

Essays in Macroeconomics and the
Role of Household Heterogeneity in the
Great Recession

Kieran P. Larkin

Submitted in fulfillment of the requirements for the degree of

Doctor of Philosophy
of
University College London

Department of Economics
University College London

August 15, 2019

Declaration

I, Kieran Patrick Larkin 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.

Abstract

This thesis studies the role that household heterogeneity plays in macroeconomic dynamics. It tackles this topic by focusing on consumption behaviour during the Great Recession. Its overarching argument is threefold: Firstly, that heterogeneity is crucial for understanding the patterns observed in the data. Secondly, that the state of the economy and preceding environment has a first order impact on how the economy responds to shocks. Thirdly, inaction and illiquidity matter and by focusing on this margin we can gain a better understanding of consumer behaviour. In the three chapters I consider the role that labor market sorting, household portfolio choice, shocks to the availability of credit, and changes in the future path of expected income played in understanding the Great Recession. The first chapter investigates the role that labor market tranquillity prior to the Great Recession played in the magnitude of the subsequent consumption decline. It highlights the interaction between labor market sorting and the household portfolio choice in the determination of the economy's response to shocks. The second chapter looks at the role of credit conditions in household consumption dynamics. It argues that much of the literature has misunderstood the effect of a decline in credit availability during the Great Recession. It instead emphasizes the inaction response generated due to a tightening of the collateral constraint, which provides a better account of the Great Recession and of the borrowing behaviour of households in the micro data. The final chapter focuses on the car market and car purchasing behaviour during the last three US recessions, documenting important unusual features of the consumption response during the Great Recession. It uses these responses as an identification mechanism to uncover the importance of the shocks hitting the economy during the crisis, finding an important role for asset price shocks and a decline in the expected growth rate of household incomes in replicating the consumption dynamics.

Impact Statement

Understanding the causes of the Great Recession is of self-evident policy relevance. The development of better economic models is also vital to the process of evaluating policy interventions. All three chapters of this thesis aim to contribute to a better understanding of economic decision making and to provide guidance on the likely effect of policy changes.

The first chapter studies the interaction between labor market sorting and the household portfolio allocation decision. It contributes by developing a model which introduces novel labor market features into a state-of-the-art general equilibrium incomplete markets model to understand how the interaction of asset and labor market decisions alter the economy's response to shocks. A key finding of the chapter is that the tranquil labor market environment prior to the Great Recession exacerbated the subsequent decline by leading households to adopt illiquid portfolios. This emphasizes the importance to policy makers of understanding how the response of the economy may differ given the prior economic context. The model also highlights the sizeable amplification mechanism generated by the restrictions on mortgage refinancing faced by unemployed households. This friction suggests an opportunity for policy intervention to improve welfare.

The second chapter argues that much of the existing literature has misunderstood the effect of a decline in credit availability during the Great Recession. It develops a framework for understanding the role of credit shocks that fits the data better than the existing models, emphasising the inaction response. It also finds that policy interventions to offset credit shocks appear to be relatively ineffective and may end up lowering welfare.

The final chapter contributes by documenting new facts regarding car purchasing behaviour during the Great Recession. These features of the data can be useful in understanding the sources of shocks to the economy during the crisis. Given that significant resources were allocated to the car market during the crisis the chapter also provides a clearer understanding as to the way future policy interventions might be better designed.

Acknowledgements

I thank my supervisors Morten Ravn and Orazio Attanasio for continued direction and support during my PhD. Morten's guidance, mentoring and encouragement in particular has been invaluable. I also thank members of the UCL Economics Department faculty that have contributed time and ideas to the development of this thesis. Special mention goes to Vincent Sterk, Wei Cui, Ralph Luetticke, Wendy Carlin and Victor Rios-Rull for helpful comments and input. I acknowledge the financial support of the Economic and Social Research Council (ESRC) and Centre for Macroeconomics (CfM). Finally, I thank my necessary and sufficient condition Dr Natalie Welsh for putting up with my prolonged studies and tolerating my commitment to forward looking rational optimisation in all aspects of our lives. No thanks to Fortran and Matlab who have made my life [error!].

Contents

Introduction	12
1 Job Risk, Separation Shocks and Household Asset Allocation	19
1.1 Introduction	19
1.2 Related literature	25
1.3 Data	28
1.3.1 Job separation rate	29
1.3.2 Job risk distribution	30
1.3.3 Asset allocation and job risk	32
1.4 Model	37
1.4.1 Model environment	37
1.4.2 Household asset choice	38
1.4.3 Job choice	40
1.4.4 Recursive formulation	42
1.4.5 Labor market	44
1.4.6 Government	45
1.4.7 Stationary recursive equilibrium definition	46
1.4.8 Calibration	46
1.4.9 Asset distribution	52
1.5 Properties of the steady state equilibrium	53
1.5.1 Policy functions	53
1.5.2 Liquidity regressions	56
1.5.3 Unemployment responses	58
1.5.4 Role of job risk heterogeneity	59
1.6 Great Recession experiment	61
1.6.1 Aggregate responses	62
1.6.2 Cross sectional responses	67
1.7 Conclusion	69
1.8 Tables	71
1.9 Figures	79

2	Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics	91
2.1	Introduction	91
2.2	Empirical evidence	97
2.2.1	Consumption in the Great Recession	97
2.2.2	Decline in adjustment rate	97
2.2.3	Measures of credit availability	98
2.2.4	Growth in home equity line of credit borrowing	100
2.2.5	Loan to value ratios by behaviour	101
2.3	Model	102
2.3.1	Utility function and idiosyncratic uncertainty	102
2.3.2	Adjustment costs	104
2.3.3	Saving, borrowing and budget constraint	105
2.3.4	Collateral constraint	105
2.3.5	Firms	108
2.3.6	Government	109
2.3.7	Recursive formulation	109
2.3.8	Equilibrium	111
2.3.9	Model intuition	112
2.4	Calibration	114
2.4.1	Collateral constraint process	114
2.4.2	Preference and technology parameters	115
2.5	Properties of the economy	117
2.5.1	Model validation: behavior of adjusters and non-adjusters	117
2.5.2	Long-run averages, distributions and moments	119
2.5.3	Conditional means and varying aggregate states	120
2.6	Responses to shocks	122
2.6.1	Great Recession experiment	122
2.6.2	Alternative models	123
2.6.3	Impulse response decomposition	124
2.6.4	State dependent responses	126
2.6.5	Non-linearities	128
2.6.6	Forecast error variance decomposition	129
2.7	Policy experiment	130
2.8	Robustness	134
2.8.1	Business cycle moments of alternative specifications	134
2.9	Conclusion	135
2.10	Tables	137

2.11	Figures	141
3	(S)Cars and the Great Recession	152
3.1	Introduction	152
3.2	Consumption dynamics and the Great Recession	159
3.2.1	The data	159
3.2.2	Descriptive statistics	161
3.2.3	Car purchases during the Great Recession	163
3.3	The model	166
3.3.1	Household problem	167
3.4	Calibration	173
3.5	Life-cycle and policy functions	178
3.5.1	Life-cycle profiles	178
3.5.2	Policy functions	179
3.6	The Great Recession	181
3.6.1	Uniform income shocks	181
3.6.2	Cohort specific shocks	183
3.6.3	Financial market shocks	185
3.6.4	Wealth shocks	186
3.6.5	Life-cycle profile income shock	188
3.7	Conclusion	192
3.8	Tables	194
3.9	Figures	198
A	Additional material	219
A.1	Appendix to Job Risk, Separation Shocks and Household Asset Allocation	219
A.1.1	Unemployment response methodology	219
A.1.2	Great Recession shocks and data	220
A.1.3	Great Recession experiment details	221
A.2	Appendix to Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics	223
A.2.1	Evidence for the existence of credit shocks	223
A.2.2	Solution method	223
A.2.3	Solving the household problem	223
A.2.4	Solving the household problem: a generalised endogenous grid method	224
A.2.5	Refinements	225
A.2.6	Solving for the aggregate law of motion	226

A.3	Appendix to (S)Cars and the Great Recession	229
A.3.1	Solution method	229
A.3.2	Estimating growth profile	232
A.3.3	Growth rate shock calibration	232
A.3.4	Estimating cohort shocks	233
A.3.5	Estimating cohort responses	234
B	Additional Tables	236
B.1	Additional Tables for Job Risk, Separation Shocks and Household Asset Allocation	236
B.2	Additional Tables for Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics	241
B.3	Additional Tables for (S)Cars and the Great Recession	250
C	Additional Figures	251
C.1	Additional Figures for Job Risk, Separation Shocks and Household Asset Allocation	251
C.2	Additional Figures for Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics	260
C.3	Additional Figures for (S)Cars and the Great Recession	262

List of Tables

1.1	Ratio of liquid assets to illiquid assets	72
1.2	Ratio of liquid assets to total assets	73
1.3	Hand-to-Mouth status	74
1.4	Externally calibrated parameters	74
1.5	Calibrated parameters	74
1.6	Targeted moments	75
1.7	Earnings process	75
1.8	Impact of job choice on equilibrium	76
1.9	Model portfolio regressions	76
1.10	Model Hand-to-Mouth status	77
1.11	Role of job risk heterogeneity	77
1.12	Period one Great Recession response	78
2.1	Stochastic process for the collateral constraint	137
2.2	Calibration of the model, Preference and technology parameters	137
2.3	Calibration of the model, wealth parameters	137
2.4	Co-movement of credit variables	138
2.5	Baseline model long-run averages	138
2.6	Business cycle moments	138
2.7	State conditional first moments	139
2.8	Policy experiment: Relaxing credit conditions	139
2.9	Policy experiment: Relaxing credit conditions	140
3.1	Comparison of Great Recession car market outcomes	194
3.2	Probability of Purchasing a Car (Avg. Marginal Effects)	195
3.3	Externally calibrated parameters	196
3.4	Calibration targets	196
3.5	Estimated Parameters	197
3.6	Non Targeted Moments	197
B.1	Ratio of liquid assets to illiquid assets, CPS	237

B.2	Ratio of liquid assets to total assets, CPS	238
B.3	Hand-to-Mouth status, Core	239
B.4	Hand-to-Mouth status, CPS	239
B.5	Great Recession shocks	239
B.6	Comparison of High average δ and Pre-Great Recession Equilibria	240
B.7	Business cycle moments: US data 1950.I-2014.I	242
B.8	Business cycle moments: US data, sub-time periods	243
B.9	Survey of Consumer Finances, 1989-2010	244
B.10	Commonly used durable adjustment costs specifications	244
B.11	Conditional responses: regressions of Θ IRFs on state space .	245
B.12	Conditional responses: regressions of small Z IRFs on state space	246
B.13	TFP and Credit shock regression results (difference)	247
B.14	TFP and Credit shock regression results (ratio)	248
B.15	Accuracy of aggregate law of motion	249
B.16	Consumer Expenditure Survey Summary Statistics	250

List of Figures

1.1	Distribution of US job risk	79
1.2	Job separation rate	79
1.3	Job separation rate: compositional effects	80
1.4	Job separation rate in Great Recession	80
1.5	Time series variation in job risk distribution	81
1.6	Portfolio choices during decline in job risk	81
1.7	Calibrated job distribution	82
1.8	Policy functions	82
1.9	Job offer indifference curve	83
1.10	Asset distributions	83
1.11	Job distributions	84
1.12	Labor market response to unemployment shock	84
1.13	Housing response to unemployment shock	84
1.14	Housing response to unemployment shock, by size of income shock	85
1.15	Marginal Welfare Gain from Insurance	85
1.16	IRF to lump sum tax shock	86
1.17	Role of job risk in response	86
1.18	Great Recession shocks	87
1.19	IRF to Great Recession shock	88
1.20	Persistence of consumption response	88
1.21	Amplification with prior increase in wage uncertainty	89
1.22	Relative group responses to Great Recession shock	89
2.1	Consumption shares during recessions, 1967.II-2010.II	141
2.2	Decline in durables transaction volumes	141
2.3	Comparison of FRB’s Senior Loan Officer Opinion Survey credit variables, 1990.III-2014.III	142
2.4	HP filtered aggregate series: US data 1966.III-2014.I	142
2.5	Growth and availability of Home equity line of credit mortgages	143
2.6	Mortgage to housing stock ratio of movers and non-movers	143

2.7	Durables policy functions under low and high collateral constraints	144
2.8	Durables policy functions under a drop in aggregate credit conditions	144
2.9	Baseline model distributions of durables and capital	145
2.10	Great Recession experiment shocks	145
2.11	Great Recession experiment	146
2.12	Great Recession experiment: Comparator models	146
2.13	Impulse response decomposition	147
2.14	Impulse response: credit shock under prior conditions	148
2.15	State dependence	148
2.16	Impulse response: Non-linear response	149
2.17	Forecast variance decomposition	149
2.18	Policy experiment: relaxing credit conditions	150
3.1	CEX cohorts: non-durable consumption	198
3.2	CEX cohorts: number of cars	198
3.3	CEX cohorts: % households with no cars	199
3.4	CEX cohorts: Value of car stock	199
3.5	CEX cohorts: Ratio of car stock to non-durable consumption	199
3.6	NIPA: car expenditure & durables growth	200
3.7	Comparing NIPA to CEX	200
3.8	CEX: share of households purchasing car	201
3.9	CEX: value of car purchased	201
3.10	CEX: Probability of purchase, year fixed effects	202
3.11	Model life-cycle fit (I)	202
3.12	Model life-cycle fit (II)	202
3.13	Model policies	203
3.14	Model policies	203
3.15	Uniform income shocks	203
3.16	Uniform income shocks: consumption	204
3.17	Uniform income shocks: extensive and intensive margin	204
3.18	Uniform versus cohort shocks	205
3.19	Interest rate spread shock	205
3.20	Interest rate spread shock	206
3.21	Asset price shock	206
3.22	Asset price shock: consumption response	207
3.23	Asset price shock: extensive and intensive margin	207

3.24	Evidence supporting decline in life-cycle growth rate (CEX)	208
3.25	Growth rate shock: extensive and intensive margin	208
3.26	CEX cohort car response	209
3.27	Cohort response: model comparison	209
C.1	Consumption drop during Great Recession	251
C.2	Job finding rate	252
C.3	Weekly earnings in Great Recession	252
C.4	Job risk distribution with alternative income controls	253
C.5	Distribution of US job risk by year	253
C.6	Portfolio choices during decline in job risk	254
C.7	Job separation rate of jobs by age	255
C.8	Value function in (ω, δ) space	255
C.9	Role of job risk in MPCs	256
C.10	Housing status response to unemployment shock	256
C.11	Other asset responses to unemployment shock	257
C.12	Group responses to Great Recession shock I: level response	258
C.13	Group responses to Great Recession shock II: level response	259
C.14	Financial Assets and Automotives Distributions, SCF, 2010	260
C.15	Financial Assets and Housing Distributions, SCF, 2010	260
C.16	Financial Assets and Housing Distributions, SCF, 2010	260
C.17	Accuracy of the aggregate law of motion	261
C.18	Interest rate: car loan response	262
C.19	Evidence supporting decline in life-cycle growth rate (CPS)	262
C.20	Growth rate shock: example paths	263

Introduction

This thesis studies the role that household heterogeneity plays in macroeconomic dynamics. It tackles this topic by focusing on consumption behaviour during the Great Recession, the preeminent economic event of our lifetimes. The decline in consumption during this period was unprecedented (DeNardi et al. (2012), Rios-Rull and Huo (2016)). In the US consumption fell by 3.7 percent from peak to trough. This compares to an average decline of 1.5 percent across all post World War II recessions. Durables consumption was hit particularly hard. Narrowly defined, durables consumption declined by 14.2 percent from peak-to-trough between 2007.IV and 2009.II, whereas the average decline across all US post-war recessions was 9.7 percent. As will be discussed below in greater detail, the car market underwent an unusual and severe contraction with fewer households purchasing cars and those households that purchased cars spending less. From 2007.III to 2010.I the average real value of a car being purchased fell by 21.4 percent. I will argue that understanding these consumption responses is fundamental for understanding the causes of the crisis.

The chapters of this thesis will study these patterns in depth drawing on important features of the economy, to provide a fuller understand of the Great Recession. Its overarching argument is threefold: Firstly, that heterogeneity is crucial for understanding the patterns observed in the data (Heathcote et al. (2009), Guvenen (2011), Kaplan and Violante (2018)). Secondly, that the state of the economy and preceding environment has a first order impact on how the economy responds to shocks. Thirdly, inaction and illiquidity matter and by focusing on this margin we can gain a better understanding of consumer behaviour. In the three chapters I will consider the role that labor market sorting, household portfolio choice, shocks to the availability of credit, and changes in the future path of expected income played in the

Great Recession. On a technical level each chapter will demonstrate how the application of different macroeconomic methodologies - whether one-time unexpected shocks, dynamics stochastic general equilibrium models or a partial equilibrium lifecycle setting - can be informative for understanding aggregate economic dynamics. Methodologically, I also emphasise the value of micro datasets in disciplining and informing macroeconomic analysis.

The first chapter, *'Job Risk, Separation Shocks and Household Asset Allocation'*, investigates the role that labor market tranquillity prior to the Great Recession played in the magnitude of the subsequent consumption decline. It motivates a model where households' choices in both the labor market and asset market respond to a change in household risk. It finds sorting in these dimensions left the economy more vulnerable to large unexpected shocks of the type experienced in the labor market between 2007 and 2009. The motivation for this setting is that prior to the Great Recession the economy can be characterised as going through a period of tranquillity. One manifestation of this is the long-run decline in the job separation rate ([Davis \(2008\)](#)). During the same period the composition of household portfolios also underwent a change. Between 1980 and 2007, households reduced liquid savings and increased illiquid asset holdings. These trends suggest that interactions between the labor and asset markets may have affected the responsiveness of the economy to the shocks experienced during the Great Recession.

The chapter develops a model which introduces novel labor market features into a state of the art general equilibrium incomplete markets model to understand how the interaction of asset and labor market decisions alter the economy's response to shocks ([Lise \(2013\)](#)). I build a model that incorporates a jobs ladder, heterogeneity in job risk and saving in liquid and illiquid assets. In the labor market, I propose a model where households receive job offers that vary in two dimensions – the wage and separation rate ([Jarosch \(2015\)](#), [Burdett and Mortensen \(1980\)](#)). A household that receives a new job offer is

able to accept or reject the new wage-job risk combination and this decision depends on its current asset portfolio. Over time a continuously employed household will move to better paid more secure employment. In the asset market households hold liquid savings, illiquid housing and long-term mortgages. The risk faced by a household, an endogenous object, will alter their portfolio allocation decision.

In the chapter's main result, I investigate the response of the economy to the labor market shocks seen in the Great Recession. In particular, I shock the economy with changes to the wage level and job separation rate. I compare the response of the aggregates when the shocks occur following a period of relative calm in the labor market - capturing the pre Great Recession environment - to the counterfactual case with higher prior job risk. The fact that the Great Recession occurred following a period of moderation and low risk in the labor market, amplified the negative response of consumption by up to 40 percent. This is due to households holding more illiquid portfolios and sorting into better jobs.

I also show that the model is able to deliver empirically realistic asset and labor market interactions. This serves to validate the findings in response to the Great Recession shocks. Firstly, the model generates key features of the income process including negative skewness and excess kurtosis as documented in administrative datasets ([Guvenen et al. \(2015\)](#), [Hubmer \(2018\)](#)). The jobs ladder feature reconciles the model with the empirical fact that unemployment shocks are often associated with a persistent decline in income ([Stevens \(1997\)](#), [Davis and von Wachter \(2011\)](#), [Krolkowski \(2017\)](#), [Huckfeldt \(2018\)](#)). Secondly, the model replicates the positive correlation between job risk and the liquidity of a household's portfolio as observed in the Panel Study of Income Dynamics (PSID). Thirdly, I show that the model captures the housing choice of households that suffer a job separation shock. I study the asset choices of households facing an exogenous job separation and show the model replicates

the magnitude of downsizing in the housing market in response to unexpected unemployment, also measured in the PSID.

A key feature of the Great Recession was the accompanying financial crisis (Brunnermeier (2009), Jermann and Quadrini (2012), Eggertsson and Krugman (2012), Hall (2014)). The second chapter, ‘*Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics*’, looks at the role of credit conditions in household consumption dynamics. It argues that much of the literature has misunderstood the effect of a decline in credit availability during the Great Recession. Rather than forcing indebted households to reduce their consumption to lower borrowing, the main mechanism by which a credit crunch operates is to lower the probability of households making large durables or housing purchases. The reason for this is that households that are already borrowing are unlikely to be immediately affected by a contraction in credit availability. However, the decline in credit availability will reduce their ability to make future purchases which now require a larger downpayment. Further, their existing credit terms secured during more favourable aggregate conditions have become more valuable, providing a reason to avoid a purchase that would result in credit renegotiation. Such an incentive is likely to have become stronger following the growth of home equity line of credit (HELOC) loans during the first half of the 2000s (Johnson and Sarama (2015)). This mechanism can help us understand the prolonged response of durables consumption and deep decline in transaction volumes following the Great Recession.

The chapter presents a general equilibrium infinite horizon heterogeneous agents model, with durables consumption, subject to realistic non-convex adjustment costs, and featuring two forms of aggregate uncertainty: productivity shocks and collateral constraint shocks. The wealth distribution is a state variable and the solution to the equilibrium requires the estimation of the aggregate law of motion (Krusell et al. (1998)). Critically, the model relaxes

the commonly made assumption that a change in aggregate credit conditions affects all agents immediately (see: [Guerrieri and Lorenzoni \(2017\)](#), [Guerrieri and Iacoviello \(2017\)](#)). Instead in this setting conditional on non-adjustment agents retain previously agreed credit terms and this introduces a powerful channel by which the aggregate shock is propagated.

In particular, the model is able to replicate the consumption response of agents when a recession coincides with a contraction in the availability of credit, such as the Great Recession. The majority of consumer credit extended to US households is for the purpose of durables purchases and is usually secured against that stock. For example, in the 2010 Survey of Consumer Finances, 83.9 percent of family debt was secured against a residential property. Therefore, it is natural to assume that if households are affected by credit availability that credit secured against durable assets will be the most important class to consider for their consumption behaviour.

The main results of the chapter shows that a contraction in aggregate credit conditions leads to a deep, persistent decline in durables consumption and the percentage of agents adjusting their durables stock. This is the case even though only a small proportion of agents are near the collateral constraint. When a credit contraction occurs concurrently with a negative productivity shock there is additional propagation and the decline in the durables consumption share resembles that of the Great Recession.

Further, the model proposed with agent specific credit terms is better able to replicate the differing behaviour of adjusters and non-adjusters. The standard specification in the literature implies that when credit conditions change it is the adjusting group of households who's leverage ratio will closely follow changes in credit availability, with their behaviour constrained by movements in aggregate conditions. In fact, in the data this is not the case. Instead, as in the model there is a compositional change between adjusted as non-adjusters when credit conditions tighten, driven by the benefits of inaction – namely

retaining a household's prior agreed terms of credit. In this dimension the model replicates the empirical correlation more successfully than the widely studied alternative. Finally, the chapter finds that perhaps surprisingly shocks to credit availability are not amenable to policy intervention. Government action to loosen the constraint reduces consumer welfare, albeit the welfare loss is small.

The final chapter, '*(S)Cars and the Great Recession*', focuses on the car market and car purchase behaviour during the last three US recessions. As alluded to before, car expenditure during the Great Recession was remarkably different from that in previous recessions. Making use of data from the Consumer Expenditure Survey I show during the Great Recession there was a decline in both the probability of purchasing a car and the size of car purchased. Previous recessions featured only the former response, indicating that the Great Recession differed not only in its magnitude but also in the types of shocks hitting households.

In keeping with a strong economic tradition, I claim consumption dynamics can be highly informative of the unobserved shocks hitting households (Blundell et al. (2008)). The chapter uses the empirical features unique to the Great Recession to identify the source of shocks during the crisis within a partial equilibrium lifecycle setting featuring a rich income process, where households purchase cars subject to a transaction cost implying an (S,s) type of durable adjustment (Bertola and Caballero (1990), Grossman and Laroque (1990), Attanasio (2000), Bertola et al. (2005)). I expose households in the model to a wide variety of shocks consistent with the data and ask which of these generates consumption choices that provide the best fit with the empirical patterns observed. I argue that asset price changes, cohort-specific income shocks and a change to the expected growth rate of future income are important for understanding the consumption dynamics during the Great Recession.

Whilst providing a good account of the previous recession dynamics and

despite their significant size, aggregate income shocks alone are unable to generate the magnitude of the response of car expenditures seen during the Great Recession. Uniform income shocks also fail to adequately capture the heterogeneity of responses across the lifecycle. Introducing cohort specific shocks provides a better account of this variation, but as the magnitude of the shocks is the same the aggregate, responses are largely unchanged.

The introduction of asset price changes and a shock to the lifecycle growth rate, in combination with the large permanent income shock, provides a much fuller explanation of the consumption dynamics. The inclusion of these shocks allows me to capture both the extensive and intensive margin responses in the aggregate and the variation observed across the lifecycle. The house price boom helps to match the growth in consumption just prior to the recession, when the labor market was beginning to slow. During the crisis the fall in house prices generates a significant decline in the wealth of households, particularly concentrated on those that make up a larger share of aggregate non-durables and durables consumption. With lower cash in hand, households planning to adjust now prefer a smaller car, especially those with a shorter planning horizon. This has a strong effect on the intensive margin. The combination of the income shocks and wealth shock is particularly strong for middle aged households, that exhibit a large intensive margin response in the data.

Chapter 1

Job Risk, Separation Shocks and Household Asset Allocation

1.1 Introduction

Consumption underwent a large decline during the Great Recession. In the US consumption fell by 3.7 percent from peak to trough.¹ Why was this decline so large and to what extent did the conditions prior to the Great Recession amplify the consumption response? Further, if the conditions prior to the crisis did exacerbate the fall which channels were responsible for the amplification?

This paper investigates the role of labor market tranquility prior to the Great Recession as a key contributing factor to the magnitude of the subsequent consumption decline. It motivates a model where households' choices in both the labor market and asset market respond to a change in household risk. It suggests sorting in these dimensions left the economy more vulnerable to large unexpected shocks of the type experienced in the labor market between 2007 and 2009.

¹This compares to an average decline of 1.5 percent across all post World War II recessions. As a share of output consumption fell by 0.7 percent during the Great Recession compared with an average decline of 0.47 percent. The decline is much larger once detrending is taken into account. A comparison of the consumption declines during the post-war recessions is provided in Appendix C.1, Figure C.1

Prior to the Great Recession the economy is often characterised as going through a period of tranquility. A large literature has studied the so called Great Moderation, with the cause variously attributed to an exogenous reduction in firm volatility, improvements in the operation of monetary policy or good luck (see: [Davis and Kahn \(2008\)](#), [Clarida et al. \(2000\)](#) and [Stock and Watson \(2003\)](#)). In the labor market the period also saw a persistent decline in the job separation rate, the transition rate from employment to unemployment. The monthly job separation rate fell from 4.2 percent in 1980 to 2.7 percent in 2007 (Figure 1.2). Concurrently, the US also saw a decline in the cross sectional variance of the job separation rate. Even accounting for this fall significant heterogeneity remains with some households facing much larger risks in the labor market than others. Figure 1.1 shows the cross sectional distribution of the monthly job separation rate in the US as estimated from data sampled from the Current Population Survey (CPS).² There is significant heterogeneity in the level of job separation faced with some individuals facing a very low risk of unemployment while others face a much larger expected separation rate. This variation in risk is often abstracted from in macro models.

Another key stylized fact for this period is the change in household portfolios. Between 1980 and 2007, households reduced liquid savings and increased illiquid asset holdings. In particular, during this period households' housing stocks increased in both value and quantity. The combination of these asset choices meant that as the risk of job separation fell the median household's asset allocation became more illiquid.

These trends suggest that interactions between the labor and asset markets may have affected the responsiveness of the economy to shocks or that state dependency matters for aggregate dynamics. It is well understood that household's exposure to risk in the labor market is a key determinant of their asset

²Full details of the construction of this measure are provided in section 1.3.2.

allocation decision.³ Modern macroeconomics has increasingly recognized that the distribution of assets is important to the aggregate response of the economy to shocks (Kaplan and Violante (2018)). Recently the savings decisions of households and the resulting wealth distribution has received renewed focus in a growing literature that has linked features of the asset distribution, in particular the illiquidity of some forms of saving, to the distribution of the marginal propensity to consume (MPC).⁴

This paper studies the joint determination of households' asset allocation decisions and labor market outcomes. Its main contribution is to introduce novel labor market features into a state of the art general equilibrium incomplete markets model to understand how the interaction of asset and labor market decisions alters the economy's response to shocks. To do this I build a model that incorporates a jobs ladder, heterogeneity in job risk and saving in liquid and illiquid assets. In the labor market, I propose a model where households receive job offers that vary across two dimensions - wage and separation rate, a setting similar to Jarosch (2015) or Burdett and Mortensen (1980). A household that receives a new job offer is able to accept or reject the new wage-job risk combination and this decision depends on its current asset position. Over time in expectation, a continuously employed household will move to better paid, more secure employment.

To capture the rich heterogeneity in the wealth distribution the model features liquid savings, illiquid housing and long term mortgages. On the illiquid asset side I focus on housing as it is the largest illiquid assets for many households, with housing equity accounting for 52 percent of the median household's illiquid asset portfolio.⁵ Housing is also important because it is

³The role that income uncertainty plays in shaping the consumption-savings decisions of households has been one of the primary areas of study in macroeconomics during the past 30 years. Prominent examples include: Zeldes (1989), Deaton (1991), Aiyagari (1994), Carroll and Kimball (2001), Castañeda et al. (2003)

⁴See for example Kaplan and Violante (2014), Kaplan et al. (2016), Lütticke (2017)

⁵In the 2016 Survey of Consumer Finance (SCF) the median net housing holdings was \$25,000. The median net worth holdings was \$48,360

financed by long term debt obligations, which leverage a household's illiquid asset position and can make them sensitive to changes in house prices.

In the main result, I show the importance of the joint determination of asset and labor market outcomes by investigating the response of the economy to labor market shocks seen in the Great Recession. I shock the economy with labor market shocks from the Great Recession to the wage level and job separation rate, estimated in the CPS, allowing for heterogeneity in the shocks across job types. I compare the response of the aggregates when the shocks occur following a period of relative calm in the labor market - capturing the pre Great Recession environment - to the counterfactual case with higher prior job risk. For the purpose of the paper the fall in the job separation rate that captures the tranquil environment is treated as exogenous. The fact that the Great Recession occurred following a period of moderation and low risk in the labor market, amplified the negative response of consumption by up to 40 percent. Households' equity holdings also suffer a larger fall, incorporating both price and quantity effects. The amplification of the response is due to households adopting more illiquid portfolios, increasing their housing stocks and sorting into higher wage jobs. More illiquid portfolios increased households' sensitivity to transitory shocks by raising MPCs, the larger housing stock holdings placed greater downward pressures on the housing market generating larger house price declines. Stronger labor market attachment allowed households to sort into "better jobs" prior to the crises, raising the subsequent cost of unemployment.

Labor market heterogeneity is also important for this result. In particular, the persistence of the consumption decline is exacerbated by job risk heterogeneity as households moving into employment from unemployment have a higher expected separation rate than the average household already in employment. Therefore, an increase in the unemployment rate raises the future separation rate, increasing the persistence of the shock. I also evaluate the

model's ability to replicate the cross sectional consumption and asset choice dynamics. The model is able to replicate key features of the asset market decisions during this period, including the relative housing response across job types compared with households in the Panel Study of Income Dynamics (PSID). This is a test of the joint covariance of the shocks hitting the economy during the Great Recession and prior allocation in both the labor and asset markets. The success in replicating this pattern is interpreted as supportive evidence in favor of sorting in the labor and asset market as being important for understanding the magnitude of the Great Recession.

I also show that the model delivers empirically realistic asset and labor market interactions. This serves to validate the findings in response to the Great Recession shocks. Firstly, the model generates key features of the income process including negative skewness and excess kurtosis as documented in administrative datasets (see: [Guvenen et al. \(2015\)](#)). The new consensus in the labor literature is that unemployment shocks can result in significant and persistent income losses. As in other work, the jobs ladder feature reconciles the model with this empirical fact, meaning that for many households unemployment shocks are far from transitory. The jobs ladder also endogenously generates a positive correlation between the wage level and job security as seen in the data.

Secondly, the model replicates the positive correlation between job risk and the liquidity of a household's portfolio as observed in the PSID. Job risk is important for how households choose to allocate their assets, as households facing less risk and requiring lower precautionary savings are able to allocate a greater share of their assets to more illiquid forms of wealth.

Thirdly, I show the model captures the housing choice of households that suffer a job separation shock. Using methods similar to those used in the labor literature, I study the asset choices of households facing an exogenous job separation and show the model replicates the magnitude of downsizing in

the housing market in response to a job separation shock as measured in the PSID.

A contribution of the paper is to allow for a feedback mechanism from the asset allocation decision to the labor market via households' choices over which jobs to accept or reject. Counterintuitively, in the model I find that poorly insured households with low liquid asset holdings are more willing to forgo increased job security for an increased wage than households with high liquid asset holdings. The reason for this is the slope of the illiquid household consumption function, favoring wage gains today over security tomorrow. As a result, if all households were to make decisions over jobs following the choices of wealthy households, the economy would feature a lower unemployment rate.

Finally, introducing job risk heterogeneity is found to raise the aggregate MPC. This provides an alternative explanation to the previous literature which has introduced lifecycle motives, a large spread between the return on illiquid and liquid assets or heterogeneous preferences (see [Kaplan and Violante \(2014\)](#) and [Carroll et al. \(2017\)](#)). However, I find that job risk only raises MPCs by a small amount (around 15 percent on a quarterly basis). A key reason for this is that households' precautionary savings motive is highly non-linear in uncertainty with only very low risk households reducing liquid savings sufficiently to become hand-to-mouth. In the data there are relatively few of these households. Further, the positive correlation between job security and earnings means that households with low job risk tend to have high earnings - strengthening their precautionary savings motive. These mechanisms weaken the relationship between job risk and wealthy hand-to-mouth status in the model.

The rest of the paper is set out as follows. Section 1.2 sets out the related literature, section 1.3 provides empirical evidence on job risk and asset allocation; section 1.4 presents the model; section 1.5 analyses households choices in the steady state equilibrium, section 1.6 presents the response of the economy

to the Great Recession shocks; while section 1.7 concludes.

1.2 Related literature

This paper brings together to two broad literatures: firstly, research that studies heterogeneous agent economies with incomplete markets and multiple assets, and secondly, research into the effect of variation in income risk or job uncertainty on household choices. In the former strand, [Kaplan and Violante \(2014\)](#) is arguably the foremost contribution. That paper introduced the concept of wealthy-hand-to-mouth agents that hold little liquid assets but substantial illiquid assets and have a high marginal propensity to consume. Introducing wealthy hand-to-mouth households into a heterogeneous agents model enabled the authors to replicate the expenditure out of tax rebates empirically observed in the literature e.g. [Souleles et al. \(2006\)](#) and [Parker et al. \(2013\)](#). [Kaplan et al. \(2016\)](#) embedded this framework in a New Keynesian sticky price model, finding the response to monetary policy shocks to be of the same magnitude as that of a representative agent model.

The role of housing in the heterogeneous agent environment has been studied by a wide number of authors. Notable recent contributions include [Berger and Vavra \(2015\)](#); [Hedlund \(2016\)](#); [Favilukis et al. \(2017\)](#); and [Kaplan et al. \(2017\)](#). Of particular relevance to this paper is [Berger et al. \(2018a\)](#) which relates the consumption response to house price shocks to the marginal propensity to consume and shows that the response is sensitive to the level of household debt in the economy. [Rios-Rull and Huo \(2016\)](#) study the large consumption decline during the Great Recession and the role of heterogeneity in the housing market in amplification but consider financial rather than labor market shocks as the primary source of disturbance. With a focus on monetary policy, [Hedlund et al. \(2017\)](#) develop a frictional housing market in a heterogeneous agent New Keynesian model and highlight the housing channel in the

transmission of monetary policy. On the interest rate channel, [Wong \(2017\)](#) documents and replicates a lifecycle pattern in the consumption response to interest rate shocks, where the heterogeneity is due to the desire to remortgage a household's mortgage in response to an interest rate cut. A related result is [Cloyne et al. \(2016\)](#), who find in the US and UK data that mortgage holders respond strongly to interest rate changes, whereas homeowners do not.

The link to the labor market in this class of models has been made by [Ravn and Sterk \(2017\)](#) where countercyclical income risk is driven by endogenous changes in the job finding rate. [Challe \(2017\)](#) studies optimal policy when unemployment risk is endogenous, finding that the setting calls for accommodative monetary policy responses to cost-push shocks, the opposite of the optimal policy in representative agent models, due to the additional precautionary savings motive. More closely related to this paper is [Bayer et al. \(2015\)](#), where the precautionary savings motive responds to time varying changes in income uncertainty. Output contracts in response to an increase in income risk with the liquidity of the average household portfolio increasing. Here, I mainly focus on the implications of cross sectional variation in income uncertainty rather than time variation. However, the Great Recession shocks analysed also led to an increase in labor market uncertainty.

A second important literature is the study of labor market risk and its effect on household choices. Recent contributions include [Lise \(2013\)](#) who studies on the job search with precautionary savings, highlighting the interactions between choices in the labor and asset market. [Hubmer \(2018\)](#) extends this model with human capital and lifecycle dynamics to show a jobs ladder model is capable of replicating the higher order incomes moments estimated in administrative data sources, documented by [Guvenen et al. \(2015\)](#). In that model the wealth of households effects their income process by determining search effort, rather than a wage-job risk trade-off studied here. [Low et al. \(2010\)](#) make progress on disentangling the exogenous and endogenous sources

of risk in a lifecycle set up with job mobility, focusing on the welfare value of different forms of partial insurance. Also relevant is [Krusell et al. \(2010\)](#) who introduce a Diamond-Mortensen-Pissarides style labor market in an otherwise standard incomplete markets model and find important implications for optimal unemployment insurance. [Eeckhout and Sepahsalari \(2015\)](#) study the relationship between asset holdings and unemployment, but focus on the precautionary job search motive rather than job risk heterogeneity. Their model is directed search but features only one asset. [Jung and Kuhn \(2017\)](#) introduce heterogeneity in job stability to explain the size of income losses, stressing the importance of the loss of good jobs at the top of the jobs ladder.

The study most related to this paper is [Jarosch \(2015\)](#). In a similar spirit to the model studied here, that paper proposed a random search model where jobs vary over two dimensions, productivity and job security. This generates ‘slippery’ lower rungs of the jobs ladder, characterized by movement in and out of unemployment. The model is able to capture the persistent negative employment and wage effects following job loss in Germany. In comparison to that paper, here a simpler labor market is studied, but the model is enriched with a consumption savings choice. This allows the implications for the wealth distribution to be analyzed. In contemporaneous work, [Krivenko \(2018\)](#) also builds a model combining unemployment scarring and the housing market during the Great Recession, emphasizing the importance of exogenous moving shocks in generating the large house price decline of the Great Recession. Compared to that paper the labor market specification here is richer featuring household choice over job risk, rather than an exogenous process.

There is also a literature that has sought to find evidence of a precautionary savings motive in the data, without full agreement.⁶ Our approach is most

⁶[Carroll and Samwick \(1998\)](#) find evidence in the PSID of higher wealth for individuals with greater income uncertainty consistent with a buffer-stock model. While [Mishra et al. \(2012\)](#) find evidence of precautionary savings for farm households. [Guiso et al. \(1992\)](#) present evidence of precautionary savings in the Italian data, but the size of the effect is fairly small. Conversely, [Fulford \(2015\)](#) finds little evidence of precautionary savings based

similar to [Carroll et al. \(2003\)](#) who find evidence of a precautionary effect for moderate and higher income households but not for low income households in the Survey of Consumer Finance. Surprisingly this effect is only found to be present when illiquid housing equity is included in the definition of wealth. In contrast in the PSID there is a correlation between risk and the liquidity of the household portfolio. [Basten et al. \(2016\)](#) present Norwegian evidence consistent with an increase in financial wealth and reallocation towards safe assets prior to a unemployment spell.

In the broader empirical literature, once controlling for sorting [Cubas and Silos \(2017\)](#) find evidence of labor earnings compensation for higher permanent risk. [Chetty et al. \(2017\)](#) consider the effect of housing on a household's portfolio choice, with higher property value reducing stock holdings while greater equity wealth has the opposite effect. [Chetty et al. \(2017\)](#), shows that job-stayers face less dispersion in earnings growth, with positive rather than negative skew and experience greater kurtosis in income shocks than job switchers.

1.3 Data

This section motivates the model presented subsequently. Firstly, it discusses the aggregate job separation rate and shows the separation rate changed for different groups during the Great Recession. Secondly, it presents the measure of job risk used in this paper and considers how it varies over time. Finally, it presents analysis linking job risk to household asset allocations, providing evidence that supports the hypothesis that households facing lower job risk keep more of their wealth in illiquid assets.

on households stated preferences. [Jappelli et al. \(2008\)](#) reject the buffer stock behavior in Italian data.

1.3.1 Job separation rate

Figure 1.2 illustrates the job separation rate (or employment exit probability) using aggregate data from the CPS and the methodology of [Shimer \(2012\)](#). The series is constructed from the level for employment, unemployment and short term unemployment (less than five weeks) published by the Bureau of Labor Statistics for the period 1976-2018.⁷ It accounts for time aggregation that causes short spells of unemployment to be otherwise missed.

The estimated job separation rate exhibits a strong downward trend. While the decline is partly explained by demographic and education changes, it is also symptomatic of declining turnover in the US labor market (see [Molloy et al. \(2016\)](#), [Kaplan and Schulhofer-Wohl \(2017\)](#), [Fujita \(2018\)](#)). Indeed, even after controlling for compositional effects, sub groups still exhibit sizeable declines in the job separation rate. Figure 1.3 shows that after constructing the same measure for males or low educated males from the CPS micro data a decline is still apparent. While there were a number of trends in the risks facing households during the period known as the Great Moderation, in the job separation dimension households experienced rising economic security ([Davis \(2008\)](#)).

The key finding of [Shimer \(2012\)](#) is the lack of cyclicality in the job separation rate, with business cycle movements in the unemployment rate instead being largely determined by changes in the job finding rate. In contrast to the separation rate the job finding rate does not display an obvious secular trend prior to the Great Recession (see Appendix C.1, Figure C.2). Despite the well known Shimer conclusion, a notable feature of the Great Recession was the significant rise in the job finding rate, which increased from 2.6 percent in December 2007 to 3.8 percent in January 2009.⁸

⁷For the period after January 1994, CPS micro data is needed to construct the short term unemployment measure. See [Shimer \(2012\)](#) for details

⁸A recent literature has also reevaluated the importance of job separation shocks. [Ahn and Hamilton \(2016\)](#) emphasis that Shimer understates the importance of job separation

The rise in the job separation rate during the Great Recession was not uniform across job types. Figure 1.4 presents the change in the job separation rate at an annual frequency for different groups of workers. Job types are broken up into high and low wage types and into jobs with a higher or lower ex ante risk of job separation (high and low risk types to be discussed in Section 1.3.2). Between 2007 and 2009, the job separation rate rose by more in percentage terms for high risk types (76 percent versus 50 percent) and high wage types (84 percent versus 56 percent).⁹ After 2009, the heterogeneity in separation rate increases dissipates for both groups. As shown in Appendix C.1 Figure C.3 the effect on weekly earnings also differed by group type. The fall in wages by risk type was fairly similar, with high risk types seeing a slightly larger fall. In comparison low wage types saw much larger declines in weekly earnings.

1.3.2 Job risk distribution

As shown in Figure 1.1 the US labor market is characterized by significant heterogeneity in the degree of job risk workers face. Job risk is a key measure used in this paper, in this section I explain how this measure is constructed. The data used is from the Current Population Survey (CPS) micro data for the years 1987-2018. Let the outcome variable u be a dummy with the value of 1 if an individual i is unemployed in period t and 0 otherwise. I estimate the Probit model for job separation:

$$Pr(u_{i,j,o,s,t+1} = 1 | \mathcal{X}) = \Phi \left(\alpha_0 + \gamma_j + \mu_o + \eta_s + \phi_t + \theta y_{i,t} + \beta \mathbf{X}_{i,t} \right) \quad (1.1)$$

shocks to unemployment variations as these shocks alter the composition of the unemployed driving changes the job finding rate. From a theory perspective, [Coles and Kelishomi \(2015\)](#) relax the free entry condition in the standard DMP model and argue the model no longer implies a small role for job separation shocks.

⁹The aggregate rise in the job separation rate is larger here as these figures are not adjusted for temporary unemployment spells

I include fixed effects for industry j , γ_j ; occupation o , μ_o ; state s , η_s ; and ϕ_t time.¹⁰ Industries and occupations are designated at the broad industry and occupation classification level generating 13 industry and 25 occupation types in the data set. The additional controls are a measure of income, $y_{i,t}$, and a set of observables including a quartic in age, education dummies, and other demographic controls, $\mathbf{X}_{i,t}$. These conditional variables are summarised in $\mathcal{X} = (j, o, s, t, y_{i,t}, \mathbf{X}_{i,t})$

The value of job risk used is the predicted probability of job separation for an industry-occupation-state cell, excluding idiosyncratic factors:

$$\hat{Pr}(u_{j,o,r} = 1|\bar{\mathcal{X}}) = \Phi\left(\hat{\alpha}_0 + \hat{\gamma}_j + \hat{\mu}_o + \hat{\eta}_r + \phi_{2007} + \hat{\theta}\bar{y} + \bar{\mathbf{X}}\hat{\beta}\right) \quad (1.2)$$

The year is set to 2007. The distribution of these jobs can then be plotted using the distribution of jobs in the dataset.

In the baseline specification shown in Figure 1.1 the log. of average weekly earnings is used.¹¹ Figure C.4 in Appendix C.1 presents the distribution using alternative income controls. The distribution produced is fairly consistent, with the income controls accounting for part of the heterogeneity and reducing the mean separation rate.

This method can also be used to investigate how job risk has changed over time. To do this I run the same regression (1.1) for five year intervals, $\{j - 4, \dots, j\}$ and then use the distribution of jobs in the last year, j , as a measure of the job risk faced. This approach captures changes in both the job risk of a given job and changes in the composition of jobs in the economy.¹²

The distribution of job risk has changed substantially over time as shown

¹⁰For industry the 1990 Census Bureau industry classification scheme is used. For occupation the 2010 Census Bureau occupation classification scheme is used.

¹¹For each individual weekly earnings is recorded twice, the observations are one year apart at the end of individuals' four month rotations. The average over these two values is taken to reduce measurement error.

¹²No income control is used for this exercise, to extend the years available in the analysis and to maximize the sample size for each year.

in Figure 1.5, Panel a. Prior to 2005, both the mean and the median job risk underwent a secular decline as previously seen in Figure 1.4. There has also been a change in the higher cross sectional moments of the the distribution. There has been a large reduction in the cross sectional standard deviation of job risk, falling by almost 50 percent between 1985 and 2000, before rising after the Great Recession (see Panel b). Similarly the cross sectional skewness of the distribution reached a minimum in 2005, before increasing during the recession.¹³ Kurtosis (not shown) follows a similar pattern to skewness. Finally, Panel d provides a cross sectional example of the extent to which the job risk distribution has changed, comparing the distribution in 2007 when the standard deviation was low with 1984 when the standard deviation was high. Figure C.5 in Appendix C.1 provides a further selection of years evidencing how the distribution of job risk has evolved over time.

1.3.3 Asset allocation and job risk

Having constructed a measure of job risk in the data, I now consider how the liquidity of a household's asset allocation is correlated with job risk it faces. Later, I will compare these results to the model to validate the ability of the model to match the joint determination of assets and labor market outcomes. For this analysis I use data from the PSID which due to its panel structure can be used to construct a measure of job risk and since 1984 also contains information on household wealth. Household wealth information is available in 1984, 1989, 1994 and then every year in the biannual PSID from 1999 onwards.

Asset time series

Before undertaking the cross sectional analysis it is worth considering the aggregate time series movements of the household portfolio. Figure 1.6 presents the PSID time series for median housing holdings and median liquid assets

¹³Pearson's moment coefficient of skewness is used

(excluding stocks) as defined below. Between the 1980s and 2007 when the job separation rate was falling, the median household increased their housing stock and reduced liquid assets. This pattern is consistent with the reallocation of the household portfolio into more illiquid forms of saving. This remains the case when we deflate the housing stock by a house price index (labeled quantity) rather than the consumer price index, indicating that the effect is not purely driven by the increase in house prices. Figure C.6 in Appendix C.1 shows these patterns are robust across a variety of alternative data definitions and consistent with portfolio reallocation.

Portfolio liquidity

We now take advantage of cross sectional variation. The idea is to regress a measure of liquidity, $L_{i,t}$, on estimated job risk, $\hat{\delta}_{i,t}$, a function of the industry, occupation and state of the individual. Household wealth is divided into liquid and illiquid assets. Liquid wealth includes checking accounts and stocks net of credit card debt, student loans, medical debt, legal and family debt, while illiquid wealth includes housing equity, other real estate, other assets and IRA accounts, following [Kaplan et al. \(2014\)](#).

To address the non-linearities in household liquidity, with the possibility of the denominator becoming very small or negative, two alternative measures of liquidity are used: 1. the liquid asset to illiquid asset ratio; 2. liquid asset to total asset ratio. An inverse hyperbolic sine transformation is applied to this ratio following [Carroll et al. \(2003\)](#).¹⁴ For both dependent variable specifications regressing the non-transformed liquidity measure on job risk does not generate significant results.

In the baseline specification job risk is estimated in the PSID. This follows

¹⁴Like a log transformation the inverse hyperbolic sine transformation $g(Y, \theta) = \log(\theta Y + \sqrt{(\theta Y)^2 + 1})/\theta$, down weights large values of Y , however, in contrast to log it admits zero and negative values of Y , as suggested by [Burbidge, Magee and Robb \(1988\)](#). In [Carroll et al. \(2003\)](#) the parameter θ is estimated, here following [Basten et al. \(2016\)](#) θ is set to one.

the same methodology as in 1.3.2, with the log of permanent family income used for the income control.¹⁵ In the second stage the equation:

$$L_{i,t} = \alpha_0 + \lambda \hat{\delta}(j, o, s)_{i,t} + \phi_t + \theta y_{i,t} + \beta \mathbf{X}_{i,t} + \epsilon_{i,t} \quad (1.3)$$

is estimated where controls for permanent income $y_{i,t}$, year fixed effects, ϕ_t , a quartic in age, education dummies, and demographic variables, $\mathbf{X}_{i,t}$, are included.¹⁶ The parameter of interest is λ .

The results for the ratio of liquid assets to illiquid assets are presented in Table 1.1. In the baseline specification, column (1), it can be seen that increased job risk is associated with a more liquid asset allocation. The relationship is highly significant. At the 10th percentile of the liquid asset to illiquid asset ratio, a one standard deviation increase in job risk implies a 5.5 percent reduction in the liquid asset to illiquid asset ratio. The 10th percentile liquidity holding is negative so this is an increase in the liquidity of the portfolio. At the median liquidity holdings a one standard deviation increase in job risk implies 74 percent increase in the liquidity ratio.

As will be seen below, the model predicts the relationship between liquidity and job risk will be particularly strong at very low job risk levels. Evidence for this in the data is provided by using dummy indicators for the lower percentiles of job risk, column (2). The effect on asset allocation is particularly strong for households in the 1st job risk percentile. After that the effect decreases although is still significant for the 5th and 10th percentile.

