Heterogeneity and Asymmetry in the Marginal Propensity to Consume

Alexander Ballantyne
ORCID: 0000-0002-6188-6613

Submitted in partial fulfilment of the requirements of the degree Doctor of Philosophy.

December, 2021

Department of Economics
THE UNIVERSITY OF MELBOURNE

Abstract

Building on the 'partial insurance' framework of Blundell, Pistaferri, and Preston (2008), I estimate households' marginal propensity to consume (MPC) out of transitory and permanent income shocks using data from the US Panel Study of Income Dynamics over 1999-2019. Following Chatterjee, Morley, and Singh (2021), I use the state space representation of the model but estimate it with Bayesian methods, which provide a natural setting for estimating high dimensional models with unobserved components. I extend the model in different directions by allowing for: i) asymmetry in MPCs with respect to the sign of transitory income shocks; ii) observable and unobservable heterogeneity in MPCs; iii) robust aggregation of population statistics; and iv) selection of the main drivers of MPC variation from a rich set of covariates. The results challenge some conventional understanding of consumption smoothing behaviour.

I find clear evidence of *positive* asymmetry – larger responses to positive transitory income shocks – for households with substantial home equity. This contrasts the existing survey evidence on 'reported preferences' and is suggestive of a behavioural interpretation. However, evidence of *negative* asymmetry for so called 'poor hand-to-mouth' households (who have low liquidity and low wealth) is consistent with the Kaplan-Violante two-asset model. Households with heads aged 55–64 are found to exhibit negative asymmetry, irrespective of their financial position, which is interpreted as a savings preference before retirement.

Under the assumption of symmetry, I look for evidence of variation in MPCs that would accord with theoretical predictions. Distributions of MPCs along balance sheet positions show some evidence of financial frictions, but demographic characteristics account for most of the variation in MPCs. I estimate *unobserved* MPC heterogeneity and find that it is large for transitory MPCs, but that observed variation explains much of the heterogeneity for permanent shocks. Population-representative dollar-for-dollar MPCs for the US are carefully aggregated from microdata elasticities using an extended dataset that includes retired households. Households increase permanent nondurables consumption by 7 cents per dollar for transitory income shocks, and 38 cents for permanent shocks.

Using a rich set of demographic and financial covariates, Bayesian shrinkage and sparsification techniques are applied to identify the key drivers of MPC heterogeneity. Households with heads who are black, female, or born in the 1980s have larger responses to permanent income shocks, while those with four or more years of college have lower responses. These four covariates drive the overwhelming majority of heterogeneity in MPCs out of permanent income shocks. The pattern of results is consistent with financial literacy being an underlying mechanism; however, incorporating a proxy for financial literacy does not remove the demographic effects.

This thesis provides a nuanced view of how different households respond to income shocks, which can help guide better targeted policy interventions. The findings are consistent with some aspects of the life-cycle literature; however, they also challenge the conventional assumption of symmetry and focus on financial drivers of MPCs. Further research into behavioural factors that affect consumption smoothing and demographic differences across the life cycle is warranted.

Declaration

This is to certify that

(i) the thesis comprises only my original work towards the degree Doctor of Philosophy,

(ii) due acknowledgement has been made in the text to all other material used,

(iii) the thesis is less than 100,000 words in length, exclusive of tables, maps, bibliographies and appendices.

Signed A.B.

Dated: December 21, 2021.

Preface

This thesis contains original research in Chapters 2–5. Chapter 3 and most of Chapter 2 are based on a working paper "Household Consumption: MPC Asymmetry and Financial Frictions" (Ballantyne (2020)).

This thesis was supported by an Australian Government Research Training Program (RTP) Scholarship.

The author was on leave from the Reserve Bank of Australia while conducting this research. Views expressed in this thesis are those of the author and not necessarily those of the Reserve Bank. Use of any results from this thesis should clearly attribute the work to the author and not to the Reserve Bank of Australia.

Impact Assessment for Net Zero Emissions: This thesis has negligible direct impact on the transition to net zero emissions. It does not analyse aggregate resource use or long-run dynamics and is agnostic on the composition of household consumption. In estimating household consumption responses to income shocks, it may have indirect relevance for transition policy, particularly if larger and/or more frequent supply-side shocks increase income volatility. The focus on asymmetry and heterogeneity could help in better planning for how adverse shocks will affect particular groups.

Acknowledgments

"No man is an island entire of itself", and this thesis is no exception. It is the product of invaluable feedback, suggestions, critique, support, friendship, love and laughter.[1] A few of these bear special mention.

I am greatly indebted to my supervisors Bruce Preston, Yong Song and James Hansen. They have pushed me to think deeply about this material, asked piercing questions, suggested ways forward, and provided crucial feedback. It can be a rough ride tumbling down the Dunning–Kruger effect, but the rigour and clarity of thought that result have made me a much better economist. I am very grateful.

Thanks to the PhD course coordinators Kalvinder Shields, John de New and Tim Robinson, who have provided me continual support, as they do for all students in the program. Thanks to Kevin Staub; working as a research assistant was a delight and provided me with microeconometrics and HPC experience that markedly enhanced this thesis. Thanks to Mei Dong for the Money & Banking tutoring gig, which was both fun and insightful. A big thanks to the Macro and Bayesian groups in the Economics department, and early assistance from Chris Edmond in particular. I am also thankful for comments from James Morley, Thijs van Rens, Michael Coelli, Tom Cusbert, participants at the 2021 IAAE conference and an RBA seminar, and to Justin Weidner for kindly providing additional replication code.

It has been grand working alongside my PhD colleagues, who have made the hard slog so much more bearable. A huge thanks to the Melbourne University Mountaineering Club who taught me the value of community and helped keep me sane.

[1]With a side of hard work and occasional isolation, of course!

Above all else, thanks to my family, friends and housemates for all their emotional support and putting up with a shittier version of me for four and a half years. In particular, thanks to Tom, Carla, Jude, therapy dog Ralph and cat-loaf Bunji. My deepest appreciation and love goes out to Ali Hall, who has been a shining light.

I would like to acknowledge and pay respect to the Wurundjeri people of the Kulin nation, the Traditional Owners of the land on which this thesis was written. Reflecting on the history of this country is an important reminder to me of the daily struggle faced by many, and the amazing privilege it has been to undertake independent research at the University of Melbourne. Global events over 2020–21 have further brought this into sharp perspective. But times of upheaval present both challenges and opportunities – the Black Lives Matter, Stop Black Deaths in Custody and #MeToo movements have added depth to my thinking about the results in this thesis.

Finally, people told me that doing a PhD was hard, but that didn't really sink in until it was too late. Nonetheless, to riff on John Maynard Keynes, my only regret in PhD life is that I didn't drink enough Champagne.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"In the past half century, economics has been enriched by vast new resources of microeconomic data. These data have opened the eyes of economists to the diversity and heterogeneity of economic life."
— James J. Heckman, 2000 Nobel Lecture

"Explicitly modeling earnings at the individual level, separating out macro from idiosyncratic effects, assessing insurance, and thinking about macroeconomics as an aggregate of heterogeneous agents has become a (or perhaps even the) central topic in macroeconomics"
— Angus Deaton, 2015 Nobel Lecture

Since James J. Heckman's address to the Nobel committee over 20 years ago, the analysis of microdata has continued to offer new insights into individual behaviour that informs our understanding of the economy. Indeed, 15 years later when Angus Deaton took to the podium, the impact on the field of macroeconomics was clear. In particular, life-cycle models provide a tractable framework for studying household income, consumption and savings decisions that can be informed directly by microdata (e.g. Browning and Crossley (2001) provide a nice overview). This facilitates a much richer appreciation of distributional effects when designing economic policy.

Consumption decisions vary from one household to another in ways that affect aggregate economic outcomes. The marginal propensity to consume (MPC) out of income shocks plays an important role in the transmission of macroeconomic policy interventions (e.g. Kaplan and Violante (2014), Auclert (2019)). The past decade has seen heightened interest in understanding the household characteristics that affect MPCs. In particular, household balance sheets vary widely across the population, which leads to different financial pressures on consumption decisions. This not only drives heterogeneity in MPCs, but can also engender asymmetric reactions to positive and negative shocks.[2]

There are large bodies of literature analysing the determinants of consumption responses to expected and unexpected income changes.[3] This thesis contributes towards the latter, estimating how household characteristics drive heterogeneity and asymmetry in MPCs out of unexpected permanent and transitory income shocks.[4] I build on the 'partial insurance' framework of Blundell et al. (2008) (BPP) and estimate various empirical models on data from the US Panel Study of Income Dynamics (PSID) over 1999-2019. The results are interpreted in light of the theoretical literature.

Chapter 2 introduces the partial insurance framework, estimation techniques and data that serve as the foundation for the following chapters. I use Bayesian methods for estimation, which provide a natural setting for estimating high dimensional models with unobserved components. I extend the model in different directions by allowing for: i) asymmetry in MPCs with respect to the sign of transitory income shocks; ii) observable and unobservable heterogeneity in MPCs; iii) robust aggregation of population statistics; and iv) selection of the main drivers of MPC variation from a rich set of covariates.

[2]For example, a credit constraint is naturally directional, limiting households' capacity to respond to negative shocks. In fact, asymmetry could be considered a more fundamental property than heterogeneity as it is also a possible feature of a representative agent model (e.g. Kocherlakota (2000)).

[3]For example, see Parker and Preston (2005) and Johnson, Parker, and Souleles (2006) on expected changes, or Jappelli and Pistaferri (2014) and Blundell et al. (2008) on unexpected changes. See Jappelli and Pistaferri (2010) for a review.

[4]There are three broad approaches to estimating MPCs out of unexpected income shocks: natural experiments like the fiscal stimulus of 2008 (e.g. Parker, Souleles, Johnson, and McClelland (2013)); surveys that directly ask households about their reactions to hypothetical or recalled shocks (e.g. Bunn, Le Roux, Reinold, and Surico (2018)); and inferring MPCs from panel data on household income and consumption – the approach taken in this thesis.

Chapter 3 extends the partial insurance framework to estimate sign asymmetry in MPCs out of transitory income shocks, and investigates the drivers of asymmetry through heterogeneity in household balance sheets. Financial frictions, such as credit constraints, can drive such MPC sign asymmetry; observation of aggregate *symmetry* may obscure asymmetry in the cross-section. Furthermore, there are relatively few studies that measure MPC asymmetry due to requiring observation of both positive and negative shocks; however, panel data offers such a setting.

I find positive asymmetry – larger MPCs for positive shocks than negative shocks – for households with substantial home equity. Explanation of this result primarily focusses on behavioural mechanisms. Households who have both limited liquidity and no other illiquid wealth ('poor hand-to-mouth') are found to have negative asymmetry, which accords with the predictions of the Kaplan-Violante two-asset model (Kaplan and Violante (2014)). Households aged 55–64 display negative asymmetry across all balance sheet positions, which is interpreted as a savings preference late in the working life.

Chapter 4 builds on the BPP framework to estimate the distribution of household MPCs out of transitory and permanent income shocks. Existing empirical methods tend to rely on subsampling along a few observable dimensions, which has limited the capacity to measure the *distribution* of MPCs in the population. The approach allows for observable household characteristics to drive MPC heterogeneity, while accommodating potential nonlinearities. Much of the variation in MPCs is not statistically different to a simple linear model and few of the relationships are robust in a model with multiple covariates. The largest source of variation in transitory and permanent MPCs comes from demographic variables.

The analysis is extended to incorporate *unobserved* heterogeneity, similar to a 'random coefficients' model. Unobserved variation appears to be an important source of heterogeneity for the response to transitory income shocks, whereas observed variation explains much of the heterogeneity for permanent shocks. In addition, I generate policy-relevant MPCs that are representative of the US population by i) extending the sample to include retired households, ii) formally modelling heterogeneity, iii) population weighting the sample, and iv) converting the microdata elasticities into dollar-for-dollar MPCs. Population average \$MPCs imply that households increase nondurables consumption by 7 cents per dollar for transitory income shocks, and 38 cents for permanent shocks.

Chapter 5 leverages the range of data available on household financial and demographic characteristics in the PSID to determine the most salient drivers of variation in MPCs. Empirical modelling of heterogeneity in MPCs has typically been conducted by testing a select few covariates with *a priori* motivation (e.g. Johnson et al. (2006), Jappelli and Pistaferri (2020)). This approach often ties in with theoretical work that has focused on financial frictions or preference heterogeneity (Heathcote, Storesletten, and Violante (2009), Parker (2017)). However, households are subject to a much wider range of social and economic factors that influence their decision making. Further extending the framework used in Chapter 4, I utilise statistical methods from the Bayesian shrinkage and sparsification literature to separate the rich set of covariates into signals that drive meaningful variation in MPCs and noise covariates that do not.

Demographic traits are found to be the main drivers of heterogeneity in MPCs. Households with heads who are black, female or born in the 1980s have larger responses to permanent income shocks, while those with four or more years of college have lower responses. These four covariates drive the overwhelming majority of heterogeneity in permanent MPCs. Financial literacy is investigated as a potential underlying mechanism and is found to play a role in consumption smoothing, but does not eliminate the effects of demographics.

The thesis presents a range of empirical evidence on the diversity of household consumption smoothing. The results support some predictions from the theoretical literature with respect to households facing acute financial constraints, but also find that financial factors may be less important than demographic traits. The estimates provide a nuanced view of how different households respond to income shocks, which can inform policy interventions. As Greg Kaplan and Giovanni Violante (2016) propose, "Government agencies could combine multiple sources of available microeconomic data to create an individual propensity to consume score that would guide the design of the program." The methods developed here go some way to achieving this goal.

Chapter 2

Empirical Framework & Data

The overarching empirical strategy used in this thesis is based on the 'partial insurance' framework presented in Blundell et al. (2008) (BPP). This chapter presents the empirical model and its interpretation, along with estimation and data details, that serve as the foundation for the following chapters. The BPP framework has been further examined in Kaplan and Violante (2010), Kaplan, Violante, and Weidner (2014), Chatterjee et al. (2021), Cho, Morley, and Singh (2021), and Commault (2021), among others; however, precise specifications and estimation methods differ. Appendix A.1 details differences across the BPP literature, and their relation to my estimates. I use Bayesian methods, which result in estimates similar to the quasi maximum likelihood estimation (QMLE) approach of Chatterjee et al. (2021).

2.1 Empirical Framework

2.1.1 BPP Model of Partial Insurance

BPP derive the empirical model as a linear approximation of "the mapping between the expectation error of the Euler equation and the income shock" for a life-cycle consumer with constant relative risk aversion (CRRA) utility. The result is a 'semi-structural' model that can be taken to panel data on household income and consumption. The structure of the model presented below follows Chatterjee et al. (2021) (CMS), but is equivalent to the specification used in the main BPP results when abstracting from distributional assumptions regarding the shocks. Log household income comprises a predictable component – based on observable life-cycle variables such as age, education and employment status, as well as time dummies and interaction terms – and an 'unexplained' component. That is, the income process is given by

$$ln(Y_{i,t}) = Z_{i,t}\Upsilon_Y + y_{i,t} \tag{2.1}$$

where Y is household labour income and Z is the set of P life-cycle variables and an intercept. Consumption follows a similar process

$$ln(C_{i,t}) = Z_{i,t}\Upsilon_C + c_{i,t}. \tag{2.2}$$

Unexplained income, y, and unexplained consumption, c, are then modelled in a second stage using an unobserved components (UC) framework, driven by permanent and transitory dynamics. Unexplained income is given by

$$y_{i,t} = \mu^y_{i,t} + \nu^y_{i,t} + \theta\nu^y_{i,t-1} \tag{2.3}$$
$$\mu^y_{i,t} = \mu^y_{i,t-1} + \zeta^y_{i,t}. \tag{2.4}$$

The unobserved, or *latent*, component μ^y reflects the permanent component of income, ζ^y is a permanent income shock and ν^y a transitory income shock. The transitory shock is assumed to follow a first-order moving-average, MA(1), process with parameter $|\theta| < 1$. Consumption follows a similar UC process; however, permanent and transitory components of income also affect consumption dynamics

$$c_{i,t} = \gamma_\mu\mu^y_{i,t} + \mu^c_{i,t} + \nu^c_{i,t} \tag{2.5}$$
$$\mu^c_{i,t} = \mu^c_{i,t-1} + \gamma_\nu\nu^y_{i,t} + \zeta^c_{i,t}. \tag{2.6}$$

As above, μ^c reflects the permanent component of consumption (independent of permanent income), and ζ^c and ν^c are permanent and transitory consumption shocks. The parameters γ_ν and γ_μ govern the response of unexplained consumption to transitory and permanent unexplained income shocks, or the 'transitory MPC' and 'permanent MPC'. These parameters are elasticities, approximately measuring the percentage change in unexplained consumption driven by a one per cent change in unexplained income.[5] This contrasts with the dollar-for-dollar

[5]The parameters are unitless since ν^y, μ^y, μ^c and $ln(C)$ are all in log-real-dollar terms.

'$MPC' concept often used in policy making; Section 4.6.1 discusses this in detail. BPP refer to the model as capturing partial insurance, to the extent that estimates of γ_ν and γ_μ are different to those implied by the underlying theoretical model. That is, estimates that are smaller than those implied by theory reflect household consumption smoothing in *excess* of self-insurance through precautionary savings (discussed further in the next section). The value $1 - \gamma_\mu$ is referred to as households' partial insurance against permanent income shocks. A small γ_μ implies larger partial insurance and more consumption smoothing.

Equations (2.1)–(2.6) form a panel state space system that is the foundation of all empirical models used in this thesis. The existing literature using this framework conducts the first-stage regressions in equations (2.1) and (2.2) independently, then takes unexplained income and consumption, y and c, as the data and estimates equations equivalent to (2.3)–(2.6) in a separate second stage. I follow this approach; however, a joint model is detailed in Appendix D.5 and has some different conceptual interpretations. The benchmark model assumes MPCs are constant across both households and time, but this thesis relaxes that assumption in various ways. Extensions to the model in Chapters 3 and 4 are focussed on capturing variation in the estimated MPCs driven by the sign of income shocks and heterogeneity across households.

Bayesian methods are used to estimate the model, which require additional assumptions about the nature of shocks. All shocks follow independent mean-zero Gaussian distributions with constant variance

$$x_{i,t} \sim \mathcal{N}(0, \sigma_x^2); \quad \forall \quad x_{i,t} \in \nu_{i,t}^y, \nu_{i,t}^c, \zeta_{i,t}^y, \zeta_{i,t}^c. \tag{2.7}$$

This approach has similarities with QMLE used in CMS; that is, estimation relies on the likelihood as the central objective function.[6] As noted in CMS, the data display some evidence of non-normality, so the estimates here carry the same caveat. The first-stage regressions strip income and consumption of much of their variation, leaving unexplained income and consumption that displays only mild

[6]Rather than maximise the likelihood, Bayesian estimation applies priors and simulates from it. Care must be taken to ensure the choice of priors does not drive the results.

departures from normality.[7] Here I take the volatility of shocks as constant over the sample, whereas BPP and CMS allow for time-varying volatility. I test for time-varying volatility and find little statistical difference over most of the sample. As a result, the estimates of elasticities are similar under both assumptions. Results and further discussion are provided in Appendix A.1.3.

2.1.2 Theoretical Background and Interpretation

The empirical model is derived from an Euler equation solution and budget constraint of a life-cycle consumer with CRRA utility. It can be considered a 'semi-structural' model (e.g. Cho et al. (2021), Commault (2021)). This section presents a brief derivation, following BPP, to highlight key aspects of its theoretical and empirical interpretation. For full details of the model derivation, see Appendix B in BPP and Appendix A in Blundell, Low, and Preston (2004). Here, I only consider the model with an MA(1) transitory income process.

Household i at age t has access to a risk free bond with real return r_t and optimises consumption subject to the budget constraint

$$A_{i,t+1} = (1 + r_t)(A_{i,t} + Y_{i,t} - C_{i,t})$$

with start-of-period assets $A_{i,t}$ and known end of life at age T, where $A_{i,T} = 0$. Optimal consumption is given by the familiar Euler equation

$$C_{i,t-1}^{-\frac{1}{\phi}} = \beta(1 + r_{t-1})E_{t-1}\left\{e^{\Delta V'_{i,t}\vartheta_t} C_{i,t}^{-\frac{1}{\phi}}\right\}$$

where β is the discount factor, ϕ is the elasticity of intertemporal substitution, Δ is the first-difference operator and $V'_{i,t}\vartheta_t$ incorporates taste shifters. A log-linear approximation is given by

$$\Delta ln(C_{i,t}) \simeq \Delta V'_{i,t}\vartheta_t\phi + \eta_{i,t} + \Omega_{i,t}$$

[7]Skewness is around -0.2 and kurtosis around 3.7, depending on the sample. The extent of kurtosis may be artificially reduced due to filtering (detailed in Section 2.3.1), particularly the removal of outliers.

where $\eta_{i,t}$ is a consumption shock that is zero in expectation and $\Omega_{i,t}$ captures the effect of interest rates, impatience and precautionary savings on the slope of the consumption path. BPP assume that variation in the consumption path due to $\Omega_{i,t}$ can be absorbed by a vector of deterministic characteristics $\Pi_{i,t}$ and a stochastic individual component $\zeta^c_{i,t}$. Defining $\Delta c_{i,t} = \Delta ln(C_{i,t}) - \Delta V'_{i,t}\vartheta_t\phi - \Pi_{i,t}$ gives

$$\Delta c_{i,t} \simeq \eta_{i,t} + \zeta^c_{i,t}. \tag{2.8}$$

Equation (2.2) implies that the deterministic life-cycle terms $\Delta Z_{i,t}\Upsilon_C$ capture variation in the consumption path due to $\Delta V'_{i,t}\vartheta_t\phi$ and $\Pi_{i,t}$. The lifetime budget constraint is subject to the income process in equations (2.1) and (2.3)–(2.4). Taking differences in expectations before and after shock realisation in period t, a log-linear approximation of the lifetime budget constraint simplifies to

$$\eta_{i,t} \simeq \pi_{i,t}(\zeta^y_{i,t} + \alpha_{t,L}\nu^y_{i,t}). \tag{2.9}$$

where $\pi_{i,t}$ is an age-decreasing self-insurance factor given by the share of future labor income in current human and financial wealth and $\alpha_{t,L}$ is an age-increasing annuitisation factor (with retirement occuring at age L). Combining equations (2.8) and (2.9) provides a linear relationship between income shocks and the expectation error of the Euler equation that can be taken to the data

$$\Delta c_{i,t} \simeq \pi_{i,t}\zeta^y_{i,t} + \pi_{i,t}\alpha_{t,L}\nu^y_{i,t} + \zeta^c_{i,t}.$$

Measurement error is included in unexplained consumption. Redefining $c_{i,t}$ as observed unexplained consumption for consistency with the empirical model and $c^*_{i,t}$ as the true value, adding measurement error $c_{i,t} = c^*_{i,t} + \nu^c_{i,t}$ gives

$$\Delta c_{i,t} \simeq \pi_{i,t}\zeta^y_{i,t} + \pi_{i,t}\alpha_{t,L}\nu^y_{i,t} + \zeta^c_{i,t} + \Delta\nu^c_{i,t} \ , \tag{2.10}$$

with unexplained income given by

$$\Delta y_{i,t} \simeq \zeta^y_{i,t} + \Delta\nu^y_{i,t} + \theta\Delta\nu^y_{i,t-1}. \tag{2.11}$$

Equations (2.10)–(2.11) are the reduced form representations of the UC models in equations (2.3)–(2.6), but with structural time-varying MPC coefficients. BPP report that $\pi_{i,t} \simeq 0.8$ for a consumer 20 years before retirement (using a simulation of the life-cycle model), although this decreases towards zero at retirement. The annuitisation factor $\alpha_{t,L}$ is quite small. For example, setting $\Delta V'_{i,t}\vartheta_t\phi = 0$, the real interest rate constant at 3 per cent, and the MA(1) parameter $\theta = 0.1$, then the annuitisation factor is approximately $\alpha_L \simeq 0.03$. From the partial insurance perspective, estimates of $\gamma_\mu < \pi_{i,t}$ and $\gamma_\nu < \pi_{i,t}\alpha_{t,L}$ would provide evidence of consumption smoothing over and above self-insurance through savings. The theory also provides a clear prediction that γ_μ should decline with age towards retirement.

It is important to highlight that the model only captures the effect of income shocks on permanent consumption. This arises from the theoretical derivation – income shocks are mapped to the expectation error of the Euler equation.[8] The model allows transitory shocks to consumption, captured by $\nu^c_{i,t}$ and interpreted as measurement error, but does not allow income to affect transitory consumption dynamics. This assumption has been brought into question by Commault (2021), who finds large consumption responses to contemporary and past transitory income shocks that imply consumption may deviate from a random walk.

From an empirical perspective, the model can be seen as simultaneous equations in log income and consumption, with an identifying restriction that unexplained consumption does not affect contemporaneous unexplained income. This interpretation implies that the life-cycle variables Z act as control variables, absorbing endogenous variation in consumption. The large set of life-cycle variables used makes this plausible, suggesting a causal interpretation of the estimated MPCs.

[8]As the Euler equation equates current and discounted future marginal utility, consumption will evolve as a random walk under rational expectations. See Hall and Mishkin (1982).

