*, 83, 13–27.
- ATTANASIO, O. AND N. PAVONI (2010): "Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption," *Econometrica* - *Forthcoming*.
- ATTANASIO, O. AND G. WEBER (1995): "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey," *Journal of Political Economy*.
- BAKER, M. (1997): "Growth-rate heterogeneity and the covariance structure of life-cycle earnings," *Journal of Labor Economics*, 338–375.
- BAKER, M. AND G. SOLON (2003): "Earnings dynamics and inequality among Canadian men, 1976-1992: Evidence from longitudinal income tax records," *Journal of Labor Economics*, 21, 267–288.
- BANKS, J., Z. SMITH, AND M. WAKEFIELD (2002): "The distribution of financial wealth in the UK: evidence from 2000 BHPS data," *IFS Working Paper*.
- BARDASI, E., S. JENKINS, AND J. RIGG (1999): "Documentation for derived current and annual net household income variables, BHPS Waves 1-7," *Institute for Social and Economic Research Working paper*, 99–25.

- BARRO, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, 823–866.
- BATTISTIN, E., R. BLUNDELL, AND A. LEWBEL (2009): "Why is consumption more log normal than income? Gibrat's Law revisited," *Journal of Political Economy*.
- BATTISTIN, E., G. WEBER, AND R. MINIACI (2003): "What do we learn from recall consumption data?" *Journal of Human Resources*.
- BLOOM, N. (2009): "The impact of uncertainty shocks," *Econometrica*.
- BLUNDELL, R. (2002): "Welfare-to-work: Which policies work and why?" *Keynes Lecture: Proceedings of the British Academy*.
- BLUNDELL, R., M. FRANCESCONI, AND W. VAN DER KLAAUW (2010): "Anatomy of Policy Reform Evaluation: Announcement and Implementation Effects," *Mimeo*.
- BLUNDELL, R., A. GOSLING, H. ICHIMURA, AND C. MEGHIR (2007): "Changes in the distribution of male and female wages accounting for employment composition using bounds," *Econometrica*.
- BLUNDELL, R. AND H. HOYNES (2004): "Has 'In-Work' Benefit Reform Helped the Labor Market?" in *Seeking a Premier Economy*, ed. by D. Card and R. Freeman.
- BLUNDELL, R., H. LOW, AND I. PRESTON (2008a): "Decomposing Changes in Income Risk Using Consumption Data," *Institute for Fiscal Studies WP08/13*.
- BLUNDELL, R. AND T. MACURDY (1999): "Labor supply: A review of alternative approaches," *Handbook of labor economics*.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2004): "Imputing consumption in the PSID using food demand estimates from the CEX," *IFS Working Paper*.
- (2008b): "Consumption inequality and partial insurance," *The American Economic Review*.
- BLUNDELL, R. AND I. PRESTON (1998): "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*, 113, 603–640.

- BLUNDELL, R., H. REED, AND T. M. STOKER (2003): "Interpreting aggregate wage growth: The role of labor market participation," *The American Economic Review*.
- BLUNDELL, R. AND T. M. STOKER (1999): "Consumption and the Timing of Income Risk," *European Economic Review*, 43, 475–507.
- BOHRNSTEDT, G. W. AND A. S. GOLDBERGER (1969): "On the exact covariance of products of random variables," *Journal of the American Statistical Association*, 64, 1439–1442.
- BREWER, M. (2007): "Welfare reform in the UK: 1997-2007," *IFS Working Paper*.
- BREWER, M., A. GOODMAN, A. MURIEL, AND L. SIBIETA (2007a): "Poverty and inequality in Briatrain," *IFS Briefing Note*.
- BREWER, M., L. SIBIETA, AND L. WREN-LEWIS (2007b): "Racing away: income inequality and the evolution of high incomes," *IFS Briefing Note*.
- BROWNING, M. AND A. LUSARDI (1996): "Household saving: Micro theories and micro facts," *Journal of Economic Literature*.
- BROWNING, M. AND C. MEGHIR (1991): "The effects of male and female labor supply on commodity demands," *Econometrica*.
- BRUMBERG, R. AND F. MODIGLIANI (1954): "Utility analysis and the consumption function: An interpretation of cross-section data," *Post Keynesian Economics*.
- CABALLERO, R. (1990): "Consumption puzzles and precautionary savings," *Journal of monetary economics*, 25, 113–136.
- CAGETTI, M. (2003): "Wealth accumulation over the life cycle and precautionary savings," *Journal of Business and Economic Statistics*.
- CAMPBELL, J. AND J. COCCO (2007): "How do house prices affect consumption? Evidence from micro data," *Journal of Monetary Economics*.
- CARROLL, C. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*," *Quarterly Journal of Economics*.