The rest of the table provides a series of alternative specifications to serve as robustness checks. Column (3) restricts the sample used in the PSID to the

¹⁵Permanent family income is the average income over all household observations. As there is a long panel for each household this measure should not be too affected by a future unemployment spell.

¹⁶Our identification shares similarities with [Carroll et al. \(2003\)](#). They estimate a logit model of the probability of unemployment and use the predicted values, instrumented by state.

Core sample, the results are significant but the coefficient is slightly reduced.¹⁷ Thus far, the evidence presented does not prove a causal link, as it could equally be the case that households with illiquid asset allocations choose low risk jobs or in the confounding case households with greater risk aversion choose safer jobs and a more liquid portfolio. While for the purpose of model validation it is sufficient to consider correlations in the data, evidence is also presented that shows a causal relationship from risk to assets appears to exist. In column (4) I include a series of controls for the household's risk preferences, including the share of liquid assets invested in stocks, expenditures on home, automotive and health insurance relative to household income, and total assets. The results remain with only a small decline in the coefficient. In column (6) I control for individual risk preferences by using a household fixed effect. Again the coefficient is a little lower but remains significant. In columns (7) and (8) I address reverse causality by instrumenting job risk by its value at two and six year lags. The results remain statistically significant, with some reduction in the coefficient in the latter case. The only robustness test which returns contrary results is reported in column (5). Here additional controls for industry and occupation in the second stage are used, essentially using only the variation across regions to identify job risk. This is a similar approach to [Carroll et al. \(2003\)](#). Here the coefficient is negative, but the effect is statistically insignificant.

I repeat the analysis with an alternative definition of liquidity: the ratio of liquid assets to total assets. The results presented in Table 1.2 are broadly similar. In contrast to the ratio of liquid assets to illiquid assets, I do find evidence of a liquidity effect when controlling for industry and occupation, column (6). However, the results when including individual fixed effects are

¹⁷The justification for not restricting the sample to the core sample in the PSID in the baseline analysis is that the unemployment rate is higher for the additional samples: Core: 0.039; SEO: 0.095; immigrant: 0.075; latin: 0.087. Including these observations provides greater variation in the job risk variable.

not significant. As a further robustness check I repeat the results using job risk estimated from the CPS. These are presented in Appendix B.1, Tables B.1 and B.2. For the liquid asset to illiquid asset ratio the results are largely supportive, although no longer significant when using individual fixed effects. For liquid assets to total assets the results are less clear with the sign changing for some of the specifications.

Hand-to-Mouth status

To relate the analysis directly to the hand-to-mouth literature, a dummy variable for hand-to-mouth status, differentiating between Poor Hand-to-Mouth (PHTM) and Wealthy Hand-to-Mouth (WHTM) is regressed on the job risk measure. Hand-to-mouth is defined as those households with positive liquid wealth, but less than one week of household income in liquid assets, or negative liquid wealth and less than one week of household income from their budget constraint.¹⁸ Poor hand-to-mouth are hand-to-mouth households with weakly negative illiquid assets.

The results in Table 1.3 show higher job risk is associated with a higher probability of being PHTM and a lower probability of being WHTM. I also present results when controlling for industry and occupation and lagged job risk.¹⁹ One explanation that fits with this pattern is that households with high job risk end up being PHTM following a period of unemployment when they consume their liquid assets. Low job risk households in contrast end up choosing to allocate their assets toward more illiquid forms, such as housing, as they require less insurance given the lower job risk faced. In Appendix B.1, Table B.3 replicates the analysis for the core PSID sample, while B.4 uses job risk estimated from the CPS. The results show the same pattern.

¹⁸The budget constraint is defined as one month of household income

¹⁹The other specifications presented in the portfolio analysis yield similar results so are not presented here. One exception is the inclusion of individual fixed effect for which the results were not significant.

1.4 Model

This section sets out the model that I will use to study the joint determination of labor and asset market outcomes. I study an infinite horizon model in which households make asset allocation decisions and sort across job types. There is a unit mass of households. In the labor market, workers' time is converted into output with a linear production function. The government collects taxes, which it spends on benefits and government expenditure.

1.4.1 Model environment

Households are infinitely lived and value consumption over non-durables, c_t and housing services, \tilde{h}_{t+1} . Households make choices to maximize the expected value of lifetime utility:

$$\mathbf{E}_t \sum_{t=0}^{\infty} \beta^t u(c_t, \tilde{h}_{t+1})$$

Households can be employed or unemployed. Employed households hold jobs that vary in two dimensions: by the wage, ω_t , and separation probability, δ_t . Job offers arrive stochastically for both the employed and unemployed. Households that receive a new offer can reject the new job and remain in their current job. If the unemployed household rejects a job offer they remain unemployed. In addition, unemployed households vary by the quality of their job draw, ϵ_t .

There is a fixed stock of owner occupied housing, \bar{H} , in the economy with aggregate house price p_t . A rental market converts the consumption good into rental units with price r^p . Each employed worker, i , is employed by a firm, j . Firm-worker pairings generate match quality, μ_t . Workers are paid their marginal product, such that $\omega_t = \mu_t$. Aggregate output is the integral over firms, j . $Y_t = \int y_t^j dj$. Finally, the government provides unemployment

insurance, collects taxes, issues bonds and undertakes government spending.

1.4.2 Household asset choice

At the beginning of each period households make a consumption savings decision. Households can choose to allocate their wealth between liquid assets, b_t , and housing, h_{t+1} . Households holding a positive housing stock may also borrow in the form of a mortgage, m_{t+1} . The three assets differ in their return and liquidity.

Liquid assets

Liquid assets are freely adjustable and earn the period return r . Households are able to borrow in liquid assets up to a borrowing constraint \underline{b} . Borrowing is more costly than saving with the spread being r^b . The interest rate schedule on liquid assets can be represented by the function $R(b)$, which features a kink at zero savings.

$$R(b_t) = \begin{cases} r & \text{if } b_t \geq 0 \\ r + r^b & \text{if } b_t < 0 \end{cases}$$

Housing

Housing choice is discrete with the choice from the set $h_{t+1} \in \mathcal{H} \subseteq [\underline{h}, \bar{h}]$. The discrete choice of housing is chosen for tractability and to capture the lumpy nature of the housing choice. The housing stock does not depreciate. In addition, housing is an illiquid store of wealth. Housing does not generate a financial return but provides a period utility flow, with the housing service flow proportional to the size of the stock $\tilde{h}_t = h_t$. Adjusting the housing stock requires the household to pay the non-convex adjustment cost $\Psi(h_t, h_{t+1})$,

which enters into the budget constraint.

$$\Psi(h_t, h_{t+1}) = \begin{cases} \Psi h_t \text{ if } h_t > 0 \ \& \ h_{t+1} \neq h_t \\ \Psi h_{t+1} \text{ if } h_t = 0 \ \& \ h_{t+1} > 0 \\ 0 \text{ if } h_{t+1} = h_t \end{cases} \quad (1.4)$$

For existing homeowners the cost is proportional to the current stock, for households with zero housing holdings the cost is proportional to the end of period choice. The cost can be thought of as a combination of realtor fees and a time cost of finding a new property. The adjustment cost is spent resources and does not flow to a financial institution.

Households can also choose to rent. I impose that renters may only rent the smallest housing stock size, $h_t^r = \underline{h}$. Renters receive a housing service flow proportional to the size they rent $\tilde{h}_{t+1} = h_{t+1}^r$ and must pay the rental price $r^h p h^r$ each period. As in [Hedlund et al. \(2017\)](#), renters do not rent housing from other households. Instead, a rental technology exist to convert the consumption good into rental services.

Mortgages

Housing can be purchased with a mortgage, m_t , which is subject to a higher interest rate $r + r^m$. There is no default and renters cannot access mortgages. Households are able to remortgage or make a mortgage repayment.²⁰ Households that remortgage take out a new mortgage worth $\Theta p_t h_{t+1}$ where Θ is the maximum loan to value ratio. They must also pay a fixed refinancing cost, Ψ^m .

Households that make the mortgage repayment must pay the interest on the mortgage today $r + r^m$ and make the minimum repayment $(1 - \gamma)m_t$. The value of the mortgage tomorrow is $m_{t+1} = \gamma m_t$. Households are restricted from paying off their mortgage early. If making a mortgage payment (rather

²⁰Restricting the mortgage choice set ensures the problem is computationally feasible

than refinancing) and moving to a smaller property a household's mortgage must not exceed the maximum loan to value ratio of the new property. In this case they must also pay the difference $\Theta p_t h_{t+1} - \gamma m_t$.²¹

Mortgage refinancing is only available to employed households with wage exceeding $\tilde{\omega}$, consistent with the requirement to provide income proof for non-self certified mortgages. This simple set up captures a variety of the features of the mortgage market - mortgage debt is a long term commitment to payments, there are restrictions on access to credit and mortgages lever a household's asset position in the face of house prices changes. Mortgages also provide a way for households to access the wealth in their illiquid asset housing and one that is cheaper than unsecured borrowing.

1.4.3 Job choice

The labor market is characterized by the stochastic arrival of job opportunities that households can choose to accept or reject. Households can either be employed or unemployed. All employed workers inelastically supply one unit of labor, with no utility cost. Unemployed households receive unemployment insurance, κ . The replacement rate does not depend on the household's previous employment.

Jobs vary across two dimensions: the wage, ω_t , and job separation probability, δ_t , which I will refer to as 'job risk'. Given this set up each job can be summarized by the state (ω_t, δ_t) . The labor market transitions take place after the asset choice for the period has been made.

Each period the unemployed receive a new job offer with probability λ_0 . The unemployed vary in the quality of job offers they receive, this is summarized by the state ϵ_t . Conditional on receiving a job offer they draw a job from

²¹Prepayment of mortgages is a feature of the US mortgage market, however in this model with a fixed mortgage rate there is relatively weak incentives for a household to pay back a mortgage early. Households are in fact able to pre-pay a mortgage in two periods, by becoming a renter and then repurchasing the same housing stock next period, although this would incur adjustment costs.

the distribution $G_\epsilon(\omega_{t+1}, \delta_{t+1})$ and may choose to accept or reject this offer. If they reject the offer the household remains unemployed and updates their job offer quality state ϵ_t with distribution $Z_\epsilon(\epsilon_{t+1})$. The quality of job offers is declining with the duration of unemployment.

After making choices in the asset market the employed first discover whether they will suffer a job separation, which occurs with the job specific probability δ_t . Households that do not lose their job then receive a job offer with probability λ_1 . The job offer arrival rates satisfy $\lambda_0 > \lambda_1$ such that job offers arrive more quickly for the unemployed, reflecting the cost of on the job search. The new job is drawn from the distribution $F_{\omega, \delta}(\omega_{t+1}, \delta_{t+1})$ capturing that the employed may receive different job offers to the unemployed, which is conditional on the current job state (ω_t, δ_t) . With probability ρ^y households can choose to accept or reject the offer. If they reject the offer they continue to the next period with their current job (ω_t, δ_t) . As jobs differ across two dimensions, wage level and job risk, the choice of whether to accept or reject a new job is not trivial and may depend on the household's other state variables. With probability $1 - \rho^y$ the household must accept the job drawn. This is interpreted as an income or contract shock rather than a job to job transition.²²

For the employed that experience a job separation they transition to the unemployed quality state with distribution $Z_\epsilon(\epsilon_{t+1})$. There is also the possibility of an immediate return to employment. With probability λ_2 they draw an in period job offer from the unemployed job offer distribution $G_\epsilon(\omega_{t+1}, \delta_{t+1})$, which they can choose to accept or reject. The result of the specification is that over time a household that remains continually employed will move towards a better paid more secure job, trading off the different attributes of new job offers against their current bundle (ω, δ) conditional on their current state.²³

²²These types of shocks are commonly used in the job ladder literature and often referred to as “Godfather shocks”, see [Moscarini and Postel-Vinay \(2018\)](#)

²³Supportive evidence for this specification can be found in the CPS. Figure C.7 in Appendix C.1 shows how job risk by job type, after controlling for age effects, still falls with age. This is purely driven by a compositional change and is consistent with a jobs ladder

1.4.4 Recursive formulation

The household choice problem can be presented in a recursive formulation. For this purpose prime notation is used to denote the next period. Households discount the future with discount factor β . There are value functions for the employed, V , and unemployed, U . The state variables for the household are liquid asset holdings, b , beginning of period housing stock, h , and current mortgage, m . Denote the household's asset holdings state as $\mathbf{s} = (b, h, m)$. The employed households have the additional state variables wage, ω , and job risk, δ . The unemployed have the additional draw quality state, ϵ .

Labor market expectation operators

Given the range of labor market outcomes it is useful to first define some expectation operators over the stochastic variables. Let the unemployed's expectation over job draw quality be: $\hat{U}_u(\mathbf{s}, \epsilon) = \int U(\mathbf{s}, \epsilon') dZ_\epsilon(\epsilon')$ and the employed's expectation over job draw quality be: $\hat{U}_e(\mathbf{s}) = \int U(\mathbf{s}, \epsilon') dZ_\epsilon(\epsilon')$. With these two objects I can now define the expectation of an unemployed household that has received a job offer:

$$\hat{V}_u(\mathbf{s}', \epsilon) = \iint \max\{V(\mathbf{s}', \omega', \delta'), \hat{U}_u(\mathbf{s}', \epsilon)\} dG_\epsilon(\omega', \delta')$$

where the first term in the max operator is the value of accepting the job and the second term is the value of rejecting the job and remaining unemployed. Similarly, the expectation of the employed household that has received a job offer can be defined as:

$$\begin{aligned} \hat{V}_e(\mathbf{s}', \omega, \delta) = & \rho^y \iint \max\{V(\mathbf{s}', \omega', \delta'), V(\mathbf{s}', \omega, \delta)\} dF_{\omega, \delta}(\omega', \delta') + \\ & (1 - \rho^y) \iint V(\mathbf{s}', \omega', \delta') dF_{\omega, \delta}(\omega', \delta') \end{aligned}$$

model where job security is positively associated with time spent climbing the jobs ladder.

where the first term in the max operator on the first line is the value of accepting the job and the second term is the value of rejecting the job and retaining the current employment state (ω, δ) . The second line is the probability of having a forced income change or contract shock.

Employed households

The value function for the employed refinancing household, $V^R(\mathbf{s}, \omega, \delta)$ is:

$$V^R(\mathbf{s}, \omega, \delta) = \max_{c, b', h' \in \mathcal{H}} u(c, \tilde{h}') + \beta \left[(1 - \delta) \left((1 - \lambda_1) V(\mathbf{s}', \omega, \delta) + \lambda_1 \hat{V}_e(\mathbf{s}', \omega, \delta) \right) + \delta \left(\lambda_2 \int \hat{V}_u(\mathbf{s}', \epsilon) dZ_e(\epsilon) + (1 - \lambda_2) \hat{U}_e(\mathbf{s}', \omega, \delta) \right) \right]$$

subject to

$$\begin{aligned} c + ph' + b' - m' &= (1 + R(b))b - (1 + r + r^m)m + ph + \omega(1 - \tau^\omega) \\ &\quad - \tau^l - \Psi(h, h') - \Psi^m - r^p h^r \mathbb{1}[h' = 0] \\ b' &\geq \underline{b} \\ m' &= \Theta ph' \end{aligned}$$

The value function for the employed repaying household, $V^P(\mathbf{s}, \omega, \delta)$, is:

$$V^P(\mathbf{s}, \omega, \delta) = \max_{c, b', h' \in \mathcal{H}} u(c, \tilde{h}') + \beta \left[(1 - \delta) \left((1 - \lambda_1) V(\mathbf{s}', \omega, \delta) + \lambda_1 \hat{V}_e(\mathbf{s}', \omega, \delta) \right) + \delta \left(\lambda_2 \int \hat{V}_u(\mathbf{s}', \epsilon) dZ_e(\epsilon) + (1 - \lambda_2) \hat{U}_e(\mathbf{s}', \omega, \delta) \right) \right]$$

subject to

$$\begin{aligned}
c + ph' + b' - m' &= (1 + R(b))b - (1 + r + r^m)m + ph + \omega(1 - \tau^\omega) \\
&\quad - \tau^l - \Psi(h, h') - r^p h^r \mathbb{1}[h' = 0] \\
b' &\geq \underline{b} \\
m' &= \min\{\gamma m, \Theta ph'\}
\end{aligned}$$

Employed households that exceed the income requirement can choose whether to repay or remortgage, otherwise they repay their current mortgage:

$$V(\mathbf{s}, \omega, \delta) = \begin{cases} \max\left\{V^R(\mathbf{s}, \omega, \delta), V^P(\mathbf{s}, \omega, \delta)\right\} & \text{if } \omega > \tilde{\omega} \\ V^P(\mathbf{s}, \omega, \delta) & \text{else} \end{cases}$$

Unemployed households

Unemployed households cannot remortgage so only have one value function. The value function for the unemployed households is:

$$U(\mathbf{s}', \epsilon) = \max_{c, b', h' \in \mathcal{H}} u(c, \tilde{h}') + \beta \left((1 - \lambda_0) \hat{U}_u(\mathbf{s}', \epsilon) + \lambda_0 \hat{V}_u(\mathbf{s}', \epsilon) \right)$$

s.t.

$$\begin{aligned}
c + ph' + b' - m' &= (1 + R(b))b - (1 + r + r^m)m + ph + \\
&\quad \kappa - \tau^l - \Psi(h, h') - r^p h^r \mathbb{1}[h' = 0] \\
b' &\geq \underline{b} \\
m' &= \min\{\gamma m, \Theta ph'\}
\end{aligned}$$

1.4.5 Labor market

In the labor market each job is a worker-firm pairing. The worker-firm pairing is characterized by a linear production technology $y = A\mu n$, where A is

aggregate productivity, n is the labor input and μ is worker-firm specific productivity. I assume competition in the labor market means that workers are paid their marginal product of labor. This assumption combined with the assumption that households inelastically supply a unit of labor means that $\omega = \mu$ and the firm production function can be written as $y = A\omega$. From this perspective a household that receives a job offer, offering a new wage and job risk, can be thought of as a household meeting a new firm which provides an alternative productivity match.

Denote the probability distribution of the employed as $\Lambda^E(\mathbf{s}, \omega, \delta)$ and the probability distribution of the unemployed as $\Lambda^U(\mathbf{s}, \epsilon)$. The distribution over assets for all households in the economy is $\Lambda(\mathbf{s}) = \iint \Lambda^E(\mathbf{s}, \omega, \delta) d\omega d\delta + \int \Lambda^U(\mathbf{s}, \epsilon) d\epsilon$. Denote the marginal distribution of the employed over job types as $\hat{\Lambda}^E(\omega, \delta) = \int \Lambda^E(\mathbf{s}, \omega, \delta) d\mathbf{s}$. Output for the economy is $Y = A \iint \omega \hat{\Lambda}^E(\omega, \delta) d\omega d\delta$.

1.4.6 Government

The government funds unemployment insurance, κ , and makes government expenditure, \mathcal{G} , which has no productivity or utility value. The government also issues bonds, B , which provide the asset in which households can purchase liquid assets. The government pays interest rate r on bonds issued and receives any interest rate wedge paid by households, via an unmodelled financial sector. To fund expenditure the government collects labor taxes τ^ω and lump sum taxes τ^l . The government budget constraint satisfies:

$$\begin{aligned} \mathcal{G} + \kappa \iint \sum_{h \in \mathcal{H}} \Lambda^U(\mathbf{s}, \epsilon) d\mathbf{s} d\epsilon + rB &= \underbrace{\tau^l + \tau^\omega \iint \omega \hat{\Lambda}^E(\omega, \delta) d\omega d\delta}_{\text{taxes}} \\ &+ \underbrace{\int (R(b) - r)b + (r^m - r)m\Lambda(\mathbf{s}) d\mathbf{s}}_{\text{interest spread revenue}} + \underbrace{B' - B}_{\text{bond issuance}} \end{aligned}$$

When the economy is hit by shocks, the government allows bonds to adjust to satisfy changes in the demand for liquid assets. Government spending adjusts to satisfy the government budget constraint.

1.4.7 Stationary recursive equilibrium definition

I can now define a stationary recursive equilibrium. An equilibrium is a value function for the employed, $V(\cdot)$, and unemployed, $U(\cdot)$, and policy functions for consumption, $c^i(\cdot)$, liquid assets, $b^i(\cdot)$, housing $h^i(\cdot)$, mortgage choice $m^i(\cdot)$ and job choice $J^i(\cdot)$ for the employed and unemployed $i \in \{E, U\}$. An interest rate schedule, $R(b)$, mortgage price, r^m and house price, p ; aggregate housing stock \bar{H} ; government policies $\{\mathcal{G}, \kappa, \tau^l, \tau^\omega, B\}$; and probability distributions for the employed $\Lambda^E(\mathbf{s}, \omega, \delta)$ and unemployed $\Lambda^U(\mathbf{s}, \epsilon)$ such that:

1. The value functions and policy functions solve household's optimum problem set out in section 1.4.4.
2. The probability distributions $\Lambda^E(\mathbf{s}, \omega, \delta)$ and $\Lambda^U(\mathbf{s}, \epsilon)$ are stationary distributions induced by the policy functions.
3. Markets clear:
 - (a) the housing market clears $\int h \Lambda(\mathbf{s}) d\mathbf{s} = \bar{H}$
 - (b) the liquid asset market clears $\int (b - m) \Lambda(\mathbf{s}) d\mathbf{s} = B$
 - (c) the government budget constraint A.1.1 holds, with $B' = B$

1.4.8 Calibration

Numerical implementation

The model does not have an analytical solution so quantitative methods are used. The household's problem is non-concave due to the housing and refinancing choices and is solved using the Generalized Endogenous Grid method

of Iskhakov et al. (2017).²⁴ This method offers substantial speed improvement relative to value function iteration, allowing for a richer and more accurate specification. I use 700 grid points for liquid assets, 5 grid points for housing, 5 grid points for mortgages, 5 grid points for wages, 7 grid points for job risk and 2 grid points for job draw quality, ϵ . When solving for the ergodic distribution I simulate the distribution rather than simulating a panel of agents, except when annualized variables are required. For the distribution the grid size over liquid assets is reduced to 140 points. The rest of this section discusses the model calibration.

Externally calibrated parameters

A subset of the model's parameters are chosen following commonly used values in the literature or based on external information. A full list of these parameters is provided in Table 1.4. The felicity utility function is Cobb-Douglas of the form shown in equation 1.5. The inverse of the intertemporal elasticity of substitution, α , is set to a standard value of 1.5. The discount factor, β is set to 0.99 reflecting the decision to model a period as one quarter.

$$u(c, \tilde{h}') = \frac{(c^\theta \tilde{h}'^{1-\theta})^{1-\alpha} - 1}{1 - \alpha} \quad (1.5)$$

The unemployment benefit level, κ , is set to $0.4\underline{\omega}$, where $\underline{\omega}$ is the lowest income realization. A number of the financial parameters are also set externally. The maximum loan to value, Θ is set to 0.8 so that the required down payment is 20 percent of the house value. The spread on borrowing in liquid assets, r^b is set to 0.011, to deliver an annual spread of 6.5 percent as in Kaplan and Violante (2014). Also following Kaplan and Violante (2014), the borrowing limit, \underline{b} is set to $0.74E[\omega]$, where $E[\omega]$ is the expected wage

²⁴I thank Giulio Fella for an implementation of this routine that built upon his previous work Fella (2014)

draw. The mortgage income constraint, $\tilde{\omega}$ is set so that the lowest income type cannot remortgage. The mortgage refinancing cost, Ψ^m is set to 0.01, to prevent constant refinancing at the maximum loan to value ratio. The mortgage repayment rate, γ , is set to 0.989 to reflect the standard 30 year mortgage typical in the US. With this repayment rate after 15 years the mortgage value will have halved.

In the housing market, the adjustment cost parameter, Ψ , in the function 1.4 is set to 0.06 a number widely used in the housing and durables literature²⁵. The price of a unit of housing in the stationary equilibrium, p , is set to 1. It is assumed that housing is perfectly elastic in the long run such that the housing stock supply, \bar{H} , adjusts to meet demand. In response to shocks the housing supply is assumed to be fixed with the price adjusting to keep the housing market in equilibrium.

Labor market calibration

Job risk distribution: The functional form of the job offer distribution is designed to capture a number of features of the data: i) heterogeneity in the wage and job separation rate, ii) persistence in the job separation rate, iii) a decline in income following job loss. I assume that the wage and job risk draws are independent. Define the primitives of the job offer distribution as the iid draws from the distributions for wage $F^1(\omega')$ and job risk $G^2(\delta')$.

The unemployed job quality is a two state distribution, $\epsilon = \{\underline{\epsilon}, \bar{\epsilon}\}$, that affects the wage draw. Unemployed in the high state, $\bar{\epsilon}$, draw the wage from the full distribution, $G^1(\omega'|\epsilon = \bar{\epsilon}) = F^1(\omega')$, unemployed in the low state, $\underline{\epsilon}$, draw the lowest wage with probability one, $G^1(\omega'|\epsilon = \underline{\epsilon}) = 1$. The unemployed remain in the high state with probability, ρ^ϵ , while the low state is absorbing. I externally set the persistence of the high state to $\rho^\epsilon = 0.5$. All unemployed

²⁵For example, see: Jose Luengo-Prado (2006), Iacoviello and Pavan (2013a), Bajari et al. (2013), Fella (2014) and Berger and Vavra (2015)

draw from the job risk distribution $G^2(\delta')$.

The employed all draw the wage from the distribution $F^1(\omega')$. The job risk distribution is persistent. In particular I assume it has the following functional form:

$$F_\delta^2(\delta') = \rho^\delta \tilde{G}^2(\delta' | \delta' \geq \delta) + (1 - \rho^\delta) \tilde{G}^2(\delta' | \delta' < \delta) \quad (1.6)$$

where \tilde{G} is a rescaled conditional distribution, such that the probability sums to 1 (e.g. $\tilde{G}^2(x|x \geq \delta) = \int_\delta^x dG^2(x) / \int_\delta^\infty dG^2(x)$).²⁶ The employed that suffer a job separation shock flow into the high state with probability ρ^ϵ . Given this set up the pdfs for the job offer distribution are: $g_\epsilon(\omega', \delta') = g_\epsilon^1(\omega')g^2(\delta')$ for the unemployed and $f_{\omega,\delta}(\omega', \delta') = f^1(\omega')f_\delta^2(\delta')$ for the employed, with $f^1(\omega') = dF^1(\omega')$, $f_\delta^2(\delta') = dF_\delta^2(\delta')$, $g_\epsilon^1(\omega') = dG_\epsilon^1(\omega')$ and $g^2(\delta') = dG^2(\delta')$.

Job risk calibration: The primitive wage offer is assumed to be normally distributed with standard deviation, σ_ω^2 , $F^1(\omega') \sim \mathcal{N}(0, \sigma_\omega^2)$. The standard deviation is chosen to match the standard deviation of annual log wage changes in the PSID, which is estimated as 0.346.²⁷ Job risk is approximated by a seven state concavely spaced grid with more points at lower values of δ . The lower bound $\underline{\delta}$ is set to 0.005, to represent very low job risk and implies a job would be expected to last 50 years. The upper bound $\bar{\delta}$ is set such that the average job separation rate is 0.043, as implied by CPS monthly separation rate of 0.0147. Using the same grid spacing as the model, the empirical distribution of job risk is divided up into the same number of bins as the job risk grid approximation. The discretized draw distributions $G^2(\delta')$ can then be calibrated so that the discretized model distribution $\int \hat{\Lambda}^E(\omega, \delta) d\omega$ matches the empirical distribution. The persistence the parameter ρ^δ is chosen

²⁶Practically, for low risk job this is similar to having a persistence parameter on the current state. However, it avoids generating too strong an incentive for unemployed households to reject higher risk jobs to avoid getting "stuck" in a bad draw.

²⁷In the PSID I regress log income on age, education and time effect, and take the standard deviation of the residual. I also exclude unrealistically large increases or decreases in the wage. For the model I simulate a panel of agents and aggregate to an annual basis.

to match the correlation in job risk at a quarterly horizon, measured in the CPS. The resulting distribution over job risk is shown in Figure 1.7 against the empirical target.

Other labor market parameters: This leaves the job offer arrival rates to calibrate. The job offer arrival rate of the unemployed, λ_0 , is set so that the unemployment to employment transition rate is 0.662, based on the monthly CPS transition rate of 0.311. In practice most unemployed accept any job offer received. The job offer arrival rate for the employed, λ_1 , is chosen to match the rate of job to job transitions. I target a transition rate of 0.076, based on the monthly transition rate of 0.022 in Fallick and Fleischman (2004). For the employed that separate, the probability of a receiving an in quarter job offer, λ_2 , is set to 0.280, based on a Markov transition of monthly CPS transition rates to quarterly rates.²⁸ The probability of being able to reject a job draw, ρ^j , is set to match the correlation between the wage and job security ($\log(1 - \delta)$) in the CPS, which is 0.375.²⁹ This gives an estimate of $\rho^j = 0.95$, implying few forced moves.

Remaining calibrated parameters

The rest of the model parameters are calibrated to hit a set of targeted moments. The values for all the calibrated parameters are shown in Table 1.5 and the targeted moments are presented in Table 1.6. The consumption share in the utility function, θ , is chosen to match the aggregate ratio of non-durable consumption to the housing stock in the US for the period 1970-2012.³⁰ The

²⁸This is the share of households that start the current and next quarter employed but are unemployed in a month between relative to the share of households that begin employed and become unemployed in any period within the quarter.

²⁹This wage refers to the average wage of an industry-occupation-state cell in the CPS, based on the same methodology as the estimation of job risk. However, for obvious reasons, for this measure of job risk the control for income is excluded. The wage measure is average weekly earnings. As such this differs from the measure of Cubas and Silos (2017) who find wage compensation for higher variance of permanent shocks, implying a positive correlation between the wage and risk.

³⁰Residential fixed asset over consumption of non-durable goods and services

interest rate on liquid assets is chosen to target the ratio of median liquid assets plus housing equity to median income as measured in the 2016 SCF. This gives an annual interest rate of 0.3 percent, similar to the rate seen since the Great Recession, but significantly lower than the historic average return. As is familiar in the literature, here there is trade off to be made between targeting median or mean asset holdings. For the current analysis it is more relevant to target the former.

The mortgage rate is chosen to target the average loan to value ratio of 0.47, conditional on owning a house. This gives a mortgage spread of 2 percent on an annual basis. The rental price r^p is chosen such that 34 percent of the households choose to rent. The housing choice is approximated by five uniformly spaced grid points. The lower bound for the housing choice, \underline{h} , is chosen to match the 10th percentile housing size to median income ratio in the SCF. This is the size of the rental house.³¹ The upper bound, \bar{h} , is chosen such that 10 percent of households choose this value.

Finally, Government expenditures are set to match the average ratio of government expenditures to GDP, which generates a target of 0.216. Taxes are chosen to ensure the government budget holds in the following way. The income tax, τ^ω is set equal to government expenditures and unemployment insurance: $\tau^\omega \iint \omega \hat{\Lambda}^E(\omega, \delta) d\omega d\delta = \mathcal{G} + \kappa \iint \Lambda^U(\mathbf{s}) d\mathbf{s}$ while the lump sum tax, τ^l , is set equal to interest payments on bonds minus revenue from the interest rate spread on borrowing: $\tau^l = rB - \int (R(b) - r)b + (r^m - r)m\Lambda(\mathbf{s})d\mathbf{s}$.

Moments of the income distribution

To assess the fit of the resulting income process, the income process in the model is compared against the higher order moments presented in [Güvenen et al. \(2015\)](#). Table 1.7 shows that the stylized income process does a fairly

³¹I experimented with allowing homeowners to own the same size property as the rental property, but in practice few households choose this housing outcome so it was dispensed with for numerical efficiency.

good job of replicating some of the key facts from the income literature. Variance is slightly lower as the target moment comes from the PSID, rather than the administrative data which does not feature top-coding. The model generates additional kurtosis versus a normal distribution particularly for 1 year changes. It also broadly captures the pattern in the distribution of annual income changes by size, although it generates too many small income changes as there are no transitory shocks. The unemployment shocks generate negative skewness in the income process. This is particularly the case at the five year horizon. This is due to the jobs ladder which allows households to reject negative income shocks, leading to an upward skew in income during a period of employment.

1.4.9 Asset distribution

Figure 1.10 presents the baseline economy's asset distribution. Panel a presents the conditional distribution over liquid assets for the employed and unemployed. The economy has a significant dispersion over assets with a mass point at zero as a result of the kink in the interest rate schedule. The distribution is shifted slightly left for the unemployed with these households running down their assets when income is low. Despite this there is a greater mass of employed households at the interest rate kink. The Gini coefficient for liquid assets is 0.62 in the model versus 0.86 in the data. The Gini coefficient for wealth is 0.50 versus 0.77 in the data.³² Panel c shows that for households that hold housing wealth the distribution is fairly equal and not as concentrated as in the data. This is partly due to the upper limit set on housing such that 10 percent of the economy choose this house size. The model does generate a reasonably good approximation of the loan to value distribution,

³²There are a number of known solutions for addressing this discrepancy such as heterogeneity in discount rates (Krusell et al. (1998)), return on assets Benhabib et al. (2017), or a superstar state of income (e.g Castaneda et al. (2003) or Lütticke (2017)). These would be interesting avenues for an extension, but would fundamentally interact with the wage-job risk trade off so are excluded from the baseline model for clarity of exposition.

particularly given the small number of grid points in this dimension (Panel d).³³

1.5 Properties of the steady state equilibrium

This section of the paper studies the key features of the steady state equilibrium. It quantifies the feedback role from assets to job choice in the joint determination of asset and labor market outcomes. It also provides model validation demonstrating the ability of the model to: i) replicate the positive relationship between liquidity and job risk and ii) capture the housing response to an unemployment shock seen in the data. Finally, it highlights the impact of heterogeneity in job risk on household choices, which alters the asset distribution and raises the average MPC.

1.5.1 Policy functions

Asset choice

The household policy functions are presented in Figure 1.8.³⁴ Panel a presents the policy choices for housing. Lower job risk households choose a larger housing stock. In particular, holding all else constant, a household facing higher current job risk requires larger current liquid assets holdings to increase its housing stock relative to a low risk household. The unemployed make a lower housing choice than the most high risk household. There is also substantial inaction due to the adjustment costs. Panel b presents the consumption policy function. As lower risk households have higher expected future income they also have a higher level of consumption.³⁵

³³In the model the maximum loan to value is capped at 0.8 whereas in the data some household exceed this value, these households are allocated to the 0.8 bin

³⁴The policies are presented for a household with 5.6 units of housing, 0.5 liquid assets, 0.47 percent equity and the highest income state.

³⁵As can be seen, the low job risk consumption function includes a non-linear section due to kinks in tomorrow's value function due to changes in the discrete choice.

Panel c shows the liquidity of next period's portfolio choice, where liquidity is next period's liquid assets over next period's total wealth (b/w). Liquidity is a concave function of current liquid assets. In general it can be seen that low job risk households adopt more illiquid portfolios, reflecting the housing choice shown in Panel a.³⁶ Panel d presents the liquidity choice as a function of the current housing stock. Liquidity is decreasing in the current housing stock. This is due to adjustment costs, which mean the choice of housing tomorrow will be higher. As in Panel c, the liquidity choice is lower for households with lower job risk. It is more clearly seen in Panel d that there is substantial non-linearity in this relationship, with a smaller reduction in liquidity when moving from a high (0.204) to mid (0.055) level of job risk than when moving from a mid to low job risk (0.005).

Job choice

In addition to making a choice over assets, households also make a choice over new job opportunities. Figure 1.9 shows the indifference curves for a household deciding whether to stay or switch following a job offer. The red dot indicates the current job pairing (ω, δ), the area above the solid line indicates the alternate job offers (ω', δ') the worker would be willing to move to if drawn.

The figure shows how this choice depends on a household's current assets holdings, where assets are end of period holdings before the job offer arrives. In each panel, the solid black line shows the indifference curve for a low liquid asset household whilst the dotted magenta line shows the indifference curve for a high liquid asset households. Panel a presents the case for the smallest equity holdings - a small housing choice and high Loan to Value (LTV) ratio. Here there is a clear difference in the job offers that will be accepted for high and low liquid asset households. Low liquid asset households have flatter indifference

³⁶The part of the policy where the lowest job risk household has a higher liquidity portfolio than the higher risk household is due to the low risk household extracting equity during this region, reducing illiquid assets and raising the liquidity of its portfolio.

curves, requiring a smaller wage increase to accept a higher risk job. This is due to the desire to move away from the budget constraint, higher income today is more important than a larger expected income that lower job risk delivers.

Panel b shows the case for a household with larger illiquid asset holdings, high housing stock and a low loan to value ratio. Similar choices are evident. The policy functions for the high liquid asset household is almost identical. The low liquid asset household retains the preference for a higher wage over security, although the effect is attenuated. Overall, current asset holdings influences the household wage-job risk trade off.

Counterintuitively, rather than preferring security, lower liquid assets result in a preference for higher job risk over lower current earnings as income is more valuable now than in the future. The cause of this is the slope of the illiquid household's consumption function. The desire to move away from the borrowing constraint increases the value of a higher wage outweighing the precautionary mechanism of a reduced probability of an unemployment spell. The trade off between wage and job risk can also be understood by looking at the derivatives of the value function. Figure C.8 shows that the low liquid asset household values improvements in both job risk and income more than the high liquid asset household. The value of the derivatives are very non-linear in job risk and large at low levels. However, the marginal rate of substitution shows that high liquid asset households value the trade off between security and wage by more than illiquid households (Panel d).

The importance of the feedback from asset to labor market can be analyzed by assessing the impact on the steady state distribution, when non-household state dependent job policy functions are imposed. Table 1.8 compares moments of the model to two cases: i) when the job decision rules are set as if all households were asset rich (Rich), ii) when households accept job offers based

on the present value of the job (PV).³⁷

Changing the job policy function alters the equilibrium. If all households use the Rich job policy functions, following the trade-offs discussed above, households now choose lower job risk resulting in a lower unemployment rate. However, as this also increases job duration the average wage actually increases. With the average employed household in a less risky job, liquid asset demand falls (-3.2 percent) and housing demand rises (+1 percent). The more illiquid allocation raises the share of hand-to-mouth households, which increase by 15 percent.

The largest change in the equilibrium is observed when households accept offers based on the present value of the job opportunity. This would be the outcome in the model with linear utility. This further shows that there is an important feedback from the asset position of households to labor market outcomes via preferences. Again in the baseline model the liquidity motive dominates the precautionary motive. When households choose jobs based on the present value of the income stream, they choose less risky jobs that have a larger long term payoff. This reduces the unemployment rate by 8.6 percent and lowers the standard deviation of annual wages. The result is a poorer economy with lower liquid asset and housing holdings. As a result households hold less liquid assets (-7.5 percent) and the hand-to-mouth share rises (+5.1 percent), though by less than in the Rich economy case.

1.5.2 Liquidity regressions

I now show that the model generates a relationship between liquidity demand and job risk consistent with that observed in the data. I do this by estimating equation 1.3 on model generated data. The results are presented in Table 1.9.³⁸

³⁷Rich means the highest liquid asset choice, the highest housing choice and lowest loan to value. PV is accept whichever job has a greater value of $\omega/(1 - \beta(1 - \delta))$.

³⁸As in the data the ratio is transformed using the inverse hyperbolic sine transformation. Housing equity is the only illiquid asset in the model. To make the magnitudes comparable,

The coefficient on job risk is positive and significant, of a similar magnitude and within the 95 percent confidence intervals of the empirical estimate from the PSID. Column (2) shows that the coefficients on the dummy variables for the lower job risk types are negative as in the data.

Columns (3)-(4) then replicate the regressions of Table 1.2, where the ratio of liquid assets to total wealth is the dependent variable. Again the coefficient is positive and close to that in the data (0.564 vs 0.450). The coefficients of the low job risk dummies are also negative as in column (2). Finally, columns (5)-(6) regress the liquid asset to housing stock ratio on job risk. The relationship is again positive as in the data.

To get a better understanding of the joint determination of assets and labor market outcomes, Figure 1.11 presents the distribution and portfolio allocation across job types. Panel a shows the distribution of households across wage, ω and job risk, δ .³⁹ There is relatively few households in the lowest job risk type and these households tend to have achieved a higher wage draw. This positive correlation between the wage and job security is both a feature of the data and generated by the model's jobs ladder. Panel b shows a measure of the average liquidity for each job type, the ratio of liquid assets to total wealth (b/w). In the job risk dimension liquidity is increasing in job risk for all but the lowest income type, who hold low illiquid asset stocks. The liquidity demand is non-linear with only above average risk households holding similarly liquid portfolios.⁴⁰

Finally, I also repeat the regressions of hand-to-mouth type on job risk. The model reproduces the negative coefficient on job risk for the probability of being a wealthy hand-to-mouth type also observed in the data albeit the elasticity in the model is lower than in the data (-0.388 vs -0.849).

average job risk is scaled to match the level in the PSID

³⁹There is a mass of households at the lowest wage as unemployed in the low job quality draw state, $\underline{\epsilon}$, initially draw the lowest wage type.

⁴⁰Figure C.9 presents complementary information on the MPCs and, tenure and expected duration of each job type

1.5.3 Unemployment responses

The new empirical consensus is that job separation shocks lead to large and persistent income losses (see for example: [Stevens \(1997\)](#), [Davis and von Wachter \(2011\)](#), [Jarosch \(2015\)](#), [Krolikowski \(2017\)](#), [Huckfeldt \(2018\)](#)). I evaluate the ability of the model’s job ladder to replicate these labor market outcomes and at the same time generate the correct housing responses.

I follow the methodology of [Stevens \(1997\)](#) and [Huckfeldt \(2018\)](#), using the PSID. Full details are available in the Appendix, Section A.1.1. The basic specification is:

$$Y_{i,t} = \mathbf{X}_{i,t}\beta + \sum_{j=-2}^{10} d_{i,t}^j \delta^j + \alpha_i + \gamma_t + \epsilon_{i,t} \quad (1.7)$$

where $Y_{i,t}$ is the labor market or housing outcome of interest, d^j is a set of dummy variable for j periods since suffering a job separation shock, $\mathbf{X}_{i,t}$ is a set of controls, including a quartic in age, education dummies and family demographics, and γ_t are year fixed effects. To control for unobservable worker characteristics an individual fixed effect, α_i , is included accounting for any systematic differences in the workers likely to lose their jobs, such as lower wages or smaller housing stocks.⁴¹

Figure 1.12 compares the results of the model to the data, the data is shown with 64 percent confidence intervals. Panel a shows that the model does a good job of capturing the income loss profile. Panel b presents the unemployment response. This is a dummy variable for being unemployed (broader than experiencing a job separation) during the past year. Again the data and model line up fairly closely, indicating a raised probability of a further spell in unemployment in the years following job loss which the model’s jobs ladder captures.

⁴¹The definition of a separation shock differs from unemployment used elsewhere in this paper and only includes company closure, layoffs or firing.

Figure 1.13 presents the main results from the housing market choices.⁴² The first observation is that the housing responses are less precisely estimated in the data than the labor market variables. Panel a shows the response of log housing, capturing the intensive margin. The model does a good job of capturing the average reduction in the housing stock size for those that remain homeowners. The average decline in the data is around 1-4 percent. The model generates a decline of 1-3 percent. The mortgage income constraint is important for this results, in its absence almost all households that remain homeowners remortgage rather than reduce their housing stock. Panel d shows the response of the housing level, capturing both the intensive and extensive margin. While the model captures the prolonged decline in the housing stock it exaggerates the magnitude. In the data the average decline is around \$4,000 whilst in the model it is around \$8,900.⁴³ For the mortgage, the model successfully replicates the data on both the intensive and extensive margin (See Figure C.11).

Finally, I consider the magnitude of the housing response conditioning on the size of the initial income shock.⁴⁴ The results are reported in Figure 1.14. The model does a good job of replicating the difference between those that suffered large or small income shocks. For the value of housing, the model captures the significantly larger fall of large wage shock households.

1.5.4 Role of job risk heterogeneity

A novel contribution of the model in this paper is the introduction of heterogeneity in job risk. The effect of the job heterogeneity on the asset distribution can be observed by comparing the model to the alternative with a single job risk level, which is referred to as the single δ economy. Table 1.11 compares

⁴²Chetty and Szeidl (2007) undertake a similar empirical analysis, but focus on housing services and the responses of those that remain homeowners or remain renters.

⁴³The model predicts too many households moving into renting (see Figure C.10)

⁴⁴I split the sample by those that income response in period 0 was more or less than the average response, including interaction dummies for the periods $j = 0, \dots, 10$

the baseline model to a recalibrated model with one job risk type. The average household in the baseline model holds more housing and more liquid assets.⁴⁵ This is caused by two mechanisms. Firstly, average wages are lower in the single δ economy as there is a lower probability of a long tenure, reducing the opportunity for a high wage to arrive. Secondly, there is higher demand for liquid assets at above average levels of job risk. I also calculate the Marginal Welfare Value of Insurance (MWG). Following Chetty and Szeidl (2007), this is given by:

$$MWG = \int \frac{\delta}{1 - \delta} \left(\frac{\mathbf{E}U_b(\mathbf{s}, \epsilon) - \mathbf{E}V_b(\mathbf{s}, \omega, \delta)}{\mathbf{E}V_b(\mathbf{s}, \omega, \delta)} \right) \Lambda(\mathbf{s}, \omega, \delta) d\mathbf{s} d\omega d\delta$$

Despite the lower precautionary savings motive, the MWG is higher in the heterogeneous job risk economy (+5.5 percent), reflecting i) the larger average income loss experienced upon unemployment and ii) the variation in MWG by job risk in the baseline economy. As shown in Figure 1.15 the MWG is steeply increasing in job risk due to the large probability of job loss for high risk households.

Impact on MPCs

Despite the greater demand for liquid assets the baseline model also features a larger share of hand-to-mouth households.⁴⁶ In the baseline model 9.7 percent of households are hand-to-mouth, whilst in the model without job risk heterogeneity this falls to 6.6 percent. As a result the aggregate MPC is also higher in the baseline model. The MPC in the baseline model is 0.096, whilst

⁴⁵The models are calibrated to target a median wealth measure.

⁴⁶Following the literature, hand-to-mouth is defined as a household having either negative liquid assets and being less than two weeks of their current wage from the borrowing constraint or having weakly positive liquid assets and having less than two weeks of their current wage of liquid assets. In the model poor hand-to-mouth is defined as a household that is renting. Wealthy hand-to-mouth households are homeowners and have positive illiquid asset holdings.

in the model without job risk heterogeneity it is 0.080.⁴⁷

The impact on MPCs can also be assessed by subjecting the economy to a lump sum tax shock. In the experiment all households receive an increase in the lump sum tax, worth 1 percent of average income. The tax then follows an autoregressive process with persistence 0.5. The present value of the shock is worth 2.4 percent of the average post-tax income.⁴⁸ Figure 1.16 presents the Impulse Response Functions for the baseline and single δ model. On impact consumption in the baseline model falls by 0.23 percent while it falls by 0.18 percent in the single δ case, representing a 22 percent larger response from the model with heterogeneous job risk (Panel a). Panel a of Figure 1.17 presents the period one response by job risk. The largest response is for the lowest risk group, only this set of jobs has a large MPC. After this the response declines and is broadly increasing in job risk.

1.6 Great Recession experiment

I now subject the economy to labor market shocks meant to replicate those that occurred during the Great Recession and consider how the initial conditions affect the size of the response. This has two purposes. Firstly, by seeing how different initial equilibria respond to equivalent shocks the importance of the joint determination of labor and asset market outcomes and role job risk heterogeneity in generating state dependent responses can be assessed. Secondly, I examine the extent to which the cross sectional variation in the responses, by wage and and job risk, matches the data. This provides information on the importance of the interaction between the pre Great Recession asset allocation and specific labor market shocks for the consumption and as-

⁴⁷MPCs in the model are at the lower bound of those estimated empirically (for example 0.12-0.30 on a quarterly basis in [Parker et al. \(2013\)](#) and 0.35-0.70 on an annual basis in [Fagereng et al. \(2018\)](#))

⁴⁸As in the single δ case the economy is poorer, the absolute size of the shock is smaller in the single δ economy

set choices during this period. Matching this cross sectional heterogeneity is strong test of the model's key mechanism.

The shocks are estimated from CPS data for four groups stratified by wage and risk: low wage, low risk; low wage, high risk; high wage, low risk; high wage, high risk. Further details on how the shocks were estimated can be found in the Appendix, Section A.1.2. The stylized fact is that the Great Recession was characterized by larger job separation shocks for the high wage types, with the low wage, low risk type experiencing a much smaller increase in job risk. The two groups that saw the largest wage falls were the low wage, high risk and high wage, low risk. The cross sectional paths of the average job separation rate and income level for each group is shown in Figure 1.18, alongside the data. The actual shocks and persistence parameters are in Table B.5.⁴⁹

1.6.1 Aggregate responses

The Great Recession is modeled as a set of one time shocks to the job separation rate and wage level. Once the shocks have realized, the households have full knowledge of their deterministic path. The economy returns to the stationary equilibrium after the effect of the shocks dies out. The aggregate housing stock is assumed to be fixed, with equilibrium in the housing market achieved through a change in the house price. Government bonds are allowed to adjust to accommodate any changes in liquid assets or mortgages to satisfy demand.⁵⁰

To assess the importance of the conditions prior to the Great Recession I undertake the following experiment. I first find the response of the economy

⁴⁹Due to the model's jobs ladder a substantial part of the fall in income is due to the rise in the separation rate which causes agents to suffer an income loss as they are required to climb the ladder again, rather than the wage falls conditional upon remaining in the same job.

⁵⁰Effectively, I do not impose market clearing in the liquid asset market. Given that the interest rate was at the zero lower bound for most of the duration of the Great Recession, this seems a reasonable assumption.

in a high average separation rate equilibrium or “*high ave. δ* ”, meant to represent the labor market of the 1980s and 1990s. Then to capture the period of low job risk prior to the Great Recession, holding all other parameters fixed I lower the separation rate and find the “pre-Great Recession” stationary equilibrium, this approach is in keeping with secular decline in the separation rate seen in Figure 1.2.⁵¹ As the average separation rate in the economy is endogenous, I reduce the expected job separation rate draw of an unemployed household ($\int \delta dG^2(\delta)$) between the two equilibria by 25 percent, matching the decline in the data of the average separation rate in 2007, relative to the average separation rate between 1980 and 1999.⁵² Table B.6 in the Appendix compares the asset distributions in these two steady states, the key feature is that households hold larger housing stocks and less liquid assets in the pre-Great Recession economy. Further, the average wage is higher as households have longer to climb the jobs ladder and are willing to trade off some of the lower job risk gains for higher wages.

I undertake two experiments to identify different channels of the response. In the first experiment households expect to remain in the pre-Great Recession equilibrium after the shocks die out (labeled *Pre-GR, fixed expectation*) in the second experiment households expect to revert to the high average separation rate equilibrium in the long run (labeled *Pre-GR, change expectation*). Full details of how an “equivalent shock” is defined in these economies is provided in the Appendix, section A.1.3. Figure 1.19 presents the aggregate response to the combined labor market shocks. Panel a shows that under the high average separation rate equilibrium consumption undergoes a 2.9 percent contraction.

⁵¹While the Great Moderation was associated with a decline in aggregate volatility, it is not clear from the data that this translated into a reduction in household uncertainty (see [Davis and Kahn \(2008\)](#)). There is a large body of evidence documenting a rise in the volatility of household earnings since the 1970s, although typically these studies find that it has been more stable since the late 1980s (see: [Heathcote et al. \(2014\)](#), [Blundell et al. \(2008\)](#) and [Gottschalk and Moffitt \(2009\)](#))

⁵²The average quarterly separation rate between 1980 and 1999 was 4.8 percent. The average quarterly separation rate in 2007 was 3.6 percent, a 25.4 percent decline

On impact there is an increase in liquid assets in response to the increase in risk, but during the course of the recession households run down liquid assets to smooth consumption (Panel b). The model generates a decline in the house price. Equity falls as households remortgage to smooth consumption and because of the decline in the value of the housing stock (Panel d).