2.2 Estimation

This section motivates the use of Bayesian techniques, discusses crucial assumptions used in estimation and outlines the estimation algorithm. Simulations to test inference are provided in Chapters 3 and 5. A full derivation of conditional posterior distributions for the benchmark model is provided in Appendix A.2. Estimates using Bayesian methods are similar to those using QMLE, detailed in Appendix A.1.

2.2.1 Motivating Bayesian Estimation

Equations (2.3)–(2.6) form a linear panel state space system. The QMLE approach developed by CMS uses the Kalman filter to estimate the likelihood for a set of parameter candidates. Bayesian estimation has core similarities with QMLE in that it relies on the likelihood as the central objective function. However, the introduction of asymmetry into the model in Chapter 3 makes the state space system *functionally* nonlinear (in the language of Harvey (1989)). As is well known, the Kalman filter is only optimal for linear Gaussian models.[9] An increasingly popular choice in economics is to use the particle filter to evaluate the likelihood.[10] This method is most often employed in a Bayesian setting given its use of simulation to evaluate the states. A Metropolis-Hastings algorithm is also an appropriate solution as it only requires evaluating the posterior density once candidates for the unobserved state parameters are drawn, which is straightforward. Consequently, a Bayesian setting is natural to estimate a model of this form, which is further exploited in the model of unobserved heterogeneity in Chapter 4 and the use of shrinkage and sparsification in Chapter 5.

A Bayesian approach therefore provides a powerful set of tools to examine the drivers of variation in MPCs using the BPP framework. The primary drawback is the large computational burden of working with microdata, which is overcome

[9]The extended Kalman filter can be applied to nonlinear models; however, this takes a linear approximation via a Taylor series expansion and so is not suitable for a discontinuous model of the type considered here (Harvey (1989)).

[10]For details on the particle filter see Gordon, Salmond, and Smith (1993), or for economic applications see (for example) Kim, Shephard, and Chib (1998) or Fernández-Villaverde and Rubio-Ramírez (2007).

by a bespoke algorithm and parallelisation. Primiceri and van Rens (2009) use Bayesian estimation in a similar framework to analyse the joint evolution of the variances of income and consumption. They also emphasise the role of heterogeneity, but with respect to expected income shocks, rather than the focus on unexpected shocks in this thesis. They note the benefit of a Bayesian setting for estimating a high dimensional model with latent variables, which also applies to the models used in this thesis.

2.2.2 Assumptions and Priors

The benchmark model assumes i) Gaussian shocks, ii) common variances across all households and time, iii) common MPC coefficients across all households and time, and iv) households are independent from one another. CMS note that the true data generating process may not be normally distributed and so consider their likelihood-based estimation under assumption i) to be quasi maximum likelihood; the same caveat applies here. Assumption ii) departs from some of the previous literature; this is discussed in Appendix A.1.3. The main contribution of this thesis is to relax assumption iii). Assumption iv) is plausible since the unexplained income and consumption variables are residuals from life-cycle regressions, so most macroeconomic shocks that affect households simultaneously have been removed from the data.

Conditional on the income shocks, the UC model for consumption requires the additional identifying restriction that there is no contemporaneous correlation between the permanent and transitory shocks $cov(\zeta^c_{i,t}, \nu^c_{i,t}) = 0$. The same restriction is imposed on the income model $cov(\zeta^y_{i,t}, \nu^y_{i,t}) = 0$. Given the short time dimension of the panel data, I opt to estimate the period zero parameters ($\nu^y_{i,0}$, $\mu^y_{i,0}$ and $\mu^y_{i,0}$) rather than setting them equal to zero as is often done.

Formal definitions of the priors are provided in Appendix A.2 and descriptions follow. The priors for $\mu^y_{i,t}$ and $\mu^c_{i,t}$ are given by independent normal distributions with variances $\sigma^2_{\zeta y}$ and $\sigma^2_{\zeta c}$. The priors for the variance of the shocks are given by inverse-gamma 2 (IG2) functions, with hyperparameters set to provide disperse priors relative to the estimates of these parameters in the literature (Figure 2.1).[11] With the same prior on each variance parameter, estimation of MPCs is not being informed by any prior information on the *relative* size of shock variances. The relative size of transitory and permanent shock variances is a key factor in identifying the latent components μ^c and μ^y, as well as the MPC parameters. The priors for the period zero parameters $\mu^y_{i,0}$ and $\mu^c_{i,0}$ are normal distributions with zero mean and a variance set to 1.4 (which is ten times the greater pooled sample variance of y or c) to provide a disperse prior.

When the MPC parameters are allowed to vary, they are written as vectors of parameters Γ_ν and Γ_μ. Priors for these parameters are taken as vectors of independent normal distributions, all with the same hyperparameters so that no prior information about the relative size of MPCs is imposed. That is, I am assuming all covariates that drive MPC heterogeneity and asymmetry have the same prior on their coefficients.[12] The theoretical and empirical literature on MPCs suggests that estimates should be in the range of $[0, 1]$, with MPCs to transitory shocks being smaller than for permanent shocks. Relative to this benchmark, the prior used is quite disperse; it has a mean $m_p = 0.35$, in between most estimates of γ_ν and γ_μ, with a variance $v_p = 0.5^2$ (Figure 2.2). This prior implies a 24 per cent probability of a negative MPC and a 10 per cent probability of an MPC greater than one, which are considerable weights to highly improbable MPCs.

2.2.3 Algorithm

The set of parameters to be estimated (over all individuals i and times t) is given by

$$\Theta = \{\Gamma_\mu, \Gamma_\nu, \theta, \mu^y_{i,t}, \mu^c_{i,t}, \mu^y_{i,0}, \mu^c_{i,0}, \nu^y_{i,0}, \sigma^2_{\nu y}, \sigma^2_{\nu c}, \sigma^2_{\zeta y}, \sigma^2_{\zeta c}\}.$$

The general form of the joint posterior distribution is detailed in Appendix A.2. Standard analytical solutions to the conditional posterior densities can be derived

[11] The inverse-gamma 2 distribution, $IG2(a, b)$, is such that its square root is an inverse-gamma distribution. See Appendix A in Bauwens, Lubrano, and Richard (2000). Here I set the scale hyperparameter $a = 0.15$ and the shape hyperparameter $b = 0.5$.

[12] This is particularly important in Chapter 3 when testing for the difference between the posterior densities of Γ^+_ν and Γ^-_ν.

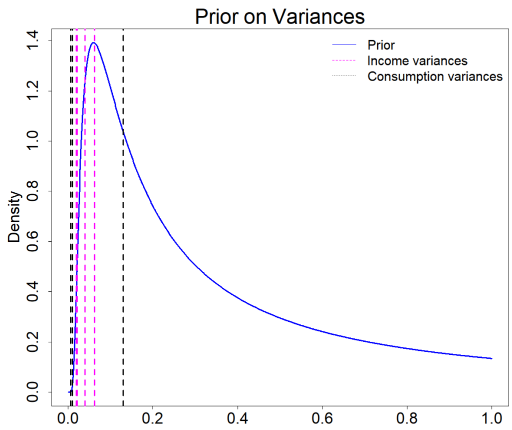

Figure 2.1: Prior on variance parameters (blue line) along with estimates from the literature (dashed vertical lines).

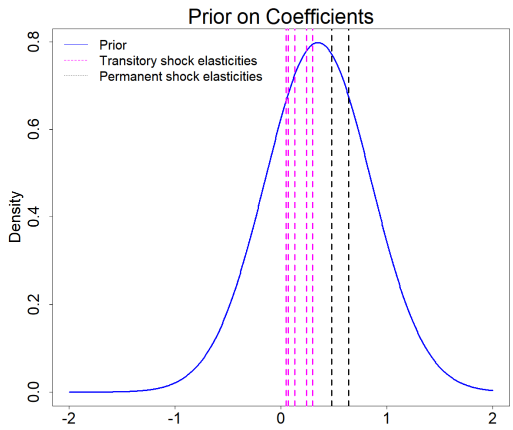

Figure 2.2: Prior on coefficient parameters (blue line) along with estimates from the literature (dashed vertical lines).

for many of the parameters, namely $\Gamma_\mu, \Gamma_\nu, \sigma^2_{\nu y}, \sigma^2_{\nu c}, \sigma^2_{\zeta y}, \sigma^2_{\zeta c}$, and so draws for these parameters are constructed from a Gibbs sampler. I use the precision sampler from Chan and Jeliazkov (2009) to sample the latent component of consumption μ^c; I also apply this routine to efficiently sample other parameters whose conditional posterior distribution is multivariate normal. Nonlinearities in the models developed in this thesis mainly affect the posterior distribution of μ^y. For Chapter 3 this is of unknown form so I implement a simple Metropolis-Hastings routine to generate draws for this vector of parameters.[13] The posterior of the MA(1) parameter θ is also of an unknown form (although this has less influence on the results compared with μ^y) and I implement a 'Griddy-Gibbs' routine to draw this parameter (see Ritter and Tanner (1992)).[14] The Griddy-Gibbs routine calculates the posterior density for each of 122 nodes on the interval $(-1, 1)$ (with a higher density of nodes between $[-0.15, 0.3]$), then constructs a discrete cumulative density function that can be used to sample θ from the set of nodes. The estimation algorithm is summarised as follows.

1. Draw $\sigma^2_{\nu y}, \sigma^2_{\nu c}, \sigma^2_{\zeta y}, \sigma^2_{\zeta c}$ using Gibbs sampling

[13]The Metropolis-Hastings routine is also used in Chapter 4, but a precision sampler is used in Chapter 5 where the conditional posterior distribution is known.

[14]A Metropolis-Hastings algorithm for θ was also tested in simulation, as in Chib and Greenberg (1994), but the results were poorer.

2. Draw Γ_μ, Γ_ν using Gibbs sampling
3. Draw θ using Griddy-Gibbs
4. Draw $\mu^c_{i,0}, \mu^y_{i,0}, \nu^y_{i,0}$ and vector $\mu^c_{i,t}$ using precision sampler
5. Draw vector $\mu^y_{i,t}$ using Metropolis-Hastings

The datasets have 20–30 thousand observations, which means the computational burden of using Bayesian methods is large. The estimation algorithm is written in `R` and run on the University of Melbourne's high performance computing environment Spartan. The Griddy-Gibbs routine is particularly intensive as it loops over 122 nodes for every observation, but its reliability justifies its use. This routine is written in `C++` using the `R` package `Rcpp`, which increases efficiency substantially. Parallelisation over four cores is used for the Metropolis-Hastings and Griddy-Gibbs routines.[15] The precision sampler is notably faster than using standard functions to sample from a multivariate normal, with the sampler for μ^c also using the `SparseM` package.

2.3 Data

Following Kaplan et al. (2014) I use data from the biennial US Panel Study of Income Dynamics (PSID), extended to include 1999-2019. This contains data on household income and non-durables consumption, along with a range of balance sheet and demographic variables. Some modifications are made to the calculation of variables relative to the Kaplan et al. (2014) dataset, detailed below. Each chapter describes the specific covariates used in more detail; here I focus on the common variables and filtering.[16] A comparison against other samples from the literature is provided in Appendix A.4.

The biennial nature of the data modifies the interpretation of the parameters. The elasticities of interest, γ_μ and γ_ν, capture the biennial response of consumption

[15] I make use of James Balamuta's `Rcpp2doParallel` package, for which I am very grateful.
[16] The main dataset is used in Chapters 3 and 5, but Chapter 4 uses an earlier 'development' dataset.

to biennial income shocks – what Commault (2021) refers to as the "biennial pass-through". The MA(1) parameter reflects the persistence of a transitory biennial income shock. Time aggregation in the data affects the identification of parameters when they are based on an underlying model with different time units. These issues are discussed in Crawley and Kuchler (2020), Commault (2021) and Cho et al. (2021) with respect to the BPP framework; here I qualify the interpretation of estimates to reflect biennial shocks and consumption responses, rather than take a stand on the true underlying time structure.

2.3.1 Income, Consumption and Filtering

The income measure is after tax labour income including government transfers. I use an updated tax adjustment via the National Bureau of Economic Research's TAXSIM32 program, using both federal and state taxes.[17] This is a much more comprehensive tax adjustment than used in Kaplan et al. (2014). Nondurables consumption includes food, utilities, gasoline, car maintenance, transportation, childcare, health expenditures and education. Income and consumption are deflated using annual headline CPI (all urban consumers).

The estimation sample removes households if missing race, education, or state of residence, and if there are fewer than three consecutive observations. Only 'labour-force' households are used to avoid shifts in income and consumption relating to the transition into retirement. Labour-force households are classified as those with heads aged 25-64 and not retired at any period.[18] Observations where the head retires in the subsequent period are omitted. Chapter 4 also considers a sample of retired households.

Again following Kaplan et al. (2014), households whose income increases by over 500 per cent, falls by over 80 per cent or is below \$100 are omitted, as well as households with zero or negative nondurables consumption. Once predictable life-cycle components Z have been removed, additional outliers are dropped from

[17] Of the 32 items used by TAXSIM32, 12 items are omitted (items 8–9, 15–16, 18, 25, 27 and 28–32).

[18] This includes a small share of households where the spouse is retired. The distribution of income of these households is slightly higher than those with both partners working, suggesting that the spouse's retirement decision is not of first order importance to household income.

the unexplained income and consumption components (absolute value greater than two), as well as observations with liquid assets over \$1 million or under –\$250,000, or illiquid assets over \$2 million or under –\$250,000. This results in a labour-force sample of 29,645 observations from 4,528 households.[19]

2.3.2 Household Covariates

The life-cycle control variables Z used in equations (2.1)–(2.2) comprise a large set of interacted dummy variables as in Kaplan et al. (2014). Year dummies are interacted with dummies for education, race, employment and region, with the uninteracted dummies of all variables retained.[20] Dummies for year of birth, family size, number of children, positive taxable income from other family members (not head or spouse) and whether the household supports a non-family unit dependent (e.g. child support). This is in excess of 190 dummies that remove around 40 per cent of the variation in log income and log consumption. The residuals, y and c, are mean zero but display mild deviations from normality with skewness around -0.2 and kurtosis around 3.7.

The focus of this thesis lies in identifying the drivers of heterogeneity and asymmetry in MPCs. Household balance sheets can be complex and diverse, potentially generating a range of financial frictions that interact with consumption smoothing. The PSID provides detailed data on balance sheets, including: checking and savings account balances; credit card balances; holdings of stocks and bonds; home equity and mortgage debt; other real estate equity; annuities and Individual Retirement Accounts (IRAs); and other assets and debt. Non-labour income is adjusted for tax (weighted by the share of asset income in total income). Balance sheet items measured as income ratios (e.g. mortgage-to-income ratio) use after tax labour income. Demographic variables also feature heavily and relate to the

[19] The unexplained income and consumption outliers result in 0.5 per cent of the sample dropped, and the asset outliers result in a further loss of 1.1 per cent.

[20] The categories for education are less than high school, high school and at least some college; for race are white, black and other person of colour; for employment status are employed, unemployed, retired and not in the labour force; and for region are North East, Midwest, South and West.

household head, such as age, sex, race and education.[21] The sample is augmented by data on life expectancy, matched by age, year and sex using the historical US Social Security Administration actuarial life tables.

Following Kaplan et al. (2014), assets and liabilities are divided into liquid and illiquid categories. Liquid assets comprise checking balances, bonds, stocks and credit card debt, while illiquid assets comprise home equity, other real estate equity, annuities and IRAs, and other assets and debt. An important difference to the Kaplan et al. (2014) data is the treatment of credit card debt.[22] Prior to the 2011 survey, the PSID asked households a single question to capture all non-mortgage debt, which includes credit card debt, student loans, medical bills, legal bills and family loans. As such, much of the sample used in Kaplan et al. (2014) included all of these items in the measure of liquid assets, whereas ideally credit card debt alone would be used. In contrast, this thesis uses the additional years of PSID data to impute an historical measure of credit card debt. This provides cleaner separation of liquid and illiquid liabilities. Details are provided in Appendix A.3. A household is designated as hand-to-mouth (HtM) if their holdings of liquid assets is less than half of one month's income, or liquid debt is more than half of one month's income. HtM households are further separated into 'poor HtM' and 'wealthy HtM' households, the latter of which have positive illiquid wealth.

[21] The difference between sex and gender is not considered in the PSID (historically the question was recorded without even asking). This thesis refers to the (binary) variable as 'sex' for consistency with the PSID data, although the data may well reflect gender.

[22] A further (minor) difference is the correction of a coding error that resulted in asset income only being captured for survey years 1999 and 2001.

Chapter 3

Asymmetry & Financial Frictions

3.1 Introduction

Household balance sheets vary widely across the population, which leads to households making consumption decisions under different financial pressures. Such 'financial frictions' have been found to influence households' marginal propensity to consume (MPC) out of income shocks. MPCs vary across households, but they may also vary with the *type* of income shock faced; the response of households to positive income shocks may be different to negative income shocks – what I deem asymmetry in MPCs.[23] There are relatively few studies that measure MPC asymmetry due to requiring observation of both positive and negative shocks.[24] This chapter estimates asymmetry in MPCs out of transitory income shocks using data from the US Panel Study of Income Dynamics (PSID) over 1999–2019, and investigates the drivers of asymmetry through heterogeneity in household balance sheets.

Empirical studies find that households with low financial means tend to be more sensitive to income shocks; that is, have larger MPCs out of transitory income (Kaplan et al. (2014), Jappelli and Pistaferri (2014), Fagereng, Holm, and Natvik (2018)). Credit constraints have long been seen as a potential mechanism to reconcile the larger empirical MPCs out of transitory income shocks relative to the almost zero response predicted by the permanent income hypothesis (PIH) (e.g.

[23] Asymmetry could be considered as a more fundamental property than heterogeneity as it is also a possible (although atypical) feature of a representative agent model (e.g. Kocherlakota (2000)).

[24] Notable exceptions are Nakajima (2018), van den Heuvel, Vandermarliere, and Schoors (2019) and Baugh, Ben-David, Park, and Parker (2021).

Deaton (1991), Huggett (1993), Carroll (2001); also see Jappelli and Pistaferri (2010)). However, households with truly binding credit constraints are relatively uncommon, and so liquidity constraints have become a dominant part of the literature as they can affect a greater share of households under certain conditions, most prominently in the Kaplan-Violante two asset model (Kaplan and Violante (2014)).

These financial frictions can also drive MPC asymmetry. For example, a household who is at a binding credit constraint must absorb negative shocks entirely via changes in consumption, whereas positive shocks may be partially saved in order to move away from the constraint. This is also true of households who are sufficiently near, but not precisely at, the constraint. To demonstrate this point (and provide theoretical motivation for investigating asymmetry) I construct MPCs for simulated households in the region of a credit constraint using a benchmark life-cycle model.[25] The strict credit constraint at zero is enough to generate asymmetry in transitory MPCs in aggregate (Table 3.1). Further splitting the sample based on wealth finds that asymmetry is stark near the constraint, but it is also present for households away from the constraint.

			Asymmetry over subsamples		
Shocks	Benchmark	Asymmetry	Zero wealth	Near zero	Unconstrained
Permanent	0.77	0.77	1.00	0.97	0.74
Transitory	0.18	–	–	–	–
Positive	–	0.12	0.28	0.22	0.10
Negative	–	0.23	0.91	0.83	0.16

Table 3.1: Simulated consumption elasticities from the Kaplan and Violante (2010) model. Near zero households are taken as those with positive wealth that is less than half of one month's income, as per hand-to-mouth households in Kaplan et al. (2014).

Observation of aggregate *symmetry* may obscure substantial *asymmetry* in the cross-section. Indeed, I find little evidence of asymmetry in aggregate, but robust asymmetry when conditioning on liquidity, home equity and age. The main

[25] I use the Kaplan and Violante (2010) incomplete-markets model with a credit constraint at zero and the same calibration as in their paper. Around 7 per cent of households are exactly on the constraint. I use a simple regression (rather than the method of moments) to construct MPCs from the true shocks, using only working age households.

findings are:

1. Positive asymmetry (larger MPCs for positive shocks than negative shocks) for households with substantial home equity.
2. Negative asymmetry for 'poor hand-to-mouth' households (who have both limited liquidity and no other illiquid wealth).
3. Negative asymmetry for households aged 55–64, irrespective of financial position.

Result (1) contrasts evidence on 'reported preferences', and also contrasts the standard theoretical models that predict wealthy households exhibit no asymmetry and very small responses to transitory income shocks. The result is intuitive from a simple heuristic perspective – households use their financial position to buffer against unexpected declines in their income – but is counter to the consumption smoothing properties of the Euler equation for such households. The result is broadly consistent with other 'revealed preferences' studies and suggests behavioural interpretations where consumption smoothing is directional, much like income insurance.

Result (2) is consistent with the predictions of the Kaplan-Violante two asset model. Theoretical predictions depend on the type of financial constraint that is binding. A credit constraint engenders negative asymmetry. However, the liquidity constraint faced primarily by *wealthy* 'hand-to-mouth' (HtM) households in the two-asset framework has ambiguous implications for MPC asymmetry. The estimates of negative asymmetry for poor HtM households and lack of asymmetry for wealthy HtM households accord.

Result (3) is unexpected. The financial drivers in results (1) and (2) still affect the degree of asymmetry for these households, but asymmetry is shifted into the negative region across the board. Large consumption declines in response to negative income shocks and very little increase in response to positive shocks implies a savings preference for these households, who are close to retirement. This could be the result of inadequate retirement savings during earlier stages of the working life, or uncertainty over the timing of retirement.

There are few estimates of MPCs out of negative income shocks since the most common identification strategy uses fiscal stimulus or tax rebates, which are far more common than unexpected temporary tax increases or levies (e.g. Johnson et al. (2006), Parker et al. (2013)). However, a strand of literature addresses this issue via estimating 'reported preferences' (Fuster, Kaplan, and Zafar (2018), Bunn et al. (2018), Christelis, Georgarakos, Jappelli, Pistaferri, and van Rooij (2017)). That is, specialised surveys are constructed that ask respondents about their reactions to either a hypothetical or recalled income shock. This literature finds substantial *negative* asymmetry across the entire sample, not solely for constrained households.

In contrast, panel data on household income and consumption can provide direct evidence on 'revealed preferences' with respect to both positive and negative income shocks, which is the approach taken here. I extend the Blundell et al. (2008) (BPP) framework detailed in Chapter 2 to allow for the MPC out of transitory income to vary conditional on the sign of the shock, and for all MPCs to vary with household characteristics. The estimation method is bespoke, overcoming nonlinearity in the state space model, and further validated by simulations. I focus on households' HtM status (which captures households who face liquidity frictions), home equity and household head age, all of which have been shown to interact with life-cycle consumption decisions (Kaplan et al. (2014), Hurst and Stafford (2004), Fagereng et al. (2018), Wong (2018), Kovacs and Moran (2021)). Arellano, Blundell, and Bonhomme (2017) also consider asymmetry in the BPP framework, but limit asymmetry to affecting the persistence of shocks, whereas here I limit it to transitory MPCs.

Baugh et al. (2021) and van den Heuvel et al. (2019) conduct empirical studies that are close to the objective of this chapter, using transaction level household data. Baugh et al. (2021) find that all households (across a range of liquidity positions) consume some portion of expected tax refunds when they are paid, but do not reduce their consumption in response to paying tax obligations. van den Heuvel et al. (2019) similarly find overall positive asymmetry for consumption responses to expected income changes, but find that households without liquid wealth do not exhibit asymmetry. I examine responses to permanent and transitory income shocks that can be interpreted as unexpected using the BPP framework. Although I also find evidence of positive asymmetry, it is driven by home equity, rather than liquidity.

3.2 Empirical Model & Data

The empirical framework is developed in Section 2.1, with the benchmark model repeated here for convenience. Working directly with unexplained income y and unexplained consumption c, the model is given by

$$y_{i,t} = \mu^y_{i,t} + \nu^y_{i,t} + \theta\nu^y_{i,t-1} \tag{2.3}$$
$$\mu^y_{i,t} = \mu^y_{i,t-1} + \zeta^y_{i,t} \tag{2.4}$$
$$c_{i,t} = \gamma_\mu\mu^y_{i,t} + \mu^c_{i,t} + \nu^c_{i,t} \tag{2.5}$$
$$\mu^c_{i,t} = \mu^c_{i,t-1} + \gamma_\nu\nu^y_{i,t} + \zeta^c_{i,t}. \tag{2.6}$$

Unless specifically refered to in levels, unexplained income and unexplained consumption are simply called income and consumption. The coefficient γ_ν is taken as the MPC out of transitory income shocks and γ_μ as the MPC out of permanent income shocks. Recall that $1 - \gamma_\mu$ is households' partial insurance against permanent income shocks.

3.2.1 Conditioning on Balance Sheet Characteristics

The empirical literature provides strong evidence that MPCs vary across households. This could be driven by a range of household characteristics, particularly the household balance sheet. The simplest method (and one commonly used) to gauge this heterogeneity is to estimate the model on subsamples of the data independently. This implicitly makes the assumption that the MA(1) coefficient and variances also vary across subsamples. In addition, households often switch status (say between HtM and non-HtM) and so subsampling has the drawback of either losing observations or introducing misclassification error by fixing household status across time.