- (2004): “Theoretical foundation of buffer stock saving,” *NBER Working Paper*.
- (2006): “The method of endogenous gridpoints for solving dynamic stochastic optimization problems,” *Economics Letters*.
- (2009): “Precautionary saving and the marginal propensity to consume out of permanent income,” *Journal of Monetary Economics*, 56.
- CARROLL, C. AND M. KIMBALL (2001): “Liquidity constraints and precautionary saving,” *NBER Working Paper*.
- CARROLL, C. AND A. SAMWICK (1997): “The nature of precautionary wealth,” *Journal of monetary Economics*, 40, 41–71.
- CASE, K., J. QUIGLEY, AND R. SHILLER (2005): “Comparing wealth effects: the stock market versus the housing market,” *Advances in Macroeconomics*, 5, 1–34.
- CUNHA, F., J. HECKMAN, AND S. NAVARRO (2005): “Separating uncertainty from heterogeneity in life cycle earnings,” *Oxford Economic Papers*.
- CUTLER, D. AND L. KATZ (1992): “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980’s,” *The American Economic Review*.
- DAVIS, M. AND F. ORTALO-MAGNÉ (2010): “Household expenditures, wages, rents,” *Review of Economic Dynamics*.
- DE NARDI, M., E. FRENCH, AND J. JONES (2010): “Why do the elderly save? the role of medical expenses,” *Journal of Political Economy*, 118, 39–75.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 1221–1248.
- (1992): *Understanding consumption*, Oxford University Press.
- DEATON, A. AND J. MUELLBAUER (1980): *Economics and consumer behavior*, Cambridge Univ Pr.
- DEATON, A. AND C. PAXSON (1994): “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 102, 437–467.

- DICKENS, R. (2000): "The evolution of individual male earnings in great britain: 1975-95," *The Economic Journal*.
- DICKENS, R. AND A. MANNING (2004): "Has the national minimum wage reduced UK wage inequality?" *Journal of the Royal Statistical Society. Series A, Statistics in Society*, 167.
- DISNEY, R. F., C. EMMERSON, AND G. TETLOW (2009): "What is a public sector pension worth?" *The Economic Journal*, 119, F517—F535.
- DYNAN, K. (1993): "How prudent are consumers?" *Journal of Political Economy*.
- EECKHOUDT, L. AND H. SCHLESINGER (2008): "Changes in Risk and the Demand for Saving," *Journal of Monetary Economics*, 55, 1329–1336.
- FARHI, E. AND I. WERNING (2009): "Capital Taxation: Quantitative Explorations of the Inverse Euler Equation," *MIT working paper*.
- FLAVIN, M. (1981): "The adjustment of consumption to changing expectations about future income," *The Journal of Political Economy*, 974–1009.
- GOLDMAN, D. AND N. MAESTAS (2005): "Medical expenditure risk and household portfolio choice," *NBER working paper*.
- GOLLIER, C. (2004): *The economics of risk and time*, The MIT Press.
- GOLLIER, C. AND J. W. PRATT (1996): "Risk vulnerability and the tempering effect of background risk," *Econometrica: Journal of the Econometric Society*, 1109–1123.
- GOLOSOV, M. AND A. TSYVINSKI (2007): "Optimal taxation with endogenous insurance markets," *Quarterly Journal of Economics*.
- GOODMAN, L. (1960): "On the exact variance of products," *Journal of the American Statistical Association*.
- GORBACHEV, O. (2010): "Did Household Consumption Become More Volatile?" *American Economic Review*, Forthcomin.