When the shocks instead hit the economy in pre-Great Recession conditions the effect is amplified substantially. The red line shows the response in the first experiment in which household expect the long run separation rate to return to the pre-Great Recession equilibria. The decline in consumption is now 26 percent larger, declining by 3.6 percent relative to 2.9 percent in the high average separation rate economy. The household equity response is also amplified. This has both a quantity and price aspect, with households taking out more equity to smooth consumption and with housing demand declining by more leading to a larger fall in the house price and decline in the value of household equity holdings. What accounts for this amplification? A key determinant is the change in the wealth and labor market distribution in the pre-Great Recession economy. Households in the pre-Great Recession equilibria hold less liquid assets, more housing and a greater share are hand-to-mouth. Due to greater labor market attachment they also have slightly higher incomes. The amplification can be broken down into three channels. Firstly, with less liquid assets and a greater share of hand-to-mouth agents, consumption is more sensitive to income shocks. Secondly, in the pre-Great Recession economy households are more dependent on selling housing to smooth consumption when unemployed, placing greater downward pressure on the housing market and generating a larger negative wealth shock which hits all homeowners. Thirdly, as households are more attached to the labor market and have been able to climb the jobs ladder further, unemployment results in a bigger average decline in household income resulting in a larger consumption response.

A similar amplification result is observed under the alternative assump-

tion on household expectations in the second experiment. The magenta lines shows the response when households expect the separation rate to return to the higher value following the Great Recession (*Pre-GR, change expectation*). In this experiment households want to return their liquid asset and housing allocations back to the high average separation rate economy. As a result there is a large increase in liquid assets. The desire to reduce the housing stock results in a 5.5 percent decline in the house price.⁵³ There is also a large reduction in household equity, falling by 11.6 percent on impact. As in the first experiment this has both a quantity and price decline element to the fall.⁵⁴ A full description of the amplification can be seen in Table 1.12 The two experiments show that there is significant state dependency in the model, with the pre-conditions mattering a great deal to the aggregate response.

Persistence

The responses in the full model can also be compared to the results when heterogeneity in job risk is switched off, to assess the role of job risk heterogeneity in amplification. The amplification of consumption upon impact is fairly similar (See Table 1.12) . If anything, the model with a degenerate job risk distribution overstates the amplification. This is because in the baseline economy households trade of some of the gains from the the secular decline resulting in a compressed distribution of job risk (as seen in the data).

While the immediate impact is similar in both economies the baseline economy generates additional endogenous persistence of the consumption response.

Figure 1.20 compares the consumption responses of the baseline and degen-

⁵³This is still below the 17.5 percent fall observed in the data between December 2007 and December 2009, measured using the Case Shiller national Index or 35 percent, between March 2006 and February 2012. [Kaplan et al. \(2017\)](#) have stressed the role of a change house price growth expectations in delivering the house price decline seen during the Great Recession

⁵⁴Consumption does not fall by much more than the fixed expectations experiment. This is because in period 0 before the shock, households are already consuming on the basis of the policy functions in the high average separation rate economy. If consumption in period 0 was held at the pre-GR equilibrium level, the decline would be -10.5 percent.

erate job risk economies in the pre-Great Recession environment. During the course of the shock the cumulative decline of consumption is 18 percent larger in the baseline economy. This is due to the jobs ladder. In the baseline households that fall off the jobs ladder, have a higher expected future job separation rate than the average employed household. In the degenerate economy these two separation rates coincide. This means that the increase in the separation rate generate a larger decline in the present expected value of future income in the baseline economy. As seen in Section 1.5.3 and in keeping with the data, some share of households will ultimately face additional unemployment spells in the future and be forced to lower consumption. This propagation mechanism is absent in the model with degenerate job risk.

Increase in income uncertainty

While the period between the 1980s and 2007 was associated with a decline in the job separation rate, it could be a concern that other changes in the income process might overturn the portfolio reallocation mechanism of households, for example the rise in the volatility of income that has been documented in the literature. I do not attempt to provide a full description of the change in the income process faced by households but show that a reasonable increase in the income risk represented by the wage offer distribution does not overturn the amplification previously highlighted. I implement this by concurrently increasing the standard deviation of the wage offer distribution, σ_w by 25 percent while reducing the job separation rate.

Figure 1.21 presents the results of this experiment. It can be seen that rather than diminishing the response, the increase in income risk faced by households actually provides additional amplification. This is because in a jobs ladder setting greater income risk has two effects. Firstly, it raises uncertainty causing households to want to hold additional liquid assets, this reduces the amplification. Secondly, as households are able to reject bad job draws it also

increases the opportunity to access better paid jobs. This raises the average wage in the economy leading households to choose greater housing holdings. The larger housing stock amplifies the decline in house prices when the Great Recession shocks hit, as seen in the second panel, as well as increasing the cost of unemployment.⁵⁵

1.6.2 Cross sectional responses

The cross sectional behavior by job type in response to the shocks faced can also be studied. This provides a test of the extent to which the joint decisions in the asset and labor markets matter, in particular it is an indicator of the importance of the interaction between the pre crisis asset allocation by job type and labor market shocks experienced during the Great Recession.

Figure 1.22 compares the relative group level model responses to panel data responses from the PSID. Details on the data definitions can be found in Section A.1.2.⁵⁶ On the left sided panels, model generated responses are presented while on the right sided panels the data from the PSID is shown. For the groups: red lines indicate low wage and blue lines indicate high wage groups. Solid lines are for low job risk, dashed lines are for high job risk. The responses are shown relative to the high wage, high job risk group that experienced the largest housing decline.

Figure 1.22, Panel a presents the housing choice. The model replicates the order of the responses across the four groups and the relative responses are of a similar magnitude to the data. The high wage, high risk group sees the largest decline in housing, followed by the high wage, low risk group as both groups experience a large increase in the job separation rate. In comparison, for the low wage, low risk type that experience the smallest increase in job risk, hous-

⁵⁵The relationship between the increase in income risk and amplification is very non-linear and alternative parametrisations can result in different results.

⁵⁶The level for the consumption and housing response are shown in Appendix C.1, Figure C.12.

ing has become relatively cheap and thus for this group the decline in housing is effectively just the price decline, while increasing their unit holdings. The high wage groups have the largest housing stock in equilibrium and experience larger job separation rate shocks. Their response is to significantly reduce their housing demand. While the total decline does not match the scale in the data, this is mainly due to the model not generating a large enough house price decline. All else being equal, a larger house price decline would scale up these response while retaining the correct ordering produced by the model.⁵⁷

Turning to consumption, Panel b shows the model does a good job of predicting the consumption responses of the high wage types that saw a large increase in the job separation rate. For households that remain employed this raises their desire for precautionary savings and for the additional workers that lose their job a large reduction in consumption is required due to the persistent income loss. An aspect of the consumption response that fits less well, is the dynamics of the low wage types, reversing the order of the low wage, low risk and low wage, high risk groups.

In Appendix C.1 Figure C.13, Panel a presents the response of liquid assets. The broad pattern is captured. Both the model and data imply a large precautionary increase in liquid assets for the low wage, low risk type and substantial declines for both high risk groups after the initial shock. In the model relative to the data all groups engage in precautionary savings.⁵⁸ Panel b presents the equity choice. Similarly to the consumption response, the model does well on the high wage groups that saw a large increase in the job separation rate, with the high wage, high risk group reducing equity by a larger amount than the high wage, low risk. Although the model over predicts the initial fall in the low wage, high risk group, it captures the faster rebound seen for this group

⁵⁷There would be additional effects due to the size of the wealth shocks and degree to which groups varied in the probability of ending up underwater following a larger house price decline

⁵⁸Given the biannual nature of the data the initial precautionary effect could be missing

in the latter part of data.

Overall, the model does a good job of matching the cross sectional responses. Effectively, this is a test of the joint covariance of the prior distribution over assets and job types and the shocks hitting the model. The fact that the model is successful in matching these dynamics is interpreted as evidence that the mechanism proposed, namely the interaction of sorting in the labor market and asset allocation decisions in the face of lower risk, is important for understanding the severity of the Great Recession.

1.7 Conclusion

This paper has investigated the importance of the joint determination of asset and labor market choices and the role of job risk heterogeneity in the emerging liquid and illiquid asset incomplete markets macroeconomic models. The model is able to replicate the relationship between job risk and portfolio liquidity and the housing responses to unemployment shocks seen in the data. Job risk heterogeneity raises the aggregate MPC and thus the aggregate response to transitory income shocks. However, the increase is fairly small. The main reason for this is that liquidity demand is highly non-linear in job risk and the data only calls for a fairly small fraction of very low risk households.

The importance of the joint determination of assets and labor market outcomes is evidenced by the response to the Great Recession labor market shocks. When the shocks hit the economy in pre-Great Recession conditions, characterized by a low job separation rate, the negative response of consumption is increased by 40 percent. This is due to the asset choices of households in equilibrium, choosing a larger housing stock and more illiquid portfolio and by households sorting into higher wage jobs. The addition of job risk heterogeneity also accounts for a rise in the amplification of the housing market responses.

The model replicates cross sectional features of the Great Recession. When feeding in the distribution of labor market shocks that occurred during the Great Recession, the model captures the ordering of housing choices across job types. This provides a strong test of the covariance of the pre-recession asset allocations and labor market shocks experienced in the Great Recession and suggests that this is an important margin for understanding the crisis. A remaining puzzle is the relatively weak response of the low wage, low risk group that experienced a small job separation shock but exhibited a strong consumption response in the data.

A limitation of this paper is that households are only able to allocate between job types by waiting for an improved offer to arrive following the stochastic arrival rate. It would be interesting to consider how more active sorting might affect the results, such as by varying search intensity or in a directed search setting. It is anticipated that these features would strengthen the interaction. This is beyond the scope of this paper and left for future research.

1.8 Tables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Job risk (δ)	0.928*** (0.168)		0.875*** (0.225)	0.890*** (0.211)	-0.052 (0.238)	0.907*** (0.418)	0.830*** (0.196)	0.536*** (0.222)
$\delta : \leq 1^{st}pc.$		-0.147*** (0.045)						
$\delta : \leq 5^{th}pc.$		-0.098*** (0.025)						
$\delta : \leq 10^{th}pc.$		-0.056** (0.023)						
$\delta : \leq 25^{th}pc.$		-0.013 (0.014)						
Income	0.195*** (0.011)	0.200*** (0.011)	0.235*** (0.014)	0.146*** (0.016)	0.200*** (0.012)	0.055*** (0.012)	0.200*** (0.012)	0.210*** (0.013)
Core			✓					
Risk pref.				✓				
Ind & Occ					✓			
FE						✓		
IV							δ_{t-2}	δ_{t-6}
N	33,932	34,629	23,717	22,212	33,932	34,814	29,342	22,429
R ²	0.046	0.046	0.048	0.090	0.052	0.010	0.050	0.052

Notes: Dependent variable is the transformation of the ratio of liquid assets to illiquid assets. δ Percentiles are dummy variables with $\delta : \leq 1^{st}pc.$ referring to the 1st percentile, $\delta : \leq 5^{th}pc.$ referring to 5th percentile < $\delta \leq 5^{th}$ percentile etc. Core refers to use of only core PSID sample, Risk pref. includes controls for stock market exposure, insurance purchases and total assets. Ind & Occ includes industry and occupation dummy in main specification. FE is panel specification with individual fixed effects. IV uses lagged job risk. *Statistically significant at 10%; **statistically significant at 5%; ***statistically significant at 1%

Table 1.1: Ratio of liquid assets to illiquid assets

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Job risk (δ)	0.450*** (0.120)		0.543*** (0.161)	0.352** (0.154)	1.367*** (0.169)	0.162 (0.276)	0.374** (0.150)	0.718*** (0.185)
$\delta : \leq 1^{st} pc.$		-0.014 (0.032)						
$\delta : \leq 5^{th} pc.$		-0.040** (0.019)						
$\delta : \leq 10^{th} pc.$		0.003 (0.015)						
$\delta : \leq 25^{th} pc.$		-0.007 (0.009)						
Income	-0.084*** (0.006)	-0.082*** (0.006)	-0.082*** (0.008)	-0.123*** (0.009)	-0.086*** (0.007)	-0.046*** (0.007)	-0.079*** (0.008)	-0.068*** (0.009)
Core			✓					
Risk pref.				✓				
Ind & Occ					✓			
FE						✓		
IV							δ_{t-2}	δ_{t-6}
N	46,664	47,680	30,439	29,872	46,664	47,885	38,207	27,526
R ²	0.052	0.050	0.043	0.055	0.056	0.048	0.043	0.035

Notes: Dependent variable is the transformation of the ratio of liquid assets to total assets. δ Percentiles are dummy variables with $\delta : \leq 1^{st} pc.$ referring to the 1st percentile, $\delta : \leq 5^{th} pc.$ referring to 5th percentile < $\delta \leq 5^{th}$ percentile etc. Core refers to use of only core PSID sample, Risk pref. includes controls for stock market exposure, insurance purchases and total assets. Ind & Occ includes industry and occupation dummy in main specification. FE is panel specification with individual fixed effects. IV uses lagged job risk. *Statistically significant at 10%; **statistically significant at 5%; ***statistically significant at 1%

Table 1.2: Ratio of liquid assets to total assets

	PHTM			WHTM		
	(1)	(2)	(3)	(4)	(5)	(6)
Job risk (δ)	0.522*** (0.056)	1.223*** (0.088)	0.659*** (0.070)	-0.849*** (0.067)	-1.263*** (0.100)	-0.995*** (0.088)
Ind & Occ		✓			✓	
IV			δ_{t-2}			δ_{t-2}
N	51,469	51,469	41,562	51,469	51,469	41,562
R^2	0.233	0.240	0.223	0.019	0.024	0.016

Notes: PHTM: poor hand-to-mouth. WHTM is wealthy hand-to-mouth.

*Statistically significant at 10%; **statistically significant at 5%;

***statistically significant at 1%

Table 1.3: Hand-to-Mouth status

α	β	Θ	γ	$\tilde{\omega}$	r^k (an.)	\underline{h}	p	Ψ	Ψ^m	ρ^ϵ
1.5	0.99	0.8	0.989	0.5	0.065	3.7	1.0	0.06	0.01	0.5

Table 1.4: Externally calibrated parameters

Labor market:							
$\bar{\delta}$	λ_0	λ_1	σ_ω	ρ^δ	ρ^y		
0.204	0.662	0.185	0.245	0.830	0.950		
Other:							
θ	r (an.)	r^m (an.)	r^p (an.)	\bar{h}	\mathcal{G}	τ^l	τ^ω
0.947	0.002	0.013	0.012	14.2	0.26	0.001	0.24

Table 1.5: Calibrated parameters

Moment	Parameter	Data	Model
<i>Labor market:</i>			
Average δ	$\bar{\delta}$	0.043	0.042
UE transition rate	λ_0	0.662	0.661
EE transition rate	λ_1	0.076	0.076
$\sigma(\Delta \log(w_{an}))$	σ_ω	0.346	0.346
$Corr(\delta, \delta')$	ρ^δ	0.912	0.916
<i>Other:</i>			
C/H	θ	0.137	0.133
med $b + eq$: med income	r	3.45	3.52
LTV	r^m	0.472	0.465
Share of renters	r^p	0.344	0.345
Share holding \bar{h}	\bar{h}	0.100	0.111
\mathcal{G}/Y	\mathcal{G}	0.216	0.225

Table 1.6: Targeted moments

Moment	Data	Model
Variance: 1yr change	0.26	0.16
Skewness: 1yr change	-1.07	-0.69
Skewness: 5yr change	-1.25	-0.14
Kurtosis: 1yr change	14.93	10.53
Kurtosis: 5yr change	9.51	5.13
Frac. 1yr change < 10%	0.49	0.62
Frac. 1yr change < 20%	0.67	0.68
Frac. 1yr change < 50%	0.83	0.86

Table 1.7: Earnings process

Moment	Baseline	% Δ	
		Rich	PV
Unemployment rate	0.044	-3.7	-8.6
E-E transition rate	0.076	+0.8	+2.7
E-U transition rate	0.030	-4.0	-8.7
Av. δ	0.042	-4.0	-8.7
Av. wage	1.206	+0.1	-0.8
$\sigma(\Delta \log(wage))$	0.346	-1.6	-4.3
$\rho(\log(\omega), \log(1 - \delta))$	0.369	-1.0	+9.1
b	1.642	-3.2	-7.5
h	5.201	+1.0	-1.7
eq	2.839	+0.2	-4.5
$\rho(b, h)$	0.291	-2.8	5.3
% renters	0.345	-1.1	+4.0
HTM	0.097	+15.0	+5.1

Table 1.8: Impact of job choice on equilibrium

Notes: Rich is model ergodic distribution using highest liquid asset, housing choice and lowest loan to value job choice policy function. PV is model ergodic distribution when all jobs of a higher present value are accepted.

	b/eq		b/w		b/h	
	(1)	(2)	(3)	(4)	(5)	(6)
Job risk (δ)	1.239*** (0.002)		0.564*** (0.001)		0.749*** (0.000)	
$\delta_1(6\%) :$		-0.535*** (0.001)		-0.367*** (0.000)		-0.317*** (0.000)
$\delta_2(27\%) :$		-0.368*** (0.000)		-0.242*** (0.000)		-0.242*** (0.000)
$\delta_3(67\%) :$		0.019*** (0.000)		-0.021*** (0.000)		-0.021*** (0.000)
Income	0.392*** (0.001)	0.555*** (0.001)	-0.535*** (0.000)	-0.421*** (0.000)	0.452*** (0.000)	0.549*** (0.000)
R^2	0.043	0.175	0.142	0.208	0.142	0.296
PSID	0.928*** (0.168)		0.450*** (0.120)		0.084 (0.076)	

Notes: Dependent variable is the transformation of the stated ratio.

*Statistically significant at 10%; **statistically significant at 5%;

***statistically significant at 1%

Table 1.9: Model portfolio regressions

	PHTM	WHTM
Job risk (δ)	-0.131*** (0.001)	-0.388*** (0.001)
R^2	0.038	0.030
PSID	0.522*** (0.056)	-0.849*** (0.067)

Notes: *Statistically significant at 10%; **statistically significant at 5%;
***statistically significant at 1%

Table 1.10: Model Hand-to-Mouth status

Moment	Baseline	Single δ	
	level	level	% Δ
Av. wage	1.206	1.112	-7.8
c	0.862	0.790	-8.3
b	1.642	1.559	-5.1
h	5.201	4.735	-9.0
$\rho(b, h)$	0.291	0.421	+44.9
PHTM	0.029	0.031	+4.5
WHTM	0.068	0.035	-48.5
HTM	0.097	0.066	-32.5
% PHTM	0.303	0.468	+54.7
MPC (+ ve.)	0.097	0.080	-16.1
MPC (- ve.)	0.085	0.080	-15.1
MPC (an.)	0.228	0.205	-10.1
MWG	0.036	0.034	-5.5

Notes: HTM is hand-to-mouth share; PHTM is poor hand-to-mouth;
WHTM is wealthy hand-to-mouth; MPC is marginal propensity to consume.
MWG is marginal welfare gain of insurance

Table 1.11: Role of job risk heterogeneity


		Heterogenous δ			Single δ		
		experiment			experiment		
		Hi $E(\delta)$	1	2	Hi $E(\delta)$	1	2
C	<i>% decline</i>	-2.9	-3.6	-3.9	-2.6	-3.5	-3.8
	<i>% amplification</i>		25.9	35.8		36.8	45.0
EQ	<i>% decline</i>	-5.5	-7.4	-11.6	-4.6	-9.0	-15.1
	<i>% amplification</i>		33.9	111.5		96.8	229.9
P	<i>% decline</i>	-2.2	-3.0	-5.5	-1.3	-3.4	-5.9
	<i>% amplification</i>		35.1	148.9		171.6	370.0

Notes: Heterogenous δ is the baseline model. Single δ is the model with no heterogeneity in job risk. *% decline* is the percentage decline in period one. *% amplification* is the additional decline in the experiment relative to the baseline.

Hi $E(\delta)$ is the baseline calibration with a higher expected separation rate. *Experiment 1* is the shock to the economy in pre-Great Recession conditions with no change in the the long run separation rate. *Experiment 2* is the shock to the economy in pre-Great Recession conditions with an increase in the expectation of the long run separation rate. See section A.1.3 for full details.

Table 1.12: Period one Great Recession response

1.9 Figures



Data: CPS (1987-2016). For details of estimation see section 1.3.1

Figure 1.1: Distribution of US job risk






Figure 1.2: Job separation rate




(a) Decomposition




(b) Rescaled

Figure 1.3: Job separation rate: compositional effects



(a) By job risk



(b) By income

Figure 1.4: Job separation rate in Great Recession

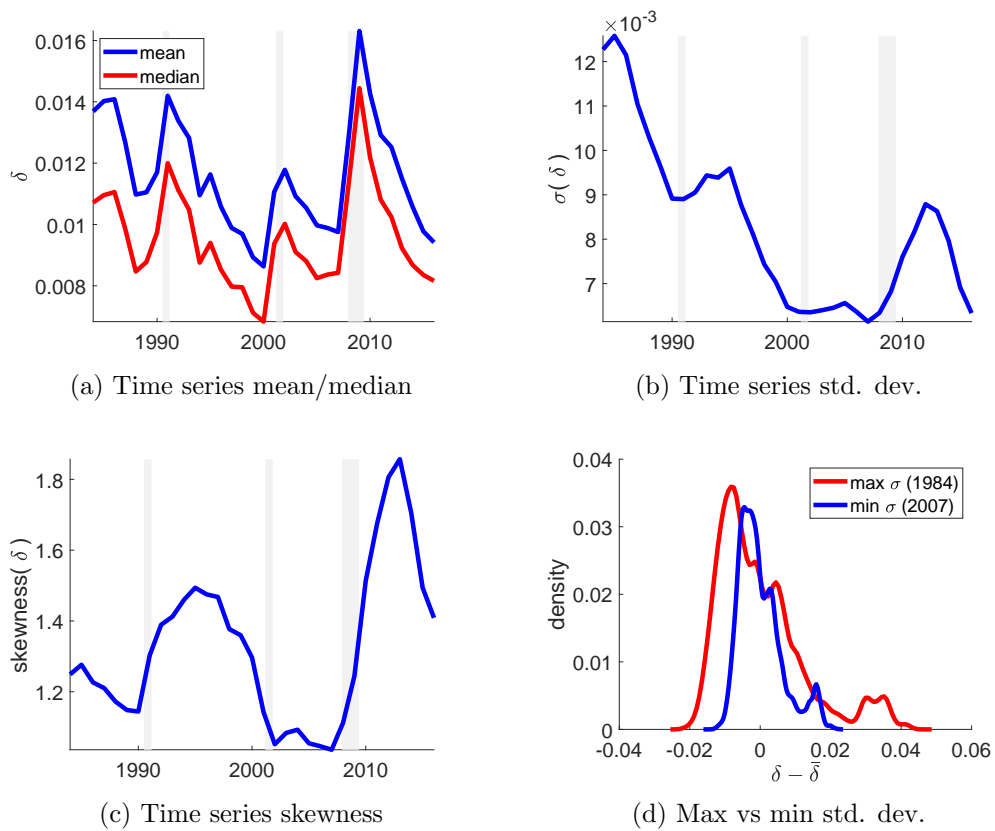


Figure 1.5: Time series variation in job risk distribution




Figure 1.6: Portfolio choices during decline in job risk




Figure 1.7: Calibrated job distribution

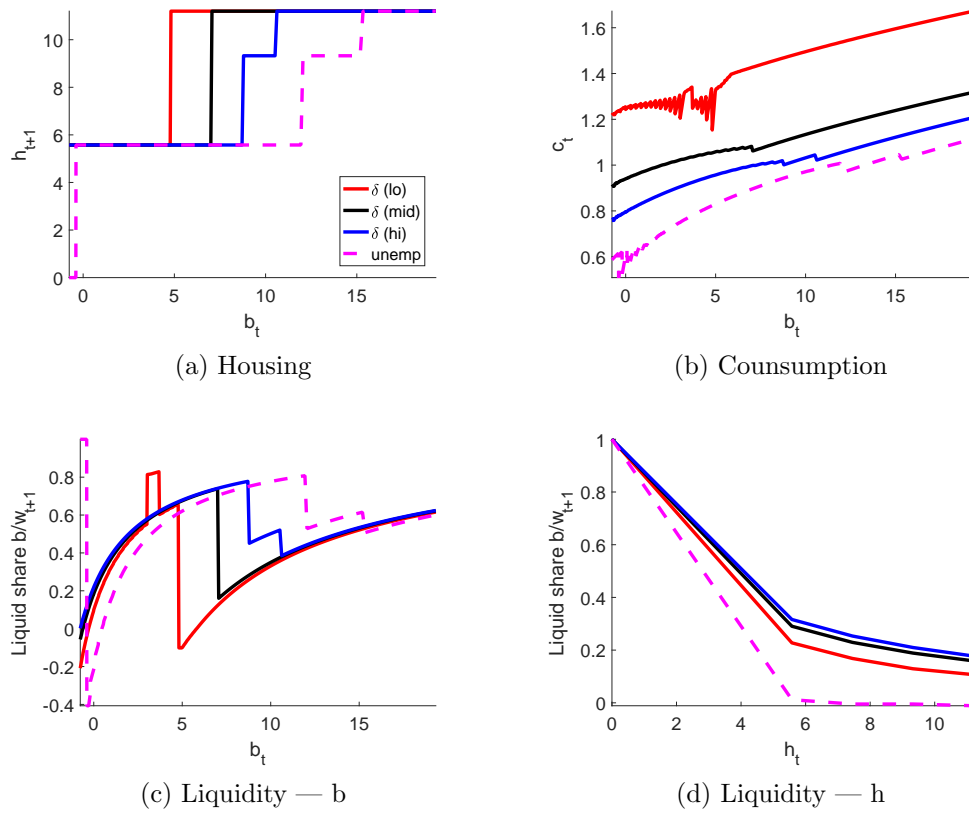


Figure 1.8: Policy functions

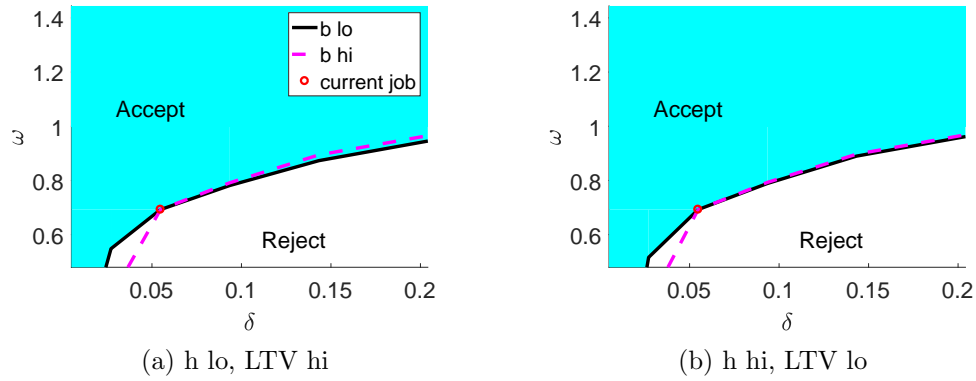


Figure 1.9: Job offer indifference curve

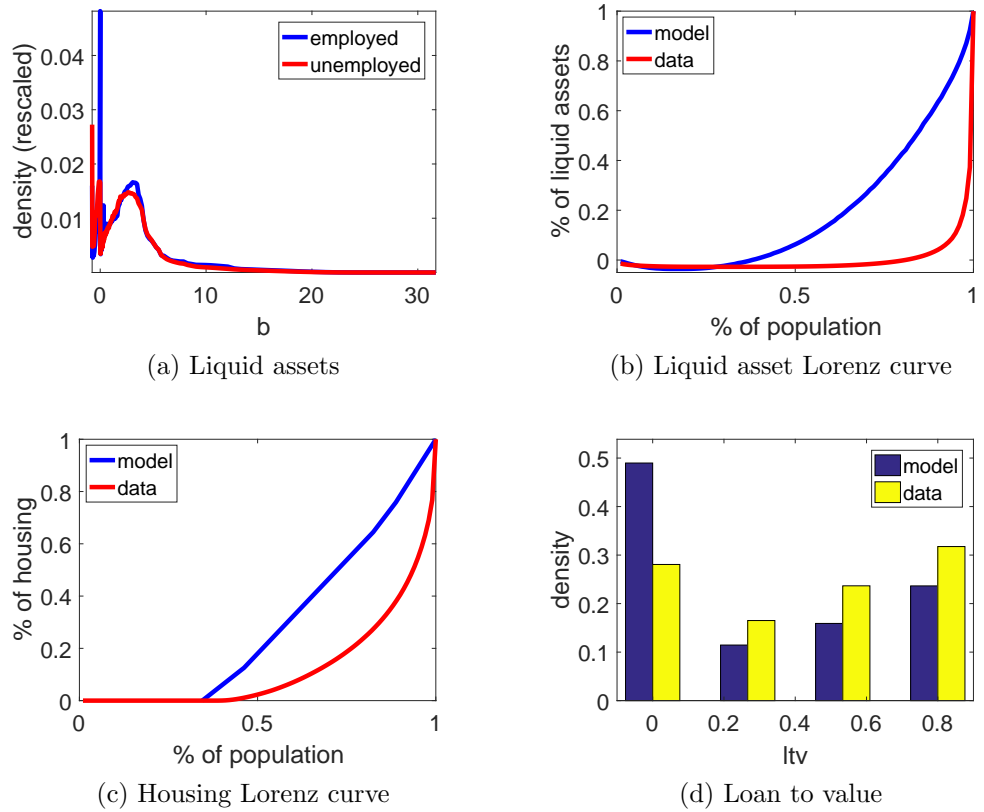




Figure 1.10: Asset distributions




(a) Density




(b) Average b/h

Figure 1.11: Job distributions




(a) Log. Income




(b) Unemployment

Figure 1.12: Labor market response to unemployment shock



(a) Housing



(b) Log. Housing

Figure 1.13: Housing response to unemployment shock

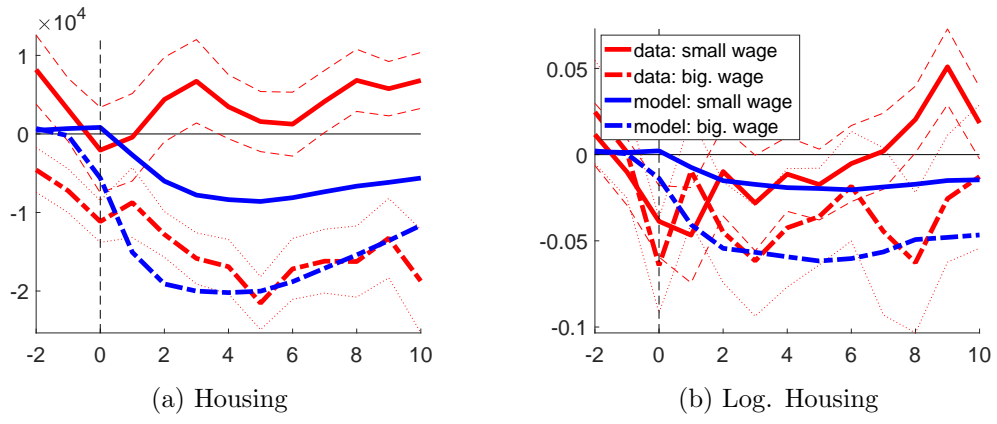


Figure 1.14: Housing response to unemployment shock, by size of income shock




Figure 1.15: Marginal Welfare Gain from Insurance

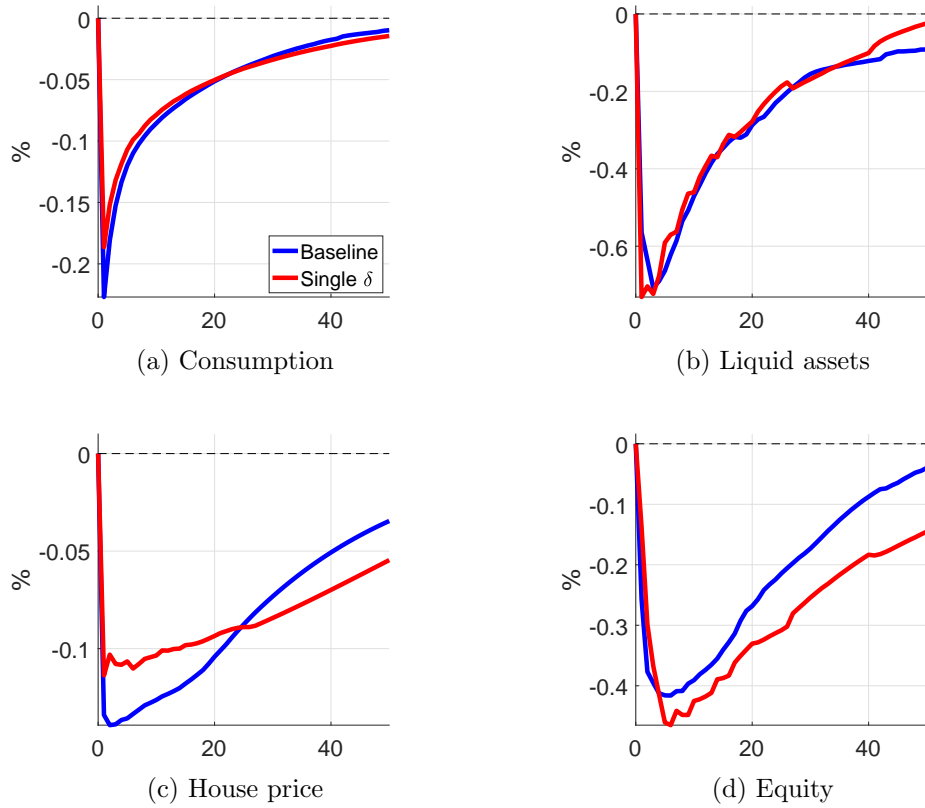


Figure 1.16: IRF to lump sum tax shock

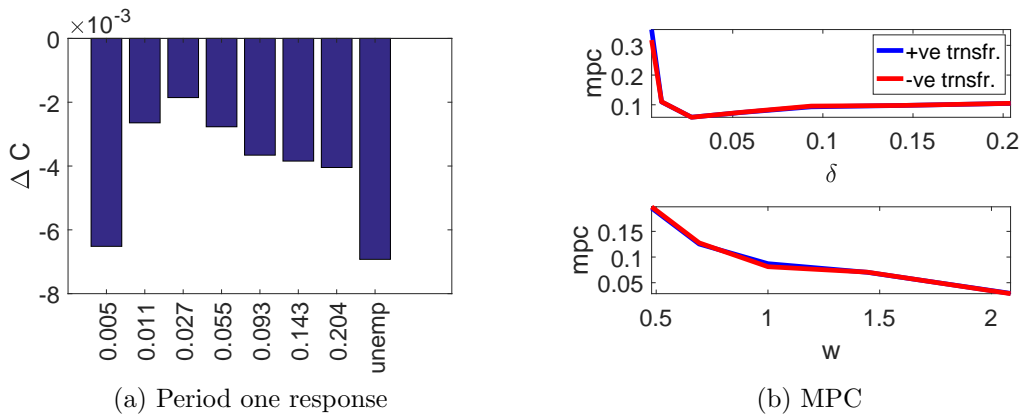





Figure 1.17: Role of job risk in response




(a) δ Model



(b) δ Data



(c) ω Model



(d) ω Data

Figure 1.18: Great Recession shocks




Figure 1.19: IRF to Great Recession shock

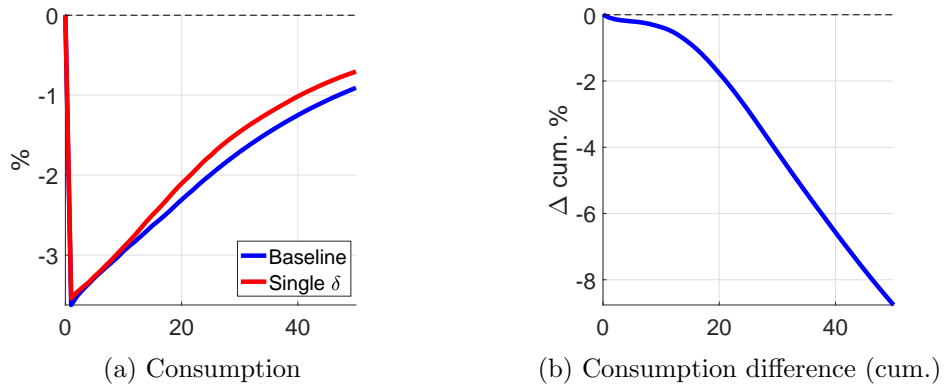


Figure 1.20: Persistence of consumption response




Figure 1.21: Amplification with prior increase in wage uncertainty

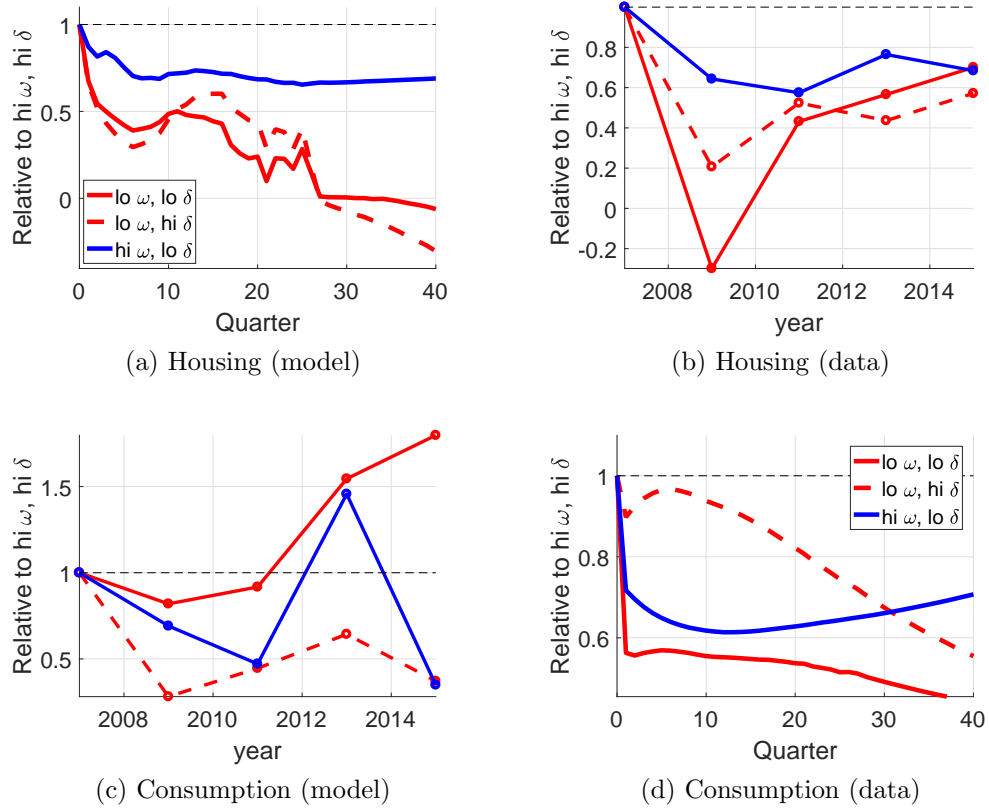


Figure 1.22: Relative group responses to Great Recession shock

Chapter 2

Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics

2.1 Introduction

Are consumer credit conditions important for aggregate economic outcomes and how do they affect consumption dynamics during recessions? The 2007 recession is widely regarded to have either been caused by, or coincided with, a severe financial crisis and contraction of credit or “credit crunch” ([Brunnermeier \(2009\)](#), [Jermann and Quadrini \(2012\)](#), [Eggertsson and Krugman \(2012\)](#), [Hall \(2014\)](#)). A more dominant strand of research since the crisis has concentrated on the effect of credit shocks on firms (for example: [Khan and Thomas \(2013\)](#)), while less work has looked at the impact on the household sector.

This paper looks at the role of credit conditions in household consumption dynamics. It argues that much of the literature has misunderstood the effect of a decline in credit availability during the Great Recession. Rather than forcing indebted households to reduce their consumption to reduce their level of borrowing, the main mechanism by which a credit crunch operates is to lower

the probability of households making large durables or housing purchases. The reason for this is that households that are already borrowing are unlikely to be immediately affected by a contraction in credit availability. However, the decline in credit availability will reduce their ability to make future purchases which now require a larger downpayment. Further, their existing credit terms secured during more favourable aggregate conditions have become more valuable, providing a reason to avoid a purchase that would result in credit renegotiation. Such an incentive is likely to have become stronger following the growth of home equity line of credit (HELOC) loans during the first half of the 2000s (Johnson and Sarama (2015)). This mechanism can help us understand the prolonged response of durables consumption and deep decline in transaction volumes following the Great Recession.

In particular, the model is able to replicate the consumption response of agents when a recession coincides with a contraction in the availability of credit, such as the 2007-09 recession compared with a more standard downturn. This is important as the recent recession saw a large adjustment on the durables consumption margin. Narrowly defined, durables consumption declined by 14.2 percent from peak-to-trough between 2007.IV and 2009.II, whereas the average decline across all US post-war recessions was 9.7 percent. The decline is much larger if a broader definition of durables, including residential investment, is used. The majority of consumer credit extended to US households is for the purpose of durables purchases and is usually secured against that stock. For example, in the 2010 Survey of Consumer Finances, 83.9 percent of family debt was secured against a residential property. Therefore, it is natural to assume that if households are affected by credit availability that credit secured against durable assets will be the most important class to consider for their consumption behaviour.¹

¹There is a long tradition of introducing durable goods into real business cycle models, as a propagation mechanism (Baxter (1996); Leahy and Zeira (2005)) and to match empirical facts, such as consumption volatility (Álvarez Parra et al. (2013)) or the co-movement of

This paper presents a general equilibrium infinite horizon heterogeneous agents model, with durables consumption, subject to realistic non-convex adjustment costs, and featuring two types of aggregate uncertainty: productivity shocks and collateral constraint shocks. Non-convex adjustment cost are an important feature of the model to ensure that the agent level behaviour matches that of the microeconomic data, where “infrequent and lumpy” purchases are observed, alongside sluggish adjustment of aggregate variables (Bertola and Caballero (1990)).² Empirical evidence for the validity of such a modeling specification has been established using automotive purchases data (Eberly (1994), Attanasio (2000)). In this class of model, the degree of agent inaction is a function of model primitives, such as the depreciation rate, wealth drift and degree of uncertainty (Caballero (1993), Carroll and Dunn (1997), Dunn (1998), Bertola et al. (2005)). In a heterogeneous agent environment, as studied here, higher adjustment costs entail a stronger precautionary savings motive (Gruber and Martin (2003)).

The model studied is similar to that of Berger and Vavra (2015), but with the addition of richly specified time varying credit conditions. Relatively few papers have looked at the importance of credit shocks for households in a fully heterogeneous environment. In the representative agent setting, Iacoviello (2005) incorporates a collateral channel into a New-Keynesian model that tightens with house prices. Campbell and Hercowitz (2005) link loosening in the collateral constraint to a decrease in the volatility in macroeconomic aggregates. Guerrieri and Iacoviello (2017) generate asymmetric responses to house price changes via a collateral channel. A contrary result is d’Albis and Iliopoulos (2013), who show that borrowing constraints can be undone by the introduction of a rental market. Buera and Moll (2015) stress the lim-

imports and exports (Engel and Wang (2011)).

²Grossman and Laroque (1990) show that in this setting optimal consumption is not a smooth function of wealth, but characterized by an optimal return point, upper and lower boundary and region of inaction.

itations of studying credit shocks in a representative agent setting. In the heterogeneous agents environment, [Jose Luengo-Prado \(2006\)](#) find that at the agent level consumption growth become more volatile with looser collateral constraints. [Silos \(2007\)](#) finds looser collateral constraints have a second-order effect on unconditional aggregate business cycle moments, but reduce the procyclicality of business investment for young, constrained agents. In a life-cycle setting, [Fernández-Villaverde and Krueger \(2011\)](#) highlight the importance of durables and endogenous collateral or borrowing constraints, in generating the hump-shaped profile of durables and non-durables consumption over the life-cycle seen in the data. While with non-convex adjustment costs, [Iacoviello and Pavan \(2013b\)](#) capture features of the Great Moderation – an increase in the level of debt and a decline in the volatility of residential investment, with amongst other changes, looser collateral constraints. Perhaps the most well known paper featuring changes in consumer credit conditions in a heterogeneous environment is [Guerrieri and Lorenzoni \(2017\)](#). They study a permanent unexpected tightening of the collateral constraint in a Bewley model with durables. In the absence of nominal rigidities when credit availability contracts the precautionary savings channel dominates leading to an increase in durables purchases. [Favilukis et al. \(2017\)](#) incorporate housing into a general equilibrium framework but focus on the role of endogenous house prices and its effect on tightening or loosening credit conditions.

This paper features shocks to credit availability that are anticipated by agents and critically it relaxes the commonly made assumption that a change in aggregate credit conditions affects all agents immediately (see: [Guerrieri and Lorenzoni \(2017\)](#), [Guerrieri and Iacoviello \(2017\)](#)). Instead in this model conditional on non-adjustment agents retain previously agreed credit terms and this introduces a channel by which the aggregate shock is propagated. [Justiniano et al. \(2015\)](#) propose an asymmetric collateral constraint similar in spirit, but in their model the tightness of the collateral constraint is linked to

changes in the durables price rather than the decision of the household and so does not feature the endogenous propagation mechanism emphasised in this paper. A few papers have linked the collateral constraint faced to the households adjustment decision. [Sterk \(2015\)](#) links the tightness of the constraint to the adjustment decision in a representative agent setting emphasising the impact on the labour market and job mobility. In a heterogenous setting [Amior and Halket \(2014\)](#) and [Halket and Vasudev \(2014\)](#) feature an asymmetric collateral constraint linked to the adjustment decision in a stationary environment not suitable for business cycle analysis. The specification used in these papers capture the long term debt nature of mortgage contracts. The constraint proposed here is looser than that suggested in the former and additionally captures the value to agents of having an agreed line of credit against their durables stock, such as HELOC loans that grew in popularity prior to the downturn. The choice aspect of the implementation of credit conditions studied in this paper is similar to [Wong \(2017\)](#) who focuses on the refinancing decision in the face of a stochastic interest rate. In a related setting, [Berger et al. \(2018b\)](#) highlight the resulting path- or state-dependency which is also a feature of this model.

Heterogeneity is also a vital feature of the model. Firstly, it allows agents to hold history specific credit terms, the distribution of which is an endogenous object. This means that whilst the available terms of credit is an exogenous variable, the effective tightness of credit availability in the model is an endogenous object. Secondly, heterogeneity also enables an accurate representation of the wealth distribution for low wealth agents. This is a key variable for the importance of credit shocks and disciplines the model's outcome. At the same time it means the model does not need to resort to imposing the condition that a fixed share of agents are at the borrowing limit as would be the case in a representative agent or related lender/borrower setting. Indeed the fraction of agents at or near the borrowing constraint is also a time varying endogenous

object.

The main results of the paper are as follows: a contraction in aggregate credit conditions leads to a deep, persistent decline in durables consumption and the percentage of agents adjusting their durables stock. This is despite the fact that only a small proportion of agents are near the collateral constraint. When a credit contraction occurs concurrently with a negative productivity shock there is additional propagation and the decline in the durables consumption share resembles that of the Great Recession. We validate the model mechanism by comparing the correlations between aggregate credit conditions and the household level behaviour of adjusters and non-adjusters in the Panel Study of Income Dynamics (PSID). The model with idiosyncratic credit terms is better able to replicate the relative leverage position of adjusters and non-adjusters than a model with a common aggregate collateral constraint shock. This correlation emphasizes that a key mechanism for the effect of credit shocks is inaction due to the option value of credit terms rather than immediate adjustment. As a consequence, the model is also able to generate the prolonged depression of the durables consumption share seen during the Great Recession, a result not produced by a uniform aggregate credit shock. In this setting the effect of credit shocks is far more persistent. Finally, these shocks are not amenable to policy intervention. Government action to loosen the constraint reduces consumer welfare, albeit the welfare loss is small.

The paper is set out as follows. Section 2.2 presents supporting empirical evidence; section 2.3 outlines the model; section 2.4 details the calibration; section 2.6 details the responses to the various shocks; section 2.5 explores the model's long-run properties and unconditional business cycle moments; section 2.7 presents the policy experiment; section 2.8 explore the model's robustness to alternative specifications, while section 2.9 concludes.

2.2 Empirical evidence

This section sets out the empirical evidence that serves to motivate focus on durable and importance of consumer credit shocks.

2.2.1 Consumption in the Great Recession

While all forms of consumption fell following the Great Recession, the decline appears to have been particularly concentrated on durables consumption expenditures. Between 2007.IV and 2009.II, non-durables consumption fell by 1.4 percent, durables consumption fell by 14.2 percent and residential investment declined by 44 percent. Further, this decline has been more persistent than in previous recessions. Figure 2.1 documents the change in the consumption share of non-durables consumption, durables consumption and residential investment during the last seven recessions. The shares are presented for ten quarters before and ten quarters after the beginning of the recession, as dated by the National Bureau of Economic Research (NBER), and normalised to one at the recessions commencement. As can be seen, the decline in the share of durable consumption and residential investment was larger in the Great Recession than in most of the previous recessions. Further, while on average this decline in the durables consumption share rebounded after ten quarters, following the 2007 recession the decline persisted. This is taken as evidence to suggest that the shocks hitting the economy during the 2007 recession differed somewhat from the typical downturn.

2.2.2 Decline in adjustment rate

The Great Recession also saw a significant decline in the transaction volumes for durable goods, with much of the adjustment taking place on the extensive margin and households postponing purchases. Figure 2.2 presents the monthly sales of existing homes, new homes and new vehicles in the period prior to and

following the 2007 recession. Also shown is an average of previous recessions' decline, depending on data availability. Existing home data is only available from 1999. New homes data is available from 1963, while new motor vehicles data is available from 1976. For these variables the average is over the 1980, 1981, 1990 and 2001 recessions. The average is calculated in percentage terms and then normalised to match the sales level in December 2007. Existing home and new home sales both began declining in advance of the recession and continued to decline following the recession's onset. The decline during this period was far more significant than that seen in previous recessions. Similarly, new car sales dropped dramatically following the recession, again greater than average decline seen in previous recessions. Also visible is the role of the Cash for Clunkers policy in 2009, which raised vehicle sales during this period.

2.2.3 Measures of credit availability

The Great Recession saw a dramatic tightening of the credit available to households. Figure 2.3 plots three series from the Federal Reserve Board's Senior Loan Officer Opinion Survey, which provide an exogenous account of the changing availability of credit in the US economy. The Credit Easing Accumulated (CEA) series reports the response to the question "Net percentage of domestic banks reporting increased willingness to make consumer installment loans". The auto downpayment series reports the net percentage of domestic banks increasing the minimum down payment for autos loans and the mortgage standards series reports the proportion of banks tightening standards for mortgage loans. The negative of these two series is taken so that a number greater than zero implies a loosening. All three series recorded highly negative value during the 2007 recession. This contrasts with the 2001 recession when the series did not deteriorate. This suggests that a lack of credit availability

was a significant factor in the Great Recession, but did not play a large role in previous recessions such as the 2001 recession. Further, all three series co-move strongly, suggesting an underlying process driving credit availability in a variety of sub-markets. This assumption is also useful for the purpose of this paper as the more general series reporting the supply of consumer installment loans is available over a much longer period, beginning in 1966.III. While the other two series provide a tighter fit to the collateral constraint specification outlined in the model below, it will be assumed that their behaviour and the tightness of exogenous credit conditions would closely mirror the consumer loans availability.