Instead, I allow MPCs (γ_ν and γ_μ) to vary across household and time according to one or several binary and mutually exclusive indicator variables. This provides estimation of group-specific MPCs, but common variance and MA(1) parameters. Equations (2.5) and (2.6) become

$$c_{i,t} = (B_{i,t}\Gamma_\mu)\mu^y_{i,t} + \mu^c_{i,t} + \nu^c_{i,t} \tag{3.1}$$

$$\mu^c_{i,t} = \mu^c_{i,t-1} + (B_{i,t}\Gamma_\nu)\nu^y_{i,t} + \zeta^c_{i,t}. \tag{3.2}$$

where $B_{i,t}$ is a $1 \times k$ selection vector indicating which of k groups the observation belongs to, and Γ_μ and Γ_ν are now $k \times 1$ vectors of MPC coefficients for the k groups. Note the model nests the base case of homogeneous MPCs across the sample ($k = 1$ and $B_{i,t} = 1$), which at times are referred to using the parameters γ_μ and γ_μ for simplicity.

3.2.2 Introducing Asymmetry

A further extension to the benchmark specification, which is the main contribution of this chapter, is to allow the effects of transitory shocks to differ depending on the sign of the shock. Thus I can observe asymmetry in MPCs out of transitory income shocks across different household characteristics. Equation (3.2) now becomes

$$\mu^c_{i,t} = \mu^c_{i,t-1} + (B_{i,t}\Gamma^+_\nu)\nu^{y,+}_{i,t} + (B_{i,t}\Gamma^-_\nu)\nu^{y,-}_{i,t} + \zeta^c_{i,t} \tag{3.3}$$

where Γ^+_ν and Γ^-_ν are the MPCs out of positive and negative transitory income shocks, while $\nu^{y,+}_{i,t} = \mathbb{1}^+\nu^y_{i,t}$ where $\mathbb{1}^+$ takes the value of one if the shock is positive and zero otherwise (and similar for $\nu^{y,-}_{i,t}$).

The extension to asymmetry is a simple modification to the framework, but it introduces a nonlinearity into the state space framework. As mentioned in Section 2.2, this makes estimation more difficult as the linear Kalman filter cannot be used. The state space system is *functionally* nonlinear, in the language of Harvey (1989). Intuitively, households respond differently to positive and negative income shocks, which creates a kink in the consumption response function at zero change in income. Mathematically, introducing asymmetry means the permanent income states μ^y have a nonlinear and discontinuous relationship with μ^c via the binary operator $\mathbb{1}^+$, as seen by rewriting equation (3.3)

$$\mu^c_{i,t} = \mu^c_{i,t-1}+(B_{i,t}\Gamma^+_\nu)\mathbb{1}^+(y_{i,t}-\mu^y_{i,t}-\theta\nu^y_{i,t-1})+(B_{i,t}\Gamma^-_\nu)(1-\mathbb{1}^+)(y_{i,t}-\mu^y_{i,t}-\theta\nu^y_{i,t-1})+\zeta^c_{i,t}.$$

The Γ^+_ν and Γ^-_ν parameters depend on the sign of *unobserved* shocks. A Metropolis-Hastings routine is used to overcome this estimation issue, as discussed in Section 2.2. Cho et al. (2021) investigate asymmetry in a very similar setting, but only with respect to the sign of *observed* income growth.

3.2.3 Model Simulation

The estimation routine can be tested by simulation to ensure it adequately identifies true parameter values. The data generating process (DGP) for a single household is given by equations (2.3)-(2.6) for the benchmark symmetry model and equations (2.3),(2.4),(3.1) and (3.3) for the asymmetry model. All shocks are assumed to follow a normal distribution. MPCs are first assumed to be homogenous across the sample. True parameter values are taken as similar to what might be expected in the dataset. The variances are set to $\sigma^2_{\nu y} = 0.06$, $\sigma^2_{\nu c} = 0.12$, $\sigma^2_{\zeta y} = 0.04$ and $\sigma^2_{\zeta c} = 0.04$, which are based on raw sample moments. The MA(1) coefficient is set to $\theta = 0.2$ based on an early version of the Chatterjee et al. (2021) (CMS) paper, and $\gamma_\mu = 0.6$ is between the BPP and CMS estimates. The choice of $\gamma_\nu = 0.3$ is inflated relative to the literature for the sake of clearer exposition, and similar rationale applies to setting $\gamma^+_\nu = 0.5$ and $\gamma^-_\nu = 0.1$. Posterior distributions are constructed using 2,000 draws after 1,000 burn in, and the statistics reported (mean and 90 per cent highest posterior density (HPD) interval) are averages over 100 replications for each set of parameters.[26]

The base results in Tables 3.2 and 3.3 use a balanced panel of $N = 500$ households and $T = 5$ time observations (the PSID dataset has around 4,500 households with an average of around 6.5 time observations each). The simulations indicate that the parameters are suitably well identified, lying within the 90 per cent confidence interval, and that the priors are not driving the results. The exercise is repeated

[26]The algorithm converges rapidly due to the high dimensionality of the data, thus 2,000 draws is sufficient for the simulation study and lowers the computational burden.

	True	Mean	5%	95%
		Average Statistic		
γ_ν	0.300	0.330	0.221	0.439
γ_μ	0.600	0.583	0.460	0.710
θ	0.200	0.175	0.095	0.249
$\sigma^2_{\nu y}$	0.060	0.059	0.052	0.067
$\sigma^2_{\zeta y}$	0.040	0.041	0.033	0.049
$\sigma^2_{\nu c}$	0.120	0.124	0.113	0.134
$\sigma^2_{\zeta c}$	0.040	0.033	0.024	0.043

Table 3.2: Benchmark symmetry model. Averages of each posterior statistic (mean, 5% & 95% HPD) over 100 replications. $N = 500$, $T = 5$.

	True	Mean	5%	95%
		Average Statistic		
γ^+_ν	0.500	0.529	0.412	0.642
γ^-_ν	0.100	0.149	0.031	0.263
γ_μ	0.600	0.597	0.471	0.729
θ	0.200	0.222	0.138	0.298
$\sigma^2_{\nu y}$	0.060	0.063	0.055	0.071
$\sigma^2_{\zeta y}$	0.040	0.038	0.030	0.046
$\sigma^2_{\nu c}$	0.120	0.124	0.113	0.134
$\sigma^2_{\zeta c}$	0.040	0.033	0.024	0.043

Table 3.3: Baseline asymmetry model. Averages of each posterior statistic (mean, 5% & 95% HPD) over 100 replications. $N = 500$, $T = 5$.

using a simulation of households with heterogeneous MPCs. Further discussion and simulation results are provided in Appendix B.1.

3.2.4 Data

Details of the PSID dataset were provided in Section 2.3. This chapter utilises only a few household covariates, all of which are discretised into groups. Specifically, I consider how HtM status (poor, wealthy or non-HtM), income (top two quintiles and bottom three quintiles), mortgage-to-income ratio (MIR) terciles, home equity terciles, and age (25–34, 35–44, 45–54, 55–64) drive MPC asymmetry in separate regressions. Summary statistics show that the sample split by the estimated sign of $\nu^y_{i,t}$ does not contain any large differences in distributions (Table 3.4).[27] There is some evidence that negative income shocks are more prevalent for lower income households with higher mortgage-income ratios, but these results are not the main focus of the chapter.

[27] A given $\nu^y_{i,t}$ is taken as positive if the share of the posterior distribution above zero is greater than or equal to 0.5, $P(\hat{\nu}^y_{i,t} \geq 0) \geq 0.5$.

	Full sample		Positive shock		Negative shock	
	Mean	Std.Dev.	Mean	Std.Dev.	Mean	Std.Dev.
y	0.001	0.495	0.004	0.495	-0.003	0.495
Δy	0.007	0.383	0.006	0.383	0.008	0.383
c	-0.006	0.431	-0.004	0.428	-0.008	0.434
Δc	0.004	0.363	0.005	0.359	0.003	0.367
HtM	0.443	0.497	0.442	0.497	0.444	0.497
Poor HtM	0.205	0.404	0.206	0.404	0.204	0.403
Wealthy HtM	0.238	0.426	0.237	0.425	0.240	0.427
Income ($000)	67.448	43.147	67.803	43.493	67.080	42.783
MIR	0.897	1.219	0.889	1.200	0.906	1.239
Home equity ($000)	72.150	117.202	73.040	119.312	71.227	114.970
Age	42.062	10.327	42.094	10.308	42.028	10.346
Observations	29645					
Households	4528					

A positive shock has $P(\hat{\nu}^y_{i,t} \geq 0) \geq 0.5$ and negative $P(\hat{\nu}^y_{i,t} \geq 0) < 0.5$, as estimated in the intercept only asymmetry model (see Table 3.6).

Table 3.4: Summary statistics.

3.3 Results

3.3.1 Benchmark Results under Symmetry

Estimates from the benchmark symmetry model are presented first to anchor the results.[28] The posterior distributions are very narrow and the results are similar to expectations, given the existing literature (Table 3.5). The two MPC parameters appear to be well identified (Figures 3.1 and 3.2). The MPC out of transitory income shocks γ_ν is small relative to that out of permanent income shocks γ_μ, and the variance of permanent income shocks $\sigma^2_{\zeta y}$ is small relative to the variance of transitory income shocks $\sigma^2_{\nu y}$.

[28]Unless stated otherwise, all results use 5,000 draws after 5,000 burn-in. The benchmark symmetry and asymmetry results use 10,000 draws and burn-in.

	Mean	5%	95%
γ_ν	0.111	0.093	0.130
γ_μ	0.533	0.497	0.569
θ	0.087	0.070	0.105
$\sigma_{\nu y}$	0.260	0.257	0.264
$\sigma_{\zeta y}$	0.153	0.148	0.159
$\sigma_{\nu c}$	0.243	0.240	0.245
$\sigma_{\zeta c}$	0.104	0.099	0.108
Observations	29645		
Households	4528		

Table 3.5: Benchmark symmetry model; statistics reflect 10,000 draws after 10,000 burn-in.

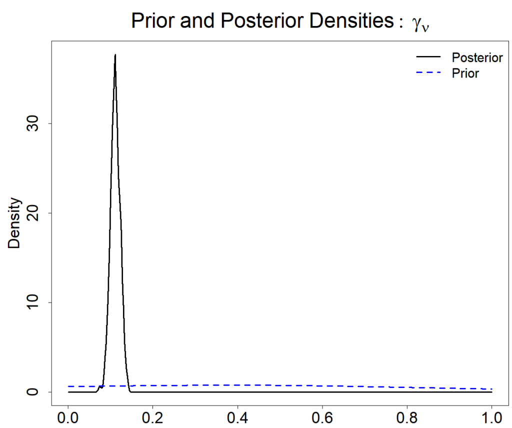

Figure 3.1: Posterior density and prior of transitory MPC γ_ν.

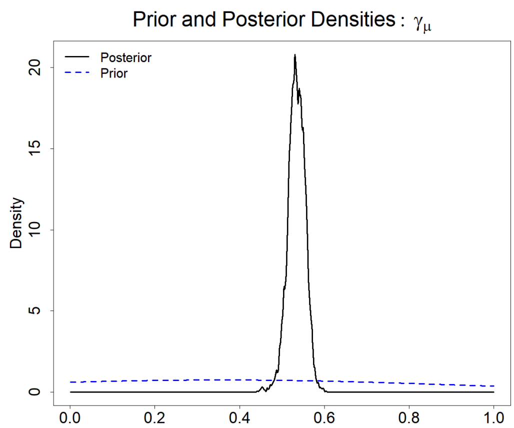

Figure 3.2: Posterior density and prior of permanent MPC γ_μ.

3.3.2 Balance Sheet Drivers of Asymmetry

Baseline results for the asymmetry model with only a common intercept are shown in Table 3.6.[29] Estimates are very similar to those obtained under symmetry and show little evidence of asymmetric MPCs. However, much like the previous literature has found with respect to heterogeneity in the relative size of MPCs, the aggregate results mask important differences that are uncovered when conditioning on household balances sheets.

[29]Estimation diagnostics are provided in Appendix B.3.

	Mean	5%	95%
γ_ν^+	0.118	0.100	0.137
γ_ν^-	0.108	0.089	0.127
γ_μ	0.534	0.501	0.566
θ	0.090	0.070	0.110
$\sigma_{\nu y}$	0.261	0.257	0.264
$\sigma_{\zeta y}$	0.152	0.148	0.157
$\sigma_{\nu c}$	0.243	0.240	0.245
$\sigma_{\zeta c}$	0.104	0.099	0.109
Observations	29645		
Households	4528		

Table 3.6: Baseline asymmetry intercept only model; statistics reflect 10,000 draws after 10,000 burn-in.

Liquidity

Much of the existing literature focusses on categorising households by their HtM status; that is, separating households who face liquidity frictions from those who do not. Table 3.7 presents asymmetry results split along this dimension. Non-HtM households exhibit a clear positive asymmetry, with little overlap of posterior distributions (Figure 3.3). This is in contrast to the reported preferences literature, which finds broad-based negative asymmetry (Bunn et al. (2018), Christelis et al. (2017), Fuster et al. (2018)), but the finding is consistent with revealed preferences estimates of Baugh et al. (2021) and van den Heuvel et al. (2019). This suggests that when unconstrained by liquidity frictions, households behave in a manner that implies a positive consumption preference relative to their lifetime income. That is, households consume a moderate amount of positive transitory income shocks but tend not to reduce consumption in response to negative shocks – showing a preference to keep consumption high and allowing savings to decline. As would be expected, non-HtM households also show some evidence of responding less to permanent income shocks than HtM households, which accords with their larger capacity to self-insure.

	Mean	5%	95%
HtM households			
γ_ν^+	0.075	0.048	0.101
γ_ν^-	0.134	0.107	0.164
γ_μ	0.547	0.514	0.586
(Group obs.	13137	)	
Non-HtM households			
γ_ν^+	0.152	0.131	0.173
γ_ν^-	0.082	0.058	0.105
γ_μ	0.502	0.470	0.539
(Group obs.	16508	)	
Observations	29645		
Households	4528		

Table 3.7: Asymmetry model, split by HtM status.

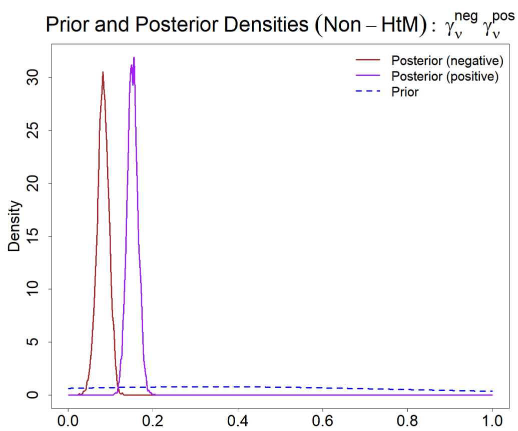

Figure 3.3: Posterior density and prior of transitory MPCs γ_ν^+ and γ_ν^- for non-HtM households.

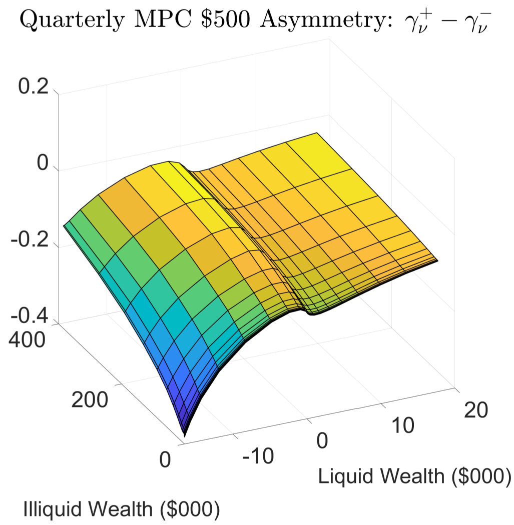

Figure 3.4: Kaplan et al. (2018) MPC asymmetry distribution across liquid and illiquid wealth generated using the replication code; difference between MPC out of positive and negative transitory income shocks; see Appendix C.2 for details.

In contrast, HtM households exhibit negative asymmetry. This is broadly consistent with the reported preferences estimates, which show that negative asymmetry is largest for households with low liquidity. However, theory does not make clear asymmetry predictions for liquidity frictions alone. The consumption response of households in the two-asset model who are near zero liquid wealth, rather than near their credit constraint, depends on a number of factors such as whether their liquid wealth is zero, slightly positive, or slightly negative, and the interest rate gap between borrowing and saving in the liquid asset. Figure 3.4 shows predicted MPC asymmetry with respect to transitory income shocks for households in the Kaplan et al. (2018) model.[30] Asymmetry is most apparent for households with no illiquid wealth and who also face credit or liquidity constraints.

	Mean	5%	95%
Poor HtM households			
γ_ν^+	0.021	-0.021	0.061
γ_ν^-	0.154	0.113	0.194
γ_μ	0.561	0.519	0.607
(Group obs.	6070	)	
Wealthy HtM households			
γ_ν^+	0.121	0.085	0.159
γ_ν^-	0.103	0.066	0.139
γ_μ	0.550	0.509	0.589
(Group obs.	7067	)	

Table 3.8: Asymmetry model, split by poor and wealthy HtM status.

Further disaggregating HtM households into the poor HtM – who have zero illiquid wealth – and the wealthy HtM results in estimates of asymmetry that are consistent with the predictions of the two-asset model (Table 3.8). Negative asymmetry is driven by poor HtM households, whereas wealthy HtM households show no clear signs of asymmetry. Poor HtM households face both liquidity fric-

[30]See Appendix C.2 and Chapter 4 for details of the relationship between the empirical model and the Kaplan et al. (2018) model.

tions and also have no other assets that can be used to smooth consumption.

Home Equity

Home equity has long been connected to household consumption smoothing over the life cycle, in retirement and also to finance consumption during the working life (e.g. Chen and Jensen (1985), Kovacs and Moran (2021)). Withdrawal of home equity, through home equity loans or mortgage refinancing, is linked to households facing liquidity issues and negative financial shocks (Hurst and Stafford (2004), Benito (2009), Agarwal and Qian (2017)). On the other hand, negative home equity can amplify households' consumption response to shocks (Disney, Gathergood, and Henley (2010), Mian, Rao, and Sufi (2013)). Examining the role of home equity provides further insight into how illiquid forms of wealth interact with MPC asymmetry.

Conditioning the model on home equity terciles results in estimates that correspond with those by HtM status (Table 3.9). Households with zero or negative home equity show clear *negative* asymmetry, like the poor HtM; households with some home equity show no asymmetry, like the wealthy HtM; and households with substantial home equity show clear *positive* asymmetry, like the non-HtM. In part, this is expected given home equity accounts for the majority of illiquid wealth and also correlates with liquid wealth. However, there is substantial overlap across the two dimensions. Roughly half of those households with no/negative equity are poor HtM, but almost all of the remainder are non-HtM households, and around one quarter of households with substantial equity are HtM. The results split by HtM indicate that negative asymmetry is driven by a combination of *both* liquidity constraints and a lack of other assets, but the drivers of positive asymmetry are unclear – the result could be driven by ample liquidity or ample illiquid home equity (examined further below).

	Mean	5%	95%
First tercile (none: ≤$0)			
γ_ν^+	0.062	0.029	0.094
γ_ν^-	0.138	0.105	0.169
γ_μ	0.551	0.507	0.600
(Group obs.	10373	)	
Second tercile (some: $1–$68,000)			
γ_ν^+	0.118	0.084	0.154
γ_ν^-	0.126	0.090	0.161
γ_μ	0.524	0.484	0.569
(Group obs.	9464	)	
Third tercile (substantial: $68,000+)			
γ_ν^+	0.158	0.129	0.190
γ_ν^-	0.047	0.019	0.078
γ_μ	0.528	0.485	0.572
(Group obs.	9808	)	

Table 3.9: Asymmetry model, split by home equity terciles.

Age

Housing decisions are closely related to the household life cycle. Fagereng et al. (2018) and Wong (2018) note that age demographics can have a substantial bearing on MPCs, but the relationship is associated with the interaction of different age profiles with their assets and liabilities, particularly housing. Younger households are less likely to own homes and older households more likely to own homes outright. This maps specific age distributions into the home equity groups (Figure 3.5). Furthermore, responses to income shocks should change over the life cycle as the incentives to save and consume change with the proximity of death. The life-cycle model predicts a clear decline in permanent MPCs towards retirement. Estimation is conditioned on household head age to assess whether the life cycle itself has a bearing on MPC asymmetry (Table 3.10).

Conditioning on age broadly supports the decline in permanent MPCs, but gen-

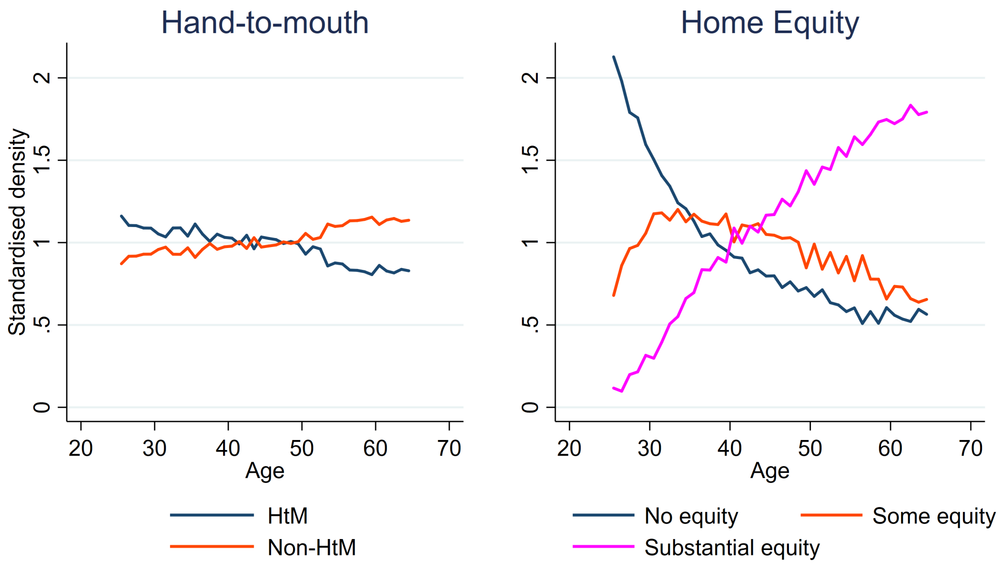

Figure 3.5: Standardised age distributions by balance sheet characteristics; a value of one means that age has the same density as in the full sample.

	Mean	5%	95%
Household head aged 25–34			
γ_ν^+	0.136	0.098	0.174
γ_ν^-	0.088	0.050	0.125
γ_μ	0.542	0.495	0.592
(Group obs.	8614	)	
Household head aged 35–44			
γ_ν^+	0.153	0.122	0.185
γ_ν^-	0.094	0.062	0.126
γ_μ	0.512	0.467	0.556
(Group obs.	8814	)	

	Mean	5%	95%
Household head aged 45–54			
γ_ν^+	0.126	0.094	0.159
γ_ν^-	0.114	0.081	0.149
γ_μ	0.526	0.482	0.572
(Group obs.	7833	)	
Household head aged 55–64			
γ_ν^+	-0.006	-0.053	0.040
γ_ν^-	0.149	0.102	0.193
γ_μ	0.489	0.440	0.538
(Group obs.	4384	)	

Table 3.10: Asymmetry model, split by household head age.

erates some unexpected results for transitory MPCs that contrast the life-cycle nature of housing. Although households with no (or negative) home equity tend to be younger, with those under the age of 30 over-represented, younger households in fact show some evidence of *positive* asymmetry, contrary to the results by home equity. This partially supports the interpretation that negative asymmetry is a result of financial frictions rather than life-cycle factors. Furthermore, the negative asymmetry observed for households with head aged 55–64 is surprising, as these households would be expected to have accumulated enough wealth to self-insure and should be least likely to face credit constraints. They are also over-represented in the substantial home equity group, which has clear positive asymmetry. Overall, the results suggest that the asymmetry found within groups based on financial measures is not driven by the age profile of these groups; however, a model that interacts groups can provide more insight.

Separating Liquidity, Equity and Age

Grouping households along a single covariate finds some overlapping results that can only be separated by interacting covariates. The conditional approach of the empirical model allows this to be done without losing observations, whereas subsampling would result in poor inference due to limitations on the number of observations in each subgroup. Of key interest is determining whether liquidity or home equity drive positive asymmetry, and interrogating the age profile of asymmetry and its connection to household finances. This section presents results from interacting subgroups, honing in on asymmetry by looking at the distribution of the difference of draws of γ_ν^+ and γ_ν^-. That is, the difference $\gamma_\nu^{+(n)} - \gamma_\nu^{-(n)}$ is taken for each draw n in the Markov chain, with statistics calculated from the resulting distribution. Statistical significance is determined by the share of the difference distribution on the opposite side of zero to the mean.

	Home Equity		
	≤\$0	\$1–\$68,000	\$68,000+
Non-HtM	−0.001	0.042	0.105***
HtM	−0.122***	−0.049*	0.122***

***, **, * represent less than 1, 5, 10 per cent of the difference distribution on opposite side of zero.

Table 3.11: MPC asymmetry by HtM and home equity; posterior means of $\gamma_\nu^+ - \gamma_\nu^-$.

The results of conditioning MPCs on both HtM status and home equity find positive asymmetry for all households with substantial home equity (Table 3.11).

That is, liquidity is not the main driver of positive asymmetry. The results also confirm the previous finding for negative asymmetry, which is contained to only households with low liquidity and a lack of other assets.