- GORODNICHENKO, Y., D. STOLYAROV, AND K. SABIRIANOVA PETER (2010): "Inequality and volatility moderation in Russia: Evidence from micro-level panel data on consumption and income," *Review of Economic Dynamics*.
- GOSLING, A., S. MACHIN, AND C. MEGHIR (2000): "The changing distribution of male wages in the UK," *The Review of Economic Studies*.
- GRUBER, J. (1997): "The consumption smoothing benefits of unemployment insurance," *The American Economic Review*.
- GRUBER, J. AND A. YELOWITZ (1999): "Public health insurance and private savings," *Journal of Political Economy*.
- GUIISO, L., T. JAPPELLI, AND M. PADULA (2009): "Pension Risk, Retirement Saving and Insurance," *EUI Working Paper*.
- GUVENEN, F. (2007): "Learning your earning: are labor income shocks really very persistent?" *The American Economic Review*, 97, 687–712.
- (2009): "An empirical investigation of labor income processes," *Review of Economic Dynamics*.
- GUVENEN, F. AND A. SMITH JR (2010): "Inferring labor income risk from economic choices: An indirect inference approach," *Mimeo*.
- HALL, R. AND F. MISHKIN (1982): "The sensitivity of consumption to transitory income: evidence from PSID households," *Econometrica*, 50, 461–81.
- HEATHCOTE, J., F. PERRI, AND G. VIOLANTE (2010a): "Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006," *Review of Economic Dynamics*.
- HEATHCOTE, J., K. STORESLETTEN, AND G. VIOLANTE (2008): "Consumption and labor supply with partial insurance: An analytical framework," *mimeo*.
- (2010b): "The macroeconomic implications of rising wage inequality in the United States," *Journal of Political Economy*, 118, 681–722.

- HRYSHKO, D. (2010): "RIP to HIP: The data reject heterogeneous labor income profiles," *Unpublished manuscript, University of Alberta.*[67].
- HUGGETT, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of economic Dynamics and Control*.
- HURST, E., A. KENNICKELL, A. LUSARDI, AND F. TORRALBA (2005): "Precautionary savings and the importance of business owners," *NBER working paper*.
- JAPPELLI, T. AND L. PISTAFERRI (2006): "Intertemporal Choice and Consumption Mobility," *Journal of the European Economic Association*, 4, 75–115.
- JENKINS, S. (2010): "The British Household Panel Survey and its Income Data," *ISER Working Paper*.
- JOHNSON, P. AND S. WEBB (1993): "Explaining the growth in UK income inequality: 1979-1988," *The Economic Journal*.
- JUSTER, F., J. LUPTON, J. SMITH, AND F. STAFFORD (2006): "The decline in household saving and the wealth effect," *The Review of Economics and Statistics*, 88.
- KAPLAN, G. (2010): "Inequality and the Lifecycle," *Mimeo*.
- KAPLAN, G. AND G. VIOLANTE (2010): "How Much Consumption Insurance Beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2, 53–87.
- KEANE, M. AND K. WOLPIN (1997): "The career decisions of young men," *Journal of Political Economy*, 105, 473–522.
- KIMBALL, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica: Journal of the Econometric Society*, 53–73.
- (1993): "Standard risk aversion," *Econometrica: Journal of the Econometric Society*, 589–611.
- KRUEGER, D. AND F. PERRI (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *The Review of Economic Studies*, 73, 163–193.