The long run business cycle dynamics of exogenous credit conditions is presented in the bottom panel Figure 2.4, as the CEA index. This is a Hodrick-Prescott HP filtered version of the data in Figure 2.3 weighted by aggregate debt to income.³ Two further credit variables are also presented, these are average loan to value series for new automotives, $LTV(1)$ and housing purchases $LTV(2)$. In the case of these series it is more likely that they also reflect endogenous responses. These series broadly move in sync with each other, albeit there are instances, such as the mid-2000s, where the loan to value variables declined, while the CEA index increased. Also apparent is that whilst there are recessions in which the CEA series declines significantly - such as 2007 recession, after which the series reached its lowest value - there are also cases where the fall is much smaller, for example the 1990s recession. This suggests that there could be recessions where credit conditions play a much more important role than others and provides a rationale for model-

³The Credit Easing Accumulated index used is as presented in [Slacalek et al. \(2012\)](#). This series seeks to provide an exogenous account of the changing availability of credit in the US economy. It is constructed by aggregating changes in response to the variable: “Net percentage of domestic banks reporting increased willingness to make consumer installment loans”, from the Federal Reserve Board’s, Senior Loan Officer Opinion Survey, weighting these changes by the aggregate debt to personal income ratio. The version of the variable used is that weighted by much more, and much less willing to make consumer loans. Debt is ‘total consumer credit’ from FRB’s G19 - Consumer Credit dataset. Personal income is personal disposable income, from the NIPA.

ing productivity and the availability of credit as different stochastic processes. The top panel presents three durables consumption series: $I^d(1)$ durables narrowly defined, $I^d(2)$ residential investment and $I^d(3)$ the composite of these two series. Clearly visible is the synchronized nature of the three durables consumption series. It is also clear that understanding durables consumption is very important for understanding consumption dynamics during recessions, with significant declines taking place during these periods.

The CEA variable is highly correlated with the durables consumption. The correlation with the broad definition $I^d(3)$ is 0.83 and this is higher than its correlation with output, 0.70, or investment, 0.56. The full business cycle moments are presented in Appendix B.2, Table B.7. This suggests credit availability might be particularly important for consumer durables. The variable is also highly persistent with a first order autocorrelation of 0.94 and a second order autocorrelation of 0.83. Similarly, the loan to value series are also more correlated with durable consumption than output or investment, although the correlation is lower. Further vector autoregression (VAR) evidence on the importance of credit shocks for durables consumption is presented in Appendix A.2.1.

2.2.4 Growth in home equity line of credit borrowing

A significant change in residential credit market has been the growth of HELOC and as such this is an important feature to capture when modelling the changing credit conditions experienced during the Great Recession. These are loans in which rather than an amortised repayment schedule, a maximum loan amount is agreed between the borrower and lender to be drawn upon in the future. This borrowing is secured against equity in the borrower's house. Typically the interest rate is floating rather than fixed. As a share of the total mortgage market HELOCs grew from 2.38 percent in 1999.I to 7.37 percent

in 2005.IV, before declining during the recession (left panel, Figure 2.5). The value of these credit arrangements is high to households because they allow for additional borrowing to be undertaken on previously agreed terms. While there were instances of lenders reducing borrowers' agreed credit limits during the recession (see New York Times, 2008), the aggregate figures suggest that there was no sudden reduction in maximum loan limits. The right panel of Figure 2.5, presents the real HELOC balances and borrowing limits. The limit declines gradually from \$1.5 trillion to \$1.0 trillion between 2007 and 2010, following the onset of the recession suggesting some renegotiations or lower renewed terms, but does not decline precipitously. The actual balance of HELOC loans declines even less substantially, suggesting a relatively small share of constrained individuals.

2.2.5 Loan to value ratios by behaviour

As shown in Figure 2.4 loan to value ratios vary over time. Microeconomic data allows us to further distinguish between the borrowing patterns of those adjusting and not-adjusting their housing stock. Figure 2.6 presents data from the PSID on the loan to value ratio for households with positive housing stock and holding a mortgage. The sample is divided by those that have and have not moved house since the last survey, this is in the previous year prior to 1997 and in the last two years from 1999 onwards, adjusters and non-adjusters respectively. As can be seen from the top panel, movers have higher loan to value ratios than non-movers, holding roughly double the share of mortgage debt to housing value.

The relative strength of the response of these two groups to changes in credit conditions is informative about the way credit shocks affect the economy. The difference between the de-meaned value of these loan to value series is presented in the lower panel, alongside the HP filtered CEA. As can be seen,

prior to the 2007 recession as credit conditions loosened the difference between the loan to value ratio of adjusters and non-adjusters increased. Once credit conditions tightened during the recession this difference fell. The correlation between these two series is 0.25 or 0.14 if using the logged value of the loan to value ratios.⁴ In Section 2.5.1 we will show that the positive correlation between these two series can be used to distinguish between models of credit shocks and supports the specification proposed in this paper with idiosyncratic credit terms.

2.3 Model

This section of the paper outlines the model. The model is an infinite horizon version of the stochastic growth model, with heterogeneous agents, incomplete markets and aggregate uncertainty. Households purchase a consumption good and a durable good and are subject to idiosyncratic employment shocks which determine their labour income. Households are able to save and borrow in non-state contingent assets, subject to a collateral constraint. The tightness of this collateral constraint is determined by aggregate credit conditions and the household's decision whether to adjust or not. Savings are transformed into a capital good which is used in production.

2.3.1 Utility function and idiosyncratic uncertainty

The environment studied here is one in which households undertake home production as in [Greenwood and Hercowitz \(1991\)](#). The households receive utility from non-durables consumption, c_t , and home production, h_t . The instantaneous utility function is given by $u(c_t, h_t)$. Home production is a function of the end of period durables stock, d_{t+1} , the disutility from working, l_t and

⁴The correlation between the difference in LTV ratios and the non-HP filtered, linearly detrended series is 0.0304 (logged -0.1501), due to a level increase in credit conditions in the 1990s, to which there was no response of the difference in LTV ratios.

home production productivity parameter ς_t , such that $h_t = H(d_{t+1}, \varsigma_t(1 - l_t))$. Following, Greenwood and Hercowitz (1991) it is assumed that:

$$u(c_t, h_t) = \frac{(c_t^\theta h_t^{1-\theta})^{1-\alpha}}{1-\alpha} - 1$$

$$H(d_{t+1}, \varsigma_t(1 - l_t)) = \{\omega d_{t+1}^\lambda + (1 - \omega)[\varsigma_t(T^{MAX} - l_t)]^\lambda\}^{1/\lambda}$$

where α is the inverse of the elasticity of intertemporal substitution, θ is the share of non-durables consumption in utility, ω is the share of durables in home production, λ is the intratemporal elasticity of durables and labour in home production, and T^{MAX} is the maximum time allocation of the agents. Each period the agent can invest in durables, i_t^d , at the relative price, p . As their name suggests durables are not fully consumed in period. The law of motion for durables is:

$$d_{t+1} = (1 - \delta^d)d_t + i_t^d$$

where δ^d is the durables depreciation rate.

Households face two idiosyncratic shocks: employment shocks and preference shocks. Starting with the employment shocks in each period an agent can either be employed, $\varepsilon_t = 1$, or unemployed, $\varepsilon_t = 0$. Agents make a decision to work on the extensive margin, such that $l_t \in \{0, 1\}$. When employed the agent earns a pre-tax wage rate, w_t , and faces a tax, τ_t , leaving an after tax wage of $(1 - \tau_t)w_t\bar{l}_t$, where \bar{l} is a productivity normalisation. When the agent is unemployed they receive the wage $\mu w_t l_t$, where μ is the unemployment benefit replacement rate. The interpretation of the decision to work, l_t , for the unemployed agents is that to receive benefits active job search must be demonstrated, thus these agents could be thought of as looking for work.

The agents also experience shocks to their preferences. The discount fac-

tor, $\beta \in \mathcal{B} \subset (0, 1)$, used to discount future utility differs across agents in the economy and is time varying. Its distribution is invariant with a fixed proportion of agents of different levels of β in every period. This feature of the model is used to capture the idea that some proportion of the population might be more patient than others. This feature is commonly used in the literature to generate more realistic wealth distributions (see Krusell and Smith, 1998). Here it is principally employed to ensure that a sufficient proportion of the agents will take on debt and remain at, or close to the collateral constraint.

2.3.2 Adjustment costs

If households do decide to adjust their durables stock they must pay a non-convex adjustment cost. Non-convex adjustment costs mean that it is optimal for an agent to not adjust their durables stock in every period, replicating the observed lumpy and infrequent behaviour of durable purchases at an agent level. The ability to differentiate between adjusters and non-adjusters is also an important part of the specification of the terms of credit faced by agents (see Section 2.3.4).

I assume that agents face a fixed cost proportional to their current durables stock when adjusting. When not adjusting agents incur a maintenance cost, χ , which is a share of current period depreciation. Not adjusting and paying the maintenance cost partially offsets this depreciation, see also [Berger and Vavra \(2015\)](#). I abstract from their specification, which includes an adjustment cost proportional to the agents labour cost. Adjustment costs of this form generate (S,s) policy rules. Such costs are most readily interpreted as either a utility cost of adjusting, e.g. time spent choosing, locating and buying new durables stock, the loss of value when selling the existing stock in the secondary market

or realtor fees. Formally, agents must pay the adjustment cost:

$$\Psi(d_t, d_{t+1}) = \begin{cases} \Psi(1 - \delta^d)d_t & \text{if } d_{t+1} \neq d_t(1 - \delta^d(1 - \chi)) \\ 0 & \text{if } d_{t+1} = d_t(1 - \delta^d(1 - \chi)) \end{cases}$$

2.3.3 Saving, borrowing and budget constraint

Agents are only able to save and borrow in one period, non-state contingent assets, k_{t+1} . Saving is ultimately in the form of capital accumulation, discussed in the next subsection, so agents receive a net rate of return of $r_t - \delta^k$ on assets, k_t , where δ^k is the depreciation rate of capital and r_t is the factor price. Agents are also able to borrow at this rate. Finally, each household also receives a share of the representative firm's profits, π_t . The agent's budget constraint can therefore be written as:

$$c_t + pd_{t+1} + k_{t+1} = (1 - \delta^k + r_t)k_t + (1 - \delta^d)pd_t - \Psi(d_t, d_{t+1}) + [(1 - \tau_t)\bar{l}\varepsilon_t + \mu(1 - \varepsilon_t)]w_t l_t + \pi_t$$

2.3.4 Collateral constraint

The approach to modelling the collateral constraint attempts to capture two salient features of the durables credit market. Firstly, that aggregate changes in credit conditions are unlikely to affect agents immediately unless they are undertaking a durables purchase. One interpretation of this is that it replicates a long-term debt contract. Secondly, capturing the rise of HELOC mortgages, agents may have an available line of credit secured against their durable assets that they are not currently taking advantage of. Borrowing is only permitted for the purpose of financing durable purchases, in the sense that loans are

subject to a collateral constraint such that:

$$k_{t+1} \geq -\zeta_t^i d_{t+1} \quad (2.1)$$

and $\zeta_t^i \in (0, 1)$. This restriction says that an agent must make some minimum down payment, $(1 - \zeta_t^i)d_t$, when purchasing a durable good, where ζ_t^i is referred to as an agent's collateral constraint. Such a restriction can be rationalized by the existence of information problems between the lender and borrower, that would restrict the ability of the lender to recover the full value of the durables stock in the case of default. Note that due to the assumption that all financial assets are one period commitments, agents can costlessly refinance up to the collateral constraint, every period, given the depreciated value of their existing durable stock.

The specification of the collateral constraint (1) means that the constraint is both idiosyncratic and time varying. I assume that the aggregate credit conditions in the economy are given by Θ_t , which follows a Markov process such that $\Theta_t \in \mathcal{X}$. The standard assumption is that $\zeta_t^i = \Theta_t$, implying that all agents are immediately affected by a change in the credit conditions (for example [Guerrieri and Lorenzoni \(2017\)](#)). However, this seems a strong assumption as usually it would only be households renegotiating their line of credit (due to a new house purchase, for example) that would be affected by the aggregate credit conditions. Indeed the evidence presented in 2.2.5 suggests a systematic difference in the LTV ratios of adjusters and non-adjusters. I assume agents enter the period with a given individual specific credit terms, Ξ_t^i , which summarises the credit agreement on their existing loan. If the agents do not adjust their durables stock it is this term that determines the collateral constraint, $\zeta_t^i = \Xi_t^i$, and they face the constraint $k_{t+1} \geq -\Xi_t^i(1 - \delta^d(1 - \chi))d_t$. Non-adjusting agents cannot hold onto credit terms indefinitely, it is assumed that the lender may be able to renegotiate credit terms. With probability ρ^Ξ

the credit terms are retained tomorrow, $\Xi_{t+1}^i = \Xi_t^i$ and with probability $1 - \rho^\Xi$ next period's credit terms will be given by next period's aggregate credit conditions $\Xi_{t+1}^i = \Theta_{t+1}$.⁵

Alternatively, if the agent adjusts their durables stock they also need to refinance their purchase taking today's aggregate credit conditions, $\zeta_t^i = \Theta_t$. They then face the collateral constraint $k_{t+1} \geq -\Theta_t d_{t+1}$. In this case the next period credit terms are given by today's aggregate state $\Xi_{t+1}^i = \Theta_t$ with probability one. This specification implements in a parsimonious way the idea that most agents in an economy will not be directly affected when aggregate credit conditions change, having already agreed their credit terms. It is predominantly those agents required to agree new credit terms upon the purchase of a new durable good, e.g. automobile or house, that must finance their purchases on the basis of the current aggregate credit conditions.

Other papers have linked the tightness of the collateral constraint to the adjustment decision. For example, [Amior and Halket \(2014\)](#) specify for non-adjusters $k_{t+1} \geq \min(k_t, -\Theta_t d_{t+1})$, such that agents must weakly reduce debt or refinance with the current credit conditions.⁶ While this captures the long term debt nature of the contract it does not feature the agreed available line of credit such as HELOC loans. Therefore, the option value of not-adjusting to retain the terms of credit is less in their specification.

Having completed the discussion of the household problem, it can now be

⁵Here the fact that agents may lose their credit terms reflects the reality that most credit contracts will not be open ended.

⁶Amior and Halket's specification actually features a location, j , varying endogenous price which tightens the collateral constraint rather than credit conditions e.g. $k_{t+1} \geq \min(k_t, -\Theta_j d_{t+1})$. While this generates different dynamics for the above discussion it can be considered broadly equivalent.

summarized denoting each agent with the superscript, i :

$$\begin{aligned} & \max_{\{c_t^i, d_{t+1}^i, k_{t+1}^i, l_t^i\}_{t=0}^{\infty}} u(c_t^i, H(d_{t+1}^i, \varsigma_t(T^{MAX} - l_t^i))) + \\ & E_t \sum_{j=1}^{\infty} \left(\prod_{m=0}^{j-1} \beta_{t+m}^i \right) u(c_{t+j}^i, H(d_{t+1+j}^i, \varsigma_{t+j}(T^{MAX} - l_{t+j}^i))) \end{aligned}$$

$$\begin{aligned} c_t^i + pd_{t+1}^i + k_{t+1}^i &= (1 - \delta^k + r_t)k_t^i + (1 - \delta^d)pd_t^i - \Psi(d_t^i, d_{t+1}^i) \\ &+ [(1 - \tau_t)\bar{l}\varepsilon_t^i + \mu(1 - \varepsilon_t^i)]w_t l_t^i + \pi_t^i \end{aligned}$$

$$k_{t+1}^i \geq -\zeta_t^i pd_{t+1}^i$$

$$\zeta_t^i = \begin{cases} \Xi_t^i & \text{if } d_{t+1} = d_t(1 - \delta^d(1 - \chi)) \\ \Theta_t & \text{if } d_{t+1} \neq d_t(1 - \delta^d(1 - \chi)) \end{cases}$$

$$\Xi_{t+1}^i = \begin{cases} \Xi_t^i \text{ with prob } \rho^{\Xi} & \text{if } d_{t+1} = d_t(1 - \delta^d(1 - \chi)) \\ \Theta_{t+1} \text{ with prob } 1 - \rho^{\Xi} & \text{if } d_{t+1} = d_t(1 - \delta^d(1 - \chi)) \\ \Theta_t & \text{if } d_{t+1} \neq d_t(1 - \delta^d(1 - \chi)) \end{cases}$$

As well as the necessary transversality conditions required for an optimal solution.

2.3.5 Firms

The goods market is competitive and characterised by a constant returns to scale production function so can be modeled as a representative firm, with the production function being concave and satisfying the Inada conditions. Denoting per capita aggregate capital, K_t , and the employment rate, L_t , with \bar{l} a normalisation, per capita output is:

$$Y_t = F(K, z_t L_t) = K_t^\eta (z_t \bar{l} L_t)^{1-\eta}$$

With prices given by:

$$w_t = (1 - \eta)z_t^{1-\eta} \left(\frac{K_t}{\bar{l}L_t} \right)^\eta$$

$$r_t = \eta z_t^\eta \left(\frac{K_t}{\bar{l}L_t} \right)^{\eta-1}$$

As $\lim_{K \rightarrow 0} = \frac{\partial F(\bullet)}{\partial K} = r = +\infty$, the equilibrium will always be characterised by a positive capital stock, with more aggregate savings than loans. The production function is affected by the level of aggregate labour augmenting productivity, z_t , which follows a Markov process, where $z_t \in \mathcal{Z}$. Household savings are transformed into capital $K_{t+1}^{AG} = \int k_{t+1}^i di$ and households receive the net return on capital, $r_t - \delta^k$, for savings.

2.3.6 Government

In the baseline model the government exists solely to provide unemployment insurance. The government collects taxes from employed agents and distributes them to the unemployed as a fraction, μ , of the current aggregate wage. The government's budget constraint is assumed to hold in each period. Therefore, taxes are given by:

$$\tau_t = \frac{\mu(1 - L_t)}{\bar{l}L_t}$$

In the policy experiment the government also levies additional taxes, to be discussed in Section 2.7, with the government budget constraint again assumed to hold in every period.

2.3.7 Recursive formulation

It is also convenient to reformulate the household's problem in the recursive formulation, where we can write the problem as one where the agent considers

the value of adjusting, $V^{A,i}$, and not adjusting, $V^{N,i}$, and maximises between these two values. Notice this adjustment decision takes into account both the difference in the end of period durables stock and the implication for the collateral constraint and next period's credit terms. Denote the idiosyncratic state $\mathbf{s} = (\beta, \varepsilon, k, d, \Xi)$, which is defined as idiosyncratic preferences, employment, assets, durables and credit terms.

At this point it is important to explicitly recognise that the interest rate and wage rate are functions of the aggregate stochastic shock, aggregate capital and aggregate employment. Additionally, households condition their expectations on the transition of the aggregate state, \mathbf{S} , with transition law $\mathbf{S}' = G(\mathbf{S})$, where the prime signifies next period's value. The aggregate state, $\mathbf{S} = (z, \Theta, \Gamma)$, is defined as the stochastic aggregate productivity shock, z , aggregate credit conditions, Θ , and the joint distribution of the idiosyncratic state, $\Gamma = \Gamma(\mathbf{s})$. Clearly this is a very complex object of infinite dimensions. Finally, define the home production productivity, ς , as a function of aggregate productivity and an agent's employment status, $\varsigma = \varsigma(z, \varepsilon)$.

The household maximises:

$$V(\mathbf{s}; \mathbf{S}) = \max \left\{ V^A(\mathbf{s}; \mathbf{S}), V^N(\mathbf{s}; \mathbf{S}) \right\}$$

Where the Bellman equation for an agent adjusting is:

$$V^A(\mathbf{s}; \mathbf{S}) = \max_{c, k', d', l} \left\{ u(c, H\{d', \varsigma(z, \varepsilon)(T^{MAX} - l)\}) + \beta E \left[V(\mathbf{s}'; \mathbf{S}') | \beta, \varepsilon, \Theta \right] \right\}$$

subject to:

$$\begin{aligned}
c + pd' + k' &= (1 - \delta^k + r(z, K, L))k - \Psi(d, d') + \\
&\quad (1 - \delta^d)pd + [(1 - \tau)\bar{l}\varepsilon + \mu(1 - \varepsilon)]w(z, K, L)l + \pi \\
k' &\geq -\Theta pd' \\
\Xi' &= \Theta
\end{aligned}$$

and the Bellman equation for an agent not adjusting is:

$$V^N(\mathbf{s}; \mathbf{S}) = \max_{c, k', l} \left\{ u(c, H\{d', \varsigma(z, \varepsilon)(T^{MAX} - l)\}) + \beta E \left[V(\mathbf{s}'; \mathbf{S}') | \beta, \varepsilon, \Theta, \Xi \right] \right\}$$

subject to:

$$\begin{aligned}
c + pd' + k' &= (1 - \delta^k + r(z, K, L))k + \\
&\quad (1 - \delta^d)pd + [(1 - \tau)\bar{l}\varepsilon + \mu(1 - \varepsilon)]w(z, K, L)l + \pi \\
d' &= (1 - \delta(1 - \chi))d \\
k' &\geq -\Xi pd' \\
Xi' &= \begin{cases} \Xi & \text{with prob } \rho^\Xi \\ \Theta' & \text{with prob } 1 - \rho^\Xi \end{cases}
\end{aligned}$$

Both are subject to the law of motion for the aggregate state:

$$\mathbf{S}' = G(\mathbf{S})$$

2.3.8 Equilibrium

An equilibrium for the economy can now be defined. Using the notation where $x(\cdot)$ is a policy for an in period choice and $x_+(\cdot)$ is a policy for an end of period choice, an equilibrium consists of a value function, $V(\mathbf{s}; \mathbf{S})$;

household policy functions: $c(\mathbf{s}; \mathbf{S}), d_+(\mathbf{s}; \mathbf{S}), k_+(\mathbf{s}; \mathbf{S}), l(\mathbf{s}; \mathbf{S})$; price functions, $r(z, K, L), w(z, K, L)$; and the joint transition law $G(\mathbf{S})$, such that:

1. The household's policy functions, $c(\cdot), d_+(\cdot), k_+(\cdot), l(\cdot)$, are the solution to the household's problem given prices.
2. Factor prices, $r(z, K, L), w(z, K, L)$, are competitive and the solution to the firm's problem.
3. The aggregate transition law, $G(\mathbf{S})$, is consistent with the individual's policy rules and the individual and aggregate transition probabilities.
4. The government budget constraint holds in each period.
5. The goods market and labour market clear and net savings is equal to capital demanded:

$$\int c_t^i di + p \int d_{t+1}^i di + K_{t+1}^{AG} + \int \Psi(d_t^i, d_{t+1}^i) di = (1 - \delta^k) K_t^{AG} + (1 - \delta^d) p \int d_t^i + Y_t$$

$$\int \varepsilon_t^i l_t^i di = L_t^{AG}$$

$$K_{t+1}^{AG} = \int k_{t+1}^i di$$

2.3.9 Model intuition

The novel feature of this model is the change in behaviour due to a tightening of the aggregate credit conditions. The mechanism can be understood by comparing the policy functions for durables consumption under alternative states. Figure 2.8 presents the policy functions, $d_+(\mathbf{s}; \mathbf{S})$ under a high and low collateral constraint for three levels of assets in hand. The agent's credit terms are set to match the aggregate credit conditions $\Xi_t^i = \Theta_t$. This shows the effect of a tightening in aggregate credit conditions under the standard assumption that agents are hit immediately by a change in the exogenous variable, Θ_t . In

this model it shows the decision for agents that lose their credit terms due to (random) renegotiation.

In the left panel, the low asset agent, who has a current allocation close to the high collateral constraint, chooses a lower next period durables choice when the collateral constraint is at its tighter level over much of the durables holding distribution. These agents are unable to finance the higher downpayment cost out of current income and are forced to reduce their durables stock. Further, the non-adjustment region shifts to the left, agents that would have adjusted up instead choose to let their stock depreciate further. Under the tighter credit conditions, these agents cannot finance the durables choice they would have previously chosen. The mid assets agent, who has a current allocation that places them close to the low collateral constraint, reduces durables consumption by a smaller amount on average, as they have less far to deleverage. Also notice, there remains a region for mid-asset case where the non-adjust option - where next period's durables equals the depreciated value of today's durables stock plus the required maintenance, along the dotted line - coincides. However, the no adjust-region shifts to the left, indicating that there is a change in the agents that would and would not choose to adjust. Finally, the high asset agent is completely unaffected by tightening of the collateral constraint.

Figure 2.8 presents the behavioral response introduced in this model by the presence of the agent specific credit terms - a function of the agent's history of choices and the aggregate states. Here the policy functions are presented for the high and low aggregate credit conditions, Θ_t , with the credit terms, Ξ_t^i , held at the highest level. These policy functions demonstrate the effect of a tightening of aggregate credit conditions - an agent adjusting faces the tighter collateral constraint of the aggregate conditions, but an agent not adjusting faces a looser constraint based on their existing terms. The non-adjusting agent no longer needs to increase their asset holdings and this enables a higher level of durables stock to be maintained. The option value to the agent of

not adjusting is larger, as by not adjusting agents are able to borrow more in the future, leaving them further from their borrowing constraint. These two effects can be seen in the expanded inaction zone in the left and central panel, including agents with higher durables stocks. The inaction zone is now wider than under the loose aggregate credit conditions. The tightening of the aggregate credit conditions reduces the probability of adjustment, which reduces durables consumption, and due to depreciation, results in a decline in the durables stock. For the high asset agent, once again the collateral constraint has no effect.

2.4 Calibration

This section discusses the calibration of the model. The model is solved non-linearly using a generalised endogenous grid method and assuming [Krusell et al. \(1998\)](#) bounded rationality for the aggregate law of motion. Full details are provided in Appendix A.2, Section A.2.2. First the calibration of the collateral constraint is dealt with, then the preference and technology parameters, and finally those affecting the wealth distribution.

2.4.1 Collateral constraint process

A key set of parameters to be calibrated in this model is the collateral constraint and the exogenous stochastic process that it follows. It is assumed that the collateral constraint follows a two-state Markov process. The highest value of the collateral constraint is set to, $\Xi^H = 0.8$, implying a down payment requirement of 20 percent. This approximates the average loan to value ratio for autos of 0.898, and mortgages of 0.757, and is in keeping with parameter choice in [Guerrieri and Lorenzoni \(2017\)](#), [Bajari et al. \(2013\)](#) and [Iacoviello and Pavan \(2013b\)](#). The lowest collateral value is set to three quarters this value, $\Xi^L = 0.6$. For comparison, [Guerrieri and Lorenzoni \(2017\)](#), who target

a reduction in the household debt-to-GDP ratio, choose 0.56, for their tighter credit constraint value.

The process for the collateral constraint is estimated from the CEA series. A linear trend is subtracted from the accumulated index, the resulting series is divided into below and above average states, and the transition probabilities between these states is computed. The resulting transition matrix is presented in Table 2.1. The transition probabilities imply the collateral constraint is a slow moving persistent process, in keeping with the observation of [Drehmann et al. \(2012\)](#) that financial cycles tend to be longer than business cycles. The expected duration of the low value constraint is 18.8 quarters. The expected duration of the high value constraint is 23.2 quarters.

Finally, the persistence of credit terms, ρ^{Ξ} , is set to 0.99, so that the expected duration of the credit terms, conditional on no adjustment, is 100 quarters or 25 years. This matches the average term of a primary mortgage, in the American Housing Survey.

2.4.2 Preference and technology parameters

The model is calibrated to target a number of long run aggregate ratios from the NIPA, using the definition of a narrow durables good. The non-durables share and durables share in home production is selected to target the durables to quarterly non-durables consumption flow: $\frac{D}{C} = 2.3$ and ratio of non-durables consumption to output: $\frac{C}{Y} = 0.68$. The elasticity of substitution is set to -1, following [Greenwood and Hercowitz \(1991\)](#). Total hours is set to 3.228, such that an employed worker, supplying labour ($l_t = 1$) works for 31 percent of the week, based on average hours worked of 34.7 in the Current Employment statistics.

The quarterly depreciation rate of durables, δ^d , is set to 0.0484, calculated from the reported BEA depreciation rates, the share of the total stock of the

various durable goods and the implied depreciation rate of cars. The fixed adjustment cost, Ψ , and maintenance parameter are selected following [Berger and Vavra \(2015\)](#).

The productivity level, unemployment rates, and associated transition matrix is a generalisation of that taken from [Den Haan et al. \(2010\)](#). In that specification is a two state Markov process, with “bad” and “good” states. Here a rare third “very bad” state is added to approximate the additional decline in productivity that occurred during the Great Recession. This very bad state features an additional one percent fall in productivity, and unemployment rises to 10 percent. The transition probabilities are computed from a three state approximation of the unemployment rate during the period. The transition process for the “bad” and “good” states is very close to that implied in [Den Haan et al. \(2010\)](#), whilst the probability of transitioning to the very bad, conditional on being to the bad state is 0.024 percent.

As almost all agents in the model supply labour during all periods, $l_t = 1$, the aggregate unemployment rate is a deterministic function of the aggregate productivity state. A multiplicative specification for the productivity of home production is used, with idiosyncratic productivity of a employed agent normalised to one: $\zeta(z, \varepsilon) = z\varsigma$ if $\varepsilon = 0$ and $\zeta(z, \varepsilon) = z$ if $\varepsilon = 1$. This leaves one parameter, ς , which is set to 0.4271, based on the time spent on home production activities (shopping, cooking and housework) of unemployed relative to employed individuals, as reported in [Krueger and Mueller \(2012\)](#). They report that employed individuals spend 82 minutes on shopping, cooking and housework, whilst the unemployed spend 178 minutes. The unemployment benefit replacement rate is set to 0.4, following [Shimer \(2005\)](#) value of leisure and the US Unemployment Insurance replacement rate, as reported by the Department of Labor. The remaining preference and technology parameters are standard in the business cycle literature. The full list of the preference and technology parameters is presented in Table 2.2.

Wealth parameters

A final set of parameters to calibrate are those governing the wealth distribution. These parameters are important for the model's predictions as they govern the proportion of agents close to, and thus affected by, the collateral constraint. A two state process for the discount rate is assumed, with agents either being patient, β^H , or impatient, β^L . It is assumed that there are equal shares of patient and impatient agents. The persistence of the discount factor type is selected to represent an agent's life, the expected duration of the discount factor is 50 years. The values of the discount factors are calibrated within an auxiliary version of the model, without aggregate shocks. In particular the high type discount factor, β^H , is selected to generate a quarterly interest rate of 1.015. The gap between the lower and upper discount rate, $\beta^{GAP} = \beta^H - \beta^L$ is selected to target the share of agents in debt: 0.20 percent.

2.5 Properties of the economy

This section sets out the properties of the baseline model and assesses its performance against the data.

2.5.1 Model validation: behavior of adjusters and non-adjusters

The baseline model performs better than a model with aggregate credit shocks when describing the behaviour of durables consumption during the Great Recession. We can also distinguish between this model and the model with only aggregate collateral constraint shocks by considering the behaviour of adjusters and non-adjuster as presented in the empirical section 2.2.5. To see why the behavior differs consider the following. In the model where all households follow the aggregate collateral constraint, when credit conditions contract it is

the highly leveraged that are forced to adjust. This results in high loan to value households that would not have adjusted in the absence of the shock, deciding to adjust. This increases the difference in the expected loan to value ratio of adjusters and non-adjusters. By contrast in a model with idiosyncratic credit terms, when aggregate conditions tighten there is a strong option value for non-adjustment. This produces the opposite compositional change, with high loan to value households that were planning to adjust deferring adjustment to retain their existing terms of credit. As a result there is a decline in the difference between the average loan to value ratio of adjusters and non-adjusters.

In Section 2.2.5, we showed that in the PSID the difference between the loan to value ratio of adjusters and non-adjusters is positively correlated with changes in credit conditions. Table 2.4 shows that following the intuition above, the baseline model generates this positive correlation while the alternative aggregate collateral constraint model generates a negative correlation. This provides empirical support for the baseline specification versus a model where all agents are affected by changes in credit conditions concurrently. This validation can also be seen in a selection of other related summary statistics. In particular, for similar reasoning the baseline model generates a lower correlation of the adjustment rate with aggregate credit conditions than the alternative specification. The relationship is negative in the data. The correlation between the expected loan to value ratios of the adjuster and non-adjuster groups is also closer to that measured in the data. Finally, the baseline model generates a ratio of the standard deviations of the loan to value of adjusters to non-adjusters that is greater than one, again in line with the data.⁷

⁷This ratio is larger in the baseline model because non-adjusters have more acyclical loan to value ratios, whereas in the aggregate collateral constraint model, they must still adjust leverage with credit conditions even if they keep the same durables stock.

2.5.2 Long-run averages, distributions and moments

Table 2.5 displays the long run averages from the model. The targeted moments: $\frac{D}{C}$, $\frac{C}{Y}$ and the percent of agent in debt, are well approximated and the model generates an average quarterly interest rate of 1.015. The proportion of agents adjusting each period also closely matches the value from the PSID.

By comparing the average and median value of capital and durables, it can be seen that distribution of assets is far more skewed than the distribution of liabilities. The average assets holding is 32.1 whereas the median holding is only 12.7. For durables, the mean and median holdings are far closer, 5.7 and 5.4 respectively. A very small proportion of agents are at the aggregate collateral constraint, but seven percent of agents are near the aggregate constraint. Near the collateral constraint is defined as when distance of an agent's chosen financial asset position, to the maximum they could borrow given their choice of durables using the current aggregate credit conditions, is less than their current labour income e.g. $k_{t+1} \leq -\Theta_t d_{t+1} + [(1 - \tau_t)\bar{l}\varepsilon_t + \mu(1 - \varepsilon_t)]w_t$.

The model predicts too large a capital to output ratio. However, this statistic is pinned down by the targeted interest rate and so to achieve this ratio would require a very high interest rate, which is considered undesirable. Both the loan to value ratio of those in debt and the durables share of wealth are a little too low in the model, implying that in reality those in debt are more indebted than the model suggests, this should work against the model's findings.

The full distribution of financial assets and durables are shown in Figure 2.9. Compared with the empirical densities presented in Figure 2.9 in Appendix C.2, Figures C.14 to C.16, the model densities provide a reasonable qualitative approximation, particularly in the difference between the distribution of durables and assets. A greater proportion of the agents in the data have no or very low durables stock. While the model features skewed distributions

for both durables and assets it is not able to match the extreme degree seen in the data. For example, in the data assets have a skew of 178.2 whereas in the model the figure is only 2.3.

Table 2.6 presents the standard business cycle moments for the baseline model. Here the model does fairly well in replicating the stylized facts. The standard deviation of consumption is a little too low, a feature not uncommon in these models. The standard deviations of durables consumption and investment are fairly closely matched, although durables consumption is not quite volatile enough and investment a little too volatile. The correlations with output are well matched and with the inclusion of home production preference the positive co-movement between durables consumption and investment is replicated (see Section 2.8). Consumption is positively correlated with both investment and durables consumption, although a little too correlated with investment and not correlated enough with durables consumption. The model does slightly less well in replicating the moment of the durables and capital stock. However, it is likely that these are less well measured in the data.

2.5.3 Conditional means and varying aggregate states

The aggregate shocks in this model also imply that, unlike in an Aiyagari model, there is no invariant distribution of the agents' holdings of capital and durables. As such the model features state dependency and the history of shocks has implications for aggregate variables. This can be seen by comparing various moments of the model following different histories. More precisely, in a simulation the average of the aggregate shocks over the past 24 quarters is calculated and the sample divided into periods when the aggregate shock has on average been above and below its mean value. These conditional moments are reported in Table 2.7. The moments with respect to the productivity level are largely as to be expected. The capital to output and durables to output ra-

tios are higher following an on average “better” history of productivity shocks. This is due to both, higher levels of investment in response to the greater return and precautionary savings motive, and that there is a lower proportion of unemployed during a positive productivity quarter, with employed agents having higher optimal durables targets. Given greater asset accumulation, the durables share of wealth is slightly lower. Following a bad set of productivity realisation more agents are in debt and more are near the collateral constraint.

The contrasts between the state dependency of the model to collateral constraint shocks are less significant than to productivity shocks. This is because as shown in Section 2.3.9, collateral constraint shocks will not affect those agents with a sufficiently large holdings of financial assets and due to the fact that the idiosyncratic credit terms allow agents to shield themselves from aggregate shocks. However, due to the reasonable proportion of agents in debt in the model differences do arise. The durables to output ratio is marginally higher following a series of looser aggregate credit conditions, as agents are able to fund larger durable purchases. Following tighter aggregate credit conditions agents’ borrowing is restricted and they choose to undertake more precautionary savings. This is clearly seen in the stock of debt of those in debt which falls from 1.118 to 1.030. The average collateralised portion of a durable good for an agent with negative financial assets also declines from 0.231 to 0.215. As a result durables share of wealth is lower. Unsurprisingly, the share of agents near the collateral constraint is higher when in recent history aggregate credit conditions have been tighter.

Finally, the proportion of agents adjusting rises when aggregate credit conditions are tight. The interpretation of this is a little subtle. As shown in Appendix A.2.6, which decomposes the probability of adjustment for agents across all possible states, agents are more likely to adjust when aggregate credit conditions are loose and when idiosyncratic conditions are tight, as this maintains their option value of favourable credit terms. They are also less likely

to adjust when their idiosyncratic credit terms are better than the aggregate conditions. However, after a period of tight aggregate credit conditions, as captured here, a substantial proportion of the agents will have switched to tighter idiosyncratic credit terms, due to eventually deciding to adjust their durables stock. At this stage the later force of increased probability adjustment with tight idiosyncratic credit terms outweighs the former of lower probability due to tight aggregate terms.

2.6 Responses to shocks

This section presents results for the aggregate responses to shocks to the stochastic variables. A simulation of the Great Recession and impulse response functions to a collateral constraint, a small and large productivity shock, and both shocks concurrently are presented. Given the heterogeneity present the model naturally generates a variety of responses to a shock, depending on the distribution of agents, with the model exhibiting rich state dependency. In particular, a longer history of favourable aggregate conditions prior to a credit tightening results in a deeper decline in the aggregate durables stock.

2.6.1 Great Recession experiment

Firstly, the impact of a Great Recession type episode is presented and compared with the response to a normal recession. A normal recession is a decline in productivity from the good state, z^g , to the bad state, z^b , lasting eight periods. For the Great Recession, the timings of the shocks are taken from the data. There is an initial decline in productivity, from the good state to the bad state in 2008.Q1, then a further decline to the very bad state, z^{bb} , which takes place in 2008.Q4. In addition to this in 2008.Q2 the aggregate credit conditions tighten. This is a 25 percent decline of the parameter Θ , from 0.8 to 0.6. The exact sequence of shocks are shown in Figure 2.10.

As can be seen from Figure 2.11, the model is able to capture the magnitudes and alternative behaviour of the non-durables and durables consumption share during a standard recession relative to the Great Recession. The durables consumption share declines more significantly during the Great Recession shock and remains depressed for a considerable period. After eight quarters durables consumption is 8.1 percent below its pre-recession level, whereas following a normal recession it is only 4.3 percent. The additional propagation is most easily visible in the response of the share of agents adjusting, in which it remains depressed for the full ten quarters following the onset of the recession. The data series plotted for percentage of agents adjusting is existing home sales normalised by the civilian population, which provides a high frequency measure of the durables adjustment decision. Finally, the plot for output shows that the productivity shocks being used to generate the responses are reasonable.

2.6.2 Alternative models

The baseline model results also fit the data better than alternative models. This suggests that the specification of the collateral constraint presented in the baseline of the model, as well as being more reasonable in its assumptions, also better explains the behaviour of consumption dynamics during the Great Recession. The top panels of Figure 2.12, presents the durables consumption share and share of agents adjusting. The magenta line shows the response of the variables in an economy where credit conditions affect all agents simultaneously. The dotted line shows the response of the model with an additional productivity shock, but no tightening in the collateral constraint. As can be seen the baseline model out performs both of these specifications. The model with only aggregate credit conditions displays a very strong response on impact of the collateral constraint tightening, but no propagation. After ten

quarters the percentage of agents adjusting their durables stocks has counterfactually returned to the pre-recession level. The addition of the collateral constraint tightening provides an additional decline and propagation of the durables consumption share.

The bottom panels of Figure 2.12, show the cumulative effect of these differences relative to the baseline model. The cumulative effect is important as the baseline model emphasises the propagation of the shock. The red bars indicate the difference relative to the model with no tightening of credit conditions, while the magenta bars show the difference relative to a model with only aggregate credit conditions. For comparison the blue bars show the cumulative difference between a small productivity shock and a large productivity shock (with no tightening of the collateral constraint).

After 10 quarters, the cumulative difference between the baseline effect of the shock and in the absence of a tightening of credit conditions is 21 percentage points. This is 67 percent of the cumulative difference between a small and large productivity shocks. The credit shock is playing a sizeable role in the dynamics. Whilst the cumulative difference between the baseline and a model with aggregate credit conditions is not that large for the durables consumption share, it is sizeable for the percentage of agents adjusting. After 10 quarters this difference is greater than the difference between a small and large productivity shock.

2.6.3 Impulse response decomposition

The previous section showed that the model is able to generate the different consumption dynamics associated with standard recessions and deeper credit constrained recessions such as the Great Recession. However, the precise mechanisms involved can be difficult to disentangle due the sequencing of the shocks. In this section a direct comparison between the shocks is pro-

vided, in a setting where the economy experiences a one time shock, lasting eight quarters. To calculate the impulse response functions, a simulation of 105,000 periods is generated, with the first 5000 periods discarded. Every 50 periods the economy is hit by the chosen shock, which is always negative - the four periods before the shock are set to the high state and the shock moves the economy to the low state for eight quarters. Otherwise, the stochastic variables are held at their pre-shock level.⁸

In Figure 2.13, the dotted blue line shows the response to a small productivity shock, the red line the response to a credit shock, the solid blue line a response to a larger productivity shock, and the black line a response to a concurrent large productivity shock and credit shock. The credit shock has a minimal impact on non-durables consumption relative to a productivity shock. It has a much larger effect on durables consumption. The on impact effect of a drop in the aggregate credit conditions from 0.8 to 0.6 has a similar magnitude to a small negative productivity shock. The larger productivity shock results in a stronger response of durables consumption. On impact, a shock to credit conditions concurrent with a deep productivity shock does not generate an additional decline in durable consumption, but it does offer additional propagation in future periods. As a result the total durables stock declines significantly more during the course of the shock period. The effect of a credit shock on investment is quite different to that of a productivity shock, with a tightening of credit conditions leading to an increase in investment as agents seek to reduce their levels of debt to meet the tigher credit conditions

⁸An impulse response for the shock, j , to the variable, X , is then calculated as the percentage deviation of a variables from the level the period before the shock hits: $IRF_t^j = \frac{X_t^j - X_{SS}^j}{X_{SS}^j}$. The level before the shock is defined as the average over the preceding year, $X_{SS}^j = \frac{1}{4} \sum_{i=-3}^0 X_i^j$. For the capital and durables stock the timing is that the initial level is at the beginning of the period, before the shock hits. Having found the response for each of the j shocks, the median in each time period following the shock is taken as informative of the model's dynamics, e.g. : $\tilde{IRF}_t = \frac{1}{j} \sum_j IRF_t^j$ or $\tilde{IRF}_t := \{IRF_t^j : P(IRF_t \leq IRF_t^j) = 0.5\}$. As the value of the variable in the initial period, $t = 0$, may differ from its average level over the past four quarters, the aggregated series is normalized so that at $t = 0$ the deviation is zero: $\tilde{\tilde{IRF}}_t = \tilde{IRF}_t - |\tilde{IRF}_0|$.

in a future period. Hence a credit shock alone can not account for recession dynamics and standard business cycle correlations in this model.⁹

Agents in debt and close to the collateral constraint are important for the mechanism discussed in this paper. Negative productivity shocks and credit tightening cause a reduction in the percentage of agents in debt. As with the response of durables consumption the combined credit and productivity shocks generates a steep decline in the indebtedness of agents - as the durable stock of households fall they face a tighter collateral constraint and must reduce indebtedness. There are a number of ways to measure the share of agents at the collateral constraint. The share of agents near to the collateral constraint (ncc), as specified by aggregate credit conditions, rises mechanically with the contraction of credit availability. However, as many agents will be able to avoid this tightening, by non-adjustment, the effective collateral constraint (ecc) falls only slowly over time. In contrast to the share of agents near the collateral constraint, the share of agents near the effective collateral constraint ($necc$) falls on impact of a negative credit shock. This is because in the absence of adjustment an agents' terms of credit will not change, but they will seek to reduce their level of debt in anticipation of needing to meet the tighter constraint in the future when they choose to adjust. In the longer run the share of agents near the effective collateral constraint begins to rise as the effective constraint tightens and agents are willing to be closer to this constraint (see Table 2.7).

2.6.4 State dependent responses

The model generates substantial heterogeneity in the responses conditional upon the distribution of durables and assets in the economy, this distribution is an important state variable. One way of seeing this is by comparing

⁹Credit shocks may be able to generate the business cycle correlations if the model was augmented with nominal frictions

the responses of the aggregates conditional upon the history of the exogenous variables. Figure 2.15 compares the effect of a credit shock conditional upon the number of periods prior to the shock that productivity and credit conditions were held at the highest level. This experiment is meant to capture the additional effect of the Great Recession, due to the fact that it occurred following a prolonged period of favourable economic conditions. The fall in the durables stock is increasing in the number of periods prior to the shock that conditions were favourable. Non-durables consumption also experiences a smaller boom, if productivity has been high and credit conditions relaxed for a longer period. The positive response of consumption in response to a credit tightening is due to the additional resources freed up by durables non-adjustment. These resources can be consumed or saved. By consuming a fraction of the resources agents are able to smooth their marginal utility in response to the shock.

A further way of understanding the cause of these state dependent responses is to look at the marginal effect of summary statistics for the distribution of durables and capital on the on impact responses. These are presented in Figure 2.15, with full details in Appendix K.1. The response of durables consumption and the share of agents adjusting is more negative if there has been a high level of adjustment recently, this is due to the fact that agents are closer to their optimal durables stock and therefore in the tradeoff between maintaining credit terms and diverging from the optimal durables holdings, the former becomes more valuable. The more leveraged (as measured by average durables over total wealth) and the higher debt to durables ratio of adjusters, the greater the negative response of durable consumption. This is a form of debt overhang in the model, with higher levels of accrued debt resulting in stronger negative responses. Finally, the higher the effective collateral constraint the more negative the response of durables consumption as the value of non-adjustment is higher when on average the economy is enjoying

looser credit conditions. In general the response of non-durables consumption is in the opposite direction to that of durables consumption for the reasons described above. However, when the debt to durables ratio of adjusters has been high consumption also experiences a more negative response.

2.6.5 Non-linearities

The responses in the model are highly non-linear in the size of the shock and with the marginal effect of a tightening of credit conditions differing depending upon the size of a concurrent shock to productivity. Figure 2.16 presents the impulse response functions to a small and large productivity shocks, when the credit conditions are fixed and when they tighten concurrently. The responses are all rescaled by the initial decline in output to emphasise the non-linearities. The rescaled response of non-durables consumption is almost 50 percent stronger in response to a deep productivity shock relative to a small shock. In comparison once rescaled the durables consumption response is of a similar magnitude in the case of a small or large productivity shock. While the concurrent effect of a credit tightening adds a significant additional decline in durables consumption when occurring at the same time as a small productivity shock, it has only a minor impact when occurring at the same time as a large productivity shock. In the former case the share of agents choosing not to adjust is lower, so there are additional agents that the tightening of credit conditions also forces to postpone their adjustment decision. When the productivity shock is larger, there are fewer additional agents that choose to not adjust for credit reasons. The tighter credit condition do still continue to propagate the shock.

2.6.6 Forecast error variance decomposition

An alternative approach to assessing the importance of the various shocks in this model is to conduct a forecast error variance decomposition. This is a measure of the variation in an aggregate variable at a given horizon that can be attributed to a particular shock. Figure 2.17, presents the decomposition in terms of the share of the variance due to productivity shocks. The left panel presents the decomposition for the baseline model, whilst the right panel present the case with aggregate credit conditions. In the long run productivity shocks are the driving force for many of the key variables. Investment and consumption are almost entirely determined by productivity shocks at all forecast horizons, with collateral constraint shocks playing only a minor role.

The collateral constraint shocks are far more important for durables consumption, accounting for the majority of the variance in the short run and about 40 percent of the variance of the cyclical behaviour in the longer run. For the percentage of agents adjusting, a variable closely related to durables consumption, a similar pattern is observed. Unsurprisingly, the variable that credit shocks are particularly important for is the percentage of agents in debt. Productivity shocks play a much smaller role in the determination of the debt level, accounting for 40 percent of the variance in the cyclical behaviour. Comparing the left to the right panel, it is clear that the history dependence introduced in the baseline model generates a much greater role for credit shocks, than the traditional specification. For example, productivity shocks account for 30 percent more of the variance of durables consumption at the 50 period horizon in the model with aggregate shocks relative to the baseline model.

2.7 Policy experiment

This section sets out a policy experiment in which the government safeguards agents in the economy from variations in the collateral constraint. In particular, I assume that the government fixes the collateral constraint at its highest value - the loosest credit conditions, and taxes agents to fund this subsidy.

This policy experiment replicates features of the UK government's Help to Buy mortgage guarantee scheme. This policy was a response to tighter credit constraints in the UK following the Great Recession and sought to expand the availability of mortgages requiring a small deposit, to a down payment requirement as low as five percent. Such products had largely ceased to be provided by the private sector.¹⁰

More precisely, agents in the economy use the same policies and have the same value functions as in the baseline model. The aggregate credit conditions, Θ_t , varies stochastically. When at its highest level there is no intervention by the government. When the collateral constraint takes its lower value, the government steps in and ensures that all agents can continue to borrow up to the highest collateral constraint. Therefore, from the agent's perspective while ex ante they face uncertainty over the level of Θ_t , and so still value having relaxed idiosyncratic credit terms, ex post they can always borrow up to $k_{t+1} \geq -\Theta^H d_{t+1}$.¹¹

The financial sector is only prepared to provide loans at the exogenously determined aggregate credit conditions, Θ_t . To ensure agents remain unconstrained the government funds any borrowing in excess of the current aggregate

¹⁰Under this policy, which has two components, the government provides loans of up to 20 percent to those buying new-build properties and provides financial institutions with mortgage guarantees of up to 15 percent of the loan value, for mortgages with low down payment requirements (Powley, 2013). Clearly, the model in this paper can only stylistically capture the Help to Buy scheme. For example, as the model does not feature default, rather than guaranteeing future losses, I assume the government instead fully subsidizes loans that exceed the true model consistent collateral constraint.