	Non-HtM	Poor HtM	Wealthy HtM
Aged 25–34	0.111***	−0.156***	0.226***
Aged 35–44	0.140***	−0.134***	0.051
Aged 45–54	0.082***	−0.132***	−0.021
Aged 55–64	−0.113***	−0.505***	−0.180***

***, **, * represent less than 1, 5, 10 per cent of the difference distribution on opposite side of zero.

Table 3.12: MPC asymmetry by age and HtM; posterior means of $\gamma_{\nu}^{+} - \gamma_{\nu}^{-}$.

	Home Equity		
	≤\$0	\$1–\$68,000	\$68,000+
Aged 25–34	−0.059**	0.070*	0.466***
Aged 35–44	−0.093***	0.086**	0.180***
Aged 45–54	−0.067**	−0.122***	0.154***
Aged 55–64	−0.349***	−0.169***	−0.105***

***, **, * represent less than 1, 5, 10 per cent of the difference distribution on opposite side of zero.

Table 3.13: MPC asymmetry by age and home equity; posterior means of $\gamma_{\nu}^{+} - \gamma_{\nu}^{-}$.

The results of conditioning on age and either HtM status or home equity find that positive asymmetry is accentuated in the youngest age group and negative asymmetry is accentuated in the oldest (Tables 3.12 & 3.13). The asymmetry identified above for poor HtM households and those with substantial home equity continues to be seen in age groups 25–54. However, all households with heads aged 55–64 show significant negative asymmetry. These households should also be the least likely to face financial constraints.[31] The household balance sheet

[31] In the sample, 10 per cent of observations of those aged 55-64 are poor HtM and 27 per cent are wealthy HtM. These shares could be upwardly biased due to self selection into retirement of households with adequate savings.

continues to play a role for these households – with the largest asymmetry seen in the poor HtM subgroup and the smallest asymmetry in the subgroup with substantial home equity – but financial characteristics are overwhelmed by a tendency for large consumption responses to negative income shocks.

The sample selection includes only households who are active in the labour force, and removes the observation prior to retirement. Nonetheless, these results may reflect the anticipation of retirement, as negative asymmetry implies a savings preference relative to lifetime income. The optimal savings profile of a life-cycle model predicts households reduce their average savings propensity (if anything) in the years prior to retirement. Yet Cagetti (2003) shows in a life-cycle model that retirement only begins to be an important motive for saving, relative to precautionary motives, after around 50 years of age. This is also consistent with the larger share of respondents aged 55–64 reporting retirement as a reason for saving in the Survey of Consumer Finances. To generate MPC asymmetry, however, current savings of these households would need to be below their optimal target. Uncertainty in the timing of retirement can generate additional precautionary motives for saving as the probability of retirement nears (Blau (2008)).

3.3.3 Comparison to Survey Evidence

The evidence on revealed preferences reported above contrasts the evidence found in research that relies on reported preferences in some key aspects.[32] Bunn et al. (2018) asked British households to recall shocks that occurred in the past year and their response to them. They find that household MPCs are much larger for unanticipated falls in income (0.64) than for rises (0.14). The result is broad-based across many cross-sectional dimensions, and regressions on observable correlates cannot fully explain the gap in MPCs. Christelis et al. (2017) surveyed Dutch households, presented respondents with hypothetical scenarios of different sized and signed income shocks, and asked respondents how they would react. They find that in the presence of liquidity constraints, MPCs are larger for negative

[32]Note that the reported preferences estimates do not explicitly measure the permanent consumption response, so the magnitude of estimates are not directly comparable to those of this thesis.

income shocks (0.50) compared with positive (0.39). In addition, asymmetry is pervasive across the entire sample for larger shocks and cannot be explained fully by observable characteristics. Fuster et al. (2018) conducted a similar survey on US households, with the mean negative MPC for a $500 shock (0.30) substantially larger than for a positive shock of the same size (0.08).

Overall, the reported preference studies find pervasive *negative* asymmetry across the whole sample, whereas I only find robust evidence of negative asymmetry for households who are most likely to face financial constraints and those aged 55–64. This suggests that all households report behaving in a financially conservative way (saving windfalls and reducing expenditure to cover losses), but only those facing genuine financial pressure follow through in action. Notably, I find that asymmetry runs in the opposite direction for non-HtM households overall (Figure 3.6). The reported preference studies rely on survey responses matching true reactions. Various studies analyse the pitfalls of recall and hypothetical questions, so it is no surprise that reported and revealed preferences differ (e.g. Bertrand and Mullainathan (2001), Browning, Crossley, and Weber (2003), Crossley and Winter (2013)). However, the extent to which they differ – a net response of the opposite sign – presents a puzzle.

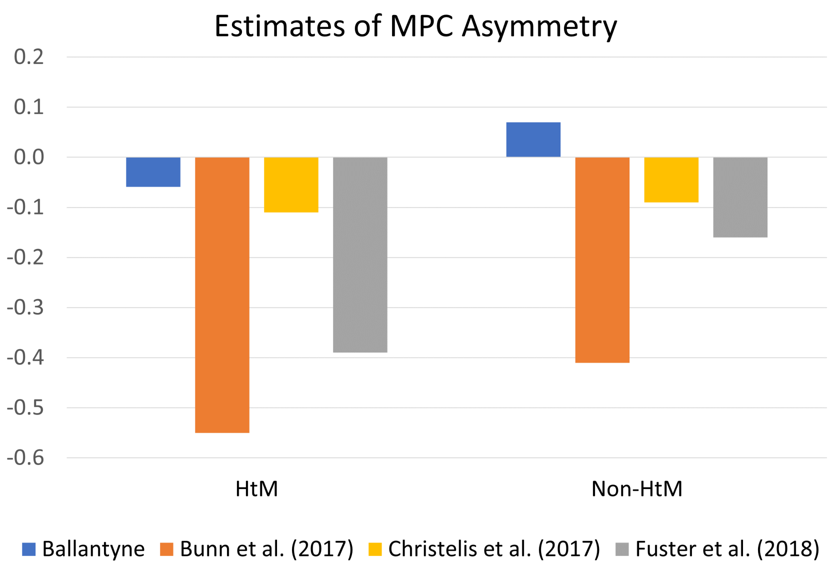

Figure 3.6: Comparison of asymmetry findings against the reported preferences literature. Bars are positive MPC minus negative MPC, with data derived from tables or graphs. HtM status for Bunn et al. (2018) and Christelis et al. (2017) is taken as the average over the bottom half of the liquid asset or cash-on-hand distribution, with non-HtM households the average over the remaining households.

3.4 Robustness

3.4.1 Asymmetry across Specifications

To establish the robustness of the findings, I estimate the model with and without the MA(1) parameter and using the original KVW sample (Table 3.14). Estimation using separate subsamples – assuming the variance and MA(1) parameters are unique to the subgroup, rather than conditioning only the MPCs – is also checked. I focus on the split by HtM status for simplicity in capturing asymmetry in both directions. In all cases and for both subgroups, the probability of asymmetry being reversed is less than 0.1 per cent.[33]

There is evidence that separate subsampling shifts the overall size of MPCs for HtM households (Table 3.14). This is likely driven by the variance of the shocks

[33]Calculated as the share of posterior draws in which $\gamma_\nu^- > \gamma_\nu^+$ for non-HtM and $\gamma_\nu^> - \gamma_\nu^-$ for HtM households.

	Non-HtM		HtM	
	γ_ν^+	γ_ν^-	γ_ν^+	γ_ν^-
MA(0) model				
KVW sample	0.179	0.058	0.114	0.191
KVW(ext) sample	0.154	0.077	0.070	0.136
MA(1) model				
KVW sample	0.173	0.061	0.114	0.183
KVW(ext) sample	0.152	0.082	0.075	0.134
Separate subsamples	0.137	0.046	0.124	0.190

Table 3.14: Mean estimates of transitory MPCs for various specifications split by HtM status.

	Non-HtM			HtM		
	Mean	5%	95%	Mean	5%	95%
θ	0.050	0.005	0.085	0.070	0.030	0.110
$\sigma_{\nu y}$	0.236	0.230	0.243	0.272	0.264	0.281
$\sigma_{\zeta y}$	0.152	0.142	0.160	0.168	0.156	0.181
$\sigma_{\nu c}$	0.202	0.197	0.206	0.274	0.269	0.280
$\sigma_{\zeta c}$	0.122	0.114	0.130	0.094	0.081	0.106

Table 3.15: Posterior statistics for estimations with separate subsamples.

being specific to the subgroup (Table 3.15). Notably, HtM households face transitory income and transitory consumption shocks with larger variance. For these shocks, there is no overlap of the 90 per cent HPD intervals for HtM and non-HtM households, as well as between these households and the common estimates in Table 3.6. This suggests a more flexible specification would be preferred, although this does not affect the results concerning asymmetry.

3.4.2 Are Priors Driving the Results?

The priors for γ_ν^+ and γ_ν^- are the same and so should not drive the asymmetry results. Nonetheless, a useful robustness check is to gauge what the priors on these parameters would have to be in order to eliminate the result. If a prior that eliminates the result is also an economically reasonable assumption, then the robustness of the result can be called into question. I conduct this analysis on the non-HtM households in two ways. First, I keep the same prior for both γ_ν^+ and γ_ν^-, but shift the prior to be centered at 0.12 (in between the previous estimates). Second, I consider different priors for each of γ_ν^+ and γ_ν^-, where I take the mean of the γ_ν^+ prior to be zero and the mean of the γ_ν^- prior to be 0.2. This is a prior that reverses the non-HtM findings but provides information similar to the survey evidence. In both methods I estimate the model for a range of prior variances (common to both parameters) decreasing from 0.25 (the baseline estimation prior variance) down to 2.5×10^{-5} (which implies an extremely strong prior).

Comparing when the posterior densities of each parameter overlap provides an indication of a prior that eliminates the baseline result (Figures 3.7 & 3.8). For the case with the same prior mean, a variance of 2.5×10^{-5} is required for substantial overlap, whereas when they have different prior means a variance of 2.5×10^{-4} is required. In the former case, this prior implies a 95 per cent certainty that both parameters lie between 0.11 and 0.13 – clearly an absurd prior (Figure 3.9). In the latter case, the prior implies a 95 per cent certainty that γ_ν^+ lies between -0.03 and 0.03, and a 95 per cent certainty that γ_ν^- lies between 0.17 and 0.23. Again an absurd prior.

3.5 Discussion

The results can be summarised by three stylised facts on MPC asymmetry:

1. Positive asymmetry for households with substantial home equity.

2. Negative asymmetry for poor HtM households.

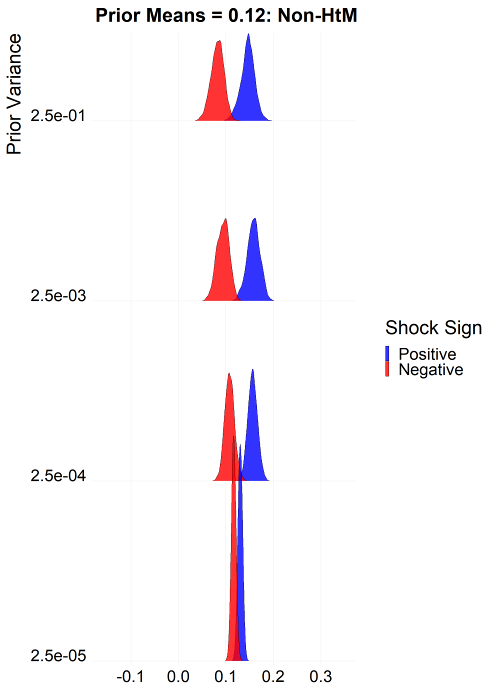

Figure 3.7: Posterior distributions of γ_ν^+ and γ_ν^- for non-HtM households under alternative priors (with the same prior mean).

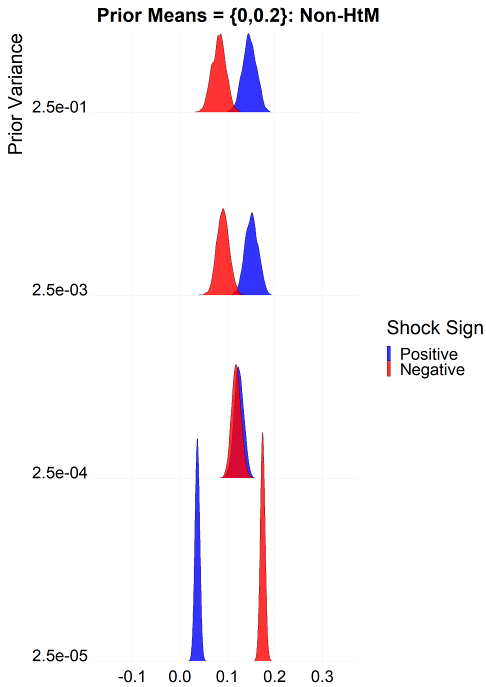

Figure 3.8: Posterior distributions of γ_ν^+ and γ_ν^- for non-HtM households under alternative priors (with different prior means).

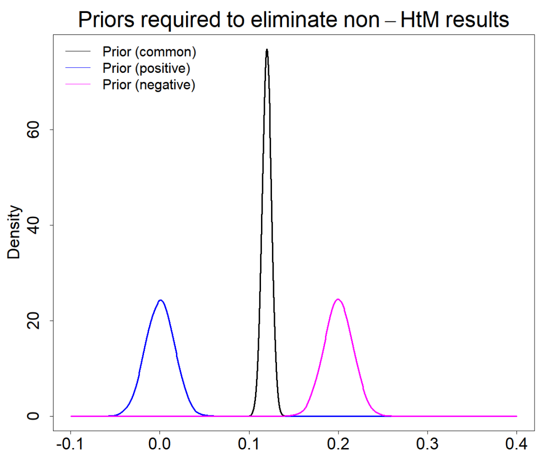

Figure 3.9: Priors on γ_ν^+ and γ_ν^- required to eliminate difference between parameters for non-HtM households.

3. Negative asymmetry for households aged 55–64, irrespective of financial position.

The results provide mixed evidence in supporting and contrasting theoretical predictions and previous empirical findings. Result (1) is broadly consistent with other revealed preferences studies, which also find positive MPC asymmetry (van den Heuvel et al. (2019), Baugh et al. (2021)). However, it contrasts the evidence on reported preferences and contrasts the standard theoretical models that predict wealthy households exhibit no asymmetry and very small responses to transitory income shocks. The result is intuitive from a simple heuristic perspective – households use their financial position to smooth through temporary declines in their income – but is counter to the consumption smoothing properties of the Euler equation for such households. This could be interpreted as households using home equity to smooth consumption in a directional manner, similar to income insurance that pays out for negative shocks (and does not penalise positive shocks).

Reproducing positive MPC asymmetry across a fairly broad range of household balance sheet positions does not appear to be consistent with known financial frictions, which only bind at particular parts of the wealth distribution. Asymmetry typically features due to curvature in the consumption function near constraints. For example, the two-asset model features a consumption policy function with curvature near the credit limit and at zero liquid wealth (Figure 3.10). This generates negative MPC asymmetry near the credit limit (as pictured), positive asymmetry just below zero and negative asymmetry just above zero.[34] However, asymmetry is negligible for all liquid wealth positions above a few thousand dollars.

Result (1) points towards behavioural interpretations. Baugh et al. (2021) posit that models of mental accounts help to explain their findings with respect to tax refunds, and conjecture that tax payments may be made out of a different mental account. Fuster et al. (2018) also suggest a model of mental accounts may help explain asymmetry in the opposite direction. Such a framework might allow separation of the household consumption function with respect to positive and

[34] This description takes changes in liquid wealth as equivalent to income shocks.

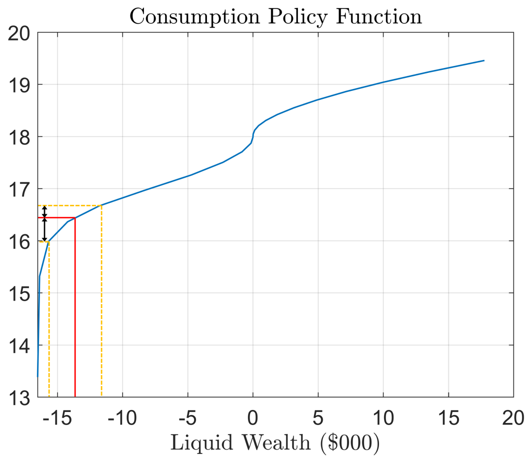

Figure 3.10: Consumption policy function for liquid wealth states in the two-asset model; average over income and illiquid wealth states; generated using replication code of Kaplan et al. (2018).

negative income shocks, generating asymmetry without relying only on curvature of the function. However, in the context of the results here, mentally separating positive and negative income shocks into different accounts seems arbitrary, and the relationship to home equity is opaque.

A mechanism such as consumption loss aversion may be able to explain the difference between the reported and revealed preferences via cognitive bias (e.g. Kahneman and Tversky (1979)). It could also affect all households and so might be able to generate positive asymmetry across a wide range of asset positions. An interesting precedent for this is Aizenman (1998), who finds that loss aversion in a precautionary savings model leads to the accumulation of additional savings of a first-order magnitude compared with the standard motive generated by the concavity of marginal utility. In addition, van den Heuvel et al. (2019) propose myopic loss aversion to explain their empirical findings of positive asymmetry for households who have liquid assets.

An empirical link is established between home equity access and smoothing negative income shocks (Hurst and Stafford (2004), Sodini, Van Nieuwerburgh, Vestman, and von Lilienfeld-Toal (2016)); however, the theoretical link does not appear to be developed. Kovacs and Moran (2021) model access to home equity under temptation, where households use commitment devices (illiquid home equity) to overcome the temptation of immediate gratification. The implications for MPC asymmetry are not clearly established; a shift to allowing home equity access decreases MPCs in the short term (as households can buffer negative shocks) but increases MPCs in the long term due to households carrying larger mortgage balances and thus being closer to financial constraints. The model may then generate positive asymmetry in the short run and negative asymmetry in the long run.

Result (3) is unexpected as pre-retirement households should have the largest financial capacity to smooth consumption. The financial drivers in results (1) and (2) still affect the degree of asymmetry for these households, but asymmetry is shifted into the negative region across the board. This is interpreted as an implied savings preference – the expected consumption response to a mean-zero transitory shock is negative, implying the household is trying to save out of transitory income.[35] Such a dynamic is consistent with increased financial uncertainty around retirement or a behavioural bias such as time-inconsistency of retirement savings.

It is plausible that a substantial share of the age group have a shortfall in wealth relative to the optimal (Lusardi and Mitchell (2007), Skinner (2007)). Myopia would lead households to undersave during the working life, as suggested by Hurst (2003). Indeed, hyperbolic discounting (a version of myopia) predicts that households delay retirement savings; however, simulations do not show a clear change in savings behaviour in the period before retirement (e.g. Angeletos, Laibson, Repetto, Tobacman, and Weinberg (2001)). A shift towards more aggressive saving late in the working life can be generated by precautionary motives such as an uncertain timing for retirement, as in Blau (2008). Households nearing the end of the working life may be caught out by early retirement and so have an incentive to accumulate additional savings. This generates a precautionary

[35]It is possible that unobserved heterogeneity could drive the result; however, there is also little evidence for this. Across all age groups, the share and size of positive shocks and negative shocks are very similar. The model further controls for household-specific fixed effects on income and consumption by estimating the time-zero parameters $\mu^{c}_{i,0}$ and $\mu^{y}_{i,0}$.

savings motive similar to that of a credit constraint, with the precautionary motive becoming sharper as the constraint is more likely to bind. Average saving rates in the PSID data accord – showing a shift towards increased saving after age 55 – which is at odds with life-cycle models in which retirement age is known (e.g. Kaplan and Violante (2010)). Ample literature addresses the gaps between observed behaviour and the life-cycle hypothesis in terms of inadequate retirement savings, consumption decline at retirement and the lack of dissaving during retirement (e.g. Davies (1981), Banks, Blundell, and Tanner (1995), Jappelli and Modigliani (1998), Skinner (2007), Hurst (2008), Love, Palumbo, and Smith (2009)), but there appears to be little work studying savings patterns leading into retirement specifically.

3.6 Conclusion

I extend the BPP framework introduced in Chapter 2 to estimate household MPCs conditional on the sign of transitory income shocks, and conditional on different household characteristics. This extension induces nonlinearity in the state space framework and so Bayesian estimation is a natural choice. Simulation studies provide evidence that inference is satisfactory.

Under the baseline specification, there is little evidence of MPC asymmetry over the full sample, but this obscures important differences when taking into account heterogeneity in household balance sheets. A central result of the chapter is clear evidence of *positive* asymmetry (larger response to positive shocks) for households with substantial home equity. For these households, the MPC out of positive transitory shocks is roughly three times that of negative shocks. This finding contrasts survey evidence on 'reported preferences' that directly ask households about their MPCs in hypothetical situations; however, it is largely consistent with other revealed preferences studies. Explanation of this result primarily focusses on behavioural mechanisms.

The results within HtM households offer support for some of the predictions of the two-asset model. Households who are most likely to face financial constraints

– the poor HtM – show marked *negative* MPC asymmetry, whereas wealthy HtM households do not. The asymmetry results by HtM household and home equity are found across households with heads aged 25–54. In contrast, those aged 55–64 display negative asymmetry across all balance sheet positions, which is interpreted as a savings preference late in the working life.

Chapter 4

Heterogeneity & Aggregation

4.1 Introduction

Consumption decisions vary from one household to another in ways that can affect aggregate economic outcomes. The marginal propensity to consume (MPC) out of income shocks plays an important role. However, MPCs can be difficult to estimate in the data, which has limited the capacity to measure the *distribution* of MPCs in the population. Existing empirical methods tend to rely on subsampling along one or two observable dimensions.[36] Building on the empirical framework in Chapter 2, I develop an approach to estimate the distribution of household MPCs out of both transitory and permanent income shocks using data from the US Panel Study of Income Dynamics (PSID) over 1999–2017. The approach allows for observable household characteristics to drive MPC heterogeneity, while accommodating potential nonlinearities, and is extended to incorporate *unobserved* heterogeneity. In addition, population-representative MPCs are carefully aggregated from microdata elasticities using an extended dataset that includes retired households.

A standard empirical approach is to capture MPC heterogeneity by separating the sample into groups along an observable characteristic like liquid wealth. Modelling the distribution of MPCs (rather than subsampling) provides a more direct connection between the empirical model and theoretical models that feature heterogeneity. This can be seen clearly in the two-asset model (Kaplan and Violante

[36]Recent exceptions are Lewis, Melcangi, and Pilossoph (2020) and Jappelli and Pistaferri (2020), who use different empirical settings to this thesis.

(2014), Kaplan et al. (2018)). It features households whose behaviour is conditional on heterogeneous income, liquid asset and illiquid asset states, and so generates a distribution of MPCs that varies over these states (Figure 4.1). In contrast, empirical models tend to separate the distribution into three stylised groups – poor 'hand-to-mouth' (HtM), wealthy HtM and non-HtM households – based on their holdings of liquid and illiquid wealth. Figure 4.2 shows how such a grouping may capture different aspects of the MPC distribution. The shaded regions are categorised as HtM households (specifically wealthy HtM households since the average illiquid asset position is positive), but it is not obvious that empirical estimates of this group's MPC are informative of the predicted MPC distribution. This chapter introduces a method to estimate the fulls distribution in an empirical setting.

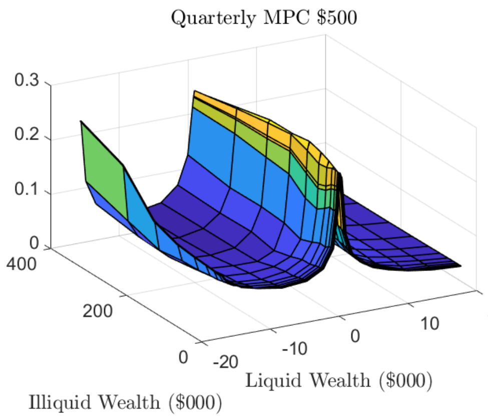

Figure 4.1: Kaplan et al. (2018) MPC distribution across liquid and illiquid wealth generated using the replication code; average over income states.

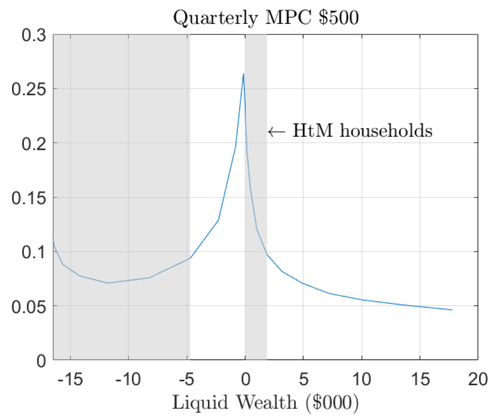

Figure 4.2: Kaplan et al. (2018) MPC distribution across liquid wealth generated using the replication code; average over illiquid asset and income states; shaded regions are HtM households.

Predictions from the literature are compared against a flexible specification of the empirical model that allows for nonlinearities in the distribution of MPCs along a single covariate. Much of the estimated variation in MPCs is not statistically different to a simple linear model and few of the relationships are robust in a setting with multiple covariates. Life-cycle dynamics are an exception; age has a robust negative relationship with MPCs out of permanent income for labour-force households. Business income and the average propensity to consume (APC) also appear to have robust correlations with MPCs. Some suggestive evidence is found in support of the predictions of the two-asset model, but liquid and illiquid wealth do not have robust correlations with MPCs in a model with multiple covariates.