- KRUEGER, D., F. PERRI, L. PISTAFERRI, AND G. VIOLANTE (2010): "Cross-sectional facts for macroeconomists," *Review of Economic Dynamics*, 13, 1–14.
- LEE, B. AND B. INGRAM (1991): "Simulation estimation of time-series models," *Journal of Econometrics*.
- LI, W. AND R. YAO (2007): "The Life-Cycle Effects of House Price Changes," *Journal of Money, Credit and Banking*.
- LOW, H., C. MEGHIR, AND L. PISTAFERRI (2010): "Wage risk and employment risk over the life cycle," *The American Economic Review*, 100, 1432–67.
- LUCAS, R. (1987): *Models of business cycles*, Basil Blackwell New York.
- LUSARDI, A. AND O. MITCHELL (2011): "Financial literacy and planning: Implications for retirement wellbeing," Tech. rep., National Bureau of Economic Research.
- MÄÄTTÄNEN, N. AND M. TERVIÖ (2010): "Income Distribution and Housing Prices: An Assignment Model Approach," *CEPR Discussion Papers*, 7945.
- MACURDY, T. (1982): "The use of time series processes to model the error structure of earnings in a longitudinal data analysis," *Journal of econometrics*.
- MEGHIR, C. AND L. PISTAFERRI (2004): "Income variance dynamics and heterogeneity," *Econometrica*.
- (2010): "Earnings, consumption and lifecycle choices," *NBER working paper*.
- MEYER, B. AND J. SULLIVAN (2003): "Measuring the well-being of the poor using income and consumption," *NBER Working Paper*.
- MOFFITT, R. AND P. GOTTSCHALK (2002): "Trends in the transitory variance of earnings in the United States," *Economic Journal*.
- MORETTI, E. (2008): "Real Wage Inequality," *NBER Working Paper*.
- MUELLBAUER, J. AND A. MURPHY (1997): "Booms and busts in the UK housing market," *The Economic Journal*.

- ORTALO-MAGNÉ, F. AND S. RADY (2006): "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints," *The Review of Economic Studies*.
- OWEN, D. (1956): "Tables for computing bivariate normal probabilities," *The Annals of Mathematical Statistics*.
- PAIELLA, M. (2007): "Does wealth affect consumption? Evidence for Italy," *Journal of Macroeconomics*.
- POSTEL VINEY, F. AND H. TURON (2010): "On-the-job search, productivity shocks, and the individual earnings process," *International Economic Review*.
- PRATT, J. W. AND R. J. ZECKHAUSER (1987): "Proper risk aversion," *Econometrica: Journal of the Econometric Society*, 143–154.
- PRIMICERI, G. AND T. VAN RENS (2009): "Heterogeneous life-cycle profiles, income risk and consumption inequality," *Journal of Monetary Economics*, 56, 20–39.
- RAJAN, R. (2010): *Fault lines: How hidden fractures still threaten the world economy*, Princeton University Press.
- SEFTON, J., J. VAN DE VEN, AND M. WEALE (2008): "Means Testing Retirement Benefits: fostering equity or discouraging savings?" *The Economic Journal*.
- SKINNER, J. (1987): "A superior measure of consumption from the Panel Study of Income Dynamics," *Economics Letters*, 23, 213–216.
- SMITH JR, A. (1993): "Estimating nonlinear time-series models using simulated vector autoregressions," *Journal of Applied Econometrics*, 8, S63–S84.
- STORESLETTEN, K., C. TELMER, AND A. YARON (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633.
- STOYE, J. (2010): "Partial identification of spread parameters," *Quantitative Economics*.
- TANNER, S. (1998): "How much do consumers spend? Comparing the FES and National Accounts," in *How Reliable is the Family Expenditure Survey? Trends in Incomes and Expenditures Over Time*, ed. by J. Banks and P. Johnson.

VAN NIEUWERBURGH, S. AND P.-O. WEILL (2010): “Why Has House Price Dispersion Gone Up?” *The Review of Economic Studies*.

VOENA, A. (2010): “Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?” *Unpublished Manuscript, Stanford University*.

WAKEFIELD, M. (2009): “Consumption and saving decisions in the face of choices about housing and pensions,” Ph.D. thesis.

ZILIAK, J. (1998): “Does the choice of consumption measure matter? An application to the permanent-income hypothesis,” *Journal of monetary Economics*, 41, 201–216.