¹¹The assumption that the policy intervention does not affect the uncertainty faced by agents can be rationalized by assuming a lack of commitment on the government's part that it will intervene during future credit constrained periods.

credit conditions. It does not fund borrowing by agents that agreed their idiosyncratic terms when $\Theta_t = \Theta^H$, but does fund borrowing by agents that benefit from future loose idiosyncratic terms $\zeta_{t+1}^i = \Theta^H$ when the aggregate credit conditions are tight $\Theta_t = \Theta^L$. The government balances its budget every period. Denoting taxes not related to unemployment transfers, T_t , in each period the following condition is satisfied:

$$T_t = \sum_i |k_{t+1}^i + \Theta_t d_{t+1}^i| \cdot \mathbf{1}[k_{t+1}^i < -\zeta_t d_{t+1}^i]$$

To fund the policy, three alternative taxes are considered. To simplify the computation the taxes do not enter the household's value functions and as such expectations are not formed on their evolution.¹² Firstly, a tax on consumption is considered. This is levied after the agents' level of consumption has been chosen, so in effect it reduces the consumption of all agents, but does not effect the agent's choices. Secondly, a proportional tax on the current assets holdings, for agents with positive asset holdings, $k_t > 0$. The timing is that this tax is levied at the beginning of the period, before the capital stock is used in production, assets accrue interest, and before depreciation, i.e. in the budget constraint the tax appear as: $(1 - \delta^k + r_t)(1 - \tau^{pol})k_t$. The third case is a tax on durables. To avoid interfering with the adjustment decision, this is levied on the agent's assets holdings, as in the previous case, but is proportional to the agent's durables stock so the distribution of taxation differs.¹³ When simulating the policy experiment, in each period the tax rate is iterated on to ensure that the government budget constraint holds.

Having calculated the new agent choices, subject to the tax, the implied

¹²To provide a more robust assessment of the policy the model could be resolved with a flexible government tax as a state variable, this could be an avenue for future research.

¹³To ensure the maximum collateral constraint is not broken, agents who's asset position is such that paying the tax would mean that they exceed the collateral constraint are not taxed.

value of each agent is recorded, based on the choices made:

$$V_t^{i,pol} = u(c_t^i, H(d_{t+1}^i, \varsigma(z_t, \varepsilon_t)(T^{MA} - l_t)) + \beta^i E_t V(k_{t+1}^i, d_{t+1}^i, \beta_{t+1}^i, \varepsilon_{t+1}^i, \Xi_{t+1}^i; \Theta_{t+1}))$$

Using the same series of aggregate and idiosyncratic shocks, the mean value across agents and across time can then be taken as a measure of welfare, $W = \frac{1}{NT} \sum_{t,i} V_t^{i,pol}$ and used for policy comparison purposes.¹⁴

The major mechanisms that prevail in this policy experiment are: 1.) the government reduces market incompleteness, by relaxing the exogenous collateral constraint in all periods of the economy in which credit is more tightly constrained; 2) there is a reduction in volatility along the credit dimension and this influences the frequency of adjustment. Recall positive collateral constraint shocks are associated with increases in the percentage of agents adjusting. Given that adjustment results in the loss of real resources, this could be an efficiency gain; 3) implementing the policy requires taxation. While the taxes considered are lump sum and do not alter incentives they do lower welfare by reducing consumption or financial asset holdings.

The results of the policy experiment are presented in Table 2.8. The key finding is that intervening to relax credit conditions in the economy on average reduces welfare, although the decline is small. For all of the taxes implemented the average agent value is lower under intervention, than without. Hence, offering additional market completeness is not offset by the loss in consumption or wealth. It can be seen that in all cases the tax rate required to implement the policy is fairly small, for the asset and durables tax the rate is less than 0.1

¹⁴Two subtleties of the set up with regard to the capital stock are worth stating. Firstly, as the government subsidy is a transfer to an agent exceeding the period collateral constraint, this borrowing is not also cross-financed by the claims of the economy's savers. Holding all else fixed this implies a slightly larger capital stock, which is now calculated: $K_{t+1} = \sum_i k_{t+1}^i \cdot \mathbf{1}[k_{t+1}^i \geq -\zeta_t d_{t+1}^i]$. This is not true of additional borrowing driven by a lower precautionary savings motive that does not exceed the true period collateral constraint, ζ_t . Secondly, the timing of the assets and durables tax, both reduce the capital stock. Both of these mechanisms have consequences for prices and are not fully consistent with the agent's expectations, captured in the aggregate law of motion.

percent. While the policy does not improve average welfare it does increase weighted average welfare, where the weights used are the marginal utility of consumption. The policy provides the greatest benefits to those that are credit constrained and have a high marginal utility.

Details of other aggregates in the economy are provided to give understanding of the outcome of the policy experiment. In the asset tax and durables tax case the total capital stock in the economy is reduced, this is a direct result of the taxation mechanism and accounts for the larger falls in the welfare measure. In these economies lower assets reduce consumption. In the consumption tax example, the welfare fall is instead the result of the direct reduction in consumption. In line with the state conditional moments in Table 2.7, the percentage adjusting is lower when the collateral constraint is permanently set to its higher value and as such fewer resources are wasted on adjustment costs. The policy encourages agents to take on more debt, this is most clearly observed in the consumption tax case, where no intervention takes place on agents' net financial assets.

The reason that the durables tax is the least effective way of implementing the policy is that the burden is shared more evenly across the distribution and falls on those with a higher marginal utility of wealth. Note from Table 2.8, that while the mean values of assets and durables are respectively 32.1 and 5.7, the median values are 12.7 and 5.4. A tax proportional to assets has far smaller impact on the available resources of the median agent than a tax proportional to their durables holdings. Further, the durables tax is also levied on those with a negative net financial assets position and these agents have a higher marginal utility of wealth.

A representative time path of the deviations of the average agent's value from that under the baseline model with no intervention is presented in Figure 2.18. As can be seen the deviations vary over time significantly. In the case of the asset and consumption tax there are many periods in which the

average agent value converges to the no subsidy case. It can be seen that the greatest negative deviations occur following a series of low collateral constraint realisations - for example, around period 900. After this series of shocks the increased requirement for the collateral constraint subsidy negatively impacts on the welfare of the average agent. The extended period of tight credit conditions leads more agents to optimise into the position where the subsidy is required. The aggregate productivity state is much less important, as it does not directly impact whether there is policy intervention in the quarter or not.

2.8 Robustness

2.8.1 Business cycle moments of alternative specifications

This section reviews the model's sensitivity to certain parameter choices and highlights important features of the baseline model. First compare the model to an alternative specification, more commonly used, where all agents take the same aggregate collateral constraint (model 1), henceforth referred to as the "aggregate" model. Most of the ratio and levels of this model are shared by the baseline specification. The model without idiosyncratic credit terms features a higher correlation of durables consumption and output and co-movement between durables consumption and investment. This is due to the propagation mechanism generated by the idiosyncratic credit terms, which makes changes in the aggregate credit conditions more important for durables consumption.

Turning to the other parameter specifications. A model with low adjustment costs (2) features more indebted households, as the cost of being indebted is lower, and a higher standard deviation of durables consumption. As there is a weaker distinction between adjusters and non-adjusters, with agent adjusting more regularly, the LTV distance is also negative. The importance of

heterogeneous discount factors for implementing a realistic wealth distribution is demonstrated by model (3), which feature a low share of agents in debt and near the collateral constraint. Finally, if standard Cobb-Douglas preference are used (model 4), the model feature a negative co-movement between durables consumption and investment. Whereas, a greater degree of complementarity between durables and leisure time would generate a higher correlation between durables consumption and investment.

2.9 Conclusion

This chapter has set out a model capable of addressing the question: “what role do time varying consumer credit conditions play in the consumption dynamics of non-durable and durable goods?” It is the first analysis to do so in a fully heterogeneous agents setting with production, suitable for the analysis of business cycle fluctuations.

The model shows that a tightening of credit conditions leads to a prolonged decline in durables consumption, and a decrease in the proportion of agents adjusting their durables stock. The setting also introduces a powerful endogenous propagation mechanism, whereby the credit terms available to agents depends on their adjustment decision. This was shown to generate an empirically realistic correlation between the different behaviour of adjusters and non-adjusters and aggregate credit conditions, which a simpler model featuring only aggregate credit shocks does not replicate. A combination of a negative productivity shock and a collateral constraint shock was proposed as a reasonable explanation for the sharp and prolonged decline in the share of durables consumption following the Great Recession, with the concurrent incidence of these shocks generating non-linear responses.

A government policy to relax credit conditions is shown to lower average agent welfare, with the costs of raising the subsidy required through taxation,

exceeding the benefits of reducing market incompleteness. The welfare losses are small reflecting the low proportion of individuals that take on highly indebted positions in this economy. Considering the welfare of low asset, credit constrained agents, on average welfare weighted by the marginal utility of consumption was increased by the policy.

Further research will seek to provide additional evidence in the micro data for the specification of the collateral constraint proposed here. It would also be interesting to introduce nominal-rigidities, whereby the fall in durables consumption demand could generate its own output effect. However, such an extension lies beyond the scope of the current chapter.

2.10 Tables

	Ξ_{t+1}^L	Ξ_{t+1}^H
Ξ_t^L	0.9467	0.0533
Ξ_t^H	0.0431	0.9569

Table 2.1: Stochastic process for the collateral constraint

Description	Parameter	Value	Source
Non-durables share	θ	0.7962	Target ratio C/Y and D/C
Inverse E.I.S	α	2.0	
Durables share	ω	0.816	Target ratio C/Y and D/C
EIS (home production)	λ	-1	Greenwood and Hercowitz (1991)
Total hours	T^{MAX}	3.228	Current Employment Statistics
Durables depreciation	δ^d	0.0484	NIPA
Adjustment cost	Ψ	0.0525	Berger and Vavra (2015)
Maintenance parameter	χ	0.8	Berger and Vavra (2015)
Capital share	η	0.36	
Capital depreciation	δ^k	0.025	NIPA
Productivity level	$z \in \{z^{BB}, z^B, z^G\}$	(0.98,0.99,0.01)	Den Haan et al. (2010)
Unemployment rate	$u \in \{u^{BB}, u^B, u^G\}$	(0.1,0.07,0.04)	Den Haan et al. (2010)
Home prod. id. productivity	ϑ	0.4271	Krueger and Mueller (2012)
Labour normalisation	\bar{l}	1/0.9	Den Haan et al. (2010)
Benefits replacement rate	μ	0.4	Shimer (2005)
Collateral constraint	$\Xi \in \{\Xi^L, \Xi^H\}$	(0.6,0.8)	Guerrieri and Lorenzoni (2017)
Persistence of credit terms	ρ^Ξ	0.99	American Housing Survey

Table 2.2: Calibration of the model, Preference and technology parameters

Description	Parameter	Value	Source
Discount factor (high)	β^H	0.9892	Quarterly interest rate $r=0.015$
Discount factor gap	β^{GAP}	0.0158	Share of agents in debt ($k < 0$) = 0.2
Share of low β types	$F(\beta^L)$	0.5	
Persistence of discount rate	ρ^β	0.995	Implied dynasty duration: 50 years

Table 2.3: Calibration of the model, wealth parameters

Statistic	Baseline	Agg CC	Data
$\rho(E[\Omega^A] - E[\Omega^N], \Theta)$	0.32	-0.37	0.25
$\rho(\%Adj., \Theta)$	0.13	-0.31	0.45
$\rho(E[\Omega^A], E[\Omega^N])$	0.50	0.35	0.89
$\frac{\sigma(\log(\Omega^A))}{\sigma(\log(\Omega^N))}$	1.93	1.26	0.84

Note: $\Omega^A = \frac{|k|}{d}|Adj., k < 0$ and $\Omega^N = \frac{|k|}{d}|NAdj., k < 0$

Table 2.4: Co-movement of credit variables

Levels		Ratios & Debt			Debt and Collat.Const.		
Variable	Value	Variable	Value	Data	Variable	Value	Data
K	32.1	$\frac{K}{Y}$	9.0	5.9	% $k < 0$	0.20	0.20
Med. k	12.7	$\frac{C}{Y}$	0.7	0.7	$k/h k < 0$	0.216	0.44
D	5.7	$\frac{D}{Y}$	1.6	1.7	$\frac{d}{d+k}$	0.552	0.78
Med. d	5.4	$\frac{D}{C}$	2.3	2.3	% near C.C	0.074	
			% adj.	0.027	% at C.C	0.003	

LTV distance is a measure of the correlation between the LTV ratios of adjusters and non-adjusters and aggregate credit conditions

$$\Upsilon = \log(E[\frac{k}{h}|adj. k < 0]) - \log(E[\frac{k}{h}|nadj. k < 0]).$$

Table 2.5: Baseline model long-run averages

	$\sigma(x)/\sigma(y)$		$corr(x, y)$		$corr(x, I^k)$		$corr(x, I^d)$	
	model	data	model	data	model	data	model	data
C	0.22	0.55	0.83	0.78	0.73	0.56	0.49	0.59
I^d	2.38	2.94	0.62	0.69	0.42	0.52	1.00	1.00
I^k	3.40	4.82	0.97	0.83	1.00	1.00	0.42	0.52
D	0.26	0.70	0.02	0.21	-0.05	0.10	-0.11	0.03
K	0.22	0.70	-0.02	-0.11	-0.13	-0.04	-0.08	-0.24

As with empirical data all statistics are logged and HP filtered with smoothing parameter 1600.

Table 2.6: Business cycle moments

	$z < z^*$	$z > z^*$	$\Theta < \Theta^*$	$\Theta > \Theta^*$
K/Y	8.969	9.068	9.021	9.022
D/Y	1.604	1.607	1.603	1.607
$\frac{d}{d+k}$	0.571	0.534	0.539	0.563
$sk(k)$	2.278	2.240	2.262	2.255
$sk(d)$	1.301	1.357	1.298	1.360
<i>per. adj.</i> (%)	0.028	0.026	0.029	0.026
% $k < 0$ (%)	0.209	0.193	0.191	0.209
$k k < 0/Y$	0.069	0.055	0.056	0.067
$k k < 0$	1.147	1.026	1.036	1.126
$k/h k < 0$	0.243	0.208	0.216	0.232
% at <i>C.C.</i> (%)	0.007	0.004	0.005	0.005
% near <i>C.C.</i> (%)	0.086	0.064	0.084	0.066

Table shows average model outcomes conditional on aggregate state history. History is defined to be the average of the state in the previous 24 quarters.

Table 2.7: State conditional first moments

	No subsidy	Consumption tax	Asset tax	Durables tax
Ave. agent value (V_i)	38.4472	38.4428	38.4324	38.4239
Ave. st.dev(value) ($\sigma(V_i)$)	13.0612	13.0711	13.0694	13.0703
W. Ave. agent val ($\varpi \cdot V$)*	4.9259	4.9375	4.9390	4.9456
Ave. st.dev(w.value) ($\sigma(\varpi \cdot V_i)$)	2.1796	2.1776	2.1769	2.1858
Tax rate (τ)**	0.000	0.1066	0.0082	0.0522
Other aggregates				
K	32.2734	32.2727	32.1655	32.1895
C	2.4791	2.4758	2.4748	2.4746
H	5.7127	5.7270	5.7169	5.7165
<i>per. adj.</i>	0.0257	0.0254	0.0254	0.0254
<i>per. k < 0</i>	0.1904	0.2102	0.2097	0.2148
$\frac{d}{w}$	0.5395	0.5601	0.5598	0.5654

*Weight for weighted average value is marginal utility of consumption:

$\varpi = u_c(c, H\{d', \varsigma(z, \varepsilon)(T^{MAX} - l)\})$ ** Tax rate presented in terms of rate i.e. $\tau \in [0, 100]$

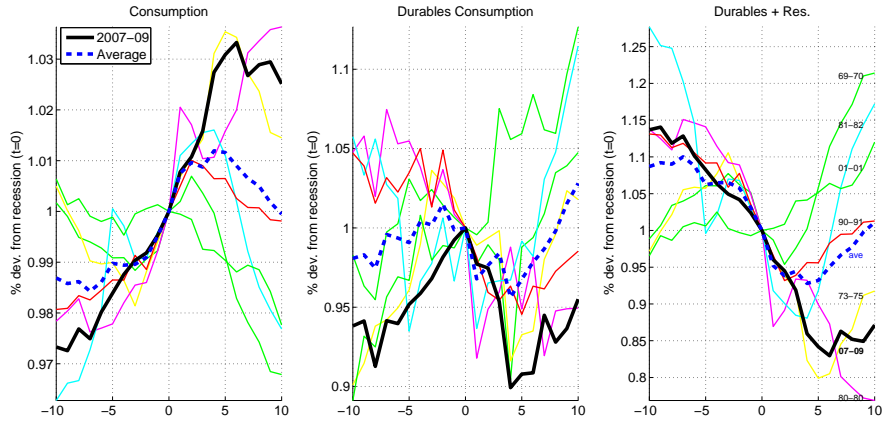
Table 2.8: Policy experiment: Relaxing credit conditions

Variable	Model						
	Baseline	1 $\zeta_t^i = \Theta_t$	2 $\Psi = 0.0005$	3 $\bar{\beta}$	4 $\omega = 1$	5 $\lambda = -2$	6 RBC
$\frac{K}{Y}$	9.021	9.023	9.001	9.049	8.976	9.037	8.992
$\frac{D}{C}$	2.310	2.310	2.414	2.361	3.709	1.783	2.489
$\frac{d}{d+k}$	0.552	0.540	0.640	0.202	0.708	0.482	0.161
% $k < 0$	0.200	0.191	0.322	0.003	0.326	0.122	-
% <i>adj.</i>	0.027	0.028	0.186	0.025	0.027	0.031	-
% <i>near C.C</i>	0.075	0.070	0.132	0.001	0.086	0.072	-
$\sigma(C)/\sigma(Y)$	0.223	0.223	0.215	0.122	0.280	0.499	0.268
$\sigma(I^d)/\sigma(Y)$	2.376	1.993	4.252	1.925	3.726	3.832	10.967
$\sigma(I^k)/\sigma(Y)$	3.399	3.331	3.024	4.138	4.708	3.197	5.541
$corr(C, Y)$	0.827	0.770	0.797	0.456	0.924	0.742	0.643
$corr(I^d, Y)$	0.619	0.807	0.656	0.302	-0.021	0.692	0.129
$corr(I^k, Y)$	0.965	0.949	0.949	0.977	0.803	0.973	0.666
$corr(I^k, I^d)$	0.422	0.743	0.407	0.110	-0.601	0.535	-0.628
$\rho(\Upsilon, \Theta)$	0.110	-0.323	-0.210	0.079	-0.032	0.187	-

*Weight for weighted average value is marginal utility of consumption:
 $\varpi = u_c(c, H\{d', \zeta(z, \varepsilon)(T^{MAX} - l)\})$ ** Tax rate presented in terms of rate i.e. $\tau \in [0, 100]$

Table 2.9: Policy experiment: Relaxing credit conditions

2.11 Figures



Source: National Income and Product Accounts

Figure 2.1: Consumption shares during recessions, 1967.II-2010.II




Figure 2.2: Decline in durables transaction volumes




Figure 2.3: Comparison of FRB's Senior Loan Officer Opinion Survey credit variables, 1990.III-2014.III

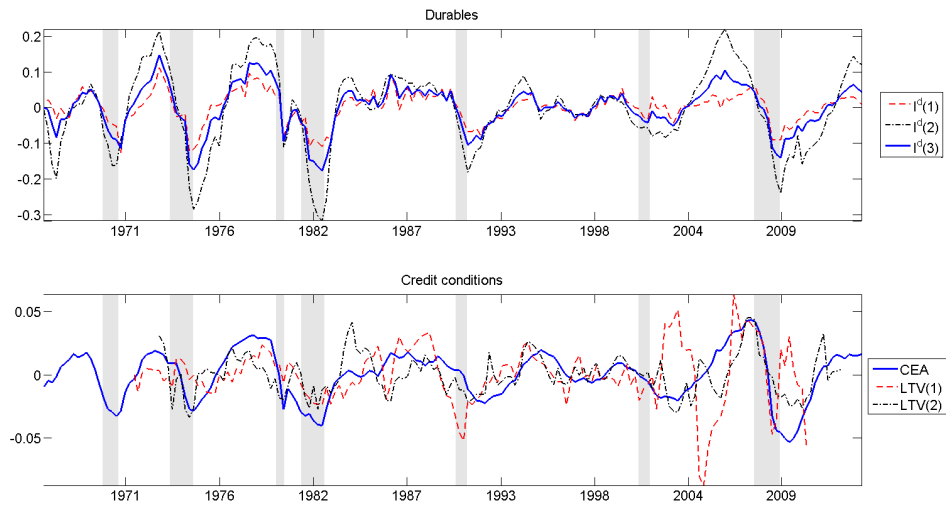




Figure 2.4: HP filtered aggregate series: US data 1966.III-2014.I




Source: Federal Reserve Board of New York, Consumer Credit Panel / Equifax

Figure 2.5: Growth and availability of Home equity line of credit mortgages




Source: Mortgage and housing data from the PSID, 1976-2013. Series have been linearly detrend. Value of housing deflated using the HPI, mortgage debt deflated using CPI. HP filtered CEA series as described above

Figure 2.6: Mortgage to housing stock ratio of movers and non-movers




As a result of the solution method (see Section A.2.2) Assets in hand are defined $a_t = k_t + \Xi^H d_t$. The asset levels are: low $a=0.04$, medium $a=1.15$, high $a=36.57$. The policies are shown for an unemployed agent in a good productivity state. Agents credit terms are set to same level as aggregate collateral constraint.

Figure 2.7: Durables policy functions under low and high collateral constraints



As a result of the solution method (see Section A.2.2) Assets in hand are defined $a_t = k_t + \Xi^H d_t$. The asset levels are: low $a=0.04$, medium $a=1.15$, high $a=36.57$. The policies are shown for an unemployed agent in a good productivity state. Agents credit terms are set to same level as aggregate collateral constraint.

Figure 2.8: Durables policy functions under a drop in aggregate credit conditions



Empirical data has been rescaled to plot against model. For durables empirical data has been rescaled to have same mean and standard deviation as model. For capital the standard deviation is rescaled. For both series, the density rescaled such that maximum density set equal to that of model maximum density for comparability. For raw empirical series see Figures C.14 to C.16 in Appendix C.2.

Figure 2.9: Baseline model distributions of durables and capital




Figure 2.10: Great Recession experiment shocks

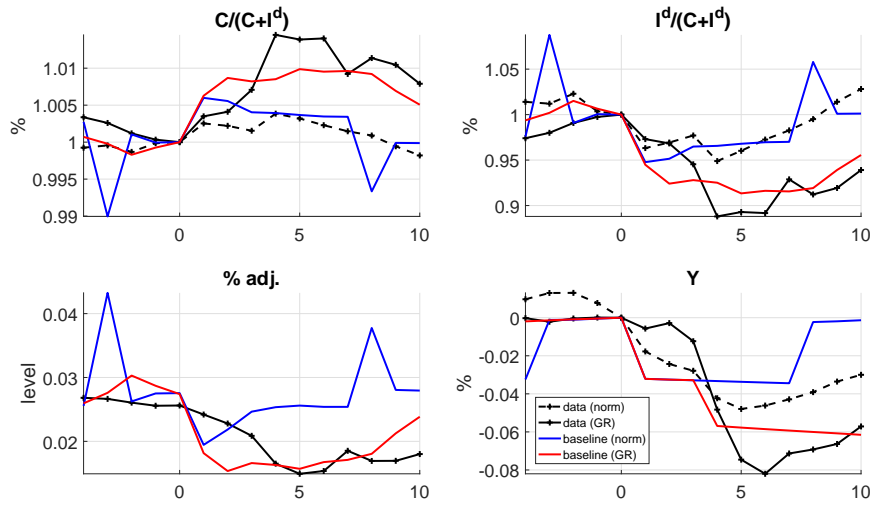


Figure 2.11: Great Recession experiment

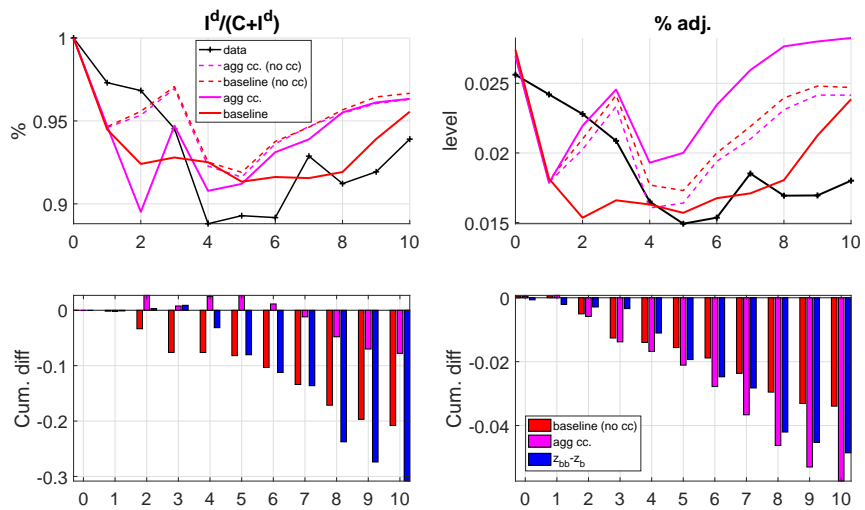
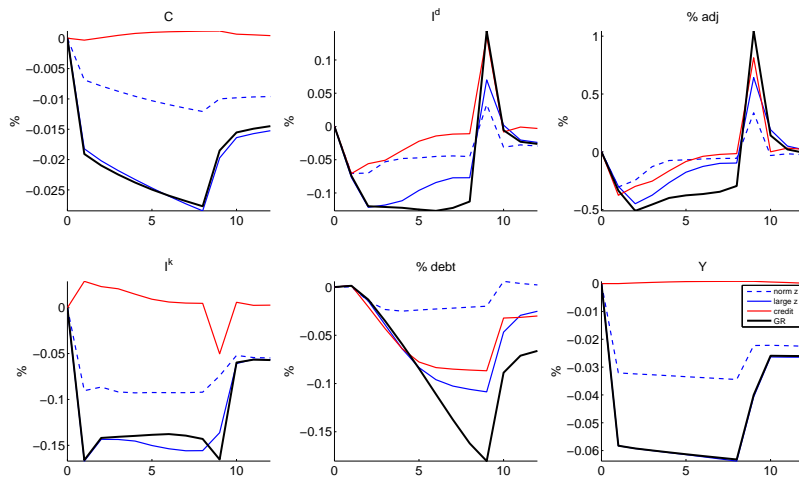
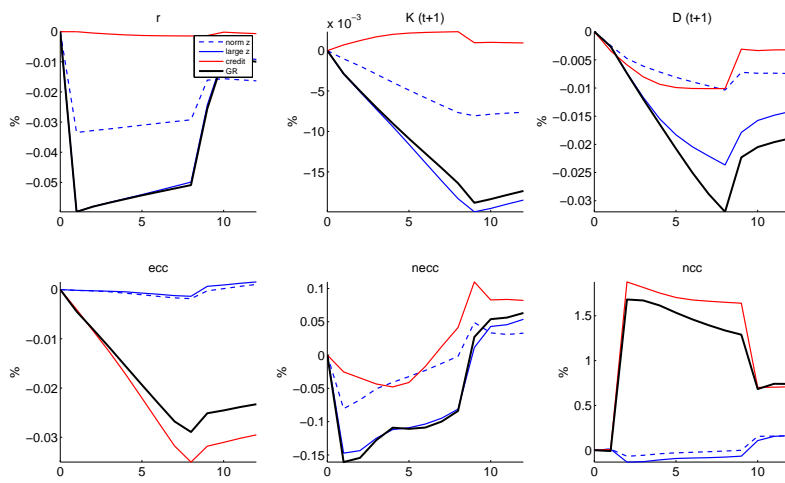


Figure 2.12: Great Recession experiment: Comparator models



(a) Quantities



(b) Credit variables

Figure 2.13: Impulse response decomposition

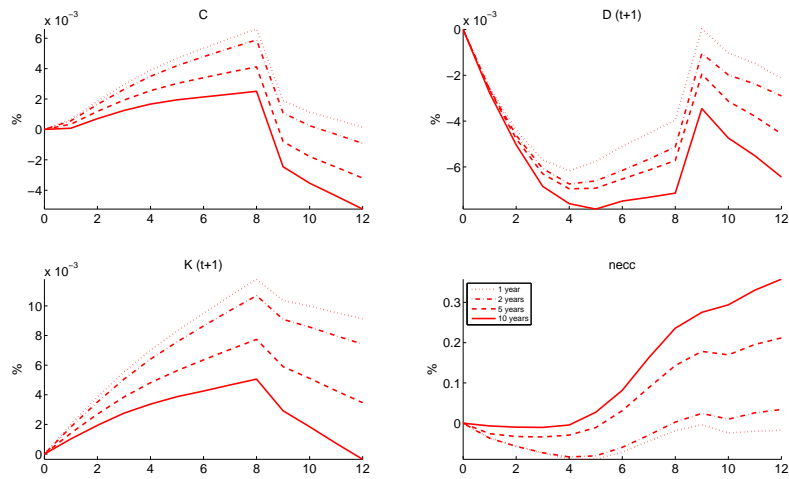



Figure 2.14: Impulse response: credit shock under prior conditions



Notes: Marginal effect is the second period IRF. Previous state is defined as average value over past 16 quarters prior to shock. The marginal effect is conditional upon all other state. Such an approach is necessary as there is a strong correlation between these conditions, which can mask the effect of an individual feature of the distribution on the responses. Precisely, let $Y = X\beta + \epsilon$, be a regression of IRF impact, Y , on state vector, X , where X is de-meaned and standard deviation normalised to one. Marginal IRF is $\hat{Y} = Y - \tilde{X}\hat{\beta}$, where \tilde{X} is vector X with conditioning state, x , set to zero. Plotted is \hat{Y} against x . Solid black lines indicate statistically significant relationships. Full regression values, $\hat{\beta}$, are presented in Table 17 of Appendix L.1.

Figure 2.15: State dependence

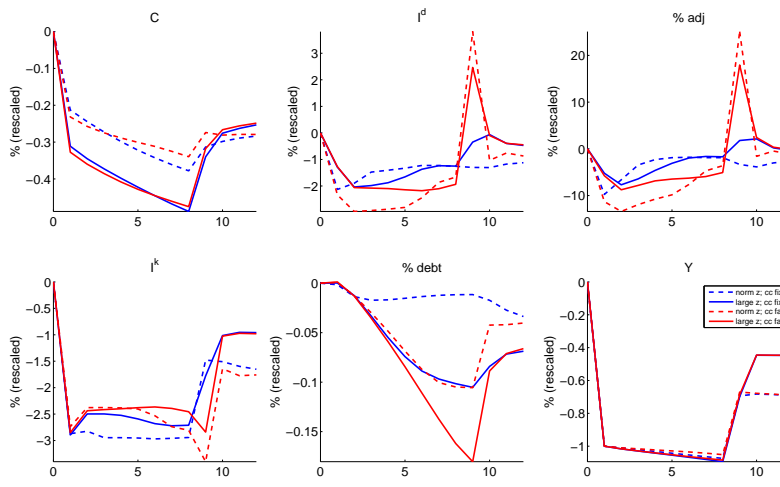


Figure 2.16: Impulse response: Non-linear response

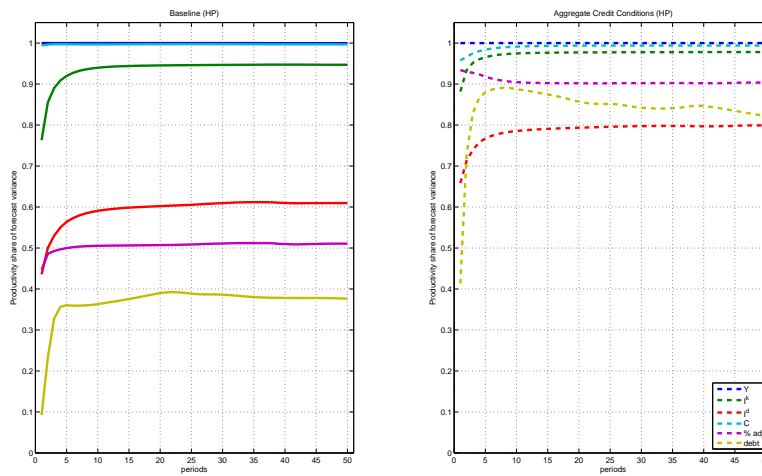




Figure 2.17: Forecast variance decomposition



(a) Deviations of agents' average value



(b) Aggregate State

Figure 2.18: Policy experiment: relaxing credit conditions

Chapter 3

(S)Cars and the Great Recession

3.1 Introduction

The financial crisis and Great Recession that followed hit the US economy and individual households hard.¹ It saw a persistent and deep decline in aggregate activity, substantial downward revisions of consumption, a significant deterioration in the state of the labor market and a large decline in house prices. Unsurprisingly, the causes of the Great Recession have been a significant focus of research in the past decade. New macroeconomic theory has been developed with the aim of addressing many of the issues brought up by this recession and how best to understand and model household behaviour.² In this paper, we study household level data with the aim of understanding aggregate outcomes. We fix attention on household consumption choices and contrast empirical evidence on how households adjusted their spending patterns with the implications of an augmented life-cycle consumption model. A

¹This chapter is based on work in a joint paper with Morten Ravn, Orazio Attanasio and Mario Padula

²This is an extensive literature. Notable examples include: [Hall \(2011\)](#), [Stock and Watson \(2012\)](#), [Christiano et al. \(2015\)](#), [Krueger et al. \(2016\)](#), [Kaplan et al. \(2017\)](#) and [Ravn and Sterk \(2017\)](#)

distinguishing feature of our analysis is that we pay close attention to consumer durables and work with a model with a richer specification of income dynamics than usually studied in the literature. In keeping with a strong economic tradition, we claim consumption dynamics can be highly informative of the unobserved shocks hitting households (for example: [Blundell and Preston \(1998\)](#), [Blundell et al. \(2008\)](#), [Blundell et al. \(2013\)](#), [Heathcote et al. \(2014\)](#) and [Olivi \(2019\)](#)). However, departing from this literature we make use of the additional information that can be extracted by distinguishing between the choices households make at the intensive and extensive margins of durable purchases. To this end, we expose households in the model to a wide variety of shocks which are all consistent with the data and ask which of these generates consumption choices that provide the best fit with the empirical patterns observed. We argue that asset price changes, cohort-specific income shocks and a change to the expected growth rate of future income are important for understanding the consumption dynamics during the Great Recession.

We first document patterns of household consumption adjustments over the life-cycle and during the Great Recession. To do this, we study household consumption data collected from the Consumer Expenditure Survey (CEX) produced by the Bureau of Labor Statistics. The CEX is particularly useful for our purposes for three main reasons. Firstly, it contains rich consumption expenditure data at the household level including expenditures on a large ticket consumer durable, cars. Aggregated over households, the CEX consumption data also match the National Income and Product Account (NIPA) consumption data quite well, especially for new car purchases. Secondly, the CEX data has been collected annually since 1984 making it useful for studying time-series as well as life-cycle dynamics. Thirdly, although individual households are only interviewed for four consecutive quarters, the cross-sectional dimension is sufficiently rich that the data can be used for cohort analysis. We use the CEX data to build a synthetic panel by grouping households by

year of birth. There is a strong life-cycle aspect of household non-durables consumption expenditure. Household expenditures on non-durables rise steeply with age for households in their 20s, increases further but at a slower rate from the early 30s until peaking when households reach their mid-40s. After that non-durables expenditures decline slowly. There is also a clear life-cycle component of household spending on cars. Around one third to a quarter of households in their mid-20s do not own a car while the corresponding number for prime-aged households is only around 7-10 percent. The household stock of cars rises slowly for young households and peaks for households in their early 50s. The maximum value of a household's car stock occurs at a later stage of the life-cycle than their maximum non-durables consumption. These patterns are remarkably consistent for different birth cohorts.

Consumption data displays strong cyclical patterns, variation which is stronger for consumer durables and in particular for cars purchases. Cars take a disproportionately large share of the cut in total consumption expenditures during a recession and the spending share on cars is procyclical in the household data. The standard deviation of motor vehicles expenditure is five time larger than the standard deviation of non-durables consumption. Indeed, as Chair of the Federal Reserve, Alan Greenspan was renowned for having kept close track of the car market ([Greenspan and Cohen \(1999\)](#)). The reasoning is that consumer choices regarding automotive expenditure can be very revealing for the state of the economy. Firstly, household spending on cars accounts for a large share of aggregate household spending on durable goods. Secondly, choices about car expenditure carries important information on the source of shocks affecting households. The durability of cars implies expectations about future income is particularly important when choosing whether to buy a car. Moreover, car purchases are large ticket items with significant adjustment costs which in combination imply (in)sensitivity to (small) large changes in permanent income. The existence of adjustment costs also means

that car spending decisions can be further decomposed into an extensive and intensive margin response. Households must choose whether to buy a car and conditional upon the decision to purchase households then choose the size or value of car to buy (Bertola and Caballero (1990)). These choices contain additional information about the state of the economy. The decision to respond on either margin will depend on both the shock hitting the households and the households' income, wealth and position in the life-cycle.

While aggregate car expenditure is typically volatile, as described above, the household level patterns during the Great Recession were particularly interesting. These features can be obscured by focusing on aggregate data. We document that while most recessions are associated with a decline in the probability of purchasing a car (an extensive margin response), the Great Recession in addition saw a large decline in the intensive margin response. Conditional on choosing to buy a car, households that did purchase bought a smaller car. In 'normal' recessions, the households that choose to purchase a new car in the midst of a recession, typically spend the same on the car as households that purchased a car just prior to the recession. This indicates either that a recession has minor impact on the optimal car size, which would be consistent with the downturn in the economy having a small effect on permanent income, or with some households being fairly unaffected by the recession. By contrast, in the Great Recession we find that households that purchased a car during the recession spent less than those that purchased just prior to the crisis. Further, the decision as to whether to purchase a car also differed in the Great Recession relative to previous episodes. Indeed, there was a larger decline in the probability of adjustment given the underlying economic conditions than is usually observed. Combined, this evidence suggests that the shocks hitting during the Great Recession were manifestly different to those experienced in previous downturns.

Important differences are also noticeable when we look at the pattern of

responses across the life-cycle. In the data we see that the response of car expenditure was larger for younger and middle-aged households, primarily driven by a low extensive margin response of older households. In addition, prime working aged households exhibited the largest response along the intensive margin. These varying responses are potentially informative for understanding the root causes of the crisis. Changes in the cost or availability of credit, for example, impact mainly on indebted households or households planning to borrow. Income shocks differ across households both due to idiosyncratic reasons but also because some households – e.g. retirees – may not be participating in the labor market, have a shorter planning horizon or as human capital makes up a smaller fraction of their lifetime wealth. Similarly, shocks to asset prices will affect those households with larger stocks of certain forms of wealth.

We interpret the consumption patterns in the CEX data through the lens of a partial equilibrium life-cycle model. Households have finite planning horizons, face a stochastic income process, and consume non-durable goods and the service stream from cars. We incorporate two salient features of car expenditure. Firstly, households face non-convex adjustment costs and therefore adjust their car stock in a discrete manner over time as their cars depreciate and in response to large shocks to their income stream. Secondly, while households cannot issue unsecured debt, they have access to car credit which we specify as collateralized loans. The availability of car credit allows households to purchase cars but because the price of car credit is high, only young households who have little savings and households who have exhausted their savings use this type of credit.

We allow for transitory and permanent income shocks and impose that the permanent shocks are composed of both an idiosyncratic and aggregate component. The partial equilibrium nature of the model means that households do not need to distinguish between idiosyncratic or aggregate shocks when solv-

ing their optimal problem, but the latter allows us to model business cycles. Motivated by the empirical observation that the response of car expenditure was different during the Great Recession, we also introduce a range of other shocks that have been commonly associated with the crisis. Firstly, we introduce fluctuations in car credit conditions to capture the fact that car credit spreads increased during the Great Recession ([Johnson et al. \(2014\)](#)). Secondly, we introduce an asset price shock to capture the large fall in house prices seen during the period ([Mian et al. \(2013\)](#)). Finally, we introduce a novel shock to the expected growth rate of future income, with the idea being that younger households perceived the Great Recession as additionally causing a deterioration in their long-run prospects. We show the data is broadly supportive of such a phenomena. Despite its salience we do not explicitly model the “Cash for Clunkers” policy. This is due to it being active for a relatively short period of time (two months) compared with the annual frequency of our model. The evidence largely suggests that the main effect of the program was to shift consumption over time, so we do not consider its absence critical ([Mian and Sufi \(2012\)](#)).

Whilst providing a good account of the previous recession dynamics and despite their significant size, aggregate income shocks alone are unable to generate the magnitude of the response of car expenditures seen during the Great Recession. Uniform income shocks also fail to adequately capture the heterogeneity of responses across the life-cycle. Introducing cohort specific shocks provides a better account of this variation, but as the magnitude of the shocks is the same the aggregate responses are largely unchanged. Somewhat surprisingly changes in the car loan interest rate spread adds relatively little to the understanding of the consumption response, although it does impact saving and loan decisions. The intuition is that due to strong life-cycle and precautionary savings motives, and in contrast to mortgage financing, relatively few households are dependent of car loan financing and those that are by and large

purchase smaller cars and have lower non-durables consumption, representing a small share of the aggregate. Hence changes in the cost of credit have second order effect on spending aggregates.

The introduction of asset price changes and a shock to the life-cycle growth rate, when in combination with the large permanent income shock, provides a much fuller explanation of the consumption dynamics. The inclusion of these shocks allows us to capture both the extensive and intensive margin durables response in the aggregate and the variation in expenditure across the life-cycle. Firstly, the house price boom helps us to match the growth in consumption just prior to the recession, when the labor market was beginning to slow. During the crisis the fall in house prices generates a significant decline in the wealth of households, particularly concentrated on those that make up a larger share of aggregate consumption. With lower cash in hand, household's planning to adjust now prefer a much smaller car, especially those with a shorter planning horizon. This has a strong effect on the intensive margin. The combination of the income shocks and wealth shock is particularly strong for middle aged households, that exhibit a large intensive margin response in the data.

Young households that hold relatively low levels of wealth and in the CEX experience a relatively mild income shocks are still observed in the data to respond strongly especially on the extensive margin. The introduction of the shock to the growth rate of future income hits these households by reducing permanent income and shifting them away from their adjustment point. With income expected to rise less quickly, constrained households also reduce the frequency with which they adjust their car stock. The inclusion of this shock enables us to capture the variation in responses across the life-cycle. Finally, a decline in the expected growth rate of future income is consistent with the large savings response seen by this cohort. We provide empirical evidence from the income data which is supportive of this type of shock. [Dupor et al. \(2019\)](#) also analyzes the collapse in the 2008 auto market, attributing somewhat less

of the decline to the house price falls and a greater role for the oil price shock.³ The structure of the paper is as follow: section 3.2 outlines the salient features of the car expenditure data at both the aggregate and household level; section 3.3 sets out the model; section 3.4 discusses the calibration; section 3.5 presents the policy functions and model fit across the life-cycle; Section 3.6 uses the model to investigate the shocks hitting households during the Great Recession; while Section 3.7 concludes.

3.2 Consumption dynamics and the Great Recession

3.2.1 The data

We examine the dynamics of households' consumption expenditures on non-durable goods and their spending on cars. The data on household consumption expenditures and characteristics are obtained from the CEX. The CEX has received a considerable amount of attention because it is the only US data set with detailed information on consumption expenditure. While the Survey has been criticized for not aggregating up to the NIPA statistics, it has become clear that when concentrating upon comparable items and populations, changes in the CEX are mirrored in changes in the NIPA data. For some items in particular (see [Garner et al. \(2009\)](#)), the CEX seems to be particularly good at reproducing the NIPA data including the expenditure on cars a feature that may follow from the salience of car purchases for most consumers.

We find it informative to examine car expenditures because of their durability. In addition, the CEX contains reliable information not only on household

³We experimented with changes in the gasoline price, but concluded a relatively small role for these shocks in the recession dynamics we are interested in. While prices spiked in 2008 they fell precipitously in 2009. If the elasticity was large the decline in 2009 implies a boom in the value of cars purchased which is strongly at odds with the data.

car purchases, but, since the mid 1980s, also on the stock of cars owned by each individual household. For each car owned by the household, the survey provides the make, model and year of the car owned, as well several characteristics of the car.⁴ Moreover, the price of cars (new or used) in the 12 months preceding the interview is also available. Finally, the CEX provides a considerable amount of detail on the financing of the car purchase, including the origin and terms of the car loan, interest rates and maturity of the loan. The survey also records information on cars sold.

We use the information contained in the CEX to estimate the value of the cars owned at the household level. We can analyze the extent to which households actively adjust their car stock by buying and/or selling a car and by how much the stock changes. Finally, we can relate this behavior both to aggregate events (such as the level of economic activity, including the recession) and to individual characteristics (such as by age and the relationship to non-durable consumption expenditure).

We will later consider a life-cycle model in which households purchase both cars and non durable consumption goods. Therefore, we find it useful to consider the individual dynamics of non durable consumption and car expenditure explicitly. As the CEX survey has only a very limited longitudinal dimension (each household stays in the survey only for four quarters), we use synthetic cohort methods to describe the life-cycle patterns in the data. In particular, we will consider cohorts defined on the basis of the year of birth of the household head. To guarantee that each cohort has a sufficiently large number of individuals in each time period considered, we consider 10-year cohorts.

⁴Unfortunately, the characteristics of the car reported in each survey are not the same over time.

3.2.2 Descriptive statistics

We start our analysis with a discussion of some descriptive statistics on consumption expenditure, which we divide into expenditure on non-durables and services and on cars.⁵ In keeping with our model specification we will focus on households aged 25 to 80. A full set of summary statistics for the CEX is provided in Appendix B.3, Table B.16. Figure 3.1 plots the average real spending on non-durable consumption goods and on consumer services against age for each of the cohorts we are considering. Consumption expenditures on non-durables and services follow a hump shaped pattern over the life-cycle with a positive trend until age 30-40 (depending somewhat on the cohort) and with a declining trend from around age 55. In real terms, household consumption increases from \$8,000 per year at age 25 to somewhere between \$11,000 and \$13,500 per year at age 45 depending on the cohort and then returns to a level around \$8,000 per year when households reach their late 60s.⁶ This is the classic hump shaped consumption profile [Attanasio et al. \(1999\)](#).

The data display a similar life-cycle pattern for the number of cars per household, see Figure 3.2, but there are some noticeable differences between how households accumulate cars over the life-cycle and how they allocate their non-durables expenditures over time. Most noticeably, the peak in the number of cars per household occurs at a later age than the peak in the expenditure on non-durables and grows very slowly at younger ages. In their early 20s, each household owns on average little more than 1 car increasing to 1.5 cars at age 30 and peaking at 2-2.2 cars for households in the range of 40-50 years of age. After this peak, household slowly decrease the number of cars that they own.

Figure 3.3 reports the percentage of households which holds no cars. Around 25 to 30 percent of households in their early 20s have no car but this share

⁵We ignore durables other than cars. One possible justification for such an approach (exclusively motivated by data considerations) is that other durables are separable in the utility function

⁶The values reported here are in in real 1984\$

declines quickly with age and stabilizes at close to 10 percent for households in the age range of 30 to 65 years of age. Car non-ownership for households above the age of 65 declines moderately. The stability of the share of car non-owners at ages 30 to 65 and the increase in the number of cars per household in the age range from 30 to 50 implies that existing car owners in this age range increase their stock of cars while some of the increase in the number of cars for younger households derives from more households becoming car owners.

Figure 3.4 illustrates the life-cycle profiles of the value of households' car stock which, although similar to that of the number of cars per household, displays much larger and persistent differences across cohorts than in any of the plots discussed above. This may indicate that aggregate shocks which impact on a large fraction of a cohort are important for understanding the dynamics of household car purchases and that these effects are persistent. For that reason the second panel of Figure 3.4 illustrates the cohort data on the value of the car stock plotted against time rather than age. This figure shows that all cohorts regardless of age experienced a decline in the value of their car stock in the early 1980s recession and during the Great Recession. It therefore follows that it takes households a long time to adjust their car portfolio after a negative shock.

Finally, Figure 3.5 reports the mean ratio of the value of the car stock to annual non-durables expenditure across cohorts plotted against age and time. There is hardly any life-cycle pattern to how households divide their expenditure between cars and non-durables consumption but this ratio does display a very strong cyclical pattern with the relative value of the car stock falling in recessions. For example, in the early to mid 1980s, the value of the households' car stock corresponded to approximately 2 to 2.5 times their annual non-durables and service consumption expenditures and it then increased markedly across cohorts until around 2005 where it peaked at around 3 to 3.5 (depending on the cohort). The Great Recession witnessed a large substitu-

tion away from cars to non-durables with the median ratio of the value of the car stock to non-durables consumption falling back to levels last seen in the early 1990s.

Thus, our results indicate strong life-cycle dynamics in household consumption but also that aggregate shocks have large and persistent effects. For that reason, the next section will further examine the cyclical patterns of the CEX data.

3.2.3 Car purchases during the Great Recession

We now document the behavior of cars expenditure during three recent recessions covered by the CEX survey. In particular, our sample period includes the 1990-1991 recession, the 2001 recession and the Great Recession of 2008-2010.

Total consumer car expenditures is one of the most volatile and procyclical expenditure series and is well-known also for being an important leading business cycle indicator.⁷ For example, NIPA data on car spending fell by almost 20 percent (annually) in the 2nd quarter of 1980, more than 18 percent in the first quarter of 1991, and almost 24 percent in the last quarter of 2008, episodes that all correspond to NBER recessions, see Figure 3.6. On the other hand, as the US economy recovered in the mid-1980's, car spending grew by 15 percent or more in each of the four quarters from 1983Q3. These patterns are also discernible in the CEX data. Figure 3.7 illustrates the time-series of expenditure on cars aggregated across all households and they follow the NIPA data fairly closely. The fit is particularly good for new cars, whereas some deviation exists for old car expenditure. This is because the NIPA series is based on the net expenditure while the CEX data is computed on the gross expenditure

This evidence still leaves open the issue of whether the intensive margin

⁷Alan Greenspan is supposed to have kept a close eye on car sales when considering the path of monetary policy, Greenspan and Cohen (1999) study patterns of car adjustments.

of car adjustment displays any cyclical patterns. For investigating this, the CEX data offers very useful information. Bar-Ilan and Blinder (1992) have previously noted that the extensive margin adjustment of car purchases moves procyclically with fewer households choosing to adjust their car stocks during recessions. To investigate this, Figure 3.8 plots (seasonally adjusted) observations on the percentage of households in the CEX buying a car in a specific quarter. There is a secular decline in the percentage of households that purchase a car. In the late 1980s, around 8-10 percent of households purchased a car every quarter while only around 5 percent of households did so in the aftermath of the Great Recession.⁸ It is evident that fewer households purchase cars in recessions. This fraction falls prior to each of the recessions in our sample and continues to remain low (or fall even further) during the recession. The end of the recessions are associated with discernible recoveries in the frequency of car adjustment. The decline in the fraction of households purchasing a car is particularly significant in the Great Recession (falling from above 6 percent per quarter in 2006 to below 5 percent per quarter during late 2008, but this may simply be due to the depth of the recession).

Figure 3.9 illustrates the average value of car purchased conditional on purchasing and it demonstrates a key difference between the Great Recession and previous downturns in the economy. Fewer households usually purchase cars in a recession but conditional upon purchasing the real value of car purchases remained unchanged during the early 1990s and early 2000s recessions while it drops significantly during the Great Recession. Thus, in “normal” recessions fewer households purchase a car but those that do buy cars, spend a similar amount to those that purchased a car immediately prior to or after the recession. In the Great Recession instead, fewer households bought a car and those that did, spent less than they did prior to the downturn. We summarise this relationship in Table 3.1, where we see that across all, new

⁸Consistent with this there is a concurrent secular increase in the average car age.

and old car purchases the decline in the probability of purchase in the Great Recession was similar to that of a “normal recession”, but on the intensive margins households responded much more strongly.⁹

Next, we estimate some simple models of car purchases. We use a probit model to relate the probability of buying a car to a number of control variables and to the current stock of cars. We control for cyclical patterns by including year fixed effects. Table 3.2 reports the estimated average marginal effects and Figure 3.10 illustrates the sequence of year fixed effects. Having taken household characteristics into account, we find that the probability of purchasing a car is decreasing in the value of the households’ stock of cars. The effect is statistically more significant and the model a better fit when we relate the probability of buying a car to the value of the stock of cars relative to non-durables expenditure rather than to just to the value of the stock of cars. The findings are consistent with the idea that households are more likely to replace cars when the stock depletes. We interpret the improved statistical fit of the model when taking non-durables spending into account as indicating the importance of non-separabilities between cars and non-durables consumption. The pattern of the time fixed effects are consistent with the evidence from Figures 3.8-3.9 showing little variation in the time-fixed effect with the important exception of the Great Recession which sees a large fall in the likelihood of a car purchase. This suggests that the decline in the probability of car purchase during the Great Recession period also exceeded that of previous recessions, once conditioning on the change in full time employment.