The largest source of variation in transitory and permanent MPCs comes from demographic variables, a result that previous studies tend to overlook. Households with black or female heads have much larger permanent MPCs, which means they smooth consumption less in response to permanent income shocks. The result is robust to controlling for household balance sheet features and so is unlikely to be driven by financial constraints. Education also plays a clear role, with higher levels of education corresponding to lower transitory and permanent MPCs. These demographic variables drive the majority of observable MPC heterogeneity.

I extend the model with multiple covariates to capture *unobserved* heterogeneity in MPCs. Lewis et al. (2020) also estimate unobserved MPC heterogeneity, using the economic stimulus payments of 2008.[37] They cluster predicted MPCs using a fuzzy C-means algorithm with no information from covariates; heterogeneity arises from statistical properties and is mapped to covariates *ex post*. My approach is similar to the well-established 'random coefficients' model (see Griffiths, Drynan, and Prakash (1979), Zeger and Karim (1991), Hsiao and Pesaran (2004)), and formally models observed and unobserved heterogeneity jointly. Unobserved variation appears to be an important source of heterogeneity for the response to transitory income shocks, whereas observed variation explains much of the heterogeneity for permanent shocks.

Heterogeneity in MPCs can have important aggregate consequences (Jappelli and Pistaferri (2014), Kaplan and Violante (2014) , Carroll, Slacalek, Tokuoka, and White (2017), Auclert (2019)). To assess this, I generate policy-relevant MPCs that are representative of the population. This involves i) extending the sample to include retired households, ii) formally modelling heterogeneity, iii) population weighting the sample, and iv) conversion of the microdata elasticities into

[37] Arellano et al. (2017) and Jappelli and Pistaferri (2020) also *control* for unobserved heterogeneity through household fixed effects, but do not estimate how large this is.

dollar-for-dollar MPCs (\$MPCs). I match the sample moments of the US Survey of Consumer Finances (SCF) and utilise the Bayesian Weighted Estimation procedure in Gunawan, Panagiotelis, Griffiths, and Chotikapanich (2020). Population weighting tightens the distributions of elasticities and shifts them towards zero. Conversion of elasticities to \$MPCs further shifts the distributions lower and has a larger impact on central moment statistics than modelling heterogeneity or population weighting. Population average \$MPCs imply that households increase permanent nondurables consumption by 7 cents per dollar for transitory income shocks, and 38 cents for permanent shocks.

4.2 Empirical Model & Data

The empirical framework is detailed in Section 2.1, with the benchmark model repeated here for convenience. Unexplained income y and unexplained consumption c, are given by

$$y_{i,t} = \mu^y_{i,t} + \nu^y_{i,t} + \theta\nu^y_{i,t-1} \tag{2.3}$$

$$\mu^y_{i,t} = \mu^y_{i,t-1} + \zeta^y_{i,t} \tag{2.4}$$

$$c_{i,t} = \gamma_\mu \mu^y_{i,t} + \mu^c_{i,t} + \nu^c_{i,t} \tag{2.5}$$

$$\mu^c_{i,t} = \mu^c_{i,t-1} + \gamma_\nu \nu^y_{i,t} + \zeta^c_{i,t}. \tag{2.6}$$

The coefficient γ_ν is the MPC out of transitory income shocks, and γ_μ the MPC out of permanent income shocks (with $1 - \gamma_\mu$ the household's partial insurance against permanent income shocks). As mentioned in Section 2.1, these coefficients regulate the consumption response to income shocks and are (approximately) income elasticities of consumption. In this chapter, the MPC concept commonly used in policy making – one that measures dollar-for-dollar responses – is derived in Section 4.6.

4.2.1 Modelling Observed Heterogeneity

A simple extension to capture heterogeneity allows households' consumption response to income shocks to vary with observable covariates through interaction terms. Observation-specific consumption responses are captured by replacing γ_μ and γ_ν with $X_{i,t}\Gamma_\mu$ and $X_{i,t}\Gamma_\nu$; where $X_{i,t}$ is a $1 \times (K+1)$ vector of K household covariates and an intercept term, while Γ_μ and Γ_ν are $(K+1) \times 1$ vectors of parameters that capture how household MPCs change with the covariates. This allows MPCs to vary across households and time, but restricts all households to respond to covariates in the same way. Using interaction parameters alone may capture undesired *direct* correlation between consumption and covariates that is unrelated to income. To control for this, the covariates are included in levels as well, with the addition of the term $X_{i,t}\Psi$. Equations (2.5) and (2.6) become

$$c_{i,t} = X_{i,t}\Gamma_\mu \mu^y_{i,t} + X_{i,t}\Psi + \mu^c_{i,t} + \nu^c_{i,t} \tag{4.1}$$

$$\mu^c_{i,t} = \mu^c_{i,t-1} + X_{i,t}\Gamma_\nu \nu^y_{i,t} + \zeta^c_{i,t} \tag{4.2}$$

The main object of interest is the *distribution* of estimated MPCs captured by the vectors $X_{i,t}\hat{\Gamma}_\mu$ and $X_{i,t}\hat{\Gamma}_\nu$; I focus on the distribution using posterior means of parameters. The priors for Γ_μ and Γ_ν are taken as $(K+1) \times 1$ vectors of independent normal distributions. The priors on the intercept terms use the hyperparameters from the homogeneous model (a mean of 0.35 and variance of 0.25), but the remaining elements are mean zero and variance four. Estimation details are provided in Appendix C.1.

The typical empirical approach in the literature separates households into groups based on covariates being in a certain range of values, such as quantiles. Although sometimes motivated by theory, the precise thresholds for subgrouping households along a continuous covariate are always arbitary in part. The approach implies that threshold effects or discrete nonlinearities are present in the distribution of MPCs. In addition, subgrouping can result in loss of observations and poorer inference. This also limits the ability to capture continuous nonlinearities, as each subgroup requires a sample size large enough to identify the full parameter set. In contrast, the approach applied here allows the vector $X_{i,t}$ to contain the level of the covariate itself, along with higher dimension power terms to give a polynomial approximation of nonlinearities.

Fitting polynomials to the distribution of MPCs is tractable and informative, but can be unreliable due to features of the data such as ill-conditioning and outliers. As such, the baseline interaction model uses b-splines rather than raw polynomials. In order to flexibly model the full distribution of MPCs along covariates, nodes at the 20th, 50th, 70th and 90th percentiles are included.[38] A cubic spline is used for covariates of continuous values, and a basic linear specification is used for binary variables, with categorical variables separated into binary terms.[39] A stochastic MPC model that accounts for *unobserved* heterogeneity is detailed in Section 4.5.

4.2.2 Data

Details of the PSID dataset were provided in Section 2.3. However, this chapter uses a preliminary version of the main dataset, which covers survey years 1999-2017 and uses the NBER's TAXSIM9 adjustment as implemented in Kaplan et al. (2014). The samples used in much of the previous literature are restricted to prime working-age households, typically where the household head is under the age of 65. I extend coverage with an additional sample of retired households. Both samples are augmented by sample weights taken from the 1998-2016 waves of the US Survey of Consumer Finances (SCF). Equivalent SCF versions of PSID covariates are created and used to generate population shares to reweight the PSID samples. The SCF provides some advantages over the PSID for population weighting (see Section 4.6.2).

4.2.3 Including Retired Households

The previous literature focusses on prime working-age households, which only account for around half of the total US population of households. Estimates of aggregate effects that do not capture the response of these omitted households are potentially misleading. Modelling older households is complicated by the changes in income and consumption that accompany retirement decisions (e.g. Browning

[38]Additional nodes at the 0.5th and 99th percentiles are used to eliminate the effect of outliers and are omitted from the graphical results. Surplus nodes are discarded for variables with a large mass at a particular value (usually zero).

[39]Appendix C.5 presents results using linear splines – the major difference is tighter highest posterior density (HPD) intervals given fewer parameters are estimated.

and Crossley (2001) discuss this with respect to the life-cycle model). Roughly half of the omitted households have a head or spouse who is retired, so retirement decisions could have a substantial bearing in the data.[40] For simplicity, I do not model the transition to retirement but instead consider retirees separately to households participating in the labour force. This allows me to capture much more of the population without misclassifying transitional retirement decisions as income and consumption shocks.

Labour-force households are classified as those with heads aged 25-64 and not retired at any period.[41] Retired households are only those where both head and spouse are retired. This extends coverage to around three-quarters of all households.[42] Observations where the head retires in the subsequent period are ommitted from the labour-force sample, and retired households are only included from one period after the last member retires. The income definition for retired households includes asset income.

The labour-force sample contains around 20,000 observations and 3,500 households, wherease the retired sample is around one-tenth as large. The small retired sample imposes limitations on inference. It thus provides less insight into the drivers of MPC heterogeneity, but still plays an important role in estimating aggregate effects.

[40] Hurst (2008) finds that work-related expenses and food expenditure fall upon retirement for the average household. In particular, households shift to home production of food so that their food expenditure falls while their food consumption is stable. Food is an important component of the nondurables consumption measure used in this thesis, but the empirical framework does not readily allow for such expenditure shifts.

[41] This includes a small share of households where the spouse is retired. The distribution of income of these households is slightly higher than those with both partners working, suggesting that the spouse's retirement decision is not of first order importance to household income.

[42] About half of the remaining households are those where the head is retired but the spouse is not. These households exhibit a distribution of income that is markedly lower than fully retired households, suggesting that continued labour-force participation might be driven by need rather than coincidence. This may imply some sample selection bias, but modelling of joint labour supply and retirement decisions is beyond the scope of this thesis.

4.2.4 Household Covariates

I limit the study to key covariates identified by the theoretical and empirical literature, supplemented with some salient demographic variables. The balance sheet covariates used are continuous and are not divided into categorical subgroups or centiles in the primary analysis. Covariates are lagged one period to avoid potential endogeneity between the covariates and income shocks.[43] This reduces the available sample by around 6,000 observations. In most cases, the results from using contemporaneous covariates are similar, with the exception of the APC, which is clearly endogenous. Following Kaplan et al. (2014), assets and debt obligations are separated into liquid and illiquid components.[44]

Additional balance sheet variables include the debt-to-income ratio, home equity, mortgage-to-income ratio, business and farm income and the average propensity to consume (APC).[45] Business and farm income reflects the asset component of family-owned businesses and farms. The APC is constructed as nondurables consumption divided by income.[46] Demographic variables include household head sex, age, life expectancy, education and race. Education is the highest number of years of completed schooling, categorised into less than high school, high school and/or nonacademic training, three or fewer years of college, and four or more years of college. Race is categorised into white, black, and other people of colour.

[43] In particular, wealth reflects the contemporaneous outcome of income and consumption decisions. If such variables are fairly persistent across time, which appears to be the case, then the results will primarily reflect cross-sectional heterogeneity rather than changes within the household. If there were particular covariates for which the timing was crucial, an IV approach could be applied.

[44] Liquid wealth comprise checking account balances, stocks and credit card debt. Illiquid wealth comprise home equity, individual retirement accounts, net value of other real estate, and all other (non-mortgage and non-credit card) debt. An important difference to the Kaplan et al. (2014) data is the treatment of credit card debt, as detailed in Appendix A.3.

[45] Total debt and mortgage remaining were also trialled, but the results were very similar to their income ratio counterparts.

[46] In constructing the APC, labour income is used for labour-force households and total income for retired households. All other income ratios are constructed in a similar manner.

	Labour-force		Retired	
Variable	Mean	Std. Dev.	Mean	Std. Dev.
Income ($000)*	74.24	45.66	36.01	21.37
Unexplained income	0.00	0.50	-0.04	0.47
Nondurables consumption ($000)	36.94	21.09	24.95	13.93
Unexplained consumption	-0.01	0.43	-0.02	0.45
Liquid assets ($000)	25.89	80.56	98.47	160.62
Illiquid assets ($000)	107.78	201.14	238.11	267.01
Home equity ($000)	66.94	112.65	138.69	146.09
Mortgage-to-income ratio	1.04	1.35	0.43	1.29
Debt-to-income ratio	1.24	1.46	0.51	1.48
Business and farm income ($000)	1.62	9.63	1.97	10.30
APC	0.56	0.33	0.78	0.48
Age of head	41.63	9.84	75.17	7.44
Life expectancy (years)	37.00	8.89	11.33	4.41
Female	0.15	0.35	0.31	0.46
Education: Less than high school	0.10	0.29	0.20	0.40
Education: High school	0.24	0.43	0.34	0.47
Education: 3 or fewer years college	0.28	0.45	0.23	0.42
Education: 4 or more years college	0.39	0.49	0.23	0.42
Black	0.08	0.27	0.06	0.23
Other person of colour	0.08	0.26	0.02	0.15

* Labour income for labour-force households, total income for retired households.

Table 4.1: Summary statistics for labour-force and retired samples.

4.3 Benchmark Estimates

The labour-force and retired households represent the most fundamental source of heterogeneity. These two samples are estimated independently – allowing for each type of household to be affected by shocks with different variances. Benchmark results using a simple intercept only model are presented in Tables 4.2 and 4.3 to anchor the analysis that follows. The posterior distributions for the MPC coefficients show that retired households have a smaller response to transitory and permanent income shocks; however, the differences are not statistically significant. Overall, both groups respond to income shocks in a similar manner.

This suggests that bias may not be large in estimates of aggregate MPCs using samples that do not include retired households.

The key difference between these samples is found in the income and consumption shocks they face. Retired households face larger transitory shocks to unexplained consumption. This result is primarily driven by the first stage decomposition of income and consumption into explained and unexplained components (the latter component being the focus of this thesis). The raw measures of income and nondurables consumption show that retired households have much *smaller* variation in income and nondurables consumption than labour-force households (Table 4.1). However, the predictable variation is also much smaller for retired households. In net terms, variation in unexplained consumption is a little larger for retired households. The predictable variation captures common time-varying components, such as business cycle effects, and so the result may reflect retirees being less sensitive to these shocks. In addition, retired households may face different idiosyncratic consumption shocks, such as out-of-pocket medical expenses as regularly used in life-cycle modelling (e.g. Hubbard, Skinner, and Zeldes (1995)). (This also questions the interpretation of transitory consumption shocks as measurement error.)

	Mean	5%	95%
γ_ν	0.136	0.116	0.156
γ_μ	0.550	0.514	0.594
θ	0.086	0.065	0.105
$\sigma_{\nu y}$	0.249	0.246	0.252
$\sigma_{\zeta y}$	0.146	0.142	0.151
$\sigma_{\nu c}$	0.230	0.227	0.233
$\sigma_{\zeta c}$	0.106	0.101	0.111
Obs.	20286		
Households	3479		

Table 4.2: Simple intercept only model posterior statistics; labour-force households.

	Mean	5%	95%
γ_ν	0.107	0.024	0.190
γ_μ	0.458	0.303	0.644
θ	0.079	-0.040	0.190
$\sigma_{\nu y}$	0.243	0.221	0.264
$\sigma_{\zeta y}$	0.161	0.133	0.188
$\sigma_{\nu c}$	0.250	0.240	0.261
$\sigma_{\zeta c}$	0.114	0.098	0.132
Obs.	1995		
Households	404		

Table 4.3: Simple intercept only model posterior statistics; retired households.

4.4 Observed Heterogeneity

The deterministic interactions model is estimated using observable covariates individually to assess the distribution of MPCs in relation to theoretical predictions and previous empirical findings. Evidence for the specific financial frictions that feature in a two-asset model is also analysed before a multiple covariate specification is used to determine which relationships are robust to controlling for a range of household characteristics. A number of themes are drawn out of the literature to anchor the analysis: wealth and liquidity; life-cycle patterns; debt and housing; business income and the APC; and demographic features. Table 4.4 summarises the literature and corresponding results; however, variation in MPCs due to demographic factors is found to be substantially larger than for other covariates. The sample of retired households is not large enough to provide adequate inference when using the flexible mapping of MPCs to covariates, so the analysis focusses on the labour-force sample.

4.4.1 Single Covariate Results

Figures 4.3 and 4.4 show how consumption responses to transitory and permanent shocks vary with covariates. The black markers are the predicted response using the posterior means of parameter estimates from a cubic spline model, while the red lines are predictions using the 5th and 95th percentiles of the posteriors to provide a measure of dispersion. These are contrasted with the predicted response from a simple linear model (grey dashed line). Many of the covariates have empirical distributions with long tails (particularly on the right side), so the graphs are truncated at the 5th and 95th percentile. Overall, responses to both shocks show some variation with covariates, but the predictions are sufficiently disperse to prevent them from being statistically different to a linear model. Responses to permanent shocks are better inferred by the data (they have tighter posterior distributions) and more closely align with the existing literature than responses to transitory shocks.

	Expected Relationship		Results	
Motivation & covariate	Transitory	Permanent	Transitory	Permanent
Wealth				
Net liquid wealth	$-$	$\sim$	$\sim$	$-$
Net illiquid wealth	$\sim$	$-$	$\sim$	$-$
Life-cycle model				
Age (labour force)	$-$	$-$	$\sim$	$-^{*}$
Debt & housing				
Debt-to-income ratio	$+$	$+$	$\sim$	$+$
Mortgage-to-income ratio	$+$	$+$	$\sim$	$+$
Home equity	$-$	$-$	$\sim$	$-$
Negative home equity	$+$	$+$	$+$	$+$
Lewis et al. (2020)				
Business income	$+$	$\sim$	$-^{*}$	$-^{*}$
APC	$+$	$\sim$	$-^{*}$	$+$
Two-asset model frictions				
Credit limit				
Large liquid debt	$+$	$\sim$	$+$	$\sim$
Budget constraint kink				
Net liquid wealth @ 0	$\sim$	$+$	$\sim$	$+$

$+$ represents a positive relationship, $-$ negative relationship, and $\sim$ no clear relationship
* robust result in the multiple covariate model (90 per cent HPD does not contain zero)

Table 4.4: The expected relationship between MPCs and covariates informed by consumption theory and previous empirical studies, compared with the empirical results of this chapter.

Wealth and Liquidity

The permanent income life-cycle model with 'buffer-stock' households predicts that permanent MPCs are declining in wealth, as lower wealth households have less capacity to smooth income shocks (Deaton (1991), Carroll (1997)). The two-asset model makes a similar prediction, but primarily for illiquid wealth, while households use liquid wealth to buffer against transitory income shocks. Previous empirical studies tend to find a negative relationship between liquid wealth and transitory MPCs (Jappelli and Pistaferri (2014), Bunn et al. (2018), Fagereng et al. (2018), Jappelli and Pistaferri (2020)), but there has been little analysis of permanent income shocks.

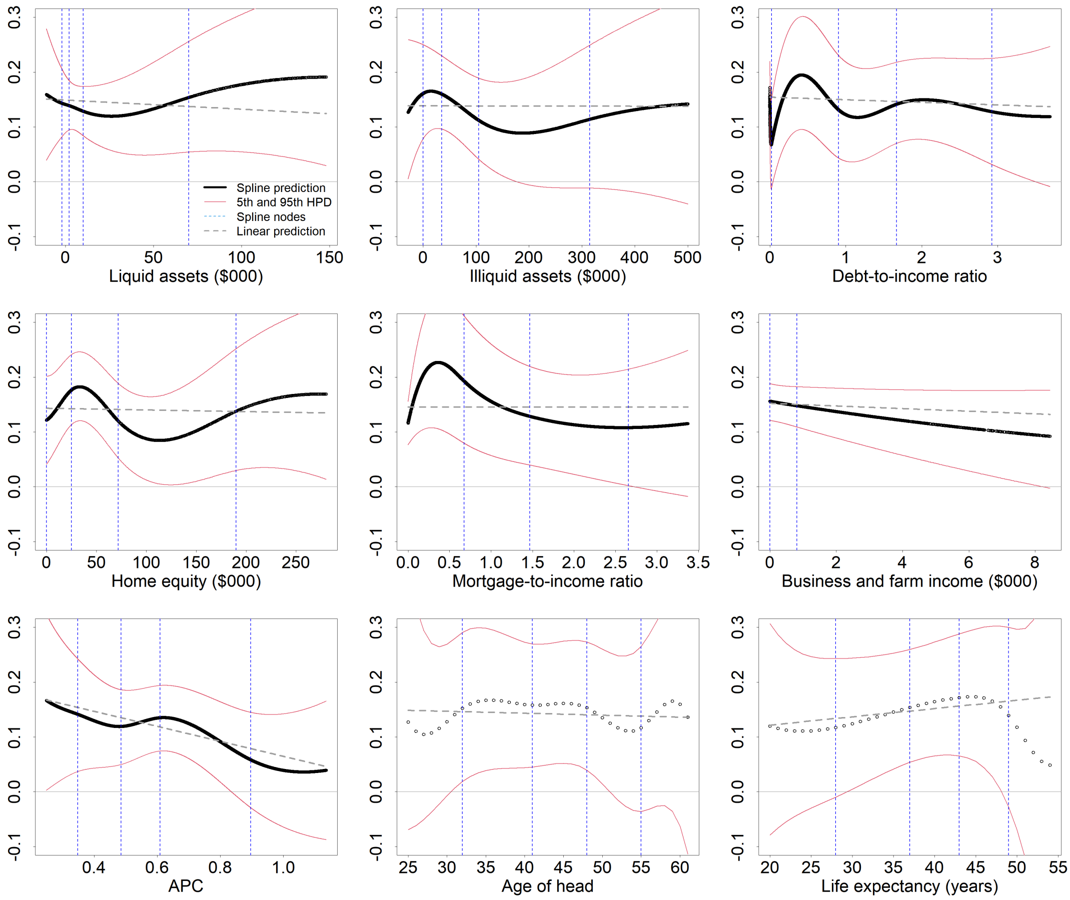

Figure 4.3: Predicted responses to transitory shocks using the posterior mean values of Γ_ν; labour-force households; covariates plotted over 5–95th percentiles, vertical dashed lines are spline nodes at the 20th, 50th, 70th and 90th percentiles.

The results show some tentative evidence that larger liquid asset positions correlate with smaller transitory MPCs for households between the 20th and 70th percentiles, consistent with the two-asset model (Figure 4.3). However, the variation is small and there is a reversal of the correlation for the top 10th percentile of the distribution. Very low (and negative) holdings of illiquid wealth are associated with larger permanent MPCs, but the relationship flattens out for the right tail of the distribution (Figure 4.4). This is broadly consistent with the two-asset model and suggests the marginal benefit (in terms of consumption smoothing) declines for large illiquid asset holdings. However, the results for liquid and illiquid wealth are similar overall, suggesting at least some fungibility between liquid and illiquid wealth.

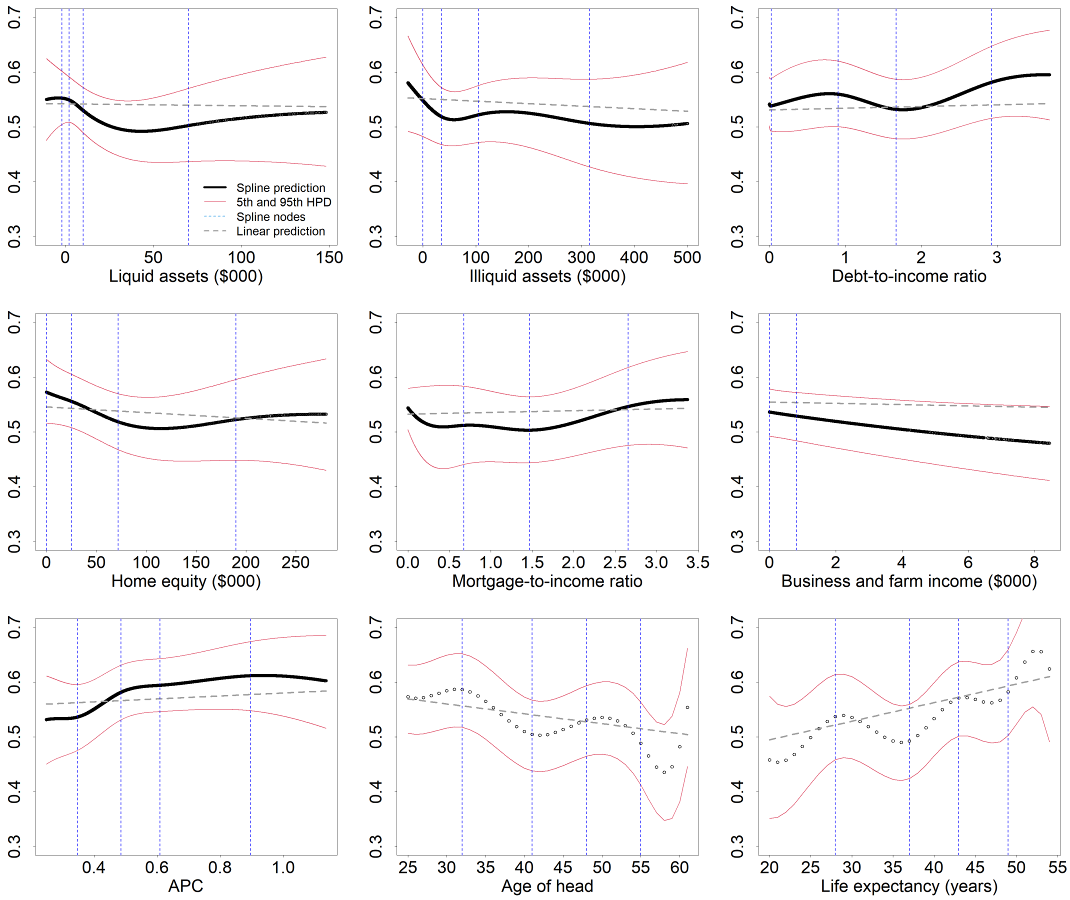

Figure 4.4: Predicted responses to permanent shocks using the posterior mean values of Γ_μ; labour-force households; covariates plotted over 5–95th percentiles, vertical dashed lines are spline nodes at the 20th, 50th, 70th and 90th percentiles.