To summarize, our analysis shows that household consumption profiles display important life-cycle dynamics and respond to business cycle conditions. In the early part of their life-cycle, the average household experiences a relatively strong growth in their expenditures on non-durables and on cars. As house-

⁹This measure somewhat understates the decline in the probability of purchase during the Great Recession which preceded the NBER recession dates which are used to calculate the average change.

holds age, their expenditure on non-durables and services peaks at around 45 to 50 years of age while their expenditure on cars peak slightly later. Older households reduce their expenditure on both non-durables and services and on cars. Most households end up owning at least one car and the incidence of car non-ownership declines fast for households over the age of 30. In recessions, households reduce their spending on cars relative to non-durables and services and are less likely to purchase a car. Thus, there is evidence that households use durable consumption goods to smooth their consumption streams. However, in the 1990 recession and in the early 2000 recession, those households who did purchase a car display no signs of reducing the size or quality of the car that they acquired relative to the typical purchase prior to these recessions. In the Great Recession instead, fewer households purchased a car and those that did, bought a smaller car relative to purchases prior to the downturn in the economy. We will now try to use these facts to gain a better understanding of the sources of the Great Recession.

3.3 The model

We study a partial equilibrium life-cycle model of household consumption smoothing. Households are subject to stochastic income shocks, face borrowing constraints, and smooth consumption over their life-cycle. We allow for common shocks in order to study aggregate fluctuations. The model includes non-convex car adjustment costs, “credit” financing of car purchases and collateral constraints.¹⁰

¹⁰Fernandez-Villaverde and Krueger (2011) and Iacoviello and Pavan (2013) also study life-cycle models of durables consumption. Fernandez-Villaverde and Krueger (2011) do not allow for non-convex adjustment costs or and assume frictionless financial markets frictions while Iacoviello and Pavan (2013) exclude idiosyncratic income shocks. We incorporate non-convex car adjustment costs, financial market frictions, and idiosyncratic income shocks which makes general equilibrium computationally “expensive.” Berger and Vavra (2015) show in a similar model to ours that the general equilibrium ramifications are moderate.

3.3.1 Household problem

The economy has a continuum of households and the total population is constant. Every period a new cohort of measure 1 is born. $\pi_a \in [0, 1]$ is the probability that a household of age a survives to age $a + 1$. We impose that at age a^{\max} , $\pi_{a^{\max}} = 0$. Households work until age $T_r < a^{\max}$, face a stochastic income stream while working and receive pensions during retirement. Each period, households make choices of how to divide their income over purchases of consumption goods and purchases of financial assets. We abstract from labor supply considerations.

Households consume non-durables and car services. The instantaneous utility function is given as:

$$u(c_{a,t}^j, d_{a,t+1}^j) = \gamma_a \frac{\left[\alpha \left(\frac{c_{a,t}^j}{\gamma_a} \right)^{1-1/\mu} + (1-\alpha) \left(\frac{\xi_{a,t+1}^j}{\gamma_a} \right)^{1-1/\mu} \right]^{(1-\varphi)/(1-1/\mu)}}{1-\varphi} - 1 \quad (3.1)$$

where a indicates age, j indicates the identity of the household, and t denotes calendar time. γ_a is a preference shock which is common across a cohort and is included to control for family size changes over the life-cycle. $c_{a,t}^j$ denotes consumption of non-durables and $d_{a,t+1}^j$ is the household's end-of-period t stock of cars. $\mu, \varphi > 0$ denote the elasticity of substitution between the service flow from cars and consumption of non-durables and the inverse of the intertemporal elasticity of substitution, respectively.

Let $i_{a,t}^j$ be the household's expenditures on car purchases. The law of motion for the stock of cars owned by the household is:

$$d_{a,t+1}^j = (1-\delta) d_{a-1,t}^j + i_{a,t}^j \quad (3.2)$$

where $\delta \in [0, 1]$ is the rate of car depreciation. We assume that households start their life-cycle with $d_{0,t}^j \in D_0 = [\underline{d}, \bar{d}]$, $\underline{d} \geq 0$.

We allow for two salient aspects of household car consumption: non-convex adjustment costs and the availability of car credit. The purchase of a car may often be associated with significant costs such as time spent looking for the preferred car, contractual costs, and costs due to differences between car purchase prices and their resale value.¹¹ Following [Attanasio \(2000\)](#) and [Eberly \(1994\)](#), we adopt a non-convex car adjustment cost function, $\Upsilon (d_{a-1,t}^j) \geq 0$:¹²

$$\Upsilon (d_{a-1,t}^j) = \begin{cases} 0 & \text{if } d_{a,t+1}^j = (1 - \delta (1 - \varsigma)) d_{a-1,t}^j \\ \psi p d_{a-1,t}^j & \text{otherwise} \end{cases} \quad (3.3)$$

p is the price of cars denominated in units of the non-durable consumption good. $\psi \geq 0$ denotes the adjustment costs as a proportion of the value of its existing car stock, $p d_{a-1,t}^j$, if it chooses to adjust its car stock beyond maintenance. $\varsigma \in [0, 1]$ is the fraction of car depreciation that can be taken care of by maintenance without inducing adjustment costs, see also [Bachmann et al. \(2013\)](#). When $\varsigma = 0$, the adjustment costs can be avoided by allowing the household's car stock to decrease at the rate of δ (the depreciation rate) over time. When $\varsigma = 1$, the household can instead maintain the standard of the car without incurring adjustment costs (thus keeping up with all the wear and tear that occurs due to depreciation). The presence of fixed car adjustment costs implies that households will in general choose not to adjust their car stocks continuously and spend $\delta \varsigma p d_{a-1,t}^j$ on maintenance during periods of inactivity.

In the US households often purchase cars on credit offered by commercial banks and by specialized Auto Finance Companies. For that reason we assume

¹¹Moreover, because of the non-convex adjustment costs, households will in general adjust their car stock in a discrete manner as is the case in the household data discussed in Section 2.

¹²[Attanasio \(2000\)](#) and [Eberly \(1994\)](#) test and estimate the parameters (S,s) rules for household automobile purchases. Others have instead looked at more indirect evidence such as the impact of uncertainty on automobile purchases, see e.g. [Bertola et al. \(2005\)](#) or [Hassler \(2001\)](#). The former of these papers also examine other types of consumer durables. [Bar-Ilan and Blinder \(1992\)](#) provide an early test of some implications of (S,s) type models.

that cars can be used as collateral for a credit account:

$$k_{a,t+1}^j \leq \eta p d_{a,t+1}^j \quad (3.4)$$

where $k_{a,t+1}^j$ denotes the end-of-period size of the credit account and $\eta \in [0, 1]$ determines the leverage. Thus, the household must have assets of at least $(1 - \eta) p d_{a,t+1}^j$ to cover the downpayments. The car credit account evolves as:

$$k_{a,t+1}^j = (1 + r_t^k) k_{a-1,t}^j + p \vartheta_{a,t}^j - \xi(k_{a-1,t}^j) \quad (3.5)$$

where r_t^k is the interest rate on car credits, $\vartheta_{a,t}^j \in [0, i_{a,t}^j]$ is the amount of new car credit issued to the household, and $\xi(k_{a-1,t}^j) \geq 0$ denotes the household's repayment on its credit account. The (log) interest rate on car loans is assumed to follow an autoregressive process with persistence ρ_r and innovation variance σ_r^2 .

We assume that household can save in a riskless bond which earns interest $r > 0$ but cannot issue uncollateralized debt:

$$b_{a,t+1}^j \geq 0 \quad (3.6)$$

The interest charged on car loans is typically a substantially higher interest rate than the return on savings and households will therefore as far as possible avoid using car loans to smooth consumption. However, the lack of access to uncollateralized debt implies that some households may have an incentive to take car credit even if r_t^k exceeds r since this allows them to purchase a car before having generated sufficient savings to finance its purchase. [Berger and Vavra \(2015\)](#) instead assume that $b_{a,t+1}^j \geq -\eta p d_{a,t+1}^j$ so that durables can be used as collateral for borrowing at the risk free rate which implies a less severe borrowing constraint than what is implied by our assumptions. However, collateralized lending for the purpose of buying durables comes at a premium

which is consistent with our set-up. Kaplan and Violante (2014) (and many others) alternatively rule out collateralized borrowing against illiquid assets altogether. However, since car loans are readily available in the U.S., our assumption appear attractive.

While working, households face a stochastic income stream. We assume that households are subject to a mix of life-cycle changes in income, persistent and transitory idiosyncratic income shocks, and persistent aggregate shocks:

$$y_{a,t}^j = p_{a,t}^j \exp(u_t^j), \quad (3.7)$$

$$p_{a,t}^j = g_{a,t} p_{a-1,t-1}^j \exp(e_{a,t}^j), \quad (3.8)$$

$$e_{a,t}^j = \varepsilon_t^j + \eta_{a,t} + v_t. \quad (3.9)$$

where $u_t^j \sim N(0, \sigma_u^2)$, $\varepsilon_t^j \sim N(0, \sigma_\varepsilon^2)$, $\eta_{a,t} \sim N(0, \sigma_\eta^2)$, $v_t \sim N(0, \sigma_v^2)$ and $g_{a,t}$ follows a two state Markov process.

According to this specification, The level of household income, $y_{a,t}^j$, is subject to transitory and permanent shocks, u_t^j and $e_{a,t}^j$ which are assumed to be mutually independent. Transitory income shocks are assumed to be idiosyncratic while the permanent income shock has an idiosyncratic, ε_t^j , and an aggregate component which may either differ across cohorts, $\eta_{a,t}$, or be common to all independently of age, v_t . Finally, there is a life-cycle component of the income process through g_a . We allow the life-cycle component g_a to be stochastic and vary between a “high growth” and a “low growth” regime.

The income process specified in equation 3.7 extends the standard approach in the life-cycle literature in several dimensions. First, we allow for aggregate shocks since we want to examine business cycle variations. The partial equilibrium aspect of our model means that, for an individual household, there is no difference whether it receives a idiosyncratic, cohort specific or common aggregate shock but the latter allows us to model fluctuations in aggregate variables. Moreover, households may not be able to tell apart idiosyncratic

and common shocks so might respond to them in the same way leaving aside any general equilibrium effects of aggregate shocks through prices.

g_a determines the expected life-cycle dynamics of household earnings. We allow this factor to be stochastic and think of shocks to $g_{a,t} \in (g_a^H, g_a^L)$ being rare with $g_a^H > g_a^L \forall a$ being an “almost absorbing” state. Hence, households are subject to rare but potentially large and persistent shocks to the life-cycle earnings profile which dictates how their expected income varies over the life-cycle. Importantly, these shocks generate not solely changes in income today but also contain information about income in the future. We therefore consider these life-cycle income profile shocks similar to the long-run risk shocks studied by [Bansal and Yaron \(2004\)](#) which have been shown to be important for understanding variations in asset prices.

At age $T_r \leq a^{\max}$, households retire and receive pension benefits. The households maximize subject to a sequence of budget constraints:

$$c_{a,t}^j + p(i_{a,t}^j - v_{a,t}^j) + \Upsilon(d_{a-1,t}^j) + b_{a,t+1}^j + \xi(k_{a-1,t}^j) \leq \quad (3.10)$$

$$(1 - \chi(a))y_{a,t}^j + \chi(a)m_{a,t}^j + (1 + r)b_{a-1,t}^j \quad (3.11)$$

where $\chi(a)$ is an indicator function:

$$\chi(a) = \begin{cases} 0 & \text{if } a < T_r \\ 1 & \text{if } a \geq T_r \end{cases}$$

$m_{a,t}^j \geq 0$ denotes the level of retirement benefits which we assume are given by a final salary scheme:

We will for now eliminate the agent subscript and consider the dynamic optimization problem for an agent of cohort a . Let $s = (b_{a-1}, d_{a-1}, k_{a-1}, a, m(T_r), x)$ be the vector of state variables. In this vector $m(T_r)$ denotes retirement benefits for a household for which $a > T_r$ while x is short hand for the exogenous

state variables $x = (r^k, p_a, g_a, u_a)$. We can then formulate Bellman's equation as:

$$\mathbb{V}(s) = \max(\mathbb{V}^p(s), \mathbb{V}^n(s)) \quad (3.12)$$

where $\mathbb{V}^p(s)$ is the value function for a household that chooses to adjust its stock of cars while $\mathbb{V}^n(s)$ is the value for a household that chooses not to adjust, i.e. households that just carry out the basic maintenance costs. The value function of a household that chooses to adjust its car stock is given as:

$$\mathbb{V}^p(s) = \max_{c_a, d'_a, k'_a, b'_a, i_a, \vartheta_a} u(c_a, d'_a) + \beta\pi(a) \mathbb{E}\mathbb{V}(s') \quad (3.13)$$

subject to:

$$\begin{aligned} c_a + p(i_a - \vartheta_a) + \psi d_{a-1} + b'_a + \xi(k_{a-1}) &\leq (1 - \chi(a)) y_a + \chi(a) m + (1 + r) b_{a-1} \\ d'_a &= (1 - \delta) d_{a-1} + i_a^d \\ k'_a &\leq \eta p d'_a \\ k'_a &= (1 + r^k) k_{a-1} + p \vartheta_a - \xi(k_{a-1}) \end{aligned}$$

and given the laws of motion of income and of the collateral constraint. The value of not adjusting the durables stock instead is determined as:

$$\mathbb{V}^n(s) = \max_{c_a, d'_a, k'_a, b'_a} u(c_a, d'_a) + \beta\pi(a) \mathbb{E}\mathbb{V}(s') \quad (3.14)$$

subject to:

$$\begin{aligned} c_a + b'_a + p(\delta \varsigma d_{a-1} - \vartheta_a) + \xi(k_{a-1}) &\leq (1 - \chi(a)) y_a + \chi(a) m + (1 + r) b_{a-1} \\ d'_a &= (1 - \delta(1 - \varsigma)) d_{a-1} \\ k'_a &\leq \eta p d'_a \\ k'_a &= (1 + r^k) k_{a-1} + p \vartheta_a - \xi(k_{a-1}) \end{aligned}$$

where $p\delta\zeta d_a$ are the maintenance costs that the household must pay when choosing not actively to adjust its car stock.

The presence of non-convex adjustment costs can prevent households from continuously adjusting their stock of cars in response to shocks to their income. The presence of depreciation of the car stock implies that households are reluctant to actively decrease their stock of cars unless they are faced with large and persistent negative income shocks. Durables consumption purchases will therefore be forward looking both for standard intertemporal smoothing reasons but also because of the durability of the stock and the presence of adjustment costs.

Finally, we define aggregate variables. Let μ_a be the measure of agents of age a as a function of the state variables s and let λ_a be the size of cohort a .¹³ Then we define a generic aggregate variable (in per capita terms as):

$$z = \frac{\sum_{a=1}^{a^{\max}} \lambda_a \int z_a(s) d\mu_a(s)}{\sum_{a=1}^{a^{\max}} \lambda_a}$$

Since we do not model in aggregate shocks to mortality (or birth rates), the normalization by the total size of the population is of no consequence.

3.4 Calibration

We solve the model by value function iteration. Further details on the solution method are provided in Appendix A.3, Section A.3.1. The externally calibrated parameters of the model are summarized in Table 3.3. One model period corresponds to a calendar year. Households enter the economy at age 25, participate in the labor market for up to 35 years and retire for up to 20 years. The survival probabilities are calibrated to match the population averages in the 2009 Life Table of the United States (US Department of Health

¹³The size of cohort of age $a^{\max} \geq a > 1$ is given as $\lambda_a = \pi_{a-1}\lambda_{a-1}$, $\lambda_1 = 1$, $\lambda_a = 0$ for $a > a^{\max}$.

and Human Services). These imply year-to-year survival probabilities above 99 percent until the age of 65, and above 98 percent until the age of 75. After that age, the survival probability decreases quickly. The life expectancy implied by the process is 75 years at age 25.

We assume that $\beta = 0.96$, which implies a rate of time preference of approximately 4 percent per year, and we set the annual real return on bonds (the return on savings) equal to 4 percent, a value that is consistent with standard estimates of long run US real interest rates. The intertemporal elasticity of substitution, $1/\varphi$, is set equal to $2/3$ which is consistent with the empirical estimates of [Attanasio and Weber \(1995\)](#), [Eichenbaum et al. \(1986\)](#), and many others who have examined either household data or aggregate time series.

The car loan rate interest rate, r^k , is calibrated to match the average auto loan rates observed in the data. In the 1970-2006 sample, the nominal interest on auto loans issued by Auto Finance Companies or Commercial Banks was 10.5 percent implying a real rate around 5.78 percent per year. Therefore, the car loan interest rate premium is 1.78 percent indicating that it is quite expensive to take out car finance.

Retirement income is assumed to be determined as:

$$m_a^j = \kappa y_{T_r}, a \geq T_r \quad (3.15)$$

where y_{T_r} denotes salary at retirement age and κ determines the replacement ratio. Following [Bernheim et al. \(2001\)](#), we assume a replacement ratio of 60 percent.

We calibrate the variances of the idiosyncratic income risk on the basis of [Blundell et al. \(2008\)](#) and [Gourinchas and Parker \(2002\)](#) and set $\sigma_u^2 = 0.246^2$ and $\sigma_\varepsilon^2 = 0.140^2$. The initial income draw is calibrated using the cross-sectional income variance of households aged 24-26 in the CEX data. We set $\sigma_{Y_{25}}^2 = 0.582^2$ and normalize the mean to unity. The life-cycle income profiles,

g_a , are estimated from CEX data by fitting a polynomial to cohort estimates of averages income controlling for demography, education, cohort and year effects. The estimated profile g_a is scaled to match the life-cycle growth of non-durables consumption.

The initial distribution of financial assets is assumed to be $b_{25} \sim N(\bar{b}_{25}, \sigma_{b_{25}})$ and is estimated in the Survey of Consumer Finances. The initial stock of cars is assumed to be $d_{25} \sim \log N(\bar{d}_{25}, \sigma_{d_{25}})$ and estimated in the CEX.¹⁴ To ensure households in the first period do not have a counterfactually large adjustment rate, we allow proportion $1 - Pr(adj|a = 25)$ to choose their optimal cars stock before their life-cycle simulation begins given the stochastic distribution of income and assets, with $Pr(adj|a = 25) = 0.323$. We additionally impose that no household breaks their collateral constraint at the initial distribution. The variance of the innovation to the car loan interest rate and the persistence of the process are fitted to match the Auto finance Company lending rate which implies an annual persistence of 0.505 and an innovation variance (in logs) of 0.295^2 .

The remaining parameters $\Phi = [\alpha, \xi, \mu, \varsigma, \psi, \delta, \sigma_v^2, g_{scale}]$ are calibrated by matching a vector of targets pertaining to US household and aggregate data for a sample period that excludes the Great Recession. Note that in the benchmark we assume that there are no shocks to the life-cycle income profile and that all aggregate shocks hit the cohorts uniformly.

The targets are summarized in Table 3.4. We use the CEX for the following targets. We target the share of households who each year purchase a car. This target is particularly informative about the size of adjustment costs. Next, we target the mean ratio of car spending to non-durables spending and include also as a target the mean ratio of the value of the car stock to annual non-durables spending in the CEX data. These numbers are helpful

¹⁴Both assets and durables the average for households aged 24-25 and are normalise by mean income aged 24-26 in the data to be consistent with the model.

when estimating preference parameters and the depreciation rate. The final targets that we match from household data is the growth in (log) non-durables expenditures from age 25 to the peak which is informative about the life-cycle pattern of consumption and the age at which the average household holds the maximum car stock. Given our focus on aggregate dynamics, we also include targets for US aggregate moments derived on the basis of annual NIPA data for the pre-Great Recession sample 1970-2006. We target the standard deviations of detrended aggregate real per capita non-durables expenditures and car purchases, and the correlation between these variables.¹⁵

Given these targets, we find the following parameter estimates, presented in Table 3.5. The preference weight on non-durables, α , is estimated to be 80.4 percent implying that non-durables consumption is the dominating component in flow utility. The parameter ξ is a normalization of the the mapping between the stock of cars and the service flow. This parameter is estimated to be 0.531. The elasticity of substitution between non-durable consumption goods and the service flow from cars is estimated to be just above one, $\mu = 1.117$, a value which is similar to [Ogaki and Reinhart \(1998\)](#), and indicates an elasticity of substitution just above the standard Cobb-Douglas unitary specification used in much of the literature including [Berger and Vavra \(2015\)](#). The higher elasticity implied by our estimates implies ceteris paribus more variation in the expenditure shares.

The car depreciation rate is estimated at 25.6 percent per year, an estimate that is similar to values adopted in the literature on car purchases, see e.g. [Attanasio \(2000\)](#). Maintenance costs are estimated to account for 11.9 percent of depreciation implying that households can avoid car adjustment costs if they are willing to let their car stock depreciate at 22.5 percent per year. It follows that the non-adjustment option is an attractive way of decreasing the car

¹⁵We detrend using the Hodrick- Prescott filter using the parameter 6.25, suitable for the annual frequency of the model studied here [Ravn and Uhlig \(2002\)](#).

stock for a household that either is subject to a negative income shock or for other reasons wants to cut back on its car portfolio. Households that instead choose to actively adjust their car stock incur a fixed transaction cost that corresponds to 6.2 percent of the value of its car stock. Given the estimates of the depreciation rate and the car adjustment costs, we impose that the collateral parameter is given as $\eta = (1 - \delta - \psi) / (1 + r^k) = 0.645$. This is the upper bound on lending that is consistent with the absence of default. The estimates of the income process imply that the standard deviation of the aggregate income shocks is 2.5 percent per year.

Given these parameters, the model matches many of the targets very well. The model implies an annual frequency of car adjustment of 27.9 percent while this statistic is 26.4 percent in the data. It is important that this target is matched well in order for the model to be consistent with the lumpiness of car purchases. The model also matches closely how consumers spread their spending over non-durables and durables and the value of their car portfolio relative to non-durables spending. The model predicts a little too much growth in non-durables spending, predicting 52 percent to the lifecycle peak relative to 34.8 percent in the data. Similarly, the maximum age of the car stock occurs later in the model than in the data (aged 61 versus 52). As far as the aggregate moments are concerned, the model matches closely the three targets (the volatility of non-durables consumption, the volatility of car expenditures, and their cross-correlation). Hence, the life-cycle model extended with aggregate shocks does a good job of accounting for key moments of aggregate consumption.

Table 3.6 contains information about the model's performance on a number of household spending moments that are not targeted when estimating the structural parameters. This includes the age at which household spending on non-durables peaks which is a bit later in the model (52 years) than in the data (44 years). However, the model does capture the fact that as in the

data the car stock peaks at a later age than non-durables consumption and by approximately the same duration (9 years in the model, 8 year in the data). Further, the growth in the car stock from age 25 to peak is very well matched (58 percent in the CEX, 60.9 percent in the model) as is the cross-sectional standard deviation of the value of the car stock (94.7 percent in the CEX, 92.1 percent in the model). We also find that conditional on adjustment the size of car purchased, relative to the current stock is a little higher in the model 1.13 relative to the data 0.83, although this difference is not that large.

3.5 Life-cycle and policy functions

3.5.1 Life-cycle profiles

Figure 3.11 illustrates the life-cycle paths of non-durables consumption expenditure and the car stock implied by our parametrization of the model. In order to produce these, we simulated the model with 27,500 agents feeding in stochastic shocks to income, wealth and interest rates as well as to mortality. We then averaged over the life-cycle profiles.

The average life-cycle path of non-durables consumption shares many salient features with the CEX estimates in Figure 3.1. As in the CEX data, non-durables consumption expenditure rises fast for younger households, flattens out when they reach their mid-40s (in the diagram, households start at age 25), and declines for households during retirement. The model implies a small reversal of the negative trend in non-durables consumption expenditure for households close to the terminal age. This feature is an artefact of the assumption that we make that the mortality risk goes to one when the household turns 80. In the periods immediately prior to this terminal age, households sell their car portfolio and use the proceeds to finance non-durables consumption spending.

The life-cycle path of the value of the household car stock is also very similar to the empirical estimates we discussed in Section 3.2. The value of the car stock rises gradually for younger households as they are likely to face binding borrowing constraints and therefore take time to build up sufficient savings before they can acquire a car. In the model, older households liquidate their car stock towards the end of the life-cycle faster than in the data but this is also an artefact of the assumption that the maximum lifespan is 80 years in the model. In Figure 3.12 we see that as in the data the probability of purchasing a car is declining in model as in the data, although not with quite the same speed.

Finally, we illustrate the life-cycle profile of car loans. Car loans are typically taken out by young households while households above the age of 30 are likely to hold financial assets. Essentially no households between the ages of 50 and 65 have car debt while there is some increased incidence in car loans for households close to the end of the life-cycle as they use this source of financing to avoid having to incur car adjustment costs. The reason for car loans are mostly accessed by the youngest households is that the car loan premium is very substantial which induces a strong savings motive.

3.5.2 Policy functions

Figure 3.13 shows various policy functions that solve the household's dynamic programming problem. Panel a shows consumption expenditures plotted against cash on hand for a young, middle-aged, and old household. There is some curvature in the consumption functions for low levels of cash on hand due to the borrowing constraint on liquid assets but mainly for young and middle-aged households (for older households, the consumption expenditure is essentially a linear function of cash on hand). It is also evident that as households age, their incentive to save declines. Panel b illustrates spending

on cars plotted against cash on hand, again for young, middle-aged and older households. These policy functions illustrate the lumpiness of investment in cars. Poor households choose not to invest in cars because of the non-convex adjustment costs and because of the need to save sufficiently for the down payment. As cash on hand rises, households eventually invest in cars but the size of the car that they own depend on their assets.

Panel a of Figure 3.14 illustrates more clearly how individual households adjust their car stock over time. This policy function shows a households' current choice of the (value of) their car stock (on the vertical axis) plotted against their beginning of period car stock holding constant cash on hand (and other state variables). The policy function displays very clearly (s,S) type behavior. Recall that households are assumed to be able to pay for maintenance without incurring adjustment costs. Hence, households let their stock of cars decline over time at the rate of 23 percent (the depreciation rate corrected for the maintenance) until the car stock hits a lower trigger point where the household makes a discrete investment in cars bringing the stock to their optimal level. The policies depend on the households' assets (and on other relevant state variables such as age). This is shown in Panel b where we illustrate the car adjustment policy for three levels of cash on hand. Poorer households delay the upward adjustment of their car stock relative to richer households and, when they adjust, choose a smaller car stock adjustment target. It follows from these policy functions that households faced with small or moderate declines in income will tend to delay their car adjustment which is consistent with the evidence that the share of households adjusting their car stock during recessions is lower than during normal times. Moreover, unless recessions are protracted, most households that do adjust their car stocks during a recession will tend to be households that despite the economy-wide contraction are doing well (experience countervailing idiosyncratic shocks) and for that reason do not delay their car adjustment. This latter observation also

implies that the size of cars purchased by households that do invest during recessions may not be very different from their pre-recession choices.

These implications imply that the Great Recession must have witnessed either different or more serious shocks than are seen in normal recessions since, as we noted in Section 3.2, the intensive margin of car adjustment did decline significantly. We will now use the estimated model to explore this empirical observation.

3.6 The Great Recession

We use the model to examine the sources of the Great Recession by simulating it for cohorts of households under alternative assumptions about the nature of the shocks impacting on the economy. We then contrast the implications of the model with the observed paths of consumption and car spending as well as with cohort level data. Importantly, the model has a rich wealth distribution and the presence of non-convex adjustment costs implies that the past history of shocks will matter for its aggregate implications when subjected to a particular sequence of shocks (in other words, the model has path-dependence). For that reason, we simulate the model for a long sequence of shocks prior to the Great Recession rather than assuming that the Great Recession started with the economy being near its ergodic distribution.

3.6.1 Uniform income shocks

We initially examine the properties of the model when assuming that households are all subject to the same sequence of aggregate income shock from 1971 onwards, $[\widehat{v}_t]_{t=1971}^{2013}$. Realized income will, however, differ across households because of idiosyncratic transitory and permanent income shocks. Moreover, households' responses to the shocks will differ due to heterogeneity of the population in the age and in their asset composition.

We derive $[\widehat{v}_t]_{t=1971}^{2013}$ by first estimating the deviations from a pre-2008 linear trend of BLS measure of labor income (the sum of wage income and proprietors income divided by the personal consumption deflator and by population) which we measure in logs. We then derive the aggregate shocks to v_t so that the implied ‘aggregate’ income fits the NIPA estimate. In either case we also draw from idiosyncratic income shocks from the calibrated distributions shown in Table 3.3. In order to allow for path dependence, we start the simulations from 1970 with households in the implied ergodic distribution over states at this point in time. We calculate the implied sample paths from 1979 to 2013, aggregating over 27,500 household in order to compute aggregate outcomes. We repeat this process 100 times to compute the average aggregate economy responses.

Figure 3.15 illustrates the aggregate shocks and the resulting aggregate income path together with its empirical counterpart. The sequence of income shocks that are needed in order to ‘fit’ the observed income process implies a longer and deeper decline in aggregate income than observed during previous recessions. Hence, the size of the recession is larger than those used for estimating the structural parameters which paves the way for the model to imply different consumption responses than during normal recessions.

This version of the model does a good job at accounting for the path of aggregate non-durables consumption prior to the Great Recession as well as from 2007 onwards, see Figure 3.16. However, it exaggerates the drop in consumption during the early 1990s recession as well as the subsequent slow recovery of consumption. It cannot account for the 2004-06 boom in consumption and it implies a counterfactual recovery of consumption in 2012-13. Nonetheless, it is remarkable how well the life-cycle model does in accounting for non-durables consumption in response to uniform aggregate income shocks.

This success, however, does not carry over to consumer durables. The model is inconsistent with the amount of variation in aggregate car expendi-

tures and, in particular, implies a very modest fall in aggregate car spending from 2007 onwards, an implication that is in stark contrast to the data. Figure 3.17 makes clear that this failure of the model applies both to the intensive and extensive margins. In the data, the share of households purchasing a car declined by 6 percentage points from 2005 to 2007, while the model implies a more modest fall of 2.5 percentage points from 28 percent to 25.5 percent. Furthermore, the model economy predicts a strong recovery in the extensive margin from 2011 onwards while no such recovery is present in the CEX data. Perhaps even starker, the intensive margin is by and large unaffected by the recession. In the data, the amount spent by households purchasing a car falls by close to 20 percent from 2006 to 2010 while the decline predicted by the model is smaller than 5 percent. It follows from this that while uniform income shocks hitting all agents in the economy allows one to account for the aggregate dynamics of non-durables expenditures, such shocks, although unusually large during the Great Recession, fail to predict the dynamics of car purchases.

Moreover, at the household level there are clear differences in car adjustments across cohorts that uniform income shocks cannot explain. In section 3.6.5 we discuss in detail the extensive and intensive margin responses across the life-cycle and show there exist a distinct empirical pattern, with a large response of overall spending by the younger cohorts and a strong intensive margin response by the middle cohort (see Figure 3.27). The uniform model is unable to reproduce these dynamics and instead generates a similar decline in car spending at all points of the life-cycle.

3.6.2 Cohort specific shocks

The last set of results above indicate that perhaps it is important to take into account that different cohorts experienced different shocks to their income. For

that reason we now expose the agents in the economy to sequences of cohort specific shocks ($\eta_{a,t}$). We compute these as follows. First, we divide households into decennial age groups. We then compute from the CEX how income has varied from 1980 onwards for these cohorts and rescale the shocks so that we match the NIPA labor income series. Thus, by construction, we constrain the “aggregate shocks” to match those of the experiment above which assumed uniform shocks across cohorts. A full description of the construction of the cohort shocks is provided in Appendix A.3, Section A.3.4.

The results are illustrated in Figure 3.18. As far as the aggregate dynamics are concerned, the model with cohort-specific shocks does not deliver results that are significantly different from those of the uniform shocks version of the model discussed above. The only slight difference between the two cases is that consumption is slightly higher during the early to mid 2000s when we allow for cohort specific shocks. Apart from this, the implied consumption and car adjustment paths are as good as identical. The reasons for this are that the policy functions i.) multiplicative in permanent income and ii.) linear in cash on hand with the exception of the very poor households. Although younger households tend to be poorer, this difference is too marginal to matter.

However, when introducing cohort specific shocks, the model’s performance in terms of accounting for the adjustment of car expenditure across cohorts does improve. In particular, this version of the model is now consistent with older households cutting their car investment less during the Great Recession than the younger cohorts and with the 35-44 year old cohort adjusting the intensive margin more than other cohorts. Nonetheless, it is clear that even after introducing cohort specific aggregate shocks, the life-cycle model does not appear to be consistent with consumption adjustments during the Great Recession.

3.6.3 Financial market shocks

In the US new car purchases are often financed through collateralized car loans issued by either specialized car loan companies or by commercial banks. The terms of these loans are likely to have been affected by the financial crisis and it is possible that these shocks are important for household consumption choices.

In order to investigate this issue, we now allow r_t^k , the interest rate on car loans to be stochastic. As when estimating the model, we assume that the car loan rate follows an autoregressive process:

$$\log r_t^k = \frac{\log \bar{r}^k}{1 - \rho_r} + \rho_r \log r_{t-1}^k + \epsilon_t^r$$

where ϵ_t^r is normally distributed with mean 0 and variance σ_r^2 . The persistence and variance of this process is estimated by fitting an autoregressive processes to an average series constructed from the Federal Reserve Board estimates of autofinance companies' New Car Average Finance rate and the commercial banks interest rate on 48 month car loans.¹⁶ We find that $\rho_r = 0.505$ and $\sigma_r = 0.295$.

We then follow estimates of ϵ_t^r for the post 1972 period. We first simulate the model including both interest rate shocks and cohort specific income shocks from 1970 to 2006. We then simulate the Great Recession assuming that either the economy was hit only by the cohort specific income shocks, or only by the car loan interest rate shock or by both at the same time. Figure 3.19 illustrates the shocks and the implied interest rate spread ($r_t^k - r_t$). Prior to the Great Recession, there is a large decline in the spread on car loans which fell by close to 140 basis points from 2001 to 2006. From 2006 to 2010, the spread rises sharply by more than 180 basis points.

¹⁶We remove a linear trend and use a sample that starts in 1971. We weight the two series using estimates from [Attanasio et al. \(2008\)](#)

Despite the large fluctuations in the car loan spread, it has very little impact on the life-cycle model, see Figure 3.20. The reason for this is that the spread is sufficiently large, combined with the life-cycle motive for saving, that very few households have car debt. In particular, car debt is concentrated on very young households who partially finance their car purchases through taking out a car loan and on older households who effectively make use of this facility to make their car stock liquid in the years before the terminal date. Younger households, however, pay off car loans as fast as they can in order to rid themselves of this costly debt instrument. The decline in the spread prior to the Recession does lead to an increase in the frequency of car adjustment but the effect is very minor as is the impact of the subsequent hike in the spread. Evidence of a change in savings and loan behaviour can be seen in Appendix C.3, Figure C.18, but these portfolio reallocations are not large enough to generate a non-trivial consumption response.

3.6.4 Wealth shocks

As we have discussed earlier, a large decline in consumer wealth does imply a downward adjustment of households' desired car stock. None of the shocks we have studied so far, however, imply sufficiently sharp adjustments in household wealth that they can account for the consumer durables dynamics witnessed during the Great Recession.

One possibility of addressing this is to allow for shocks to wealth directly (on top of income shocks and changes in the car loan spread). Indeed, a much discussed feature of the Great Recession was the sharp reduction in real estate prices which reduced the (perceived) wealth of real estate owners (Mian et al. (2013), Berger et al. (2018a)). We do not have real estate in the model but it is well-known that the dominating share of household saving is invested in real estate. Thus, we will now introduce stochastic “asset price shocks” by

allowing the return on assets for households with positive wealth to vary with house prices. We write the household budget constraint as:

$$\begin{aligned} & c_{a,t}^j + d_{a,t}^j + \Omega_{a,t}^j + \Upsilon (d_{a-1,t-1}^j) \\ & \leq y_{a,t}^j + (1 - \delta) d_{a-1,t-1}^j + (1 + r_t^\Omega) \Omega_{a-1,t-1}^j \end{aligned}$$

where $\Omega_{a,t}^j$ denotes net financial assets. We then allow for wealth shocks if net assets are positive:

$$\begin{aligned} & c_{a,t}^j + d_{a,t}^j + \Omega_{a,t}^j (1_{\Omega < 0} + p_t 1_{\Omega \geq 0}) + \Upsilon (d_{a-1,t-1}^j) \\ & \leq y_{a,t}^j + (1 - \delta) d_{a-1,t-1}^j + (1 + r_t^\Omega) (1_{\Omega < 0} + p_t 1_{\Omega \geq 0}) \Omega_{a-1,t-1}^j \end{aligned}$$

where

$$\log p_t = \log p_{t-1} + \varepsilon_{p,t}, \quad \varepsilon_{p,t} \sim \mathcal{N}(0, \sigma_p^2)$$

The asset price is assumed to follow a random walk. We estimate σ_p^2 to match the variance of the log change of the linearly detrended FHFA house price index deflated by the CPI and uncover the implied sequence of shocks. Figure 3.21 illustrates the linearly detrended real house price index and the sequence of shocks to this series. This indicates a very strong boom-bust cycle in the house price index which rises by 30 log points from 1997 to 2006 and thereafter implodes reaching its trough in 2012.

Figure 3.22 illustrates the impact of introducing such asset price shocks. We find that they are important for accounting for the observed consumption dynamics. First, the pre-recession boom in asset prices helps accounting for the early 2000 boom in spending on non-durable consumption. We now fit this series very well apart from the early 1990s recession which has too lasting an impact on non-durables spending. The decline in house prices from 2007 onwards has an impact on non-durables spending that is approximately half as large as the impact of the decline in labor income. In combination, income

shocks and asset price shocks now imply that we can account almost exactly for the path of non-durables spending. This is an interesting finding in itself which highlight how asset price movements impact on consumption (see: [Kaplan et al. \(2017\)](#)).

The asset price shocks also matter for the consumer durables dynamics. First, once we introduce the asset price dynamics, the combined impact of lower incomes and falling house prices imply that aggregate spending on cars in the model falls by around 20 log points from 2006 to 2010, a decline that is a little smaller than in the data but much larger than in the simulations considered above. As can be seen from Figure 3.23, the house price shocks are more important for explaining the decline in the intensive margin of car purchases than the extensive margin. This is for three reasons. Firstly, due to the life-cycle dynamics, wealthy households were also the ones buying larger cars. When this group is hit by the asset price shock it has a larger effect on the value of the aggregate average car purchased. Secondly, older households are also more sensitive to cash in hand shocks (see policy function in Figure 3.14) and the retired households were not previously exposed to an income shock. Thirdly, because the shock to wealth has passively rebalanced the household portfolio (between durables and assets) in addition to reducing lifetime income, households are less willing to accept a deviation from the optimal car choice and respond more strongly on the intensive margin rather than delaying their purchase.

3.6.5 Life-cycle profile income shock

Finally, we show that adding a shock to the life-cycle growth rate can complete our description of the consumption dynamics during the Great Recession. The shock reduces the expected deterministic life-cycle income profile from when it hits in 2009. We model the shocks as a persistent but non-permanent change

in the profile. Clearly, such a shock will impact heavily on households early in the life-cycle, while having no effect on those at the end of their working lives. One interpretation of such a shock is the introduction of long-run risk [Bansal and Yaron \(2004\)](#).

Empirical evidence

Our assumption is that young households perceive a decline in the future growth rate of income, not that they necessarily will experience one for the duration of their careers. We now present empirical evidence that is consistent with this type of shock. Given that there are a limited number of years of data available, it is challenging to distinguish income shocks from a change in the profile in the cross section, so we interpret these patterns as informative rather than definitive evidence of a decline in the profile. However, other authors have also pointed to empirical support for a flattening of the life-cycle profile ([Kong et al. \(2018\)](#)).

In the CEX we use two methods to estimate the growth rate of income using cross sectional variation. Firstly, we estimate the change in the growth rate directly, aggregating over age-year cells. Secondly, we estimate the life-cycle profile before and after the recession and assess the implied change in the growth rate. The years for sample before the recession are 1989 to 2012, the years for the sample after the growth rate shock are 2009 to 2012. Full details of the estimation are provided in Appendix A.3, Section A.3.2. Figure 3.24 provides the results using Financial Income Before Tax in the CEX.¹⁷ As can be seen whether estimated directly on the growth rate or on the life-cycle profile basis, the growth rate of income appears to fall after the Great Recession hits. We find this pattern is broadly replicated in other income series in the CEX and in the Current Population Survey (See Appendix C.3,

¹⁷While, Financial Income Before Tax is broader than our preferred measure of household labor income, there are more available observations. For young households with low wealth likely to be affected by the growth shock this difference should not be that problematic

Figure C.19). The decline measured in the data was used to calibrate the growth shock imposed during the Great Recession. Given the specification $g_{a,post} = \max\{g_{a,pre}^\gamma, g_{a,pre}\}$, this implied a growth rate shock of $\gamma = 0.58$.¹⁸ The implied effect on the life-cycle profile can be seen in Appendix C.3, Figure C.20.

Aggregate response

Figure 3.25 presents the extensive and intensive margin responses following the addition of the growth rate shock. In the aggregate the shock has a small additional effect, with the dynamics being fairly similar to those in the previous experiment with asset price changes. The growth shock provides an additional response of the extensive margin in 2009 when it is realised. This is because households affected by the decline choose to delay car purchase for longer for the same reasoning as with an income shock. On the intensive margin we see a small persistent additional effect from 2009 onwards. As the households predominantly affected by this shock are younger, and thus purchase smaller cars, the marginal aggregate effect is reasonably small.

Life-cycle responses

Thus far we have focused on the aggregate response, but the response across cohorts is also highly informative about the shocks hitting the economy. To take advantage of this information we compare the difference between the predicted and actual path of a cohort's consumption decisions, by estimating a life-cycle profile on pre-2008 data and comparing it to the actual consumption decisions of the cohorts. Appendix A.3, Section A.3.5 provides full details of the methodology. Example paths for the consumption of different cohorts during the recession against the life-cycle prediction are presented in Figure 3.26.

¹⁸The max operator is imposed to ensure the growth rate does not increase for older households with a declining income profile

The grey area indicates the deviation from the expected path of consumption.¹⁹ A key feature from the data is that the Great Recession was associated with a large response on the extensive margin for young and middle aged cohorts, while the middle aged cohorts responded most strongly on the intensive margin.

Figure 3.27 compares the response of the model to the data. To account for measurement error, we focus on the change between 2007 and the average of 2008-10.²⁰ We see that the fully specified model, with asset price, interest rate and growth shocks, does a much better job of matching the data than the model with uniform income shocks. As mentioned earlier, the model with uniform income shocks exhibits fairly invariant behaviour across the cohorts a pattern at odds with the data. In contrast the fully specified model replicates these dynamics. The younger and middle cohorts reduced their car expenditure by significantly more than the older cohort, whose main response was on the intensive margin. The addition of the growth shock is critical here as it raises the inaction of the youngest cohort. The full model also generates the U-shaped response on the intensive margin, with the middle cohort reducing the size of purchase most strongly. This cohort is hit by a combination of a sizeable income shock, holds reasonably significant assets so are affected by the wealth decline, and for the younger member of the cohort are not immune to the effect of the growth shock. This combination results in a strong intensive margin response replicating the dynamics observed in the data.

¹⁹It is important to take into account the predicted life-cycle path, as at different points of the life-cycle cohorts will have different expected future changes in consumption. Ignoring the likely path of consumption growth will give an inaccurate description of the effect of a shock.

²⁰Due to the small sample of car purchasers with a cohort, for the value of car conditional on adjusting we aggregate over 2006-07 in the data for the baseline

3.7 Conclusion

This chapter has sought to use household consumption data to provide a better understanding of the shocks hitting the economy during the Great Recession. Its main argument is that information about the relative importance of the often-unobserved shocks hitting the economy can be uncovered by carefully studying consumption dynamics. In contrast to the literature thus far, the analysis has made greater use of the additional information contained in the differential response along the extensive and intensive adjustment margins of durables expenditure and of the variation in responses across cohorts, to uncover the key sources of disturbance during the Great Recession. This approach was motivated by the empirical fact documented here that whereas in previous recessions households primarily responded by reducing the probability of car purchase, during the Great Recession households also purchased smaller cars.

The analysis has highlighted an important role for cohort specific income shocks, asset price shocks and a decline in the expected growth rate of income for understanding consumption dynamics during the Great Recession. We attribute a small role to changing financial conditions perhaps due to the fact that in the car market households are relatively unconstrained financially. However, we do not interpret this as evidence that financial shocks are not important for understanding the Great Recession, per se. For example, housing decisions are likely to be more dependent on the availability of credit and we attribute a sizeable role to the wealth shock which could be considered to have its roots in an unmodelled financial shock. The combination of the remaining shocks provides a good representation of consumption decisions at both the aggregate and cohort level.

While the model here is richly specified there are of course abstractions from reality. For computational purposes we have chosen to stress the richness

across the household dimension at the expense of general equilibrium effects. This could be important if it resulted in substantial endogenous and expected changes in car prices absent from the model. Further, while households in the data decreased purchasing new and used cars, there is evidence that some households also switched from purchasing new to purchasing old cars during the Great Recession.²¹ This substitution offers a further interesting margin of adjustment during the crisis that could also be informative about household's expectations. However, incorporating both new and used cars would be a significant extension and as such is left for future research.

²¹Example of papers that have include a new and used car trade-off include: [Adda and Cooper \(2000\)](#) and [Gavazza et al. \(2014\)](#)

3.8 Tables

	Normal recession	Great Recession
Probability of purchase		
All	-0.609	-0.870
New	-0.218	-0.359
Old	-0.433	-0.523
Value of car purchase		
All	0.301	-7.954
New	0.145	-4.484
Old	2.645	-8.606

Value is change average in the four quarters preceding the recession to average during the NBER recession period relative. “Normal” recession is average of 1981-82, 1990-91 and 2001. Probability of purchase is change in percentage points. Value of car purchase is percentage change.

Table 3.1: Comparison of Great Recession car market outcomes

Variable	(1)	(2)	(3)	(4)
Stock (\$10,000)	-0.024*** (0.0064)	-0.179*** (0.03201)		
Stock:ndur			-0.093*** (0.0016)	-0.068*** (0.0089)
Black	-0.031*** (0.0014)	-0.031*** (0.0014)	-0.031*** (0.0014)	-0.031*** (0.0014)
Black	-0.031*** (0.0014)	-0.031*** (0.0014)	-0.031*** (0.0014)	-0.031*** (0.0014)
Education	-0.002*** (0.0003)	-0.0025*** (0.0004)	0.004*** (0.0004)	-0.004*** (0.0004)
Sex	-0.031*** (0.0011)	-0.031*** (0.0010)	-0.030*** (0.0010)	-0.030*** (0.0010)
Family size	0.014*** (0.0003)	0.0138*** (0.0003)	0.012*** (0.0003)	0.012*** (0.0003)
Full time	0.009*** (0.0010)	0.009*** (0.0010)	.009*** (0.0010)	.009*** (0.0010)
Weeksp.	0.0003*** (0.0000)	0.0003*** (0.0000)	0.0003*** (0.0000)	0.0003*** (0.0000)
year F.E	✓		✓	
stock x year		✓		✓
R^2	0.0256	0.0258	0.0331	0.0334
N	500,018	500,018	499,984	499,984

Notes: Probit Estimation of probability of adjustment. *Statistically significant at 10%; **statistically significant at 5%; ***statistically significant at 1%. Standard errors in parentheses. Stock:ndur is the ratio of the car stock to non-durables consumption.

Table 3.2: Probability of Purchasing a Car (Avg. Marginal Effects)

Parameter		Value
a_{\max}	maximum lifespan (life starts at age 25)	55 years
T_r	retirement age	35 years
π_a	survival probability	match 2009 Life Table
β	subjective discount factor	0.96
$1/\varphi$	intertemporal elasticity of substitution	2/3
r	annual real return on savings	4 percent
r^k	annual car loan interest rate	5.78 percent
κ	pension replacement rate	60 percent
σ_u^2	variance of transitory idiosyncratic income shock	0.246 ²
σ_ε^2	variance of persistent idiosyncratic income shock	0.140 ²
$\rho_{r^k}^2$	persistence of car loan spread	0.505
$\sigma_{r^k}^2$	variance of car loan spread	0.295 ²
$\sigma_{Y_{25}}^2$	cross-sectional variance of initial log. income	0.582 ²
\bar{b}_{25}	mean initial assets	0.086
$\sigma_{b_{25}}^2$	cross-sectional variance of initial assets	1.036 ²
\bar{d}_{25}	mean initial log. car	-1.39
$\sigma_{d_{25}}^2$	cross-sectional variance of initial log. car	1.04 ²
g_a	life-cycle income factor	matched to CEX data
γ_a	household equivalent size	matched to CEX data
σ_p^2	variance of asset price	0.031 ²
g_a^{post}	life-cycle growth shock	0.581

Table 3.3: Externally calibrated parameters

Moment	Data	Model
Share of households purchasing a car	26.4	27.9
Ratio of car to non-durables spending	19.2	16.9
Ratio of car stock to non-durables spending	67.1	72.7
Growth in non-durable from age 25 to peak	34.8	52.2
Age at peak of car stock	52	61
Std dev. of aggregate non-durables	0.77	0.94
Std dev. of aggregate car purchases	5.91	5.29
Correlation of aggregate non-durables and car purchases	72.1	68.3

Table 3.4: Calibration targets

Parameter		Value
α	weight on non-durables in utility function	0.804
μ	elasticity of substitution	1.117
ξ	service flow from durables	0.531
ψ	car adjustment cost parameter	0.062
ς	car maintenance cost parameter	0.112
δ	car depreciation rate	0.256
σ_u^2	variance of aggregate permanent income shock	0.025 ²
g_a	Scaling of life-cycle income	0.711

Table 3.5: Estimated Parameters

Moment	Data	Model
Ratio of purchase value to stock adjust	0.83	1.13
Std dev. of value of car purchase adjust	2.63	1.68
Age at peak of non-durables spending	44	52
Growth to peak in car stock	58	60.9
Cross sectional std dev of non-durables consumption	58.8	87.3
Cross-sectional standard deviation of value of car stock	94.7	91.8

Table 3.6: Non Targeted Moments

3.9 Figures




Figure 3.1: CEX cohorts: non-durable consumption




Figure 3.2: CEX cohorts: number of cars




Figure 3.3: CEX cohorts: % households with no cars




Figure 3.4: CEX cohorts: Value of car stock




Figure 3.5: CEX cohorts: Ratio of car stock to non-durable consumption




Figure 3.6: NIPA: car expenditure & durables growth

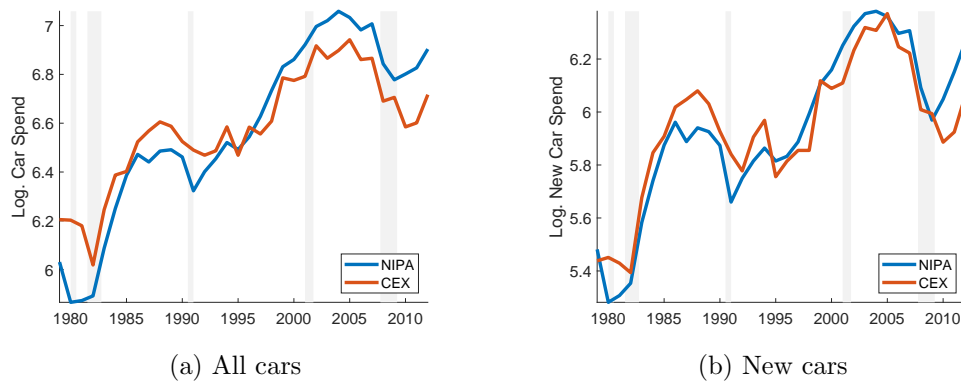


Figure 3.7: Comparing NIPA to CEX




Figure 3.8: CEX: share of households purchasing car




Figure 3.9: CEX: value of car purchased




Figure 3.10: CEX: Probability of purchase, year fixed effects




Figure 3.11: Model life-cycle fit (I)




Figure 3.12: Model life-cycle fit (II)




Figure 3.13: Model policies

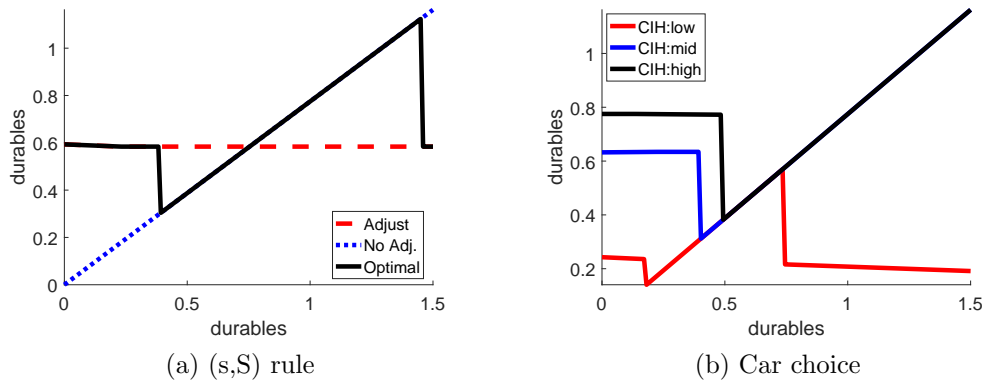


Figure 3.14: Model policies




Figure 3.15: Uniform income shocks




Figure 3.16: Uniform income shocks: consumption




Figure 3.17: Uniform income shocks: extensive and intensive margin




Figure 3.18: Uniform versus cohort shocks




Figure 3.19: Interest rate spread shock




Figure 3.20: Interest rate spread shock




Figure 3.21: Asset price shock




Figure 3.22: Asset price shock: consumption response




Figure 3.23: Asset price shock: extensive and intensive margin

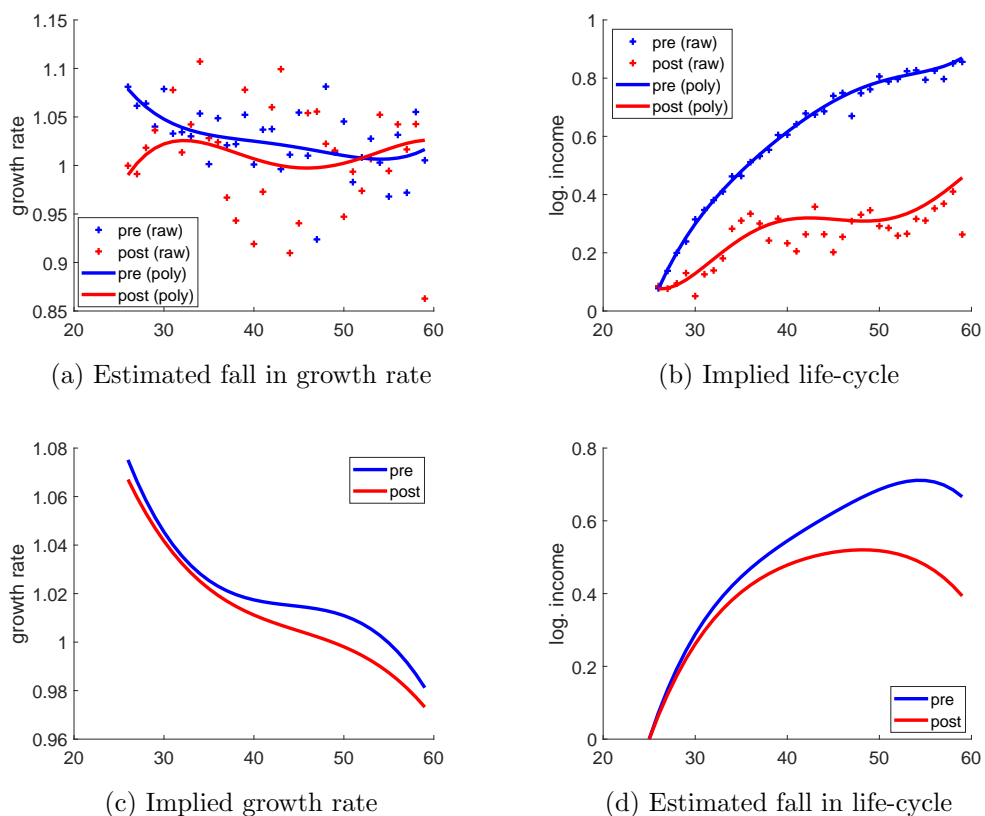


Figure 3.24: Evidence supporting decline in life-cycle growth rate (CEX)




Figure 3.25: Growth rate shock: extensive and intensive margin




Figure 3.26: CEX cohort car response

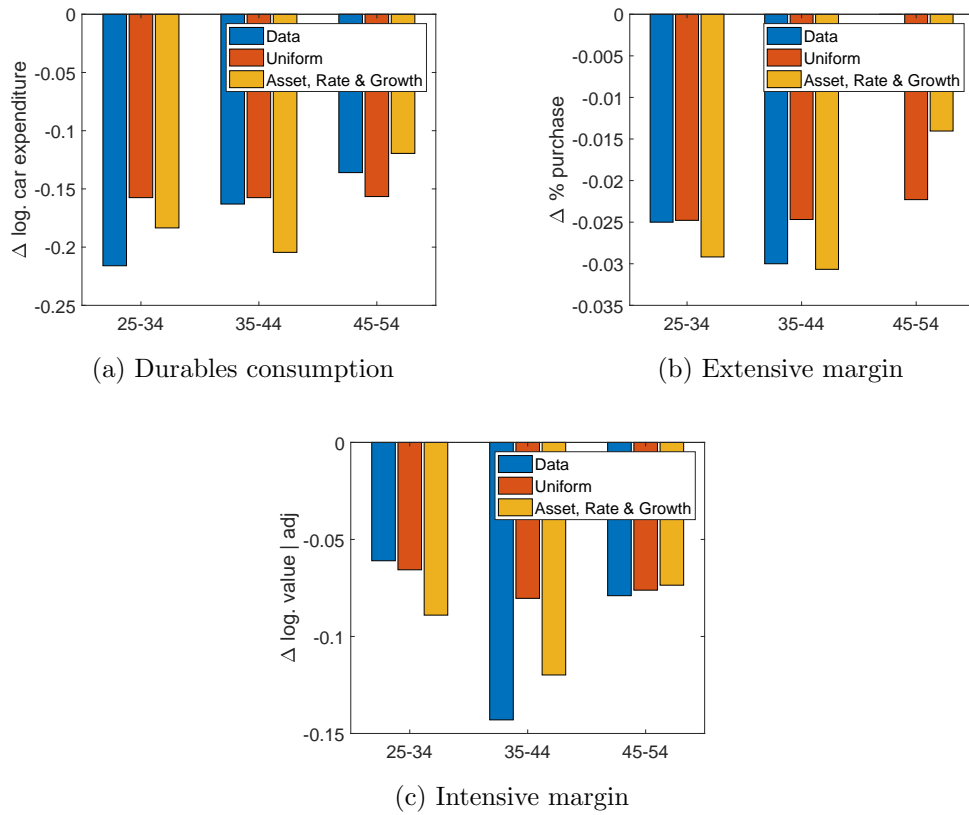


Figure 3.27: Cohort response: model comparison

Bibliography

- Abiad, A., Detragiache, E., and Tressel, T. (2010). A New Database of Financial Reforms. *IMF Staff Papers*, 57(2):281–302.
- Adda, J. and Cooper, R. (2000). Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies. *Journal of Political Economy*, 108(4):778–806.
- Ahn, H. J. and Hamilton, J. D. (2016). Heterogeneity and Unemployment Dynamics. (22451).
- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.
- Amior, M. and Halket, J. (2014). Do households use home-ownership to insure themselves? Evidence across U.S. cities. *Quantitative Economics*, 5(3):631–674.
- Attanasio, O. P. (2000). Consumer durables and inertial behaviour: Estimation and aggregation of (s, s) rules for automobile purchases. *The Review of Economic Studies*, 67(4):667–696.
- Attanasio, O. P., Banks, J., Meghir, C., and Weber, G. (1999). Humps and Bumps in Lifetime Consumption. *Journal of Business & Economic Statistics*, 17(1):22–35.
- Attanasio, O. P., Goldberg, P. K., and Kyriazidou, E. (2008). Credit Constraints In The Market For Consumer Durables: Evidence From Micro Data On Car Loans. *International Economic Review*, 49(2):401–436.
- Attanasio, O. P. and Weber, G. (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, 103(6):1121–1157.
- Bachmann, R., Caballero, R. J., and Engel, E. M. R. A. (2013). Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model. *American Economic Journal: Macroeconomics*, 5(4):29–67.
- Bajari, P., Chan, P., Krueger, D., and Miller, D. (2013). A Dynamic Model Of Housing Demand: Estimation And Policy Implications. *International Economic Review*, 54(2):409–442.
- Bansal, R. and Yaron, A. (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance*, 59(4):1481–1509.
- Bar-Ilan, A. and Blinder, A. S. (1992). Consumer Durables: Evidence on the Optimality of Usually Doing Nothing. *Journal of Money, Credit and Banking*, 24(2):258–272.
- Barillas, F. and Fernandez-Villaverde, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, 31(8):2698–2712.

- Basten, C., Fagereng, A., and Telle, K. (2016). Saving and portfolio allocation before and after job loss. *Journal of Money, Credit and Banking*, 48(2-3):293–324.
- Baxter, M. (1996). Are Consumer Durables Important for Business Cycles? *The Review of Economics and Statistics*, 78(1):147–155.
- Bayer, C., Lütticke, R., Pham-Do, L., and Tjaden, V. (2015). Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk. CEPR Discussion Papers 10849, C.E.P.R. Discussion Papers.
- Benhabib, J., Bisin, A., and Luo, M. (2017). Earnings inequality and other determinants of wealth inequality. *American Economic Review*, 107(5):593–97.
- Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2018a). House Prices and Consumer Spending. Technical Report 3.
- Berger, D. and Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, 83(1):101–154.
- Berger, D. W., Milbradt, K., Tourre, F., and Vavra, J. (2018b). Mortgage Prepayment and Path-Dependent Effects of Monetary Policy. NBER Working Papers 25157, National Bureau of Economic Research, Inc.
- Bernheim, B. D., Skinner, J., and Weinberg, S. (2001). What Accounts for the Variation in Retirement Wealth among U.S. Households? *American Economic Review*, 91(4):832–857.
- Bertola, G. and Caballero, R. J. (1990). Kinked Adjustment Costs and Aggregate Dynamics. In *NBER Macroeconomics Annual 1990, Volume 5*, NBER Chapters, pages 237–296. National Bureau of Economic Research, Inc.
- Bertola, G., Guiso, L., and Pistaferri, L. (2005). Uncertainty and Consumer Durables Adjustment. *Review of Economic Studies*, 72(4):973–1007.
- Blundell, R., Low, H., and Preston, I. (2013). Decomposing changes in income risk using consumption data. *Quantitative Economics*, 4(1):1–37.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, 98(5):1887–1921.
- Blundell, R. and Preston, I. (1998). Consumption Inequality and Income Uncertainty. *The Quarterly Journal of Economics*, 113(2):603–640.
- Brunnermeier, M. K. (2009). Deciphering the Liquidity and Credit Crunch 2007-2008. *Journal of Economic Perspectives*, 23(1):77–100.
- Buera, F. J. and Moll, B. (2015). Aggregate Implications of a Credit Crunch: The Importance of Heterogeneity. *American Economic Journal: Macroeconomics*, 7(3):1–42.
- Burdett, K. and Mortensen, D. T. (1980). Search, Layoffs, and Labor Market Equilibrium. *Journal of Political Economy*, 88(4):652–672.
- Caballero, R. J. (1993). Durable Goods: An Explanation for Their Slow Adjustment. *Journal of Political Economy*, 101(2):351–384.
- Campbell, J. R. and Hercowitz, Z. (2005). The Role of Collateralized Household Debt in Macroeconomic Stabilization. NBER Working Papers 11330, National Bureau of Economic Research, Inc.
- Carroll, C. and Dunn, W. (1997). Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets. In *NBER Macroeconomics*

- Annual 1997, Volume 12*, NBER Chapters, pages 165–230. National Bureau of Economic Research, Inc.
- Carroll, C., Slacalek, J., Tokunaka, K., and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.
- Carroll, C. D., Dynan, K. E., and Krane, S. D. (2003). Unemployment risk and precautionary wealth: Evidence from households’ balance sheets. *The Review of Economics and Statistics*, 85(3):586–604.
- Carroll, C. D. and Kimball, M. S. (2001). Liquidity Constraints and Precautionary Saving. NBER Working Papers 8496, National Bureau of Economic Research, Inc.
- Carroll, C. D. and Samwick, A. A. (1998). How Important Is Precautionary Saving? *The Review of Economics and Statistics*, 80(3):410–419.
- Castaneda, A., Diaz-Gimenez, J., and Rios-Rull, J.-V. (2003). Accounting for the U.S. Earnings and Wealth Inequality. *Journal of Political Economy*, 111(4):818–857.
- Castañeda, A., Díaz-Giménez, J., and Ríos-Rull, J. (2003). Accounting for the u.s. earnings and wealth inequality. *Journal of Political Economy*, 111(4):818–857.
- Challe, E. (2017). Uninsured Unemployment Risk and Optimal Monetary Policy. Working Papers 2017-54, Center for Research in Economics and Statistics.
- Chetty, R., Sandoràndor, L., and Szeidl, A. (2017). The effect of housing on portfolio choice. *The Journal of Finance*, 72(3):1171–1212.
- Chetty, R. and Szeidl, A. (2007). Consumption Commitments and Risk Preferences. *The Quarterly Journal of Economics*, 122(2):831–877.
- Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2015). Understanding the Great Recession. *American Economic Journal: Macroeconomics*, 7(1):110–167.
- Clarida, R., Galí, J., and Gertler, M. (2000). Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *The Quarterly Journal of Economics*, 115(1):147–180.
- Cloyne, J., Ferreira, C., and Surico, P. (2016). Monetary policy when households have debt: new evidence on the transmission mechanism. Bank of England working papers 589, Bank of England.
- Coles, M. and Kelishomi, A. (2015). Do Job Destruction Shocks Matter in the Theory of Unemployment? (14462).
- Cubas, G. and Silos, P. (2017). Career choice and the risk premium in the labor market. *Review of Economic Dynamics*, 26:1 – 18.
- d’Albis, H. and Iliopoulos, E. (2013). Collateral constraints and rental markets. *Economics Letters*, 121(3):436–439.
- Davis, S. J. (2008). The Decline of Job Loss and Why It Matters. *American Economic Review*, 98(2):263–267.
- Davis, S. J. and Kahn, J. A. (2008). Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro

- Levels. *Journal of Economic Perspectives*, 22(4):155–180.
- Davis, S. J. and von Wachter, T. M. (2011). Recessions and the Cost of Job Loss. NBER Working Papers 17638, National Bureau of Economic Research, Inc.
- Díaz, A. and Luengo-Prado, M. J. (2010). The Wealth Distribution With Durable Goods. *International Economic Review*, 51(1):143–170.
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248.
- Den Haan, W. J., Judd, K. L., and Juillard, M. (2010). Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):1–3.
- DeNardi, M., French, E., and Benson, D. (2012). Consumption and the Great Recession. *Economic Perspectives*, (Q I):1–16.
- Drehmann, M., Borio, C., and Tsatsaronis, K. (2012). Characterising the financial cycle: don’t lose sight of the medium term! BIS Working Papers 380, Bank for International Settlements.
- Dunn, W. E. (1998). Unemployment risk, precautionary saving, and durable goods purchase decisions. Finance and Economics Discussion Series 1998-49, Board of Governors of the Federal Reserve System (US).
- Dupor, B., Li, R., Saif, M. M., and Tsai, Y.-C. (2019). The 2008 U.S. Auto Market Collapse. *St. Louis Federal Reserve Working Paper*.
- Eberly, J. C. (1994). Adjustment of Consumers’ Durables Stocks: Evidence from Automobile Purchases. *Journal of Political Economy*, 102(3):403–436.
- Eeckhout, J. and Sepahsalari, A. (2015). Unemployment Risk and the Distribution of Assets. *Unpublished manuscript*.
- Eggertsson, G. B. and Krugman, P. (2012). Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. *The Quarterly Journal of Economics*, 127(3):1469–1513.
- Eichenbaum, M. S., Hansen, L. P., and Singleton, K. J. (1986). A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice Under Uncertainty. (1981).
- Engel, C. and Wang, J. (2011). International trade in durable goods: Understanding volatility, cyclicalities, and elasticities. *Journal of International Economics*, 83(1):37–52.
- Fagereng, A., Holm, M. B., and Natvik, G. J. (2018). MPC Heterogeneity and Household Balance Sheets. CESifo Working Paper Series 7134, CESifo Group Munich.
- Fallick, B. C. and Fleischman, C. A. (2004). Employer-to-employer flows in the U.S. labor market: the complete picture of gross worker flows. Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System (U.S.).
- Favilukis, J., Ludvigson, S. C., and Nieuwerburgh, S. V. (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):140–223.
- Fella, G. (2011). A Generalized Endogenous Grid Method for Non-concave Problems. Working Papers 677, Queen Mary University of London, School of Economics and Finance.

- Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, 17(2):329–344.
- Fernández-Villaverde, J. and Krueger, D. (2011). Consumption And Saving Over The Life Cycle: How Important Are Consumer Durables? *Macroeconomic Dynamics*, 15(05):725–770.
- Fujita, S. (2018). Declining Labor Turnover and Turbulence. Working Papers 18-6, Federal Reserve Bank of Philadelphia.
- Fulford, S. L. (2015). The surprisingly low importance of income uncertainty for precaution. *European Economic Review*, 79(C):151–171.
- Garner, T. I. ., McClelland, R., and Passer, W. (2009). Strength and Weaknesses of the Consumer Expenditure Survey from a BLS Perspective.
- Gavazza, A., Lizzeri, A., and Roketskiy, N. (2014). A Quantitative Analysis of the Used-Car Market. *American Economic Review*, 104(11):3668–3700.
- Gottschalk, P. and Moffitt, R. (2009). The rising instability of u.s. earnings. *Journal of Economic Perspectives*, 23(4):3–24.
- Gourinchas, P.-O. and Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1):47–89.
- Greenspan, A. and Cohen, D. (1999). Motor Vehicle Stocks, Scrappage, And Sales. *The Review of Economics and Statistics*, 81(3):369–383.
- Greenwood, J. and Hercowitz, Z. (1991). The Allocation of Capital and Time over the Business Cycle. *Journal of Political Economy*, 99(6):1188–1214.
- Grossman, S. J. and Laroque, G. (1990). Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods. *Econometrica*, 58(1):25–51.
- Gruber, J. W. and Martin, R. F. (2003). Precautionary savings and the wealth distribution with illiquid durables. International Finance Discussion Papers 773, Board of Governors of the Federal Reserve System (U.S.).
- Guerrieri, L. and Iacoviello, M. (2017). Collateral constraints and macroeconomic asymmetries. *Journal of Monetary Economics*, 90(C):28–49.
- Guerrieri, V. and Lorenzoni, G. (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap. *The Quarterly Journal of Economics*, 132(3):1427–1467.
- Guiso, L., Jappelli, T., and Terlizzese, D. (1992). Earnings uncertainty and precautionary saving. *Journal of Monetary Economics*, 30(2):307–337.
- Guvenen, F. (2011). Macroeconomics With Heterogeneity: A Practical Guide. NBER Working Papers 17622, National Bureau of Economic Research, Inc.
- Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2015). What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk? NBER Working Papers 20913, National Bureau of Economic Research, Inc.
- Halket, J. and Vasudev, S. (2014). Saving Up or Settling Down: Home Ownership over the Life Cycle. *Review of Economic Dynamics*, 17(2):345–366.
- Hall, R. E. (2011). The Long Slump. *American Economic Review*, 101(2):431–469.
- Hall, R. E. (2014). Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis. In *NBER Macroeconomics Annual 2014, Volume 29*, NBER Chapters, pages 71–128. National Bureau of Economic Research, Inc.
- Hassler, J. (2001). Uncertainty and the Timing of Automobile Purchases.

- Scandinavian Journal of Economics*, 103(2):351–366.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2009). Quantitative Macroeconomics with Heterogeneous Households. NBER Working Papers 14768, National Bureau of Economic Research, Inc.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and Labor Supply with Partial Insurance: An Analytical Framework. *American Economic Review*, 104(7):2075–2126.
- Hedlund, A. (2016). The cyclical dynamics of illiquid housing, debt, and foreclosures. *Quantitative Economics*, 7(1):289–328.
- Hedlund, A., Mitman, K., Karahan, F., and Ozkan, S. (2017). Monetary policy, heterogeneity, and the housing channel. *Unpublished manuscript*.
- Hubmer, J. (2018). The Job Ladder and its Implications for Earnings Risk. *Review of Economic Dynamics*, 29:172–194.
- Huckfeldt, C. (2018). Understanding the Scarring Effect of Recessions. 2018 Meeting Papers 1207, Society for Economic Dynamics.
- Iacoviello, M. (2005). House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle. *American Economic Review*, 95(3):739–764.
- Iacoviello, M. and Pavan, M. (2013a). Housing and debt over the life cycle and over the business cycle. *Journal of Monetary Economics*, 60(2):221–238.
- Iacoviello, M. and Pavan, M. (2013b). Housing and debt over the life cycle and over the business cycle. *Journal of Monetary Economics*, 60(2):221–238.
- Iskhakov, F., H., J. T., John, R., and Bertel, S. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365.
- Jappelli, T., Padula, M., and Pistaferri, L. (2008). A Direct Test of The Buffer-Stock Model of Saving. *Journal of the European Economic Association*, 6(6):1186–1210.
- Jarosch, G. (2015). Searching for job security and the consequences of job loss. *Unpublished manuscript*.
- Jermann, U. and Quadrini, V. (2012). Macroeconomic Effects of Financial Shocks. *American Economic Review*, 102(1):238–271.
- Johnson, K. W., Pence, K. M., and Vine, D. J. (2014). Auto Sales and Credit Supply. Finance and Economics Discussion Series 2014-82, Board of Governors of the Federal Reserve System (US).
- Johnson, K. W. and Sarama, R. F. (2015). End of the Line: Behavior of HELOC Borrowers Facing Payment Changes. Finance and Economics Discussion Series 2015-73, Board of Governors of the Federal Reserve System (US).
- Jose Luengo-Prado, M. (2006). Durables, nondurables, down payments and consumption excesses. *Journal of Monetary Economics*, 53(7):1509–1539.
- Jung, P. and Kuhn, M. (2017). Earnings Losses and Labor Mobility Over the Life Cycle. (6552).
- Justiniano, A., Primiceri, G., and Tambalotti, A. (2015). Household leveraging and deleveraging. *Review of Economic Dynamics*, 18(1):3–20.
- Kaplan, G., Mitman, K., and Violante, G. L. (2017). The housing boom and bust: Model meets evidence. Working Paper 23694, National Bureau of Economic Research.

- Kaplan, G., Moll, B., and Violante, G. L. (2016). Monetary Policy According to HANK. NBER Working Papers 21897, National Bureau of Economic Research, Inc.
- Kaplan, G. and Schulhofer-Wohl, S. (2017). Understanding The Long-Run Decline In Interstate Migration. *International Economic Review*, 58:57–94.
- Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.
- Kaplan, G. and Violante, G. L. (2018). Microeconomic Heterogeneity and Macroeconomic Shocks. *Journal of Economic Perspectives*, 32(3):167–194.
- Kaplan, G., Violante, G. L., and Weidner, J. (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity*, 45(1 (Spring)):77–153.
- Khan, A. and Thomas, J. K. (2013). Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity. *Journal of Political Economy*, 121(6):1055–1107.
- Kong, Y.-C., Ravikumar, B., and Vandenbroucke, G. (2018). Explaining cross-cohort differences in life-cycle earnings. *European Economic Review*, 107(C):157–184.
- Krivenko, P. (2018). Unemployment and the US housing market during the Great Recession. 2018 Meeting Papers 579, Society for Economic Dynamics.
- Krolikowski, P. (2017). Job Ladders and Earnings of Displaced Workers. *American Economic Journal: Macroeconomics*, 9(2):1–31.
- Krueger, A. B. and Mueller, A. I. (2012). Time Use, Emotional Well-Being, and Unemployment: Evidence from Longitudinal Data. *American Economic Review*, 102(3):594–599.
- Krueger, D., Mitman, K., and Perri, F. (2016). *Macroeconomics and Household Heterogeneity*, volume 2 of *Handbook of Macroeconomics*, chapter 0, pages 843–921. Elsevier.
- Krusell, P., Mukoyama, T., and Şahin, A. (2010). Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations. *Review of Economic Studies*, 77(4):1477–1507.
- Krusell, P., Smith, A. A., and Jr. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):867–896.
- Leahy, J. V. and Zeira, J. (2005). The Timing of Purchases and Aggregate Fluctuations. *Review of Economic Studies*, 72(4):1127–1151.
- Lise, J. (2013). On-the-Job Search and Precautionary Savings. *Review of Economic Studies*, 80(3):1086–1113.
- Low, H., Meghir, C., and Pistaferri, L. (2010). Wage Risk and Employment Risk over the Life Cycle. *American Economic Review*, 100(4):1432–1467.
- Lütticke, R. (2017). Transmission of monetary policy with heterogeneity in household portfolios. *Unpublished manuscript*.
- Álvarez Parra, F., Brandao-Marques, L., and Toledo, M. (2013). Durable goods, financial frictions, and business cycles in emerging economies. *Journal of Monetary Economics*, 60(6):720–736.
- Mian, A., Rao, K., and Sufi, A. (2013). Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics*, 128(4):1687–1726.
- Mian, A. and Sufi, A. (2012). The Effects of Fiscal Stimulus: Evidence from

- the 2009 Cash for Clunkers Program. *The Quarterly Journal of Economics*, 127(3):1107–1142.
- Mishra, A. K., Uematsu, H., and Powell, R. R. (2012). Precautionary wealth and income uncertainty: a household-level analysis. *Journal of Applied Economics*, 15(2):353 – 369.
- Molloy, R., Smith, C. L., Trezi, R., and Wozniak, A. (2016). Understanding Declining Fluidity in the U.S. Labor market. *Brookings Papers on Economic Activity*.
- Moscarini, G. and Postel-Vinay, F. (2018). The Cyclical Job Ladder. *Annual Review of Economics*, 10(1):165–188.
- Ogaki, M. and Reinhart, C. M. (1998). Measuring Intertemporal Substitution: The Role of Durable Goods. *Journal of Political Economy*, 106(5):1078–1098.
- Olivi, A. (2019). Revealed Preferences and Beliefs from Consumption Savings Decisions: Identification Results and an Empirical Investigation. *Unpublished manuscript*.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–2553.
- Ravn, M. O. and Sterk, V. (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125 – 141.
- Ravn, M. O. and Uhlig, H. (2002). On adjusting the Hodrick-Prescott filter for the frequency of observations. *The Review of Economics and Statistics*, 84(2):371–375.
- Rios-Rull, J.-V. and Huo, Z. (2016). Financial Frictions, Asset Prices, and the Great Recession. Staff Report 526, Federal Reserve Bank of Minneapolis.
- Shimer, R. (2005). The Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review*, 95(1):25–49.
- Shimer, R. (2012). Reassessing the Ins and Outs of Unemployment. *Review of Economic Dynamics*, 15(2):127–148.
- Silos, P. (2007). Housing, portfolio choice and the macroeconomy. *Journal of Economic Dynamics and Control*, 31(8):2774–2801.
- Slacalek, J., Sommer, M., and Carroll, C. (2012). Dissecting saving dynamics: measuring wealth, precautionary and credit effects. Working Paper Series 1474, European Central Bank.
- Souleles, N. S., Parker, J. A., and Johnson, D. S. (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, 96(5):1589–1610.
- Sterk, V. (2015). Home equity, mobility, and macroeconomic fluctuations. *Journal of Monetary Economics*, 74(C):16–32.
- Stevens, A. H. (1997). Persistent Effects of Job Displacement: The Importance of Multiple Job Losses. *Journal of Labor Economics*, 15(1):165–188.
- Stock, J. H. and Watson, M. W. (2003). Has the Business Cycle Changed and Why? In *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pages 159–230. National Bureau of Economic Research, Inc.
- Stock, J. H. and Watson, M. W. (2012). Disentangling the Channels of the 2007-2009 Recession. NBER Working Papers 18094, National Bureau of

Economic Research, Inc.

Wong, A. (2017). Population aging and the transmission of monetary policy to consumption. *Unpublished manuscript*.

Zeldes, S. P. (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, 97(2):305–346.

Appendix A

Additional material

A.1 Appendix to Job Risk, Separation Shocks and Household Asset Allocation

A.1.1 Unemployment response methodology

To compare the model to the data I make use of the PSID, following [Stevens \(1997\)](#) and [Huckfeldt \(2018\)](#). These papers seek to find the impact of an involuntary job loss on future income. The PSID includes a question asking respondees whether they started their job in the last year. As with the previous literature I define a job loss as a separation due to company closure, layoff or firing. I also include those unemployed that report having lost their last job to have finished due to the same criteria (company closure, layoff or firing) and that report having worked in the last year.

The sample is restricted to the pre-1999, Core sample of the PSID. I focus on head of household aged between 19-64 and drop self employed. Households that report a job loss in the past 10 years in the the first year of the PSID are also dropped, as the year of separation is not determined. As in [Huckfeldt \(2018\)](#) I include households not present throughout the entire study.

To estimate the effect of a job separation dummy are used for the period pre- and post- the separation event. More precisely, if the separation takes place in year t , let d_t^j be a dummy variable is the household experience a job separation j periods ago. Separation dummies are included for $j = -2, \dots, 10$.¹ The empirical specification is then:

$$Y_{i,t} = X_{i,t}\beta + \sum_{j=-2}^{10} d_{i,t}^j \delta^j + \alpha_i + \gamma_t + \epsilon_{i,t}$$

where, X is a set of controls, including a quartic in age, education dummies and family demographics, and γ_t are year fixed effects. To control for unobservable worker characteristics an individual fixed effect is included,

¹Pre separation dummies are important to include as the methodology does not allow precise identification of the timing of separation

accounting for any systematic differences in the workers likely to lose their jobs such as lower wages or smaller housing stocks.

To estimate the response in the model the same equation is estimated on simulated data.² A panel of workers is simulated at the quarterly model frequency and aggregated to annual observations. For housing I use end of period housing stock, following the PSID design. As the minimum period for the unemployed that do not immediately find a job is one quarter, before aggregating to annual observations I set income during the quarter unemployed to be 50 percent of the wage next period. For households not employed in the following quarter this is zero. To compare absolute values to the data the housing, equity and mortgages are rescaled so that the ratio of mean labor income in the PSID to mean labor income in the model. In the model a job separation is any household that was employed at the beginning of period $t - 1$, but was separated (with probability δ) during the period. This includes households that immediately find employment in period and start period t employed.

A.1.2 Great Recession shocks and data

The Great Recession shocks are estimated from the CPS. As in equation 1.1, a Probit regression is run to estimate the job separation risk of an industry-occupation-state cell. In this case I do not control for income as the groups and stratified by weekly earnings. A second regression is run to find the predicted log. average weekly earnings of each industry-occupation-cell. For this exercise I use data for 2000-2007 to capture the pre-Great Recession distribution.

The job cells are separated into low wage and high wage groups, with low wage job cells being those with a predicted wage below the median predicted wage. Within the wage groups the jobs are further separated by low risk and high risk with low risk being jobs with a job separation rate below the median job separation rate conditional on being in the given wage group.

Therefore, each group accounts for 25 percent of the sample.

Having assigned an ordering of jobs I follow the outcome of these jobs during the recession period. I look at the average monthly job loss rate of individuals assigned to a given group and calculate the effect of the recession relative to a baseline of the average job loss rate for the group between 2005 and 2007. To reduce measurement error and noise I calculate the effect at an annual frequency.

For the wage I look at the average log weekly earnings of each sub-group. I estimate the effect of the recession as the deviation from a group specific linear trend between 2000 and 2008. In the CPS the decline in wages occurs some time after the NBER recession date.

For the consumption and asset data I use the PSID. I use the job risk and wage estimates from the CPS, but recalculate the groups based on the distribution of jobs in the PSID, using the the years 2001-2007. Whereas for

²The are no age, demographic or time controls in the model estimating equation as these features are absent from the model

the shocks in the CPS cross sectional estimates are used, in the PSID I make use of the panel structure. I assign households to a group based upon the job they held in 2007 and follow the average for those households over the course of the recession. Responses are taken relative to the average group value in 2007. The assets are as described in the paper. For consumption food, utilities, transport, education, childcare, repairs, furniture, clothing, trips and entertainment is included. This definition of consumption is available since 2005.

To calculate the shocks in the model I match the cross sectional increase in job risk and maximum decline in earnings to the data for each group. I also match the persistence of the shocks. The shocks are shown in Table B.5. It is necessary computationally to do this in an auxiliary model that does not feature asset choices. For the low income groups in the model, a significant fraction of the wage decline is accounted for by a change in the composition of the groups with the increase in recently unemployed households starting at the bottom of the jobs ladder reducing the average wage of this group. For the consumption and asset response I replicate the data by following a panel of workers allocated to a group in the steady state. I calculate the response as deviations from the path of asset and consumption choices in the absence of a shock.

A.1.3 Great Recession experiment details

The definition of what constitutes “equivalent Great Recession shocks” is not entirely straightforward in the two stationary equilibrium, so I undertake two exercises that capture different assumptions about the household’s expectations.

Experiment 1: Firstly, I feed the shocks into the pre-Great Recession equilibrium. However, because the unemployment rate is lower in this economy if exactly the same shock were fed, in the size of the income decline would be smaller in the pre Great Recession economy. Therefore, in *experiment 1* I rescale the size of the job separation shocks to achieve the same percentage aggregate income decline. These responses are labeled “Pre-GR, fixed expectation”. The shocks to the job separation rate are essentially scaled by the decline in the job separation rate. This makes the percentage point increase in the unemployment rate the same in both cases. In this experiment households expect to return to the lower pre-Great Recession separation rate once the shocks has died out.

Experiment 2: Secondly, I take the distribution of agents from the pre-Great Recession equilibrium feed in exactly the high average separation rate economy shocks and use the policy functions from the high average separation rate economy. One way to think about this experiment is that households wake up understanding job risk has risen permanently and are then additionally hit by a further temporary shock. However, in fact in period 0 the high average separation rate policy functions are used, so a more accurate way of thinking about this is experiment is that it replicates the additional effect of the out of equilibrium asset choices. In this experiment it

is also necessary to rescale. In particular, I rescale the share of employed and unemployed to match the high average separation rate equilibrium, but leave the pre-Great Recession distribution over job types - in (ω, δ) space. In effect I rescale the unemployment rate so that in the absence of the Great Recession shocks there would not be a large increase in the unemployment rate, although there is some adjustment as households return to the high average separation rate (ω, δ) distribution. These responses are labeled “Pre-GR, change expectation”. As the aggregate housing demand is lower in the high average separation rate equilibria than the pre-Great Recession equilibria in this experiment I allow the housing stock to linearly decline over the transition period (250 periods). In this experiment households expect to return to the higher job separation rate once the shocks have died out.

A.2 Appendix to Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics

A.2.1 Evidence for the existence of credit shocks

Evidence for the role of credit constraints influencing consumption dynamics in way that differs from a productivity slowdown induced recession can be demonstrated more formally. Tables B.13 and B.14 present the results of running regressions of the the three durables definitions on a TFP shock series and either a CEA shock series or the level of the HP filtered CEA series as presented in Figure 2.3. Results are shown for both the change in the log variable (Table B.13) and the consumption to stock ratio (Table B.14). For all but one of the durables consumption series, the CEA shock is statistically significant, often having a similar magnitude to the TFP shock. For the durables series that the CEA shock is not significant, $\Delta I^d(1)$, the HP level series is significant. There is less evidence that changes in credit conditions as summarized by the CEA series are important for investment, with the level being statistically significant, but not the shocks. The reliability of these results would obviously be improved by using an instrumental variable strategy as endogeneity of the *CEA* cannot be ruled out.³

A.2.2 Solution method

The solution to the problem is non-trivial due to the combination of the non-convex adjustment costs and aggregate uncertainty. As a result the algorithm is computationally intensive. To find the solution I use a combination of an endogenous grid method and iteration on the aggregate law of motion.

A.2.3 Solving the household problem

The household's value function and policy functions are solved using a generalization of [Carroll \(2006\)](#) endogenous grid method (EGM) for non-concave problems. The original method is not suitable for the household problem studied here due to presence of the adjustment cost which means that the maximisation problem is no longer concave. However, a generalization has been proposed by [Fella \(2011\)](#), which is significantly faster than the alternative value function iteration algorithm. The solution uses 1250 grid points over assets, 50 grid point over durables and 5 grid points over the aggregate capital stock

³ [Slacalek et al. \(2012\)](#) provide results for the regressions of the savings rate on the level of the non-HP filtered CEA index. They find an IV estimation strategy using [Abiad et al. \(2010\)](#) Financial Liberation Index has little effect of the estimated coefficients.

A.2.4 Solving the household problem: a generalised endogenous grid method

First, redefine the assets choice variable: $a_{t+1} = k_{t+1} + \Xi^H d_{t+1}$, where Ξ^H is the highest, or fixed, realisation of the collateral constraint and define available resources, contingent on the next period's durables choice as: $x(d_{t+1}; a_t, d_t, y_t) = y_t + (1+r)(a_t + \Xi^H d_t) + (1-\delta^d)d_t - (1-\Xi^H)d_{t+1} - \Psi(d_t, d_{t+1})$, where y_t summarizes the period labour income and a constant interest rate is assumed. The redefinition of the assets choice variable means that for the highest collateral constraint, the choice of assets no longer depend on durables, the constraint is $a_{t+1} \geq 0$. For lower realisations of the collateral constraint, the constraint is now: $a_{t+1} \geq (\Xi^H - \Xi_t)d_{t+1}$.

The insight of the Fella solution is that the problem can be broken in two, and that in some areas of the problem the first order condition, as used in the standard EGM method will be sufficient for an optimum. Concretely, the actual value can be thought of as the max over a set of values contingent of next period's durables choice:

$$V(a_t, d_t, y_t) = \max_{d_{t+1}} V^{d_{t+1}}(a_t, d_t, y_t)$$

In fact, Fella's original solution was proposed for the case where d_t is also discrete. The basic idea of the algorithm is that $V^{d_{t+1}}$ is differentiable in a , giving the standard first order condition, away from the constraint:

$$-u_c(x_t - a_{t+1}, d_{t+1}) + \tilde{V}_a(a_{t+1}, d_{t+1}, y_t) = 0$$

where $\tilde{V}(a_{t+1}, d_{t+1}, y_t) = \beta EV(a_{t+1}, d_{t+1}, y_{t+1})$. Due to the adjustment costs and any discreteness in durables, $\tilde{V}(a_{t+1}, d_{t+1}, y_t)$ is a non-smooth, non-concave function. However, under a set of regularity conditions it is still a necessary condition for an interior local maximum, and outside the non-concave region of \tilde{V} it will also be sufficient. The algorithm works as follows: given tomorrow's state (y_t, d_{t+1}) , for each a_{t+1} , and for an initial \tilde{V} , the first order condition implies an associated x_t . Examine, the associated \tilde{V} : outside its non-concave region the pair $\{x_t, a_{t+1}\}$ is an optimum, inside the region one must check whether $V^{d_{t+1}} = \max_{a_{t+1}} u(x_t - a_{t+1}, d_{t+1}) + \tilde{V}(a_{t+1}, d_{t+1}, y_t)$ for the given x_t yields the solution a_{t+1} . If so $\{x_t, a_{t+1}\}$ is also a global maximum and the pair is retained, if not the pair is discarded. Once the conditional values are found, the upper envelope can be found by taking the maximum over d_{t+1} , as in equation (8). Given that the problem is specified as an infinite horizon set up the true $V(a_t, d_t, y_t)$ is a fixed point. After the implied V is found, as is standard, check for convergence: $\sup |V^{j+1} - V^j| / \sup |V^j|$, if this is smaller than a defined criterion, the algorithm finishes otherwise use the new V and begin the algorithm again.

A.2.5 Refinements

In addition to using the Fella algorithm, a number of refinements have been made to make the solution applicable to the problem studied in this paper. In particular, an adjustment is made for the fact that in this problem the no-adjustment case is to let durables depreciate and that durables choice is a continuous variable. A addition is also made to incorporate the stochastic collateral constraint. Consider now an adjust and no adjust solution to the algorithm stated above: $V(a_t, d_t, y_t) = \max\{V^A(a_t, d_t, y_t), V^N(a_t, d_t, y_t)\}$. Now given the initial value conditional on the choice of next period's durable stock, $V^{d_{t+1}}$, the no adjust case is immediately implied because no maximisation over d is required. However, it will be on a set of grid points that differs to that originally defined and so must be interpolated before it can be compared with V^A .

Secondly, consider the solution for V^A , which approximates a continuous choice as the number of discrete states in d is increased. One resulting problem is that as d increases, the number of elements in the global maximisation in equation (8) rises reducing the speed of the algorithm. The proposed solution is to avoid undertaking this step every iteration, this is in the same spirit as Barillas and Fernandez-Villaverde (2007), where a labour choice is only calculated intermittently. In the first iteration maximisation over the full durables space is undertaken, giving $V^A(a_t, d_t, y_t|d_{t+1}^*)$. On the next iteration it is assumed that for a given state combination a_t, d_t, y_t , if adjusting, the same choice of d_{t+1} is made, this is the value used and the maximisation $V(a_t, d_t, y_t) = \max\{V^A(a_t, d_t, y_t|d_{t+1}^*), V^N(a_t, d_t, y_t)\}$ is calculated. After a fixed number of iterations the full maximisation over d_{t+1} is once again undertaken. Notice that the adjust/no adjust problem is solved every iteration. To ensure the accuracy of the solution, in the final iteration before convergence a full maximisation over d_{t+1} is required.⁴

This refinement significantly speeds up the solution as the size of the durables grid increases and produces the same policy functions, so the solution found is the true one.

The incorporation of the stochastic collateral constraint is relatively straight forward. When the collateral constraint is fixed the redefinition of assets ensures that the lower bound feasible choice set $a_{t+1} \in [0, x(d_{t+1}; a_t, d_t, y_t)]$ does not depend on d_{t+1} . In the stochastic collateral constraint case this is also true when the collateral constraint is at its highest value. It is then ensured that the lowest feasible value of a_{t+1} when the collateral constraint takes a lower value, for each choice of d_{t+1} is present in the asset grid. Then, when the EGM step is taken the grid of assets is only evaluate for the subset of a_{t+1} such that $a_{t+1} \geq (\Xi^H - \Xi_t)d_{t+1} > 0$. This point is guaranteed to be on the assets grid.

⁴More precisely, there is a double convergence criterion 1.) that the tolerance is below a specified criterion and 2.) that the current iteration recalculated the optimum d_{t+1}^* for each state.

A.2.6 Solving for the aggregate law of motion

Krusell and Smith

Having solved the household's problem conditional on a law of motion for the aggregate states: $\mathbf{S}_t = G(\mathbf{S}_t)$, I then use the [Krusell et al. \(1998\)](#) algorithm, to solve for the true aggregate law of motion. The stochastic laws of motion for productivity, z , and the collateral constraint, Ξ , are exogenously defined. As in the Krusell and Smith solution I assume bounded rationality, whereby, a limited number of moments are a sufficient statistic for summarising the distribution, Γ . In particular, I assume the mean of the capital stock, K , provides all the additional information an agent requires to form expectations over future prices, $r(z_{t+1}, K_{t+1}, L_{t+1})$ and $w(z_{t+1}, K_{t+1}, L_{t+1})$. The aggregate law of motion is thus:

$$\log(K_{t+1}) = \xi_{z,\Theta}^{K,0} + \xi_{z,\Theta}^{K,1} \log(K_t)$$

where the subscripts emphasise that there is a law of motion for each aggregate current state combination (z, Θ) . Having defined the aggregate laws of motion and solved the household's problem, I simulate the economy with 20,000 agents for 4,150 periods, discarding the first 150 periods. When simulating the model, I use linear interpolation for the household's policy functions in the space (a_t, d_t, K_t) .⁵ To avoid any approximation errors interpolating across non-continuous areas of the policies as a result of adjustment costs the policies are computed in two steps. First, interpolation is carried out for the adjust, V^A , and no-adjust, V^N , value functions and then the policies are interpolated conditional on the optimal adjustment decision e.g. $d_+(\bullet|V^* = V^{A/N})$ and $a_+(\bullet|V^* = V^{A/N})$. Using the simulated series, I recompute the aggregate laws of motion and compare the coefficients to those specified in the model, using the criterion: $(\sum_{z,\Theta,j,i} (\hat{\xi}_{z,\Theta}^{j,i} - \xi_{z,\Theta}^{j,i})^2)^{1/2}$. If this is sufficiently small, the loop exits and the aggregate law of motion is considered to have solved the model. Else, I update the coefficients used in the aggregate law of motion in the household's problems, using a smoothing parameter, resolve the household's problem and continue.

Accuracy of the solution

Critical to the validity of the inference drawn from a model solved using the Krusell and Smith method is the accuracy of the resulting aggregate law of motion. Table B.15 presents summary statistics on the accuracy of the aggregate law of motion in the baseline model and for a version of the model where all agents are subject to the aggregate credit condition. Both models generate high R^2 values for all four exogenous state combinations. As has been commented in the literature, while widely use the R^2 may be a poor statistic for assessing the law of motion accuracy. Also reproduced is the forecast error of using the law of motion for a single period and for 1,000

⁵Notice in this section the asset space is defined over $a_t = k_t + \bar{\Theta}d_t$, see Appendix A.2.5

periods into the future. The one period deviation statistics show the percentage error for predicting the capital stock in the next period. For the baseline model, the average error is very small and the maximum error is around 0.4 percent. Even at the 1,000 period the percentage errors are generally small, with the average error being 0.05 percent and the maximum error being 3.2 percent.

Figure 2.8 graphically illustrates the accuracy of the solution, presenting a sample path of the capital stock along with the laws of motion from the baseline model. Also indicated is the position of the grid points used to approximate the aggregate moments. As can be seen, while the laws of motion do not perfectly coincide with the economy aggregate stocks, the path closely follow each other and there are few significant deviations. The reason for why the capital moment suffices can be seen in Section 2.6.6. In the model variations in aggregate investment are almost entirely due to the productivity shock and in this respect the model resembles the Krussell and Smith result.⁶

The inclusion of a greater number of moments was experimented with, including the durables stock, cross sectional variance of capital and durables across agents, and the share of agents adjusting each period. None of these additions improved the accuracy of the solution sufficiently to justify the higher computational burden.

Determinants of time to adjust

In this model, as Table 2.7 shows, agents are more likely to adjust having experienced a recent history of negative productivity realisations and having experienced a series of tighter aggregate credit conditions. Note this is not the same as the immediate response to a change in the productivity level, discussed in the impulse response section. To understand these relationships in more depth, I undertake the following experiment. For the model implied cross section of assets and durables holdings and every other state in the model's state space $((k, d), \epsilon, \beta, \Xi, z, K, \Theta)$, I use the agents' policy functions to assign their optimal asset and durables choice, conditional on the fact they adjust today $(k^*, d^*)|V = V^A$. Then given this initial starting point, $(k^*, d^*, \epsilon, \beta, z, K, D, \Xi)$, I construct their choices from the following period onwards, to find out the agent's time to adjust, $T^{\check{A}DJ}$, (the period they next do not let durables depreciate) measured in quarters. The stochastic aggregate and idiosyncratic states are held fixed, while the aggregate capital stock follows the agent's expected law of motion. From the time to adjust statistic a "pseudo-probability" of adjustment is calculated:

$P^{\check{A}DJ} = 1/\sqrt{T^{\check{A}DJ}}$. Here there is no variability as the policy choices are deterministic, conditional on the fact that the agent remains in the same idiosyncratic and aggregate state. I then regress the probability of adjustment on the initial conditions to better understand what factors lead an agent to favour non-adjustment over adjustment. The resulting equation

⁶This is not necessarily true when the calibration implies a large stock of durables, in which case an additional moment may be required.

is:

$$\begin{aligned}
P^{\check{A}DJ} = & 0.4189 - 0.0006k^* \cdot (k^* \geq 0) - 0.0071k^* \cdot (k^* < 0) + 0.0056d^* + 0.0822d^*/(d^* + k^*) \\
& - 0.1351\varepsilon - 0.0046[\beta = \beta^H] - 0.0017K + 0.0230[z = z^g] \\
& - 0.0633[\Xi^i = \Xi^H] + 0.0246[\Theta = \Theta^H] + 0.0026[k^* = -\Xi^i d^*] \\
& - 0.0663[k^* = -\Theta d^*] - 0.0063[\Xi^i > \Theta]
\end{aligned}$$

The R^2 for the resulting regression is 0.44, suggesting a reasonably good fit and that relationships should be fairly consistent across a large proportion of agents.⁷

The probability of adjustment is reduced by an agent's current asset position and durables stock and is lower for patient and employed agents. Notice also that being at the non-adjust collateral constraint, Ξ_t^i , implies an increase in the probability of adjustment as agents are forced by the constraint to act, whereas being at the adjust collateral constraint, Θ , reduces the possibility. The probability of adjustment is further reduced when an agents credit terms are more favourable than the aggregate credit conditions. Considering the aggregate states, being in the high productivity state increases the probability of adjustment. Thus, it is the smaller average size of agents' durables stock that is driving the higher probability of adjustment following a history of below average shocks. The looser aggregate credit conditions reduce the probability of adjustment, following the results presented in the conditional moments. Combined with the result for whether the agent began at the collateral constraint, this implies that there is a subset of agents that do not begin at the collateral constraint, that then find themselves constrained and adjust earlier than they would under looser aggregate credit conditions.

⁷Such a specification implicitly assumes monotonic relationships, which is unlikely to be the case, but should be a guide to average behaviour.

A.3 Appendix to (S)Cars and the Great Recession

A.3.1 Solution method

For ease of exposition and with no loss of generalisation we condense the decision whether to adjust or not into the choice of d' . Further, given the positive interest rate spread $r_k > 0$ it is not optimal for a household to hold both savings and a car loan. Therefore, we can consider a single asset k , with a kink in the interest rate schedule, and shocks to the return that are only present when $k > 0$.