Life-cycle Dynamics

Life-cycle models have long been used to explain observed patterns in income and consumption (e.g. Browning and Crossley (2001) provide a nice treatment). Life-cycle predictions of MPCs are influenced by credit constraints, the persistence of income shocks and the household's current wealth relative to their total expected future income. The basic life-cycle model predicts that households' permanent MPCs decline with age as permanent shocks have a smaller total effect on expected income closer to retirement and death, and larger buffers of relative wealth have been accumulated (Clarida (1991), Gourinchas and Parker (2002), Blundell et al. (2008)). Younger households are predicted to act as 'buffer-stock' consumers in the presence of credit constraints, with small relative wealth and large responses to transitory income shocks, whereas the accumulated savings of middle-aged households allows them to smooth transitory shocks (Carroll (1997), Gourinchas and Parker (2002)).

The results show that ageing is associated with moderate declines in permanent MPCs, consistent with life-cycle dynamics (Figure 4.4). Transitory MPCs do not appear to vary systematically with age, but they do show a small positive correlation with expected remaining life over the majority of the distribution (Figure 4.3). This latter result is consistent with the life-cycle model, although the predictions are quite disperse relative to other covariates. The life expectancy data are matched on age, sex and year, so the difference between the results for age and life expectancy could be driven by the small increase to life expectancy (1-2 years) over the sample, or the different behaviour between male and female headed households.

Business Income and the APC

This chapter is close in objective to Lewis et al. (2020), who analyse the economic stimulus payments of 2008. They find that the APC and non-labour income have statistically significant positive relationships with transitory MPCs. My approach finds negative relationships for the APC and business income (Figure 4.3).[47] A negative correlation for business income is intuitive, as it is consistent with larger financial resources allowing households to smooth consumption. However, only around 15 per cent of observations report business income, so any relationship only holds for a small subset of households.

The APC result is one of the most robust results for transitory MPCs in this chapter.[48] The negative relationship contrasts Aguiar, Bils, and Boar (2020), who show that preference heterogeneity generates a positive relationship between

[47] I use business and farm income rather than non-labour income as the result of an inherited coding error.

[48] A caveat is that the empirical definition of the APC in this thesis uses nondurables consumption, whereas total consumption might be prefered.

MPCs and the APC. However, the result is consistent with Miranda-Pinto, Murphy, Walsh, and Young (2020), who find that lagged high expenditure is associated with low MPCs (recall that the covariates are also lagged in this chapter). They propose a heterogeneous agent model with idiosyncratic consumption thresholds to explain this feature – households with high thresholds are 'saving-constrained' and less responsive to income shocks. Their model further accounts for consumption volatility being as large as income volatility, which is seen in Tables 4.2–4.3 but is not predicted by standard Euler equation consumption dynamics.

Debt and Housing

Several empirical studies find that debt is associated with an increased sensitivity to shocks that affect household finances (Price, Beckers, and La Cava (2019), Cho et al. (2021), Baker (2018), Nakajima (2018)). Although debt obligations are captured in the net measures of liquid and illiquid wealth, they may also affect household decisions independently. In particular, Price et al. (2019) argue that the composition of assets and liabilities on households' balance sheets matters for their spending, rather than only net measures of wealth.[49] The largest asset and liability on household balance sheets is housing; home equity and mortgages tend to drive illiquid wealth and total household debt.

The results suggest that households with large debt or mortgage burdens have somewhat larger permanent MPCs, which is broadly consistent with previous findings (Figure 4.4). There is no obvious relationship for transitory MPCs. The influence of debt burdens on transitory MPCs is often linked to liquidity or credit constraints (such as in Boar, Gorea, and Midrigan (2017)), but the same may be true of permanent income shocks. It would be expected that large debt burdens are eventually paid off by households, but the long time frame of mortgages (20-30 years) may be sufficiently persistent to affect households' ability to self-insure

[49] Around 30 per cent of PSID household observations over 2011–17 hold *both* positive checking account and negative credit card balances. At least half of these households have enough funds in their checking account to fully pay off their credit card debt. This 'credit card puzzle' suggests some households treat their assets and liabilities separately (e.g. Telyukova (2013)).

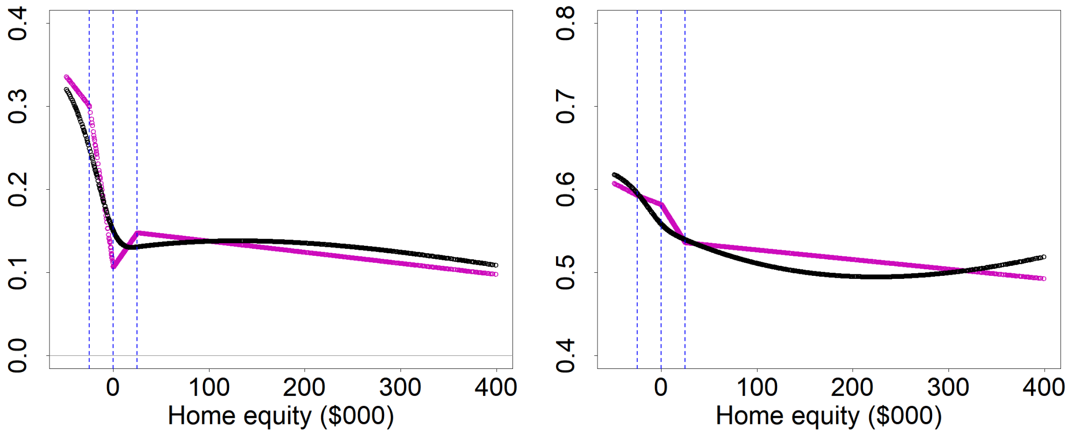

Figure 4.5: Predicted transitory (left panel) and permanent (right panel) responses over home equity using the posterior mean values of Γ_ν and Γ_μ; labour-force households; vertical dashed lines are spline nodes.

against long-lasting shocks. The results for home equity are similar to those of illiquid wealth for permanent shocks, as would be expected given the high correlation of 0.78.

Disney et al. (2010) and Mian et al. (2013) find that low or negative home equity is associated with larger MPCs out of *wealth* shocks, which might also affect MPCs out of income shocks. A specification that targets the region around zero is used, with nodes at $-\$25,000$, $\$0$ and $\$25,000$ (and the 0.5th and 99th percentiles to eliminate outliers). Households with negative home equity have larger responses to both types of income shocks, consistent with previous findings for wealth shocks (Figure 4.5). Furthermore, the estimated variation in transitory MPCs for negative home equity is the largest out of all covariates, suggesting a strong effect on household consumption decisions.

4.4.2 Constraints and Kinks

A feature of the two-asset model is that large transitory MPCs occur at two locations in the distribution of liquid wealth; near the credit constraint and near zero liquid wealth.[50] Credit constraints are a bedrock of standard precautionary savings models – households near the constraint are less able to smooth consumption and tend to consume a greater share of changes in their income (e.g. Deaton (1991), Aiyagari (1994), Carroll (2001)). Behaviour near zero liquid wealth is different; rather than facing a hard constraint, households are at a kink on their budget constraint driven by the wedge between the interest rates on borrowing and saving in the liquid asset. These households face a trade-off between the costs and benefits of i) smoothing consumption, ii) saving and borrowing in the liquid asset, and iii) saving and transaction costs in the illiquid asset.

The predicted relationships between MPCs out of income shocks and liquid and illiquid wealth from the Kaplan et al. (2018) model are shown as blue lines with markers in Figure 4.6.[51] The presence of larger MPCs at the credit limit and budget kink are muted compared to the definition presented in Kaplan et al. (2018) (Figure 4.1), which uses MPCs out of liquid wealth shocks (see Appendix C.2 for details).[52] This suggests that the empirical strategy used here, based on income shocks, may be less well suited to capturing these specific frictions.

Nonetheless, to assess the credit limit and budget constraint kink in the empirical model, linear and cubic spline specifications are used with nodes at $-\$5,000$, $\$0$ and $\$5,000$ for liquid wealth, and $-\$25,000$, $\$0$ and $\$25,000$ for illiquid wealth.[53] Including nodes at and either side of zero helps to capture a shift at the budget kink, but still imposes continuity (and differentiability for powers greater than one) at the node locations, which limits the ability to identify a discontinuity.

Figure 4.6 shows the responses to transitory (top row) and permanent (bottom row) shocks over liquid wealth and illiquid wealth. The broad model prediction

[50] These features can exist in a model with a single liquid asset, but the second illiquid asset – that incurs transaction costs – allows for a sizeable share of households to be affected by them.

[51] Note that the two-asset model predictions use a positive \$500 shock, whereas the average positive shock in the empirical estimates is roughly 15 per cent of annual income for transitory shocks and 5 per cent for permanent. Comparison of MPC variation across the distribution of wealth is the focus, not the overall magnitude of MPCs.

[52] However, consistent with Chapter 3, larger MPCs at the credit limit are much clearer for negative shocks.

[53] Nodes at the 0.5th and 99th percentiles are retained to eliminate the effect of outliers.

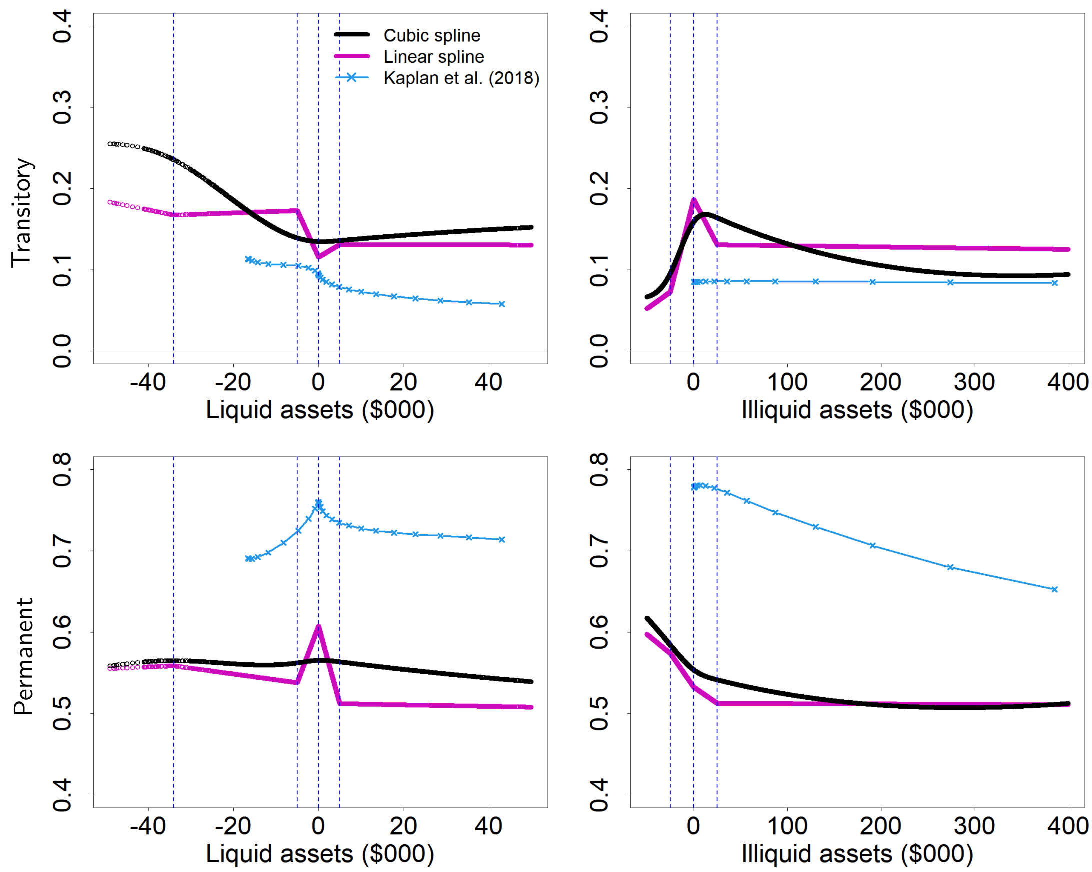

Figure 4.6: Predicted transitory (top row) and permanent (bottom row) responses using the posterior mean values of Γ_ν and Γ_μ; blue lines with markers are predictions from the Kaplan et al. (2018) model (see Appendix C.2); labour-force households; vertical dashed lines are spline nodes.

that liquid wealth is used to smooth transitory shocks and illiquid wealth for permanent shocks features in the results. However, the evidence for the specific credit constraint and budget kink frictions is mixed. The linear spline model shows a spike at zero liquid wealth, consistent with the budget kink, whereas the cubic spline model does not. There is some evidence for households with large liquid debt to have larger responses to transitory shocks, consistent with a credit limit, but again the results differ across specifications.

A specification that uses both liquid and illiquid wealth as additive cubic splines provides a joint distribution, showing similar patterns (Figure 4.7). Overall, the distributions are quite similar to those of the Kaplan et al. (2018) model. The main difference is the hump shape in transitory MPCs along illiquid wealth (also seen in the top right panel of Figure 4.6), with MPCs declining over the negative region. The measure of net illiquid wealth includes home equity and various other assets, along with student loans, medical bills, legal bills and family loans. Disaggregated data on debt are available for the 2011–17 survey years, which show that student loans dominate the sample who have negative illiquid wealth.[54] This is an interesting feature, but outside the scope of the two-asset model and this thesis.

4.4.3 Demographics

Many demographic variables are naturally categorical, so sets of binary dummies are used to capture how these variables affect consumption responses. Table 4.5 presents the results of a multiple covariate estimation using race, sex and education. The sex and education of the household head have a strong relationship with permanent MPCs. Female heads have permanent MPCs over 50 per cent larger than their male counterparts. Households with heads who did not complete high school have much larger permanent and transitory MPCs than those who completed college. Black households have very large consumption responses to permanent shocks – effectively no consumption insurance on average – and show no response to transitory shocks. Overall, the variation in MPCs related to these demographic features is much larger than that of balance sheet covariates.

There is a substantial literature documenting the income, wealth and consumption inequality along race, gender and education dimensions.[55] Furthermore, there are a range of social, economic and financial factors that may drive these patterns, not least of which is the historic and continued marginalisation of fe-

[54]For sample years 2011–17, 75 per cent of the value of other loans (non- mortgage and vehicle debt) are student loans and 73 per cent of households with negative illiquid wealth have student loans. In addition, transitory MPCs vary in opposite directions for illiquid wealth and home equity below zero, despite their high correlation, further suggesting a composition issue drives the results.

[55]See Oliver, Shapiro, and Shapiro (2006), Hardy, Morduch, Darity, and Hamilton (2018), Kuhn, Schularick, and Steins (2020) and Bhutta, Chang, Dettling, and Hsu (2020) on race; Ruel and Hauser (2013), Kassenboehmer and Sinning (2014) and Blau and Kahn (2017) on gender; and Gregorio and Lee (2002), Boshara, Emmons, and Noeth (2015), Lindley and Machin (2016) and Bartscher, Kuhn, and Schularick (2020) on education.

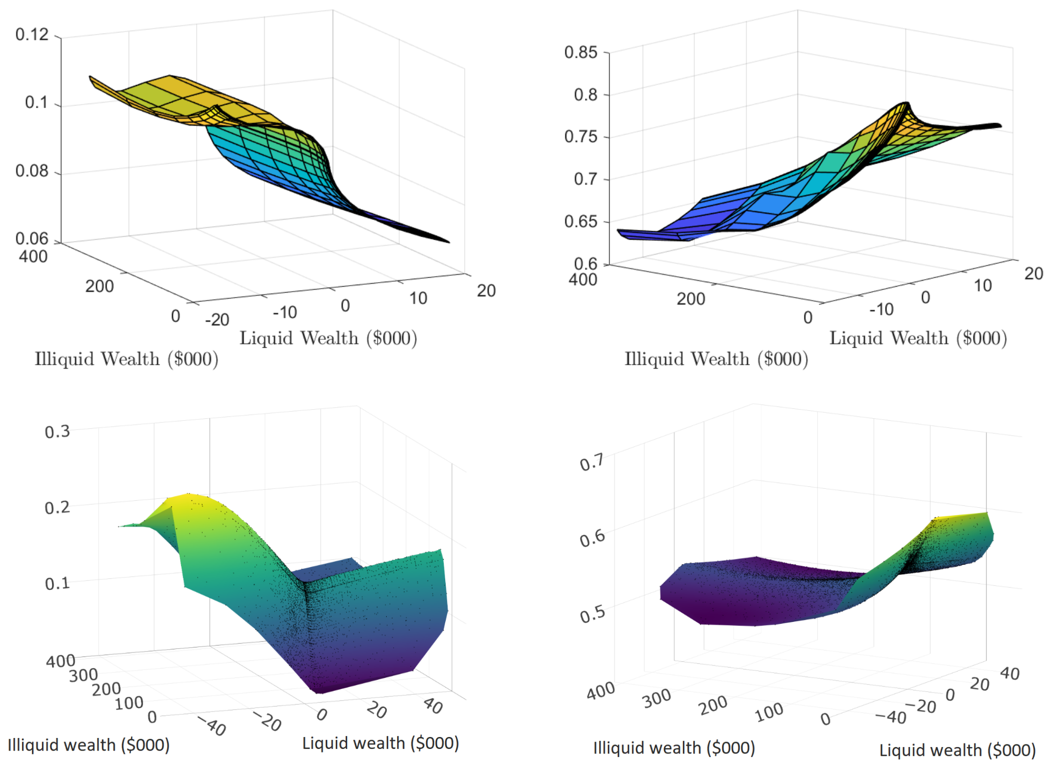

Figure 4.7: MPC distributions over liquid and illiquid wealth out of transitory income (left column) and persistent/permanent income (right column); Kaplan et al. (2018) predictions (top row) and empirical estimates (bottom row); empirical estimates use an additive cubic polynomial spline with nodes at -\$34, -\$5, \$0 and \$5 thousand for liquid wealth, and -\$25, \$0 and \$25 thousand for illiquid wealth, black dots are observation-specific predictions.

male and black workers in labour and financial markets (Darity and Mason (1998), Herring and Henderson (2016), Darity and Hamilton (2017)). In particular, financial literacy has been proposed as a mechanism to explain wealth inequality and correlates with race, gender and education in a manner consistent with the results here (Lusardi and Mitchell (2014), Lusardi et al. (2017)). However, there is comparatively little research on households' consumption smoothing specifically.[56] The findings are flagged in this chapter, with detailed discussion reserved for Chapter 5.

[56] Ganong, Jones, Noel, Greig, Farrell, and Wheat (2020) and De Giorgi, Gambetti, and Naguib (2020) are notable exceptions that address racial differences in consumption smoothing in the US. They both find that black households are substantially more responsive than white households to income shocks, but that differences in wealth explain almost all of the gap.

	Γ_ν	Γ_μ
Intercept	0.209***	0.658***
Black	−0.233***	0.444***
Other person of colour	0.036	−0.019
Female	0.035	0.299***
High school	−0.068*	−0.063
3 or fewer years college	−0.074*	−0.116
4 or more years college	−0.110***	−0.206**

Table 4.5: Demographic covariate interaction model; the intercept reflects a white male with less than high school education; labour-force households; ***, ** and * indicate that the 99, 95 or 90 per cent HPDs do not include zero.

4.4.4 Multiple Covariate Model

There is a substantial degree of correlation among some of the covariates, so a multiple covariate analysis is crucial to determine which relationships are robust. A simple linear model is used for parsimony and due to the lack of clear evidence of curvature in the single covariate models. The model includes all of the covariates examined above, except for the mortgage-to-income ratio (which has a correlation with debt-to-income of 0.9) and the other person of colour dummy (which was found to have little correlation with MPCs in the previous section). Figure 4.8 shows the posterior distributions of the MPC parameters for the continuous covariates. Each posterior draw has been multiplied by one standard deviation of the relevant covariate for comparison across drivers of heterogeneity. The graphs can be interpreted as showing how increasing a covariate by one standard deviation would shift the distribution of consumption responses.[57]

The liquid asset, illiquid asset and debt-to-income parameters have large mass on either side of zero, suggesting they are not statistically significant drivers of MPC heterogeneity.[58] For the continuous covariates that do have a statistically

[57] Sample standard deviations are reported in Table 4.1. All multiple covariate results use 5,000 posterior draws after 3,000 burn in.

[58] Simple linear, quadratic and cubic models with no splines and only liquid and illiquid wealth were also tested. The linear model did not result in statistically meaningful posteriors for the two covariates, and the quadratic and cubic models were mostly driven by variation in the

meaningful relationship, their effect on the overall elasticities are typically small. The relationship between the APC and response to transitory shocks is the main exception; an increase in the APC of one standard deviation (around 0.3) corresponds to a reduction in the MPC of around 0.05–0.08. Home equity, business income and age reduce the response to permanent shocks, but only by a small amount. Business income has a similar small effect on transitory responses.

Figure 4.9 shows the posterior distributions of the MPC parameters for binary variables. The distributions of the intercepts are much more disperse than in the model with no covariates, which is to be expected.[59] The demographic covariates retain the large and statistically significant effects found above – demographic factors appear to shift household MPCs by an order of magnitude more than the balance sheet covariates.

To gauge overall heterogeneity, the posterior means of the Γ_ν and Γ_μ parameter vectors are used to create a distribution of observation-specific predicted MPCs (that is, $X_{i,t}\hat{\Gamma}_\nu$ and $X_{i,t}\hat{\Gamma}_\mu$). A parsimonious multiple covariate specification is used, comprising a simple linear model of home equity, APC, business income and age, along with dummies for black and female heads, and education groups.[60] This results in distributions of MPCs with substantial amounts of heterogeneity (Figure 4.10), although comparable to that of Lewis et al. (2020). The distribution of responses to permanent shocks has a long right tail, indicating that a large share of households have little consumption insurance against permanent shocks, and little mass below 0.3, indicating a limit in households capacity or willingness to self insure. The long right tail is primarily driven by demographic factors. There is also little mass in the transitory distribution above 0.3, suggesting that all households are able to smooth these shocks to a large extent.

tails of the covariate distributions.

[59] The intercepts represent the response of a household with a 42 year old white male head who has less than high school education and a sample average balance sheet.

[60] Posterior distribution statistics for the covariate parameters are provided in Appendix C.6.

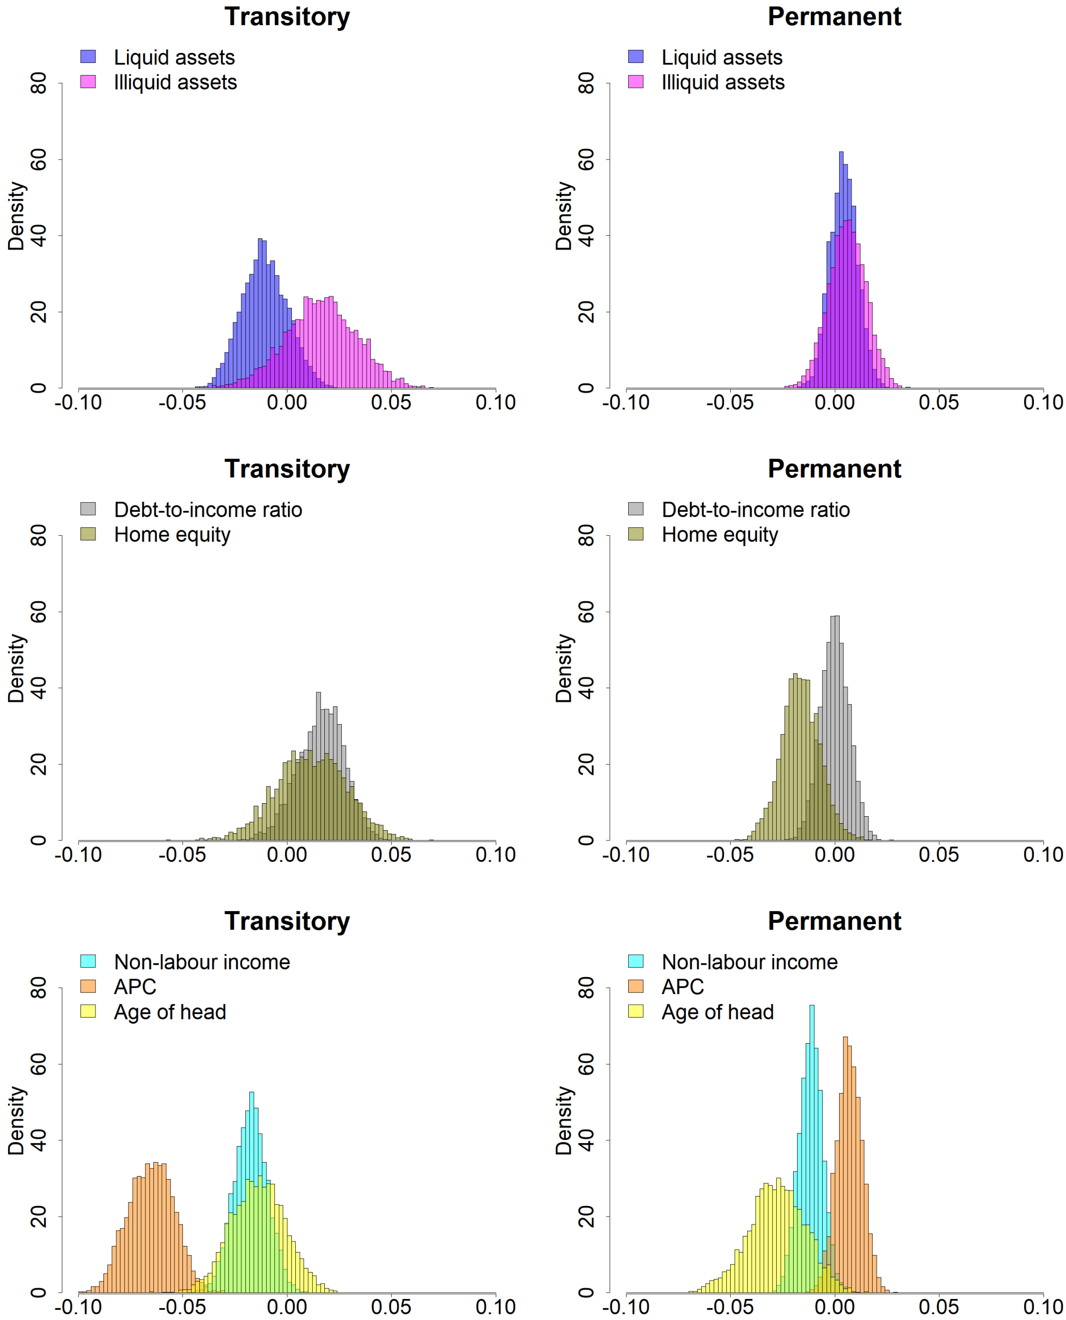

Figure 4.8: Posterior densities of Γ_ν and Γ_μ parameters for continuous variables; draws of each parameter are multiplied by one standard deviation of covariate; multiple covariate specification; labour-force households.

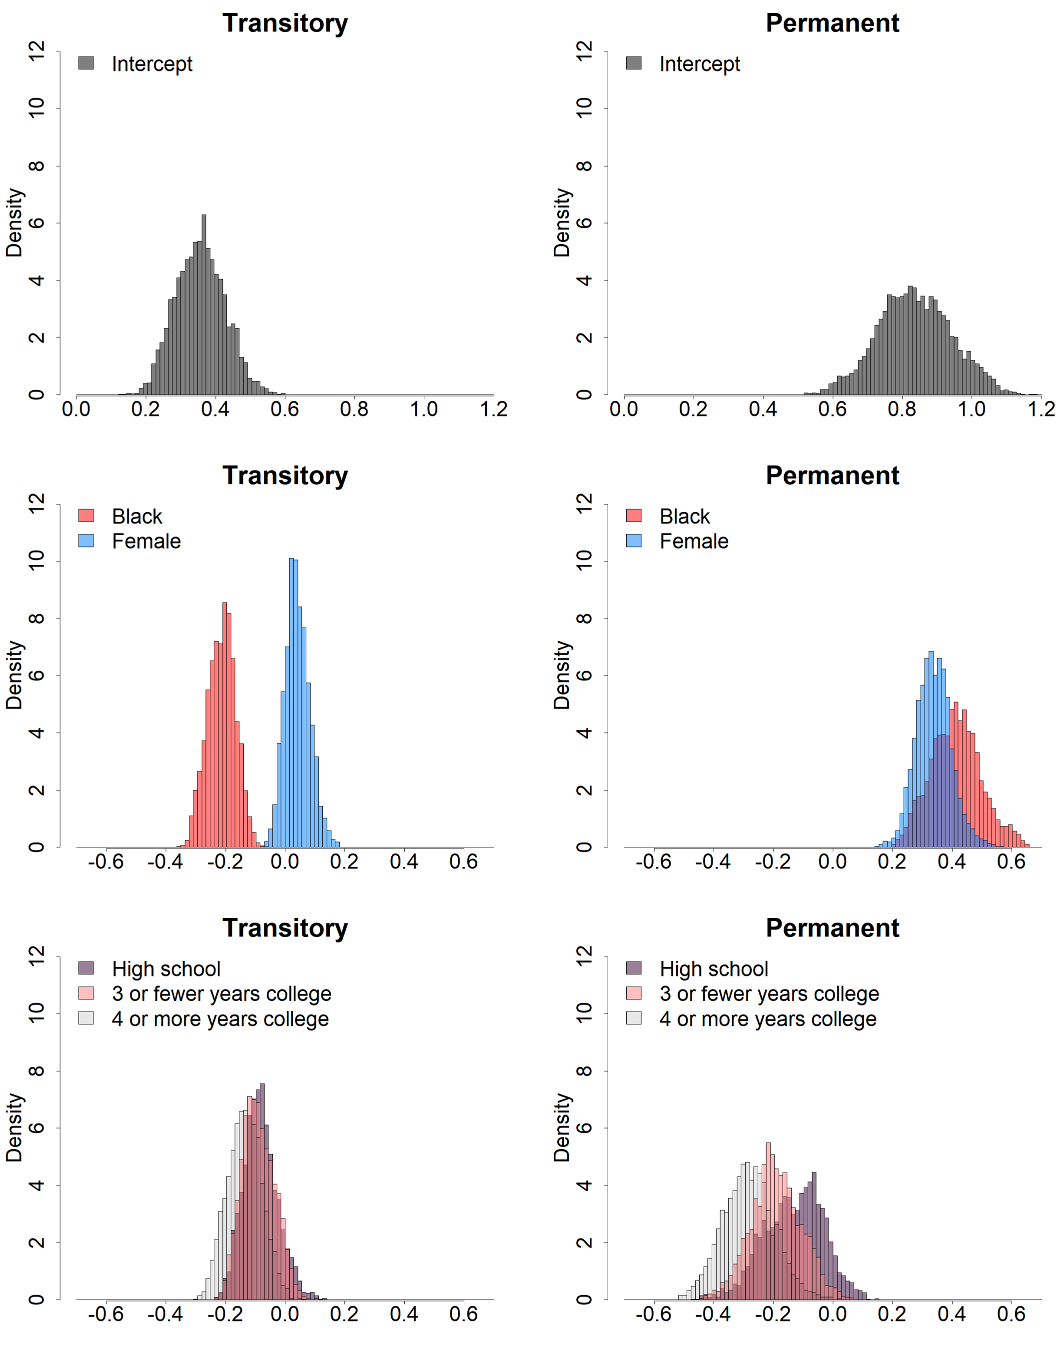

Figure 4.9: Posterior densities of Γ_ν and Γ_μ parameters for binary variables; multiple covariate specification; labour-force households.

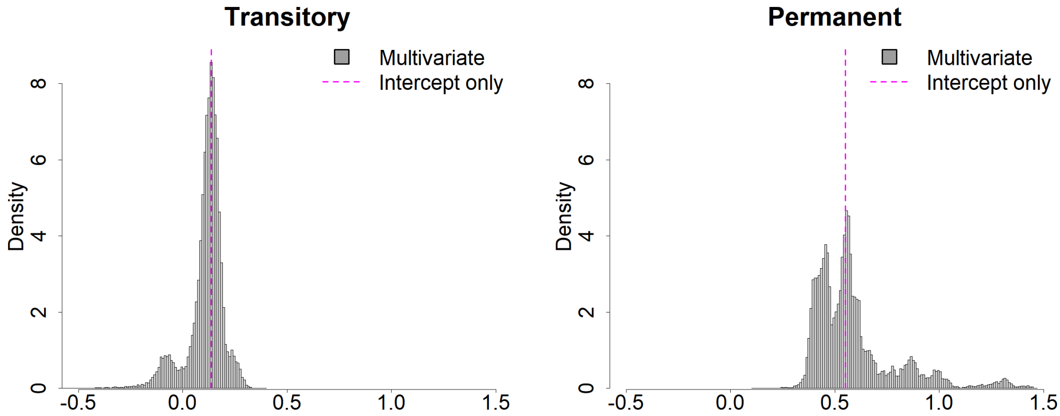

Figure 4.10: Distribution of predicted elasticities using posterior means of Γ_ν and Γ_μ; parsimonious multiple covariate specification; labour-force households; vertical dashed lines are the posterior means of the intercept only model with no covariates.

4.5 Unobserved Heterogeneity

The preceding specifications modelled MPCs as having a deterministic relationship with covariates; that is, covariates affect all households in the same way, so predicted MPCs vary deterministically with covariates. Stochastic models provide a more flexible framework that help to capture *unobserved* heterogeneity in MPCs, by allowing them to vary stochastically with covariates. For each observation, the consumption response is drawn from a distribution with a central moment given by covariate interactions and a variance informed by the whole sample. This is an application of the 'random coefficients' model (see Griffiths et al. (1979), Zeger and Karim (1991), Hsiao and Pesaran (2004)). The extension allows me to gauge how much household responses are driven by observable characteristics versus unobserved heterogeneity.[61] This is distinct from controlling for unobserved heterogeneity in consumption itself, as in Arellano et al. (2017) and Jappelli and Pistaferri (2020).[62]

[61] An alternative 'random intercept' approach does not use covariates and instead draws household-specific MPCs from a distribution with common mean and variance. See Appendix C.4 for details.

[62] In fact, the benchmark model in Chapter 2 effectively controls for this through estimating the period zero latent components $\mu^c_{i,0}$.

Equations (4.1) and (4.2) are replaced with

$$c_{i,t} = \gamma^{\mu}_{i,t}\mu^{y}_{i,t} + X_{i,t}\Psi + \mu^{c}_{i,t} + \nu^{c}_{i,t} \tag{4.3}$$

$$\mu^{c}_{i,t} = \mu^{c}_{i,t-1} + \gamma^{\nu}_{i,t}\nu^{y}_{i,t} + \zeta^{c}_{i,t} \tag{4.4}$$

$$\gamma^{\mu}_{i,t} = X_{i,t}\Gamma_{\mu} + \varepsilon^{\mu}_{i,t} \tag{4.5}$$

$$\gamma^{\nu}_{i,t} = X_{i,t}\Gamma_{\nu} + \varepsilon^{\nu}_{i,t} \tag{4.6}$$

$$\varepsilon^{\mu}_{i,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon\mu}) \tag{4.7}$$

$$\varepsilon^{\nu}_{i,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon\nu}) \tag{4.8}$$

where $X_{i,t}$, Γ_{μ} and Γ_{ν} are defined as previous and the innovations $\varepsilon^{\mu}_{i,t}$ and $\varepsilon^{\nu}_{i,t}$ are taken as independent for simplicity.[63] The same covariates are used as in the parsimonious multiple covariate specification from Section 4.4.4. Estimation details are provided in Appendix C.3. This model has large memory demands as it generates draws of $N \times T$ parameters for $\gamma^{\mu}_{i,t}$ and $\gamma^{\nu}_{i,t}$ at each step in the Markov chain.[64] To conserve memory, recursive posterior distributions of all $\gamma^{\mu}_{i,t}$ and $\gamma^{\nu}_{i,t}$ are generated using bin counting over a discretised grid, as well as recursive average values of $\gamma^{\mu}_{i,t}$ and $\gamma^{\nu}_{i,t}$ for each observation.

In the benchmark model and model with observed heterogeneity, identification of coefficients is via the conditional mean, while the unobserved components structure separates the shocks. In the model with unobserved heterogeneity, identification of the random coefficient parameters is obtained via the conditional variance, as is standard in simpler random coefficient models (e.g. Swamy and Mehta (1975)). That is, the coefficient shock $\varepsilon^{\mu}_{i,t}$ enters the error term of equation (4.3), but multiplication of $\varepsilon^{\mu}_{i,t}$ by $\mu^{y}_{i,t}$ introduces heteroscedasticity into the error term, allowing for separate identification of the variance parameters $\sigma^2_{\varepsilon\mu}$ and $\sigma^2_{\nu c}$.

The posterior distributions of the model parameters show reasonable inference (Tables 4.6 and C.3).[65] The model results in slightly different inference for co-

[63] Correction for heteroscedastic or serially correlated errors could be examined in future work, but here the main objective is to assess the overall distribution.

[64] The same issue applies to $\mu^{c}_{i,t}$ and $\mu^{y}_{i,t}$, but the chains for these parameters are thinned as they are not of economic interest.

[65] Results use 10,000 posterior draws after 10,000 burn in.

variate parameters compared with the deterministic model – posteriors tend to be closer to zero, and a little more disperse for transitory shocks, but a little less so for permanent shocks. Innovations of the transitory MPC have much larger variance than those of the permanent MPC, indicating that unobserved heterogeneity is more apparent for transitory MPCs. This is also seen in the large dispersion of the posterior distribution of all $\gamma^{\nu}_{i,t}$ draws compared to the observable predictions $X_{i,t}\hat{\Gamma}_{\nu}$, suggesting the model finds unobserved heterogeneity to be the main driver of variation in transitory MPCs (left panel Figure 4.11). This finding is similar to that of Lewis et al. (2020) for the economic stimulus payments of 2008. In constrast, the information from covariates explains much of the heterogeneity in permanent responses (right panel Figure 4.11).

	Mean	5%	95%
$\sigma_{\varepsilon\nu}$	0.549	0.514	0.592
$\sigma_{\varepsilon\mu}$	0.196	0.178	0.213
θ	0.051	0.025	0.080
$\sigma_{\nu y}$	0.235	0.230	0.240
$\sigma_{\zeta y}$	0.185	0.178	0.191
$\sigma_{\nu c}$	0.181	0.175	0.186
$\sigma_{\zeta c}$	0.088	0.082	0.094
Obs.	20286		
Households	3479		

Table 4.6: Stochastic model posterior statistics for shock processes; labour-force households.

A simple variance decomposition can be used to assess the share of the overall variance in MPCs driven by observed and unobserved heterogeneity. The ratio of the variance of the observed predictions $var(X_{i,t}\Gamma_{\nu})$ over the sum of the predicted and residual variance $var(X_{i,t}\Gamma_{\nu}) + var(\varepsilon^{\nu}_{i,t})$ is calculated for each loop of the Markov chain, then averaged.[66] This measure captures the relative con-

[66]This takes $cov(X_{i,t}\Gamma_{\nu}, \varepsilon^{\nu}_{i,t}) = 0$ as given and is equivalent to an (average) R-squared statistic for a regresion $\gamma^{\nu}_{i,t} = X_{i,t}\Gamma_{\nu} + \varepsilon^{\nu}_{i,t}$, but here $\gamma^{\nu}_{i,t}$ is unobserved and the model simulates $\varepsilon^{\nu}_{i,t}$ given a draw of $\sigma^2_{\varepsilon\nu}$. This is different to a decomposition of variance in the posterior draws of $\gamma^{\nu}_{i,t}$, since the posterior density of $\gamma^{\nu}_{i,t}$ includes additional terms from conditioning the joint

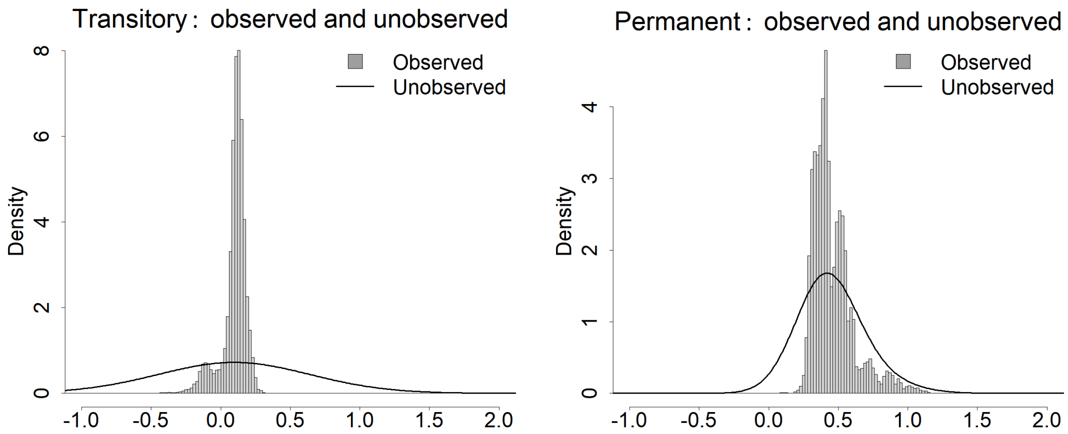

Figure 4.11: Stochastic model: distribution of observation-specific responses to transitory (left panel) and permanent (right panel) shocks; observed predictions given by $X_{i,t}\hat{\Gamma}_\nu$ and $X_{i,t}\hat{\Gamma}_\mu$, and unobserved predictions given by the posterior distribution of all γ^ν_{it} draws; 20,286 observations of labour-force households.

tribution of the covariates and accounts for parameter uncertainty. Observable heterogeneity accounts for only 3 per cent of variation in transitory MPCs, but 39 per cent in permanent MPCs.

A 'random intercept' model that uses household-specific intercepts instead of covariates is also estimated, detailed in Appendix C.4. Comparing the stochastic, deterministic and random intercept models, the distributions of predicted responses are very similar for transitory shocks, but show some differences for permanent shocks (Figure 4.12). Relative to the deterministic and random intercept models, the stochastic model's predicted permanent MPCs are shifted to the left and exhibit a thinner right tail. The permanent MPC innovation $\varepsilon^\mu_{i,t}$ is able to absorb some of the larger responses that are attributed to covariates in the deterministic model. This is consistent with smaller estimates of covariate parameters in the stochastic model and suggests that the effect of covariates on MPCs varies across the sample.

The results are potentially suggestive of weak identification of the transitory MPC variance $\sigma^2_{\varepsilon\nu}$. Substituting equation (4.5) into (4.3) and (4.6) into (4.4), there are four shocks in two equations. The structure of the unobserved components model

posterior that inform unobserved variation.

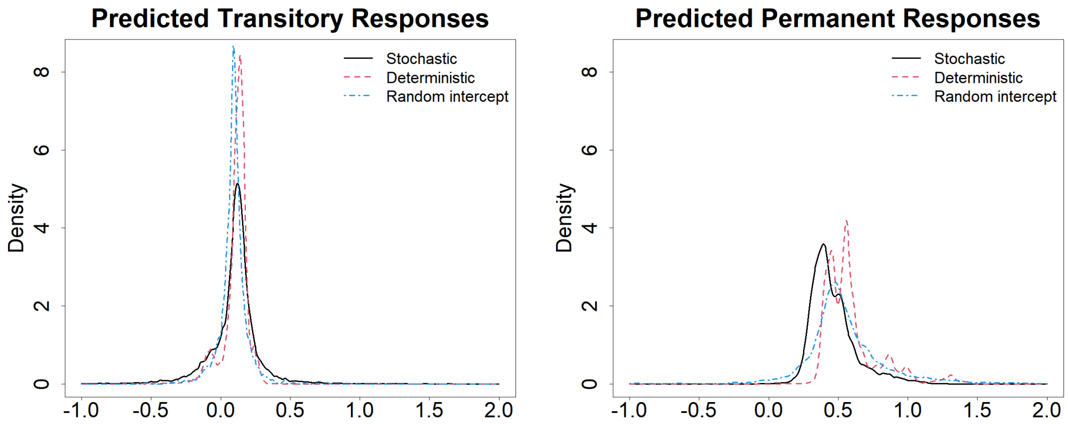

Figure 4.12: Model comparison: densities of predicted responses using posterior means of parameters; random intercept model detailed in Appendix C.4; labour-force households.

identifies $\varepsilon^{\mu}_{i,t}\mu^{y}_{i,t} + \nu^{c}_{i,t}$ from $\varepsilon^{\nu}_{i,t}\nu^{y}_{i,t} + \zeta^{c}_{i,t}$, while heteroscedasticity then separates shocks within these two terms. The heteroscedasticity induced by $\mu^{y}_{i,t}$ is larger than that of $\nu^{y}_{i,t}$.[67] In contrast, the random intercept model considered in Appendix C.4 finds smaller unobserved variation in transitory MPCs. This model restricts variation in MPCs to be driven by the cross-section only, which may help to distinguish between the equivalent $\varepsilon^{\nu}_{i}\nu^{y}_{i,t}$ and $\zeta^{c}_{i,t}$ shocks. Future work could explore this further.

Nevertheless, a greater role for unobserved heterogeneity in transitory MPCs is perhaps to be expected. The field has long sought to explain larger empirical estimates relative to theory and our understanding of MPCs out of transitory income continues to evolve. One potential interpretation is that the unobserved heterogeneity reflects model restrictions that do not allow transitory consumption responses to transitory income shocks. Cho et al. (2021) find that this plays a key role in patterns before and after the great recession. Other interpretations could include heterogeneity in preferences, such as myopia, that are not directly related to observable charactersitics. Overall, the result suggests that policy makers should expect greater uncertainty if designing policy to target transitory

[67] Recall that $\mu^{y}_{i,t}$ is the permanent component of income given by a random walk, so its variance increases linearly with time observations. Additional cross-sectional variation comes from estimating the household-specific period zero values $\mu^{y}_{i,0}$.

MPCs of a particular group. This is an important caveat to using sample or group average estimates when predicting the effects of policy.

4.6 Aggregate Effects of Heterogeneity

Simple transformations are often used to convert estimated microdata elasticities into aggregate statistics; however, correctly generating policy-relevant statistics that are representative of the population is more nuanced than it may seem. Two key factors require attention; i) using the correct functional form to generate a dollar-for-dollar MPC; and ii) weighting the estimation to correctly aggregate sample heterogeneity to a population descriptive statistic. Once these issues have been addressed, partial equilibrium predicted responses to policy interventions – cents consumed per dollar of additional income – can be obtained.

4.6.1 Converting Elasticities into MPCs

The above text refers to the parameters that govern household consumption responses to income shocks as MPCs, elasticities or consumption responses. The parameters (γ_ν and γ_μ) and predicted estimates ($X_{i,t}\Gamma_\nu$ and $X_{i,t}\Gamma_\mu$) are income elasticities of consumption; that is, they approximately reflect the percentage change in unexplained consumption in response to a one percent change in unexplained income.[68] The MPCs discussed in policy making often refer to the dollar change in consumption per dollar change in income – I label this concept $MPC for clarity. Further, there is a conceptual difference in the models that policy makers and econometricians are using. The $MPC concept assumes a constant dollar-for-dollar response across the distribution of income, whereas the regression is specified as having a constant (percentage-for-percentage) elasticity. Care is needed in transforming the regression concept into the $MPC concept.

The common approach to converting elasticities to dollar-for-dollar statistics is to

[68] Unexplained income and consumption, $y_{i,t}$ and $c_{i,t}$, are in log-real-dollar terms as these residuals have the same units as the dependent variables $ln(Y_{i,t})$ and $ln(C_{i,t})$. The parameters γ_μ and γ_ν then capture the change in log-real-dollar permanent unexplained consumption resulting from a (permanent or transitory) shock to log-real-dollar unexplained income.

multiply estimated elasticities by the ratio of sample-average consumption over sample-average income (e.g. Nakajima (2018), La Cava, Hughson, and Kaplan (2016)). I call this the quick \$MPC approximation. In the similar case of estimating fiscal multipliers, Ramey (2019) notes that this produces unreliable results. Here I tease out issues with this approach, which can be particularly salient when incorporating heterogeneity. Consider a simple regression for a single household

$$ln(C_t) = \alpha + \beta ln(Y_t) + \varepsilon_t$$

taking a first-order approximation at the sample averages $\bar{C}$ and $\bar{Y}$ gives

$$\beta = \frac{\partial C_t}{\partial Y_t}\frac{\bar{Y}}{\bar{C}}$$

and so multiplying β by $\frac{\bar{C}}{\bar{Y}}$ will return an approximation to the \$MPC, $\frac{\partial C_t}{\partial Y_t}$, at the sample average. However, this quick approximation does not necessarily provide the desired statistic, which is the sample average \$MPC.[69] Rather, taking the first derivative of the model gives

$$\beta = \frac{\partial C_t}{\partial Y_t}\frac{Y_t}{C_t}$$

and so the sample-average \$MPC is given by

$$\$MPC = \beta\frac{1}{T}\sum_{t=1}^{T}\frac{C_t}{Y_t}.$$

In effect, the quick approximation violates Jensen's inequality.[70] Potential bias in using the quick approximation may be exacerbated in a model with heterogeneous effects and sample weighting. Furthermore, the scaling factor $\frac{C_t}{Y_t}$ is, in fact, the empirical definition of the APC. Results in previous sections find that the transitory elasticity varies with households' lagged APC, which again implies the quick approach is misspecified.

[69]This is well understood in the microeconometrics literature, as it is equivalent to using marginal effects at the mean rather than average marginal effects.

Instead, I calculate the observation-level statistics required for the sample average \$MPC at each loop of the Markov chain. In the case of the benchmark empirical model used in this thesis, recall that unexplained income and consumption are the residual from a preliminary regression

$$ln(Y_{i,t}) = Z_{i,t}\Upsilon_Y + y_{i,t} \tag{2.1}$$
$$ln(C_{i,t}) = Z_{i,t}\Upsilon_C + c_{i,t}. \tag{2.2}$$

Using equation (2.2) to derive the observation-level \$MPC gives

$$\frac{\partial C_{i,t}}{\partial Y_{i,t}} = C_{i,t}\frac{\partial c_{i,t}}{\partial Y_{i,t}}.$$

Taking the case of an exogenous transitory income shock, such as might arise from one-off policy interventions (rather than structural reforms), a shock to $Y_{i,t}$ will affect $\nu^y_{i,t}$ but leave $\mu^y_{i,t}$ unchanged. Conditioning on $\mu^y_{i,t}$ and substituting in equations (2.1), (2.3), (2.5) and (2.6) gives

$$\frac{\partial C_{i,t}}{\partial Y_{i,t}} = \gamma_\nu\frac{C_{i,t}}{Y_{i,t}}.$$

For every household i and time t, the algorithm stores $\hat{\gamma}_\nu\frac{C_{i,t}}{Y_{i,t}}$ at each loop of the Markov chain. Here $\hat{\gamma}_\nu$ is a draw of the predicted elasticity, which can be generalised to also vary over households and time as in the preceding models of heterogeneity. The distribution of \$MPCs is given by the combined sample of these statistics, which can be used for posterior inference.