Redefine variables in terms of permanent assets and income

$$\begin{aligned}
 V_a(k, d, r_k, q, P, U, G_{a,j}) &= \max_{c, d', k'} u(c, d') + \beta \mathbf{E}V_{a+1}(k', d', r'_k, q', P', U', G_{a+1,j'}) \\
 & \text{s.t.} \\
 c + qk' \cdot \mathbb{1}[k' \geq 0] + k' \cdot \mathbb{1}[k' < 0] + d' &= (1+r)qk \cdot \mathbb{1}[k \geq 0] + (1+r+r^k)k \cdot \mathbb{1}[k < 0] + \\
 & \quad (1-\delta)d - \Psi d \cdot \mathbb{1}[\text{adj}] + Y \\
 k' &\geq -\eta d' \\
 Y &= PU \\
 P &= G_{a,j} P_{-1} V \\
 q &= q_{-1} W
 \end{aligned}$$

Given the restriction $q > 0$ we can define⁸:

$$\mathbf{k}' = \begin{cases} qk' & \text{if } k' \geq 0 \\ k' & \text{if } k' < 0 \end{cases}$$

And therefore if $k' > 0$:

$$\begin{aligned}
 qk &= \frac{q}{q_{-1}} \mathbf{k} \\
 &= W\mathbf{k}
 \end{aligned}$$

⁸If the interest rate premia is not impose for $k < 0$ we cannot make this simplification as $\mathbf{k} > 0$ does not guarantee $k > 0$

We can rewrite the problem, removing the asset price q as a state variable:

$$\begin{aligned}
V_a(\mathbf{k}, d, r_k, P, U, G_{a,j}) &= \max_{c, d', \mathbf{k}'} u(c, d') + \beta \mathbf{E} V_{a+1}(\mathbf{k}', d', r'_k, P', U', G_{a+1, j'}) \\
& \text{s.t.} \\
c + \mathbf{k}' + d' &= (1+r)W\mathbf{k} \cdot \mathbb{1}[\mathbf{k} \geq 0] + (1+r+r^k)\mathbf{k} \cdot \mathbb{1}[\mathbf{k} < 0] + \\
& \quad (1-\delta)d - \Psi d \cdot \mathbb{1}[\text{adj}] + Y \\
\mathbf{k}' &\geq -\eta d' \\
Y &= PU \\
P &= G_{a,j} P_{-1} V
\end{aligned}$$

Now redefine assets in terms of the maximum collateral constraint,
 $b' = \mathbf{k}' + \eta d'$:

$$\begin{aligned}
V_a(b, d, r_k, P, U, G_{a,j}) &= \max_{c, d', b'} u(c, d') + \beta \mathbf{E} V_{a+1}(b', d', r'_k, P', U', G_{a+1, j'}) \\
& \text{s.t.} \\
c + a' + (1-\eta)d' &= (1+r)W(b-\eta d) \cdot \mathbb{1}[(b \geq \eta d)] + (1+r+r^k)(b-\eta d) \cdot \mathbb{1}[(b < \eta d)] + \\
& \quad (1-\delta)d - \Psi d \cdot \mathbb{1}[\text{adj}] + Y \\
a' &\geq 0 \\
Y &= PU \\
P &= G_{a,j} P_{-1} V
\end{aligned}$$

Define variable in current permanent income: $\tilde{c} = c/P$ and yesterday's permanent income: $\hat{b} = b/P_{-1}$ and $\hat{d} = d/P_{-1}$. Divide through by permanent income, following

$$\begin{aligned}
\left(\frac{1}{P}\right)^{1-\rho} V_a(\tilde{b}, \tilde{d}, r_k, U, G_{a,j}) &= \max_{\tilde{c}, \hat{d}', \hat{b}'} u(\tilde{c}, \hat{d}') + \beta \mathbf{E} \left(\frac{1}{P}\right)^{1-\rho} V_{a+1}(\hat{b}', \hat{d}', r'_k, U', G_{a+1, j'}) \\
& \text{s.t.} \\
\tilde{c} + \hat{a}' + (1-\eta)\hat{d}' &= (1+r)W(\tilde{b} - \eta\tilde{d}) \cdot \mathbb{1}[(\tilde{b} \geq \eta\tilde{d})] + \\
(1+r+r^k)(\tilde{b} - \eta\tilde{d}) \cdot \mathbb{1}[(\tilde{b} < \eta\tilde{d})] &+ \\
& \quad (1-\delta)\tilde{d} - \Psi\tilde{d} \cdot \mathbb{1}[\text{adj}] + U \\
\hat{a}' &\geq 0 \\
\tilde{b}' &= \frac{P}{P'} \hat{b}', \quad \hat{d}' = \frac{P}{P'} \hat{d}'
\end{aligned}$$

Let $\tilde{V}_a(\cdot) = \left(\frac{1}{P}\right)^{1-\rho} V_a(\cdot)$. Then as $P/P' = G_{a+1,j'}V'$:

$$\begin{aligned} \tilde{V}_a(\tilde{b}, \tilde{d}, r_k, U, G_{a,j}) &= \max_{\tilde{c}, \hat{d}', \hat{b}'} u(\tilde{c}, \hat{d}') + \beta \mathbf{E} (G_{a+1,j'}V')^{1-\rho} \tilde{V}_{a+1}(\tilde{b}', \tilde{d}', r'_k, U', G_{a+1,j'}) \\ & \quad s.t. \\ & \quad \tilde{c} + \hat{a}' + (1 - \eta)\hat{d}' = (1 + r)W(\tilde{b} - \eta\tilde{d}) \cdot \mathbb{1}[(\tilde{b} \geq \eta\tilde{d})] \\ & \quad + (1 + r + r^k)(\tilde{b} - \eta\tilde{d}) \cdot \mathbb{1}[(\tilde{b} < \eta\tilde{d})] + \\ & \quad \quad (1 - \delta)\tilde{d} - \Psi\tilde{d} \cdot \mathbb{1}[\text{adj}] + U \\ & \quad \tilde{a}' \geq 0 \\ & \quad \tilde{b}' = \frac{\hat{b}'}{G_{a+1,j'}V'}, \quad \tilde{d}' = \frac{\hat{d}'}{G_{a+1,j'}V'} \end{aligned}$$

Finally, rewrite problem in terms of cash in hand, \tilde{x} :

$$\begin{aligned} \tilde{V}_a(\tilde{x}, \tilde{d}, r_k, G_{a,j}) &= \max_{\tilde{c}, \hat{d}', \hat{b}'} u(\tilde{c}, \hat{d}') + \beta \mathbf{E} (G_{a+1,j'}V')^{1-\rho} \tilde{V}_{a+1}(\tilde{x}', \tilde{d}', r'_k, U', G_{a+1,j'}) \\ & \quad s.t. \\ & \quad \tilde{c} + \hat{a}' + (1 - \eta)\hat{d}' = x - \Psi\tilde{d} \cdot \mathbb{1}[\text{adj}] \\ & \quad \tilde{a}' \geq 0 \\ & \quad \tilde{d}' = \frac{\hat{d}'}{G_{a+1,j'}V'} \\ & \quad \tilde{x}' = \begin{cases} (1 + r)W(\hat{b}' - \eta\hat{d}')/(G_{a+1,j'}V') + \\ (1 - \delta)\hat{d}'/(G_{a+1,j'}V') + U & \text{if } \hat{b}' \geq \eta\hat{d}' \\ (1 + r + r^k)(\hat{b}' - \eta\hat{d}')/(G_{a+1,j'}V') + \\ (1 - \delta)\hat{d}'/(G_{a+1,j'}V') + U & \text{if } \hat{b}' < \eta\hat{d}' \end{cases} \end{aligned}$$

Computation

The model is solved by Value Function Iteration. We use 200 grid points for cash in hand, x , 200 grid points for assets a' , 150 grid points for cars, d and 5 grid points for the interest rate spread, r_k . Expectations are taken over future shocks, using 5 grid points for permanent income, V , 4 grid points for transitory shocks U and 5 grid points for the asset price shock, W .

As the model is partial equilibrium, the household does not need to distinguish between aggregate and idiosyncractic shocks to permanent income, V .

We then simulate a panel of households with 500 household born age 25 each period, for 2,000 periods to calculate the aggregate properties of the economy. We also simulate the lifecycle of a panel of households without aggregate shocks to uncover the lifecycle properties.

Finally, we feed in a series of shocks estimated from the data to income $\{V_t\}_{t=1971}^{2013}$, the interest rate spread $\{r_k\}_{t=1972}^{2013}$, the asset price $\{W_t\}_{t=1976}^{2013}$ and deterministic growth rate of economy to replicate the behaviour of the economy in the period 1981-2013

A.3.2 Estimating growth profile

Growth rate measurement

The primary measure is Financial Income Before Tax. We drop household who have income below \$10,000 or work less the 20 hours a week. For the estimation of the growth rate annual household income aggregated to year, t , age, a cell: Y_t^a . Then for each year age income growth is calculated:

$$dY_t^a = Y_t^a - Y_{t-1}^{a-1}$$

For each age the average of these growth rates across years was calculated for the pre-recession (1990-2005) and post-recession (2010-2012) period.

$$d\hat{Y}^{a,\text{pre}} = \frac{1}{T^{\text{pre}}} \sum_{t=1990}^{2006} dY_t^a$$

$$d\hat{Y}^{a,\text{post}} = \frac{1}{T^{\text{post}}} \sum_{t=2010}^{2012} dY_t^a$$

Having calculated the average growth rate at each age in the pre- and post-recession period we fit a polynomial $f(a)$ across all ages to smooth the pattern:

$$d\hat{Y}^{a,\text{x}} = g^{\text{x}}(a) + \epsilon_a \quad (\text{A.1})$$

The implied lifecycle can be calculated by cumulating the estimated growth rate polynomials $\hat{g}^{\text{pre}}(a)$ and $\hat{g}^{\text{post}}(a)$.

Lifecycle measurement

Lifecycle estimation is directly estimating the lifecycle by regression of income on age polynomial and controls, using household level data. The pre-recession period is 1989-2005 and the post recession period is 2009-12. Controls are included for demography, education (e), cohort (j) and year. The estimation equation is:

$$y_{it} = \alpha_0 + f(a) + \sum_j^M \gamma^j + \sum_e^E \beta^e + \alpha_1 \text{race}_{it} + \phi t + \epsilon_{it} \quad (\text{A.2})$$

Having estimated the lifecycle polynomials $\hat{f}^{\text{pre}}(a)$ and $\hat{f}^{\text{post}}(a)$ the implied growth rate is then computed, $\hat{g}_{LC}^{\text{pre}}(a)$ and $\hat{g}_{LC}^{\text{post}}(a)$.

A.3.3 Growth rate shock calibration

Given model specification:

$$G^{\text{post}} = \min\{(G^{\text{pre}})^\gamma, G\} \quad (\text{A.3})$$

and growth rate polynomials, we estimate $\hat{\gamma}$ to find model that fits decline: $\hat{g}^{\text{pre}}(a)$ to $\hat{g}^{\text{post}}(a)$. We estimate $\hat{\gamma}$ for both the growth rate estimated series and implied growth from the lifecycle estimated series and take the average of these two measures.

A.3.4 Estimating cohort shocks

Our income measure of choice is family labor earnings. We use a sample in the CEX of households aged 24-60, for the years 1980-2013. We generate shocks for 10 year cohorts, such that a household is a member of cohort (s) if born in the 10 year period 19s0 to 19s9.

We first normalise the aggregate CEX income series to match with NIPA income. If log NIPA income is x_t and log family earnings in the CEX is \tilde{y}_{it} , we find β_t to satisfy:

$$\begin{aligned} y_{it}^{nipa} &= \beta_t \tilde{y}_{it} \\ s.t. \\ x_t &= \frac{1}{N_t^i} \sum_i^{N_t^i} y_{it}^{nipa} \end{aligned}$$

We then construct a “year of birth” synthetic cohort, j , such that year of birth earnings is:

$$y_{jt} = \frac{1}{N_t^j} \sum_{i \in j} y_{it}^{nipa}$$

To reduce measure error we smooth this income series, with the exception of 2008-10 to match the depth of the Great Recession.

$$\bar{y}_{jt} = \frac{1}{3} \sum_{i=-1}^1 y_{jt+i}$$

To find the permanent income component we regress the synthetic cohort income data on a lifecycle age polynomial, cohort (s) dummies, and linear trend, using data for ≤ 2007 :

$$\bar{y}_{jt} = \alpha + f(\text{age}_{jt}) + \sum_s \gamma^s + \phi t + \xi_{jt}$$

We then treat the residual as the year of birth income level: $\xi_{jt} \forall t$. Given the random walk in permanent income a shock for a year of birth cohort is:

$$\epsilon_{jt} = \xi_{jt} - \xi_{jt-1} \forall t$$

We then average over these year of birth cohort shocks to find the cohort

shocks that we feed into the model.

$$\hat{\epsilon}_{st} = \frac{1}{N_t^s} \sum_{j \in s} \epsilon_{jt}$$

Averaging minimize measurement error due to the mismeasurement of year of birth shocks on small sample sizes. It also enables us to draw more consistent conclusions about the effects of the Great Recession on groups of a similar life-cycle position.

A.3.5 Estimating cohort responses

To estimate the deviations in the consumption response of households during the Great Recession, we regress the log of a consumption variable (x_{it}) on a lifecycle age polynomial, cohort (s) dummies, and linear trend, using data for ≤ 2007 :

$$x_{it} = \alpha + f(\text{age}_i) + \sum_s \gamma^s + \phi t + \epsilon_{it}$$

For each household we can now predict consumption in each year during the Great Recession, in the absence of the crisis as:

$$\hat{x}_{it} = \hat{\alpha} + \hat{f}(\text{age}_{it}) + \sum_s \hat{\gamma}^s + \hat{\phi} t$$

We divide the sample up into three cohorts j based on their age in 2007: $j = 1$ if $25 \leq \text{age}_{2007} < 34$, $j = 2$ if $35 \leq \text{age}_{2007} < 44$ and $j = 3$ if $45 \leq \text{age}_{2007} < 54$. We measure average consumption for a cohort in year t as:

$$X_{jt} = \frac{1}{N_t^j} \sum_{i \in j} x_{it}$$

Actual consumption growth for year t in the recession is, growth relative to the 2007 baseline year⁹:

$$\Delta X_{jt} = X_{jt} - X_{j,2007}$$

While predicted consumption growth is the analogue:

$$\Delta \hat{X}_{jt} = \hat{X}_{jt} - \hat{X}_{j,2007}$$

Again to reduce measurement error, we also tend to focus on the average value of consumption during the recession. i.e.:

$$\Delta \hat{X}_{j,08:10} = \frac{1}{3} \sum_{t=2008}^{2010} \hat{X}_{jt} - \hat{X}_{j,2007}$$

⁹For the value of car purchase conditional on adjustment we use a baseline year of 2006-07, to increase the available observations

The reported consumption deviation is then measured as:

$$\Omega_x^j = \Delta X_{j,08:10} - \Delta \hat{X}_{j,08:10}$$

We undertake the same operation on model generated data. The measures are comparable as we are controlling for cohort specific level effects and trend growth in the data.

Appendix B

Additional Tables

B.1 Additional Tables for Job Risk, Separation Shocks and Household Asset Allocation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Job risk (δ)	3.904*** (0.966)		5.154*** (1.163)	5.911*** (1.225)	-0.075 (3.250)	0.565 (1.687)	3.736*** (1.296)	3.339*** (1.545)
$\delta : \leq 1^{st}pc.$		-0.058 (0.042)						
$\delta : \leq 5^{th}pc.$		-0.040* (0.024)						
$\delta : \leq 10^{th}pc.$		-0.125*** (0.024)						
$\delta : \leq 25^{th}pc.$		-0.026* (0.014)						
Income	0.207*** (0.011)	0.207*** (0.011)	0.247*** (0.014)	0.159*** (0.016)	0.201*** (0.012)	0.056*** (0.012)	0.215*** (0.012)	0.213*** (0.013)
Core			✓					
Risk pref.				✓				
Ind & Occ					✓			
FE						✓		
IV							δ_{t-2}	δ_{t-6}
N	33,953	34,629	23,742	22,234	33,953	34,834	28,726	22,005
R ²	0.046	0.046	0.0485	0.090	0.052	0.012	0.049	0.051

Notes: Dependent variable is the transformation of the ratio of liquid assets to illiquid assets. δ Percentiles are dummy variables with $\delta : \leq 1^{st}pc.$ referring to the 1st percentile, $\delta : \leq 5^{th}pc.$ referring to 5th percentile < $\delta \leq 5^{th}$ percentile etc. Core refers to use of only core PSID sample, Risk pref. includes controls for stock market exposure, insurance purchases and total assets. Ind & Occ includes industry and occupation dummy in main specification. FE is panel specification with individual fixed effects. IV uses lagged job risk. *Statistically significant at 10%; **statistically significant at 5%; ***statistically significant at 1%

Table B.1: Ratio of liquid assets to illiquid assets, CPS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Job risk (δ)	-1.396** (0.670)		-1.385** (0.786)	-1.554** (0.883)	10.593*** (2.220)	0.261 (1.417)	-1.980** (0.968)	-0.876*** (1.284)
$\delta : \leq 1^{st} pc.$		-0.067** (0.038)						
$\delta : \leq 5^{th} pc.$		-0.018 (0.017)						
$\delta : \leq 10^{th} pc.$		0.023 (0.015)						
$\delta : \leq 25^{th} pc.$		0.006 (0.010)						
Income	-0.082*** (0.006)	-0.082*** (0.006)	-0.080*** (0.008)	-0.122*** (0.009)	-0.082*** (0.007)	-0.046*** (0.007)	-0.074*** (0.008)	-0.057*** (0.009)
Core			✓					
Risk pref.				✓				
Ind & Occ					✓			
FE						✓		
IV							δ_{t-2}	δ_{t-6}
N	46,699	47,680	30,477	29,904	46,669	47,919	37,050	26,859
R ²	0.052	0.050	0.043	0.055	0.055	0.048	0.041	0.034

Notes: Dependent variable is the transformation of the ratio of liquid assets to total assets. δ Percentiles are dummy variables with $\delta : \leq 1^{st} pc.$ referring to the 1st percentile, $\delta : \leq 5^{th} pc.$ referring to 5th percentile < $\delta \leq 5^{th}$ percentile etc. Core refers to use of only core PSID sample, Risk pref. includes controls for stock market exposure, insurance purchases and total assets. Ind & Occ includes industry and occupation dummy in main specification. FE is panel specification with individual fixed effects. IV uses lagged job risk. *Statistically significant at 10%; **statistically significant at 5%; ***statistically significant at 1%

Table B.2: Ratio of liquid assets to total assets, CPS

	PHTM			WHTM		
	(1)	(2)	(3)	(4)	(5)	(6)
Job risk (δ)	0.541*** (0.073)	1.118*** (0.120)	0.580*** (0.091)	-0.765*** (0.092)	-1.101*** (0.146)	-0.773*** (0.121)
Ind & Occ		✓			✓	
IV			δ_{t-2}			δ_{t-2}
N	31,794	31,794	26,472	31,794	31,794	26,472
R^2	0.185	0.190	0.169	0.017	0.022	0.018

Notes: *Statistically significant at 10%; **statistically significant at 5%;
***statistically significant at 1%

Table B.3: Hand-to-Mouth status, Core

	PHTM			WHTM		
	(1)	(2)	(3)	(4)	(5)	(6)
Job risk (δ)	1.034*** (0.365)	10.289*** (1.139)	1.148*** (0.541)	-1.333*** (0.410)	-10.945*** (1.285)	-2.031*** (0.654)
Ind & Occ		✓			✓	
IV			δ_{t-2}			δ_{t-2}
N	51,508	51,508	40,058	51,508	51,508	40,058
R^2	0.232	0.238	0.210	0.016	0.023	0.014

Notes: *Statistically significant at 10%; **statistically significant at 5%;
***statistically significant at 1%

Table B.4: Hand-to-Mouth status, CPS

Moment	lo ω , lo δ	lo ω , hi δ	hi ω , lo δ	hi ω , hi δ
Job risk (δ)				
<i>shock %:</i>	0.226	0.718	1.292	0.769
<i>persistence:</i>	0.947	0.939	0.777	0.924
Wage (ω)				
<i>shock %:</i>	-0.005	0.000	-0.042	-0.016
<i>persistence:</i>	0.885	0.984	0.984	0.63

Table B.5: Great Recession shocks

Moment	Hi Av. δ	Pre GR	
	level	level	$\Delta\%$
E-U transition rate	0.034	0.025	-25.4
Av δ	0.047	0.035	-25.4
$E(\delta U)$	0.106	0.077	-27.4
Av. wage	1.191	1.233	3.5
$\sigma(\delta)$	0.051	0.034	-32.2
c	0.850	0.886	4.3
b	1.661	1.523	-8.4
h	4.619	5.478	18.6
eq	2.899	2.985	3.0
LTV	0.376	0.466	24.1
% renters	0.403	0.313	-22.2
HTM	0.097	0.116	19.9
MPC (+ve)	0.093	0.101	8.6
$\rho(\omega, \delta)$	0.370	0.351	-5.2

Pre GR is the pre-Great Recession equilibria implemented by reducing the separation rate.

Table B.6: Comparison of High average δ and Pre-Great Recession Equilibria

B.2 Additional Tables for Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics

Variable	$\frac{St.dev(x)}{St.dev(y)}$	$corr(x, y)$	$corr(x, I^k)$	$corr(x, I^d(3))$	1 st autocorr	2 nd autocorr
C	0.009	0.55	0.78	0.56	0.59	0.84
I^k	0.078	4.82	0.83	1.00	0.48	0.73
$I^d(1)$	0.048	2.94	0.69	0.52	0.93	0.73
$I^d(2)$	0.103	6.35	0.61	0.41	0.95	0.90
$I^d(3)$	0.063	3.89	0.68	0.48	1.00	0.86
Y	0.016	1.00	1.00	0.83	0.68	0.84
CEA	0.025	1.56	0.70	0.56	0.83	0.94
$LTV(1)$	0.023	1.40	0.19	0.11	0.24	0.74
$LTV(2)$	0.016	0.97	0.46	0.45	0.50	0.73

Source: All series except LTV from Bureau of Economic Analysis, tables 1.1.5 and 2.3.5. LTV from Board of Governors of the Federal Reserve, Consumer Credit, G.19. All series except CEA are logged and HP filtered, with smoothing parameter 1600. CEA is HP filtered. Aggregate series are per capita and have been deflated using GDP deflator.

$I^d(1)$: consumer durable goods, $I^d(2)$: residential investment, $I^d(3)$: durable goods and residential investment. $LTV(1)$: new auto purchases loan to value, $LTV(2)$: housing loan to value.

Table B.7: Business cycle moments: US data 1950.I-2014.I

Variable	1950.I-1982.IV			1983.I-2014.I				
	$\frac{St.dev(x)}{St.dev(y)}$	$corr(x, y)$	$corr(x, I^k)$	$corr(x, I^d)$	$\frac{St.dev(x)}{St.dev(y)}$	$corr(x, y)$	$corr(x, I^k)$	$corr(x, I^d)$
C	0.49	0.75	0.44	0.63	0.67	0.84	0.57	0.53
I^k	4.61	0.80	1.00	0.45	5.29	0.91	1.00	0.56
$I^d(1)$	3.01	0.67	0.50	0.94	2.79	0.75	0.57	0.90
$I^d(2)$	6.10	0.59	0.35	0.95	6.94	0.64	0.51	0.95
$I^d(3)$	3.86	0.67	0.45	1.00	3.93	0.73	0.56	1.00
CEA	1.39	0.69	0.38	0.89	1.92	0.74	0.67	0.80
$LTV(1)$	0.60	0.62	0.35	0.66	2.08	0.11	0.07	0.16
$LTV(2)$	0.81	0.40	0.08	0.65	1.27	0.54	0.62	0.44

Table B.8: Business cycle moments: US data, sub-time periods

	Full dataset			Model consistent		
	$d(1)$	$d(2)$	$d(3)$	$d(1)$	$d(2)$	$d(3)$
Net financial assets (k)						
<i>median : mean ratio</i>	0.01	0.00	-0.02	0.02	0.00	-0.01
$\frac{s.t dev}{mean}$	11.6	28.5	32.2	11.0	16.3	17.8
<i>skewness</i> 110.3	90.3	90.2	105.4	98.7	100.8	
Durables (d)						
<i>median : mean ratio</i>	0.55	0.38	0.42	0.56	0.36	0.43
$\frac{s.t dev}{mean}$	4.3	4.2	3.9	4.3	4.0	3.8
<i>skewness</i>	145.7	122.3	114.7	140.0	113.6	109.0
Av. dur. loan:value $loan > 0$	0.56	0.59	0.58	0.48	0.49	0.49
% in debt ($k < 0$)	0.22	0.41	0.51	0.20	0.40	0.50
Durables share of wealth $\frac{d}{w}$	1.32	3.12	2.02	0.76	1.03	1.31
Observations	142,659	412,659	42,659	132,990	124,134	129,180

Source: Survey of Consumer Finances 2010, Federal Reserve Board. Model consistent restricts the permissible durables share of wealth ratio so all observations obey the relaxed model collateral constraint, i.e. $0 \leq d/w < 1/(1 - \Xi)$. Figures are for households with head aged between 18 and 65, and are presented on a per household member basis to aid comparability with the model.

Table B.9: Survey of Consumer Finances, 1989-2010

Type	Paper	Specification	Parameter values
Current stock	Berger and Vavra (2015)	$0 \text{ if } d_{t+1} = d_t(1 - \delta(1 - \chi))$ $\Psi(1 - \delta)d_t + \Psi_2 w \varepsilon_t \text{ else}$	$\Psi = 0.0525$ $\Psi_2 = 0.001$ $\chi = 0.8$
	Jose Luengo-Prado (2006)	$0 \text{ if } d_{t+1} = d_t(1 - \delta)$ $\Psi(1 - \delta)d_t \text{ else}$	$\Psi = 0.05$
Maintain level	Iacoviello and Pavan (2013b)	$0 \text{ if } d_{t+1} = d_t$ $\Psi d_t \text{ else}$	$\Psi = 0.05$
	Díaz and Luengo-Prado (2010)	$0 \text{ if } d_{t+1} = d_t$ $\Psi(1 - \delta)d_t \text{ else}$	$\Psi = 0.05$
Future stock	Bajari et al. (2013)	$0 \text{ if } d_{t+1} = d_t$ $\Psi d_{t+1} \text{ else}$	$\Psi = 0.06$
Downward adjustment	Guerrieri and Lorenzoni (2017)	$0 \text{ if } d_{t+1} > d_t$ $\Psi(d_t - d_{t+1}) \text{ else}$	$\Psi = 0.15$

Table B.10: Commonly used durable adjustment costs specifications

	I^d	$I^d adj.$	C	I^k	$\%adj$	$\%adj u$	$\%adj d$	$D(8)$	$C(8)$	$K(8)$
$\%debt$	0.58 (3.48)	2.6 (4.54)	-0.12 (3.74)	-0.35 (1.17)	3.49 (6.41)	3.22 (5.62)	3.42 (1.42)	-0.41 (5.75)	-0.25 (4.07)	-0.42 (5.77)
$\%adj.$	-1.12 (10.93)	-3.83 (10.97)	0.2 (10.41)	2.01 (11.13)	-4.45 (13.35)	-4.22 (12.01)	-9.31 (6.33)	0.36 (8.25)	0.45 (12.20)	0.6 (13.55)
K/D	1.01 (4.41)	5.09 (6.48)	-0.15 (3.54)	-1.69 (4.15)	4.85 (6.48)	5 (6.34)	7.03 (2.13)	0.65 (6.59)	-0.32 (3.80)	-0.42 (4.28)
$d/(d+k)$	-1.19 (2.99)	-5.07 (3.72)	0.77 (10.15)	3.66 (5.18)	-9.52 (7.31)	-8.12 (5.93)	-34.32 (5.98)	1.82 (10.63)	1.63 (11.25)	2.31 (13.43)
$med(k)$	0.2 (1.17)	0.46 (0.80)	0.17 (5.33)	0.49 (1.62)	-0.3 (0.53)	-0.03 (0.05)	-9.02 (3.69)	0.44 (6.09)	0.37 (6.02)	0.54 (7.37)
$med(d)$	1.01 (4.42)	4.11 (5.26)	-0.32 (7.39)	-3.06 (7.56)	4.69 (6.30)	4.35 (5.54)	20.85 (6.34)	-0.68 (6.96)	-0.64 (7.72)	-0.82 (8.30)
r	2.37 (5.35)	10.64 (7.02)	-0.35 (4.19)	-4.8 (6.11)	10.86 (7.52)	10.72 (7.05)	23.17 (3.64)	0.41 (2.13)	-0.65 (4.02)	-0.68 (3.56)
Z	-3.25 (7.76)	-12.62 (8.80)	0.58 (7.26)	3.93 (5.29)	-14.32 (10.47)	-14.26 (9.90)	-21.89 (3.63)	0.86 (4.77)	1.26 (8.31)	1.93 (10.67)
Θ	5.36 (11.01)	15.5 (9.30)	-0.81 (8.75)	-5.56 (6.45)	21.03 (13.24)	21 (12.55)	30.07 (4.29)	-1.86 (8.92)	-1.86 (10.57)	-2.91 (13.87)
ncc	2.22 (5.92)	5.71 (4.45)	-0.59 (8.25)	-3.43 (5.16)	9.33 (7.62)	9.2 (7.14)	25.26 (4.68)	-1.46 (9.07)	-1.32 (9.72)	-2.04 (12.58)
ecc	-1.25 (4.65)	-4.96 (5.39)	-0.04 (0.82)	0.69 (1.45)	-5.01 (5.72)	-5.01 (5.43)	-6.11 (1.58)	-0.22 (1.87)	-0.08 (0.80)	-0.06 (0.55)
$n.ecc$	0.39 (1.78)	1.53 (2.06)	-0.13 (3.23)	-1.06 (2.75)	2.35 (3.33)	1.86 (2.50)	5.98 (1.92)	-0.26 (2.78)	-0.24 (3.03)	-0.24 (2.58)
$debt/d adj$	-0.54 (4.75)	-1.4 (3.58)	-0.08 (3.73)	-0.25 (1.22)	-1.25 (3.35)	-1.03 (2.63)	3.33 (2.02)	-0.11 (2.30)	-0.16 (3.86)	-0.21 (4.28)
$debt/d n.adj$	-0.1 (0.41)	1.15 (1.42)	-0.17 (3.71)	-1.86 (4.43)	2.24 (2.89)	1.27 (1.56)	14.58 (4.26)	-0.17 (1.65)	-0.27 (3.10)	-0.2 (1.95)
R^2	0.47	0.4	0.2	0.18	0.48	0.47	0.18	0.4	0.29	0.47
N	1997	1997	1997	1997	1997	1997	1997	1997	1997	1997

Notes: Results from 100,000 periods of simulation of the model. Prior conditions are specified as average value in the previous 16 quarters, to negative shock to aggregate credit conditions or productivity. Prior states are de-meanded and the standard deviation set to 100 for readability. T-statistic in parenthesis.

Table B.11: Conditional responses: regressions of Θ IRFs on state space

	I^d	$I^d adj.$	C	I^k	$\%adj$	$\%adj u$	$\%adj d$	$D(8)$	$C(8)$	$K(8)$
$\%debt$	0.42 (2.19)	1.99 (2.62)	0.04 (5.72)	-0.01 (0.12)	1.49 (1.85)	1.56 (1.98)	-24.27 (5.25)	0.1 (2.66)	0.1 (8.51)	0.06 (5.68)
$\%adj.$	-1.54 (16.32)	-5.53 (14.77)	-0.02 (5.57)	0.72 (13.80)	-5.97 (14.98)	-5.69 (14.65)	-7.2 (3.17)	0 (0.23)	0 (0.70)	-0.01 (1.34)
K/D	1.04 (4.43)	4.17 (4.48)	-0.01 (0.99)	-0.36 (2.83)	3.47 (3.50)	3.88 (4.02)	-7.47 (1.32)	0.88 (19.65)	-0.03 (2.16)	-0.07 (5.30)
$d/(d+k)$	-0.21 (0.53)	-3.3 (2.07)	-0.04 (2.25)	-0.94 (4.25)	-3.33 (1.96)	-3.21 (1.95)	11.44 (1.18)	-0.73 (9.54)	0 (0.11)	0.09 (4.06)
$med(k)$	0.31 (1.94)	0.64 (1.03)	0 (0.29)	-0.44 (5.11)	0.73 (1.11)	0.72 (1.12)	0.84 (0.22)	-0.14 (4.82)	0 (0.05)	0 (0.42)
$med(d)$	1.02 (4.88)	3.96 (4.79)	0.02 (2.83)	-0.41 (3.62)	3.75 (4.26)	3.47 (4.04)	5.29 (1.05)	-0.18 (4.61)	0.05 (4.12)	0.06 (5.31)
r	2.06 (5.97)	7.94 (5.79)	0.05 (3.99)	-1.02 (5.39)	7.56 (5.18)	7.72 (5.43)	-1.75 (0.21)	0.64 (9.68)	0.15 (7.14)	0.27 (14.18)
Z	-1.22 (3.29)	-6.61 (4.50)	-0.16 (11.01)	-0.26 (1.26)	-5.28 (3.37)	-5.28 (3.46)	30.9 (3.46)	-0.79 (11.20)	-0.33 (15.11)	-0.43 (21.37)
Θ	-0.3 (0.61)	2.26 (1.15)	-0.02 (1.05)	1.87 (6.85)	2.41 (1.15)	2.29 (1.12)	6.04 (0.50)	1.64 (17.27)	-0.18 (6.06)	-0.3 (11.30)
ncc	-1.25 (3.34)	-1.88 (1.27)	-0.01 (0.75)	1.67 (8.18)	-1.92 (1.22)	-2.25 (1.47)	18.59 (2.07)	0.94 (13.16)	-0.09 (3.91)	-0.17 (8.44)
ecc	-0.56 (2.09)	-1.97 (1.84)	0.01 (0.93)	0.51 (3.44)	-1.5 (1.32)	-1.79 (1.62)	8.6 (1.33)	-0.18 (3.42)	0 (0.18)	0.02 (1.64)
$n.ecc$	0.96 (4.61)	4.34 (5.27)	0.02 (2.53)	-0.46 (4.04)	5.49 (6.27)	4.49 (5.26)	1.11 (0.22)	-0.07 (1.69)	0.01 (1.07)	0.02 (2.17)
$debt/d adj$	-0.42 (3.80)	-1.8 (4.11)	0.03 (6.05)	0.43 (7.10)	-2.45 (5.26)	-1.8 (3.95)	-24.86 (9.35)	0.11 (5.11)	0.04 (6.58)	0.02 (2.56)
$debt/d n.adj$	1.24 (5.11)	4.88 (5.09)	-0.02 (1.86)	-0.86 (6.46)	5.39 (5.28)	5.16 (5.18)	22.86 (3.92)	0.12 (2.65)	-0.03 (1.76)	-0.02 (1.77)
R^2	0.28	0.17	0.67	0.41	0.2	0.2	0.32	0.7	0.83	0.93
N	1997	1997	1997	1997	1997	1997	1997	1997	1997	1997

Notes: Results from 100,000 periods of simulation of the model. Prior conditions are specified as average value in the previous 16 quarters, to negative shock to aggregate credit conditions or productivity. Prior states are de-meanded and the standard deviation set to 100 for readability. T-statistic in parenthesis.

Table B.12: Conditional responses: regressions of small Z IRFs on state space

Dependent Variable	$\Delta I^d(1)$	$\Delta I^d(2)$	$\Delta I^d(3)$	ΔI^k
<i>lag. dep var</i>	-0.192 (0.002)	0.394 (0.072)	0.127 (0.048)	0.159 (0.091)
<i>TFP shock</i>	0.944 (0.120)	0.766 (0.200)	0.930 (0.188)	1.853 (0.390)
<i>CEA shock</i>	0.139 (0.126)	1.895 (0.375)	1.591 (0.191)	0.402 (0.414)
<i>CEA HP level</i>		0.029 (0.389)	0.290 (0.266)	1.044 (0.379)
\bar{R}^2	0.364	0.504	0.482	0.164
<i>N</i>	191	191	191	191

Heteroskedastic autocorrelation robust standard errors in parenthesis. Dependent variable is defined $\Delta X = 100 * (\log(X) - \log(X_{-1}))$. Independent variables are rescaled such that the standard deviation of the series is one. *TFP shock* is estimate TFP residuals. *CEA shock* is the residual of an AR(1) process with a linear trend. *CEA HP level* is level of the HP filtered series.

Table B.13: TFP and Credit shock regression results (difference)

Dependent Variable	$I_D^d(1)$	$I_D^d(2)$	$I_D^d(3)$	I_K^k
<i>lag. dep var</i>	0.929	0.982	0.973	0.778
(0.018)	(0.027)	(0.024)	(0.041)	(0.041)
<i>TFP shock</i>	0.062	0.013	0.025	0.042
(0.015)	(0.018)	(0.007)	(0.008)	(0.008)
<i>CEA shock</i>	0.078	0.025	0.037	-0.003
(0.009)	(0.003)	(0.005)	(0.009)	
<i>CEA HP level</i>	0.045	0.007	0.017	0.173
(0.017)	(0.009)	(0.007)	(0.007)	
\bar{R}^2	0.912	0.972	0.972	0.576
N	191	191	191	191
			192	192

Heteroskedastic autocorrelation robust standard errors in parenthesis. Dependent variable the investment ratio rescaled such that $I^X/X \in (0, 100)$. Independent variables are rescaled such that the standard deviation of the series is one. *TFP shock* is estimate TFP residuals. *CEA shock* is the residual of an AR(1) process with a linear trend. *CEA HP level* is level of the HP filtered series.

Table B.14: TFP and Credit shock regression results (ratio)

	<i>Ave.</i>	$\{b, l\}$	R^2		Forecast error (%)				
			$\{b, h\}$	$\{g, l\}$	$\{g, h\}$	One period		1,000 period	
						<i>mean</i>	$ max $	<i>mean</i>	$ max $
Baseline	0.9996	0.9998	0.9995	0.9997	0.9994	0.0011	0.4209	0.0472	3.2099
Agg. model	0.9999	0.9998	0.9998	0.9999	0.9999	0.0014	0.1987	0.0554	3.2385

Table B.15: Accuracy of aggregate law of motion

B.3 Additional Tables for (S)Cars and the Great Recession

	1981-2007		2007-12	
	25-59	60-80	25-59	60-80
Non durables	33,750.3	29,568.6	33,755.2	32,909.8
Car purchases	3,581.6	2,206.3	3,439.6	2,583.6
Car purchase adj	16,906.3	18,026.4	21,653.2	22,210.2
New car purchases adj	27,838.2	26,195.2	36,416.4	31,489.8
Old car purchases adj	11,198.8	11,347.7	15,132.7	14,575.5
% car purchase	0.212	0.122	0.159	0.116
% new car purchase	0.067	0.053	0.045	0.050
% old car purchase	0.153	0.073	0.119	0.070
Car stock	15,950.1	14,081.9	15,406.8	14,802.1
Number of cars	1.154	1.087	0.877	0.913
Age of cars	7.935	8.265	8.646	9.193
Employed	0.922	0.354	0.904	0.414
Hours worked	40.4	12.4	39.0	15.2
Family income (before tax)	72,763.6	46,892.7	75,822.9	56,010.4
Family income (after tax)	67,059.4	44,061.6	72,802.4	53,851.3
Family labor earnings	62,390.3	14,898.1	64,983.7	21,557.9
Head labor earning	50,510.1	12,191.5	51,831.9	17,137.2
Age	40.2	68.9	42.1	68.3

Table B.16: Consumer Expenditure Survey Summary Statistics

Appendix C

Additional Figures

C.1 Additional Figures for Job Risk, Separation Shocks and Household Asset Allocation




Figure C.1: Consumption drop during Great Recession




Figure C.2: Job finding rate




Figure C.3: Weekly earnings in Great Recession




Figure C.4: Job risk distribution with alternative income controls

The alternative specification are no income control; lagged total income; and weekly earnings over the more recent time period, 1997-2016. Lagged total income is total pre-tax total income. This variable is available annually in the Annual Social and Economic Supplement (ASEC) of the CPS and refers to income accrued over the past year, with each individual featuring in two ASEC surveys. To avoid the income measure being affected by periods of unemployment, only monthly observations that occur in the same month or after the ASEC are included.

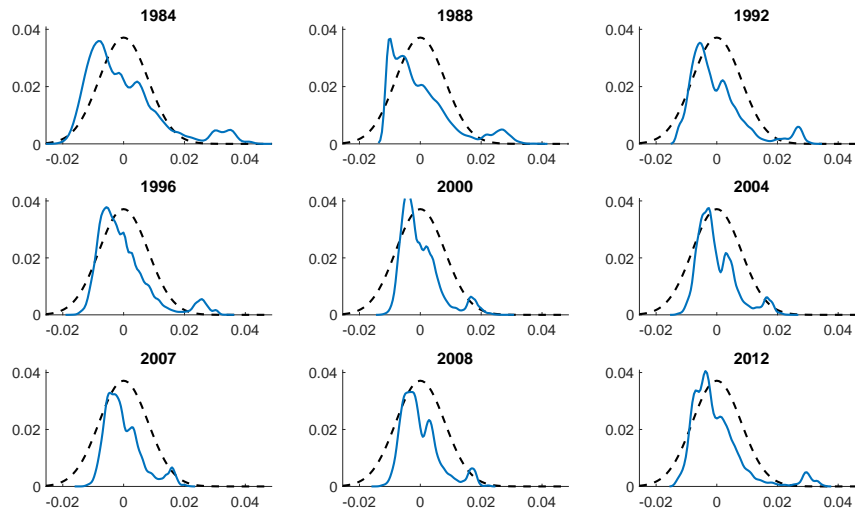





Figure C.5: Distribution of US job risk by year




(a) Housing




(b) Liquid assets*




(c) House equity



(d) Liquid assets



(e) liquid/illiquid



(f) liquid/wealth

Figure C.6: Portfolio choices during decline in job risk




Figure C.7: Job separation rate of jobs by age

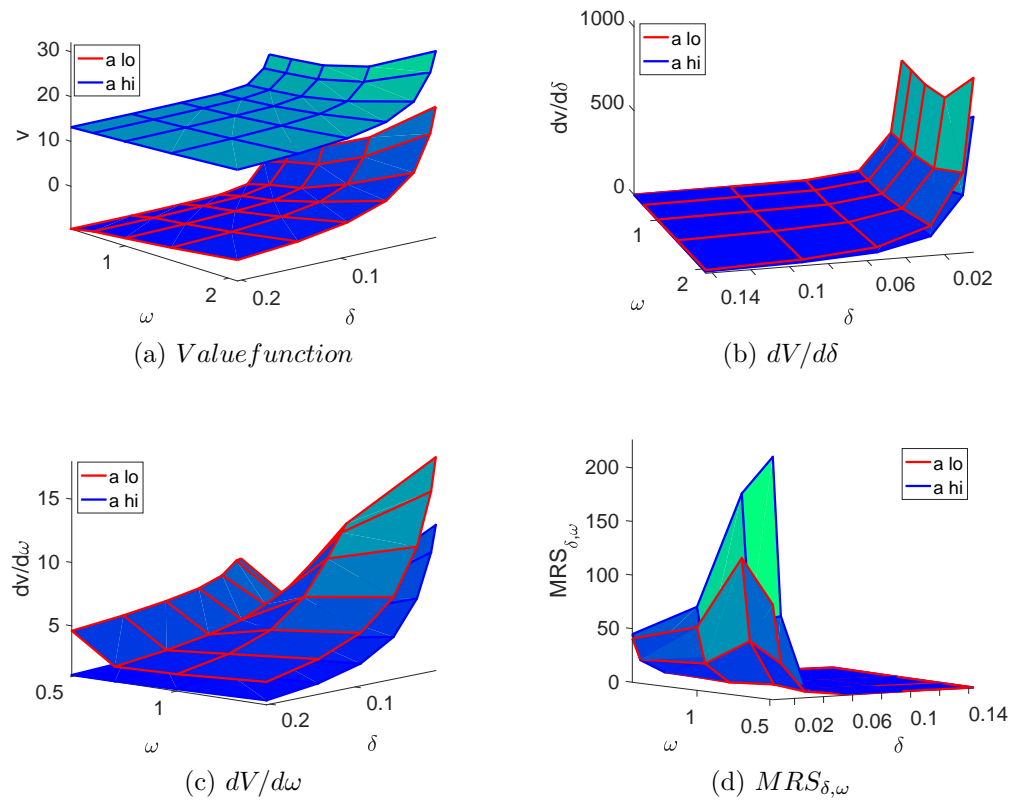





Figure C.8: Value function in (ω, δ) space




(a) MPC by job



(b) MPC by asset




(c) Log. average tenure




(d) Expected log. duration

Figure C.9: Role of job risk in MPCs




(a) Housing status


Figure C.10: Housing status response to unemployment shock




(a) Mortgage



(b) Log. Mortgage




(c) Equity




(d) Log. Equity


Figure C.11: Other asset responses to unemployment shock




(a) Housing (model)



(b) Housing (data)




(c) Consumption (model)




(d) Consumption (data)


Figure C.12: Group responses to Great Recession shock I: level response




(a) Liquid asset (model)



(b) Liquid asset (data)



(c) Equity (model)



(d) Equity (data)

Figure C.13: Group responses to Great Recession shock II: level response

C.2 Additional Figures for Aggregate Consumer Credit Uncertainty, Propagation and Consumption Dynamics

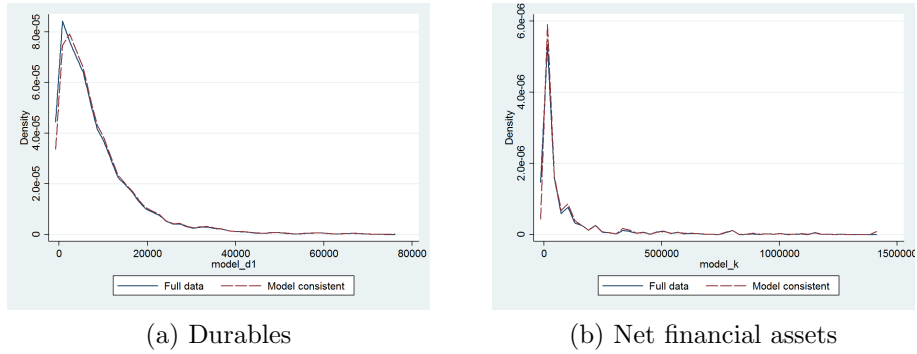


Figure C.14: Financial Assets and Automotives Distributions, SCF, 2010

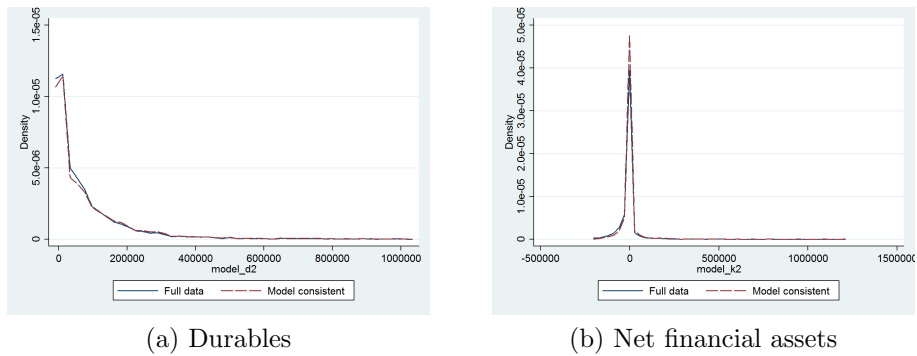


Figure C.15: Financial Assets and Housing Distributions, SCF, 2010

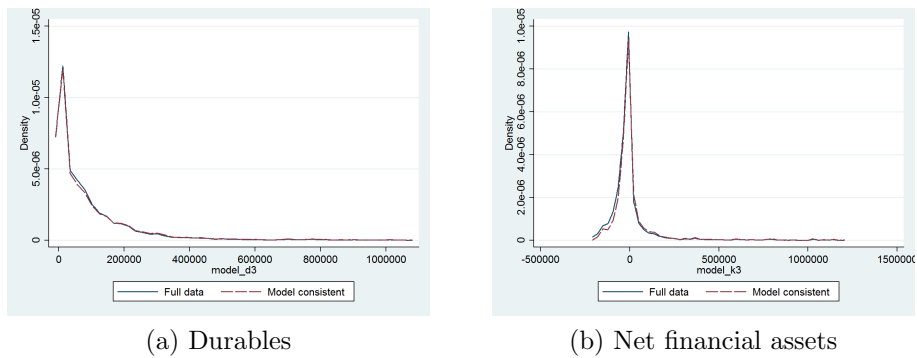


Figure C.16: Financial Assets and Housing Distributions, SCF, 2010




Figure C.17: Accuracy of the aggregate law of motion

C.3 Additional Figures for (S)Cars and the Great Recession




Figure C.18: Interest rate: car loan response

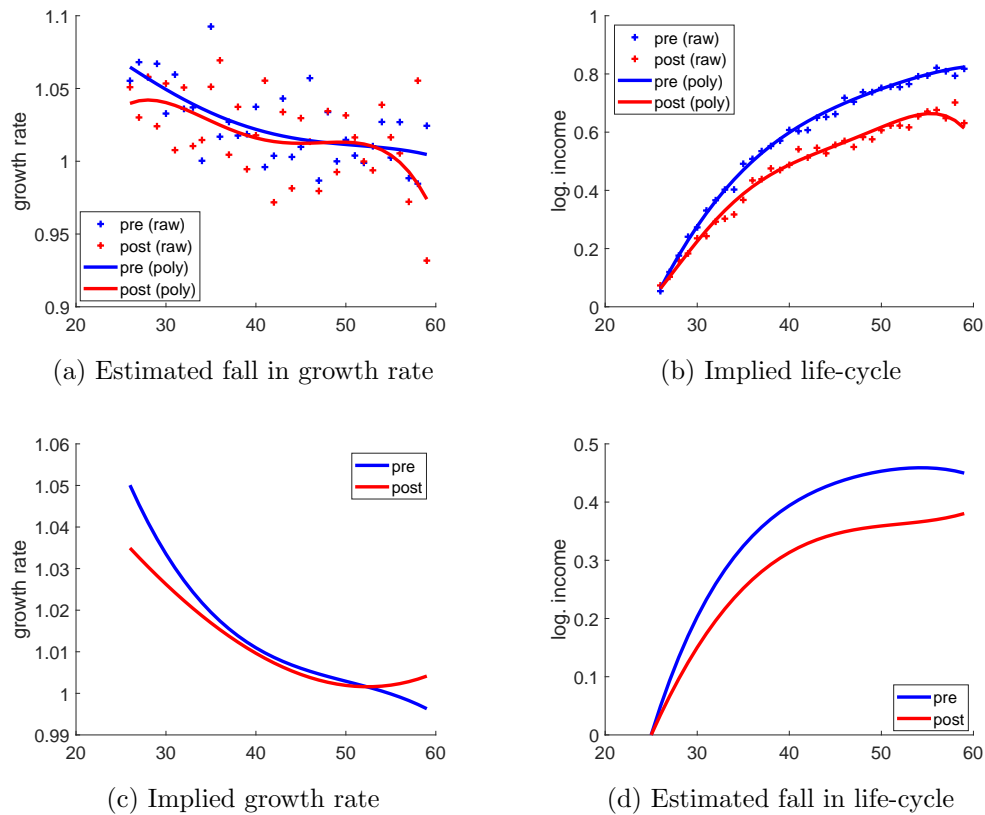


Figure C.19: Evidence supporting decline in life-cycle growth rate (CPS)




Figure C.20: Growth rate shock: example paths