[70] An alternative approach is to avoid the logarithmic transformation and estimate a model in first differences divided through by a common factor, similar to the simple regressions in Hall (2009). The equivalent for the regression above would be

$$\frac{\Delta C_t}{C_{t-1}} = \alpha + \beta\frac{\Delta Y_t}{C_{t-1}} + \varepsilon_t.$$

Hall (2009) notes that such a model implies a dollar-for-dollar interpretation of β. However, this approach introduces bias into the estimated coefficient (see Kronmal (1993)).

4.6.2 Heterogeneity in the Population

Survey weighting is commonly disregarded in economic applications where identification of causality is the driving motive; however, it is essential in the inference of population statistics. Solon, Haider, and Wooldridge (2013) argue that when a pseudo-random sample exhibits heterogeneous effects, weighting the sample can correct for the non-representative sample, but we must also model heterogeneity for inference on the population average partial effect (in this context the average \$MPC for the US population). They provide a clear exposition of this issue. The crux is heterogeneity – neglecting to formally model heterogeneity will result in inconsistent population average effects. Estimates under homogeneity do not account for different variation in the regressor across groups, irrespective of whether the sample correctly reflects the population. In the case at hand, there is good reason to think that the variance of income shocks could vary across different groups (e.g. Meghir and Pistaferri (2004)).[71]

This section both models heterogeneity and weights the sample. Previous literature using the same empirical framework has not examined the aggregate implications of the estimated consumption responses, so this is a novel contribution.[72] This aspect of the chapter can be thought of as an empirical complement to advances in the theoretical literature that match population moments of the wealth distribution and calculate aggregate MPCs (e.g. Kaplan and Violante (2014)). However, the empirical setting allows for more dimensions of heterogeneity to be matched to population moments, beyond those that have formalised microfoundations, such as sex and race.

I implement the Bayesian Weighted Estimation (BWE) procedure in Gunawan et al. (2020) to generate population representative statistics. The baseline model without weighting is numerically intensive, taking approximately one day to run using four cores, so I opt to use the BWE algorithm based on the empirical distribution function to further take advantage of parallelisation. This method is intuitive as a form of Bayesian bootstrapping. First a set of survey weights

[71]However, the results suggest that this plays only a minor role (discussed below).

[72]Lewis et al. (2020) is a notable paper that accounts for both heterogeneity and weighting, using an empirical case study of the economic stimulus payments of 2008.

(in this case inverse probability weights) is used to draw J pseudo representative samples (PRSs) from the weighted empirical distribution of the data. Parameters are estimated using each PRS as they would be on the full sample, resulting in J replications of S draws of the full parameter set Θ. The resulting combination of all $J \times S$ draws can be used for posterior inference of the parameters. As the PRSs are independent, estimation can take place in parallel.

The labour force and retired samples are treated independently for PRS construction and estimation. This preserves uncertainty about parameters for retired households, given the smaller sample, and assumes that all model parameters for retired households are distinct to those of labour force households. Due to the panel state space empirical model, a PRS is difficult to construct at the observation level. Instead, the average weight (across time) is calculated for each household and these are used to construct PRSs at the household level.[73] Each PRS has the same number of households as the original sample (3,479 for labour force and 404 for retired), but may have a different number of total observations. The number of labour force and retired PRSs used in estimation are weighted such that the ratio of labour force to retired households matches population statistics. The results use 75 labour force PRSs and 135 retired PRSs, which translates to 82.7 per cent labour force and 17.3 per cent retired households. Inference uses 5,000 posterior draws after 5,000 burn in for each PRS.

The labour-force and retired samples are subsets of the full PSID sample and are not representative random samples of the population. For example, the labour-force sample comprises around 55 per cent HtM households, which is substantially higher than the roughly 30 per cent that Kaplan et al. (2014) find in the Survey of Consumer Finances (SCF). There are subtle differences between the coverage of the two surveys, but the SCF is specifically designed to capture the very high wealth households that the PSID lacks (Pfeffer, Schoeni, Kennickell, and Andreski (2016)). The SCF is considered the benchmark for wealth surveys; for example, Kaplan et al. (2018) calibrate their model to moments of the SCF

[73]Because the panel is unbalanced, the information used is not accurately weighted on a per observation basis. Adjusting weights by the inverse of the number of household observations would correct this, but lower the average sample length length. I chose not to do this due to the short nature of the panel.

wealth distribution rather than the PSID. As such, I match the PSID sample to SCF moments along crucial dimensions of heterogeneity. This is achieved by raking the original PSID inverse probability longitudinal family weights for each year separately.[74]

Matching the sample along the entire continuous distribution of multiple drivers of MPC heterogeneity is impractical, so weighting is conducted using discrete categories. This is natural for sex, race and education, which are the largest drivers of MPC variation. Age is grouped into 6 categories: 18-34, 35-44, 45-54, 55-64, 65-74 and 75+ years old. Total net worth is chosen as the final dimension for raking due to its broad coverage, capturing both home equity and assets that generate non-labour income, and precedence for survey comparison in Pfeffer et al. (2016). I follow the same definitions of net worth in each survey as in Table A1 of Pfeffer et al. (2016) with some minor adjustments.[75] Distributions of the weights by covariate are provided in Appendix C.7.

A simple linear deterministic MPC specification is used, with the same covariates as Section 4.5; that is, home equity, APC, business income and age, along with dummies for black and female heads, and education groups.[76] The MPC specification used on the retired sample reduces covariates by dropping APC, business income and the black dummy and combining education categories into i) high school or below and ii) at least some college.

The predictions of interest can be constructed in various ways using the set of all $J \times S$ draws. I take the central object of interest to be the average predicted re-

[74] To do so, I calculate the relevant (weighted) SCF population statistics, then utilise the Stata command `ipfraking()`. As the surveys are conducted with different frequencies, each PSID survey is matched to the closest SCF year.

[75] I include the variables OTH_INST (other installment loans, which are mostly consumer and other property loans) and OTHLOC (other lines of credit) in the SCF net worth definition. These appear to be missing in Pfeffer et al. (2016) yet are covered by the PSID. Net worth comprises business assets, checking and savings accounts, stocks and mutual funds, Individual Retirement Accounts and annuities, vehicles net worth, home net worth, other real estate net worth, other assets and other debt.

[76] The predictions of this model are similar to the stochastic models, while estimating far fewer parameters. Stochastic models do not appear to be well identified using the much smaller retired sample.

sponse for each household $\overline{X_i\Gamma_\nu}$ and $\overline{X_i\Gamma_\mu}$ constructed using the posterior means of coefficients in each PRS. That is, for each PRS, the posterior means of Γ_ν and Γ_μ are used to construct response predictions, which are then averaged for each household, resulting in $75 \times 3,479$ labour force and 135×404 retired estimates. This is consistent with the application of weighting at the household level and allows for some parameter uncertainty (in contrast to using only the posterior mean of all $J \times S$ draws).

4.6.3 Aggregate Responses to Shocks

This section analyses the effects of heterogeneity, population weighting and $MPC conversion on partial equilibrium aggregate responses to income shocks. These adjustments tend to have a larger effect on permanent responses than transitory (Table 4.7). Both modelling heterogeneity and population weighting increase the mean and dispersion of permanent responses. For transitory responses, these adjustments increase dispersion, but this results in lower means. Population weighting places larger mass on the right tail of distribution of permanent responses and on the region just below zero for transitory responses (Figure 4.13). Both labour-force and retired samples respond in similar ways to population weighting, but overall it does not drastically change the distribution of predictions.

The scaling factor for the $MPC conversion, $\frac{C_{i,t}}{Y_{i,t}}$, has a skewed distribution with a mean of 0.62 and a long right tail. As around 90 per cent of the mass is between zero and one, the $MPC conversion shifts the distribution towards zero (Figure 4.13). This also reduces heterogeneity for the permanent $MPCs. However, the retired sample has a larger mean and variance of the scaling factor, which results in the distribution of transitory $MPCs becoming positively skewed with a larger mean (Figure 4.14).

Overall, the results find a population average $MPC of around 7 cents per dollar for transitory income shocks, and 38 cents for permanent shocks.[77] Compared with the simple intercept only elasticity using the labour force sample, this is a

[77]Recall that the framework only captures changes in nondurables consumption.

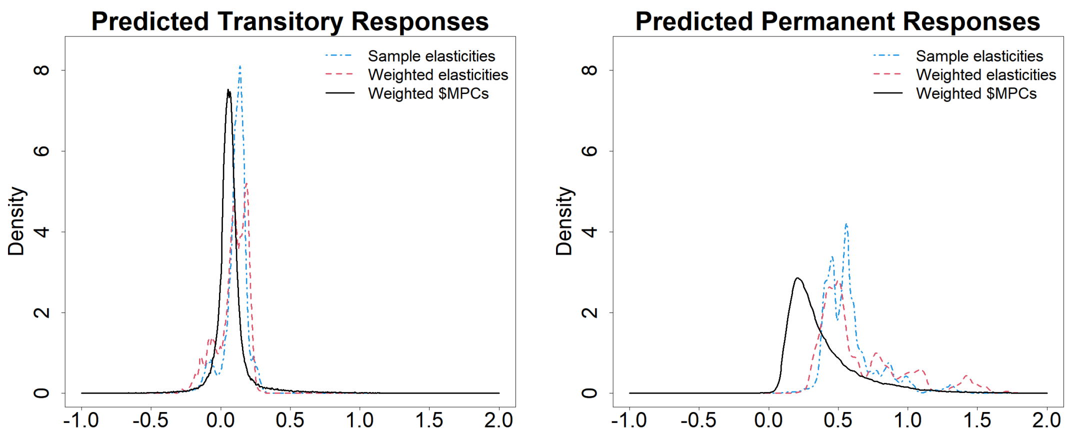

Figure 4.13: Densities of predicted responses for each step in aggregation; sample elasticities reflect observation-level predictions from labour force and retired households; weighted densities reflect household-average predictions.

	Transitory shocks			Permanent shocks		
		Heterogeneity interval			Heterogeneity interval	
	Mean	5%	95%	Mean	5%	95%
Intercept only model (elasticities)						
Labour force	0.136	–	–	0.550	–	–
Retired	0.107	–	–	0.458	–	–
Multiple covariate model (elasticities)						
Labour force	0.109	-0.088	0.223	0.591	0.389	0.999
Retired	0.090	0.021	0.164	0.538	0.221	0.824
Population weighted multiple covariate (elasticities)						
Aggregate population	0.086	-0.139	0.214	0.665	0.342	1.413
Labour force	0.096	-0.141	0.213	0.665	0.324	1.400
Retired	0.076	-0.134	0.219	0.665	0.375	1.414
Population weighted multiple covariate ($MPC)						
Aggregate population	0.074	-0.063	0.244	0.385	0.122	0.927
Labour force	0.053	-0.050	0.145	0.391	0.127	0.944
Retired	0.176	-0.153	0.767	0.356	0.092	0.832

Table 4.7: Various predicted responses to income shocks; heterogeneity interval represents percentiles of the distribution of predicted responses; multiple covariate model reflects observation-level predictions, whereas weighted statistics reflect household-average predictions; all models use deterministic interactions.

large reduction. However, the quick \$MPC approximation on this simple elasticity bridges most of the gap, resulting in estimates of 8 and 32 cents per dollar. Although there are methodological concerns in using the quick approximation, in this particular framework it produces reasonable estimates. Compared to the \$MPC conversion, heterogeneity and population weighting play a minor role when calculating the central moment.[78] This is in large part due to the PSID sample shares being close to the population shares along the key dimensions of heterogeneity, and the comparatively small effect of modelling heterogeneity on the average elasticity (Table 4.7).

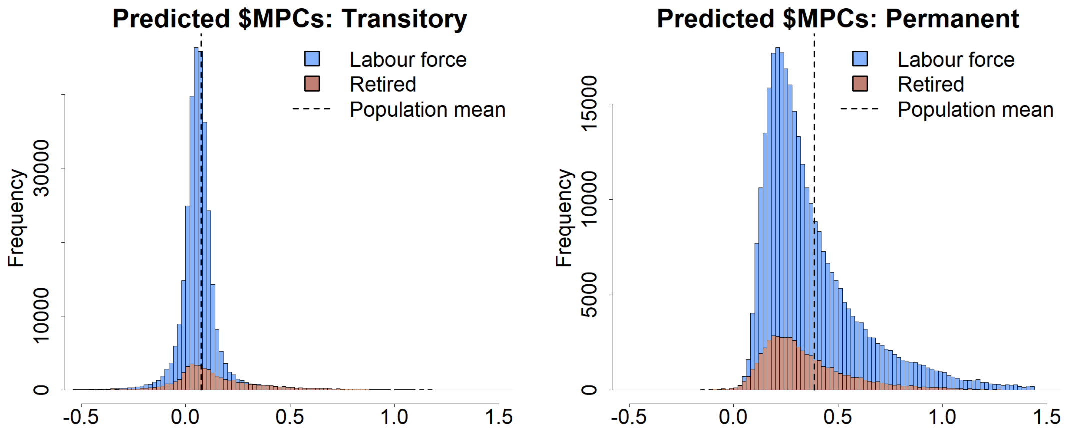

Figure 4.14: Histogram of predicted household \$MPCs; population representative sample of 75 labour force PRSs and 135 retired PRSs.

4.7 Conclusion

Building on the empirical framework in Chapter 2, I develop an approach to estimate the distribution of household MPCs out of transitory and permanent income shocks. The approach allows for observable household characteristics to flexibly drive MPC heterogeneity through the use of b-splines to capture potential nonlinearities, although the results suggest a simple linear model is usually sufficient. Single covariate results offer suggestive evidence in support of the

[78]This is in contrast to Lewis et al. (2020) analysing the 2008 tax rebate. They find that the mean heterogeneous \$MPC is much larger than the homogeneous \$MPC, particularly when calculating aggregate partial equilibrium effects.

two-asset model, life-cycle dynamics and previous empirical findings regarding housing assets and debt, but only life-cycle dynamics are robust in a multiple covariate setting. The results also suggest business and farm income and the average propensity to consume are relevant drivers of MPCs. However, demographic variables – sex, race and education – account for the majority of observable MPC heterogeneity, even when controlling for household balance sheet characteristics. This suggests that further research into broader socio-economic frictions is warranted, as they may be more important than financial frictions.

The analysis is extended to incorporate *unobserved* heterogeneity, similar to a 'random coefficients' model. Observed heterogeneity explains 39 per cent of the variation of permanent MPCs, but only 3 per cent of transitory MPCs. Thus, unobserved heterogeneity plays an important role for transitory MPCs, which is consistent with the findings in Lewis et al. (2020).

I construct policy-relevant MPCs that are representative of the US population by i) extending the sample to include retired households, ii) formally modelling heterogeneity, iii) population weighting the sample, and iv) converting the microdata elasticities into $MPCs. Population weighting tightens the distributions of elasticities and shifts them towards zero. Conversion of elasticities to $MPCs further shifts the distributions lower; population average $MPCs imply that households increase nondurables consumption by 7 cents per dollar for transitory income shocks, and 38 cents for permanent shocks. The $MPC conversion has a larger impact on central moment statistics than modelling heterogeneity or population weighting.

Chapter 5

Sparsification & Demographics

5.1 Introduction

Over the past decade, there has been heightened interest in understanding what household characteristics drive consumption decisions. This has been accentuated by theoretical developments that suggest important roles for consumer heterogeneity, and as the increasing use of household microdata facilitates much richer empirical evidence. I leverage the range of data available on household financial and demographic characteristics in the US Panel Study of Income Dynamics (PSID) over 1999–2019 to determine the most salient drivers of variation in the marginal propensity to consume (MPC) out of income shocks. Building on the framework used in Chapter 4, I utilise statistical methods from the Bayesian shrinkage and sparsification literature to separate the rich set of covariates into signals that drive meaningful variation in MPCs and noise covariates that do not. Demographic factors are found to be the main drivers of heterogeneity in consumption smoothing. Financial literacy is investigated as a potential underlying mechansim and is found to play a role in consumption smoothing, but does not eliminate the effects of demographics.

Empirical modelling of heterogeneity in MPCs has typically been conducted by testing a select few covariates with *a priori* motivation (Johnson et al. (2006), Parker et al. (2013), Fagereng et al. (2018), Lewis et al. (2020), Jappelli and Pistaferri (2020)). This approach often ties in with theoretical work that has emphasised financial frictions as crucial to consumption decisions (e.g. Kaplan and Violante (2014)). However, household balance sheets can be complex – as evidenced by the raft of questions regarding assets and liabilities included in the PSID – and empirical studies often aggregate individual items into measures such as net worth or liquid and illiquid assets. I use disaggregated balance sheet data in a number of functional forms to allow for a variety of possible effects on MPCs.[79] As such, the data is allowed to flexibly determine relationships with MPCs and models that use aggregate measures are nested (as they are simply linear combinations of the items). MPC heterogeneity is modelled using covariate interaction terms, which are extended to incorporate shrinkage and sparsification to help identify covariates that have the largest effect on MPCs.[80]

Households are subject to a much wider range of social and economic factors that influence their decision making than financial considerations alone. Demographic characteristics, such as race and gender, often receive specific attention in research concerned with income and wealth inequality, but the literature on consumption smoothing has been primarily focused on financial and preference heterogeneity (Heathcote et al. (2009), Parker (2017), Gelman (2021)). There are important exceptions, however. Ageing and cohort effects have been well studied, motivated by the life-cycle model (Browning and Crossley (2001), Gourinchas and Parker (2002), Blundell et al. (2008)). Two recent studies, Ganong et al. (2020) and De Giorgi et al. (2020), have also honed in on race. They both find that black households are substantially more responsive than white households to income shocks, but that differences in wealth explain almost all of the gap. The PSID contains a rich source of demographic information that is incorporated into the empirical model alongside financial characteristics. The results find that race, sex, education and birth cohort play key roles in MPC heterogeneity, whereas financial factors do not.

The *ex ante* expectation is that most of the covariates used have no systematic relationship with MPCs, and it is not known which are signals and which

[79] A range of factors could affect such relationships. For example, transaction costs, regulatory constraints and behavioural influences such as mental accounting can vary widely across asset classes, and there is some evidence that households respond directly to debt in addition to net positions (Misra and Surico (2014), Price et al. (2019), Cho et al. (2021)).

[80] I further develop the model to estimate all parameters jointly, including the life-cycle control variables that are estimated in a separate first stage in the previous literature. This has a different conceptual interpretation based on full information of the correlation between income shocks and the deterministic component of consumption.

are noise. The dataset provides a large number of observations for estimation; however, the inclusion of a wide (and potentially correlated) set of household covariates may lead to poor inference. As such, Bayesian shrinkage and sparsification techniques are incorporated into the estimation algorithm. Shrinkage pulls estimates of coefficients towards zero for covariates with little explanatory power, while sparsification separates covariates into discrete groups of signals and nulls. Shrinkage is implemented via the Makalic and Schmidt (2016) representation of the Horseshoe prior (Carvalho, Polson, and Scott (2010)). The Signal Adaptive Variable Selector (SAVS) algorithm of Ray and Bhattacharya (2018) is used to sparsify the coefficient vector; that is, to select which covariates drive variation in MPCs. Following the Huber, Koop, and Onorante (2020) implementation, the SAVS estimates provide posterior inclusion probabilities (PIPs) that reflect the probability that a covariate is a signal, and so are used to select a statistically robust sparse model. Estimation on simulated data guides the approach, and finds that selection of covariates is more reliable for permanent income shocks.

The results identify a few key observable drivers of MPCs. These are primarily demographic factors. In particular, households with heads who are black, female or born in the 1980s have larger responses to permanent income shocks, while those with four or more years of college have lower responses. These four covariates drive the overwhelming majority of heterogeneity in permanent MPCs. The result is robust to various specifications that allow for further flexibility in balance sheet covariates; however, the results are in part driven by these households facing shocks with different variances.

The demographic pattern of the results is consistent with financial literacy being an underlying driver of consumption smoothing. The empirical model is tested on simulated data from a life-cycle model with endogeneous financial literacy developed by Lusardi et al. (2017), confirming that financial literacy drives variation in household MPCs. Following Bialowolski, Cwynar, Xiao, and Weziak-Bialowolska (2021), a proxy for financial literacy is constructed using additional data from the 2016 PSID supplemental Wellbeing and Daily Life survey. The proxy is found to contribute to MPC heterogeneity in a reduced sample of households for which it is available. However, demographic characteristics continue to drive the majority of variation in MPCs.

5.2 Empirical Model & Data

The empirical framework is detailed in Section 2.1; here I use the extension from Chapter 4

$$y_{i,t} = \mu^y_{i,t} + \nu^y_{i,t} + \theta\nu^y_{i,t-1} \tag{2.3}$$

$$\mu^y_{i,t} = \mu^y_{i,t-1} + \zeta^y_{i,t} \tag{2.4}$$

$$c_{i,t} = X_{i,t}\Gamma_\mu\mu^y_{i,t} + X_{i,t}\Psi + \mu^c_{i,t} + \nu^c_{i,t} \tag{4.1}$$

$$\mu^c_{i,t} = \mu^c_{i,t-1} + X_{i,t}\Gamma_\nu\nu^y_{i,t} + \zeta^c_{i,t} \tag{4.2}$$

where Γ_ν, Γ_μ and Ψ are vectors of coefficients and $X_{i,t}$ is a set of K household characteristics, which may vary over household and time, and an intercept. The predictions $X_{i,t}\hat{\Gamma}_\nu$ and $X_{i,t}\hat{\Gamma}_\mu$ give the estimated MPC out of transitory and permanent income shocks.[81]

This chapter extends the model to incorporate shrinkage and sparsification over a rich set of household characteristics, without prior information on which drive variation in MPCs and which are noise.[82] Shrinkage pulls estimates of coefficients towards zero for covariates with little explanatory power, while sparsification separates covariates into discrete groups of signals and nulls. This provides a statistical approach to determining which specific household characteristics drive heterogeneity in MPCs. In contrast, previous studies tend to examine only a handful of covariates motivated by theory, potentially neglecting other meaningful sources of variation.

In Appendix D.5, I jointly estimate the full model in equations (2.1)-(2.4) and (4.1)-(4.2). That is, I jointly estimate the first-stage regressions that remove

[81] Recall that these predictions are elasticities, but are referred to as MPCs for simplicity. They are not dollar-for-dollar measures.

[82] This chapter implements the precision sampler of Chan and Jeliazkov (2009) to sample μ^y instead of the Metropolis-Hastings algorithm used in previous chapters.

predictable life-cycle variation, rather than using the residuals y and c as dependent variables. This provides true full-information estimates where households have knowledge of the correlation between predictable life-cycle variation and the stochastic income and consumption process. This results in some interesting features regarding the response to transitory income shocks, but is consistent with the main findings regarding demographics.

5.2.1 Sparsification

Sparsification is introduced following the Huber et al. (2020) implementation of the Signal Adaptive Variable Selector (SAVS) of Ray and Bhattacharya (2018). At every loop (n) of the Markov chain, auxiliary $(K+1) \times 1$ vectors $\mathcal{S}_{\nu}^{(n)}$ and $\mathcal{S}_{\mu}^{(n)}$ are calculated, according to the SAVS formula

$$\mathcal{S}_{z,j}^{(n)} = sign\left(\Gamma_{z,j}^{(n)}\right) \|X_j\|^{-2} \left(|\Gamma_{z,j}^{(n)}| \ \|X_j\|^2 - \frac{1}{|\Gamma_{z,j}^{(n)}|^2}\right)_+ \quad ; \quad \forall \quad z \in \nu, \mu$$

where X is the covariate data stacked by time and individual, X_j is the jth column of X, $sign(x)$ returns the sign of x, $\|x\|$ is the Euclidean norm and $(x)_+ = max(x, 0)$. The SAVS formula is an approximate solution to a sparsifying optimisation problem related to the adaptive lasso. It finds a sparse version of Γ_z based on minimising the Euclidean distance between the dense prediction $X\Gamma_z$ and sparse prediction $X\mathcal{S}_z$, subject to a covariate-specific penalty (Ray and Bhattacharya (2018)). Covariates that provide little additional predictive power are censored to be zero. The full set of SAVS estimates from the Markov chain provide a sparsified equivalent of the dense posterior draws of Γ_{ν} and Γ_{μ}, where estimates in each draw are separated into signals and nulls. The SAVS estimates are purely auxiliary parameters and do not affect the Markov chain directly. They can be used to construct sparse predictions of household MPCs, but are primarily used here to construct posterior inclusion probabilities (PIPs). These provide a measure of how likely each covariate is to be a signal (and hence important for MPC heterogeneity).

5.2.2 Shrinkage

Shrinkage on the coefficient vectors Γ_ν and Γ_μ is conducted using the Horseshoe prior of Carvalho et al. (2010). This posits priors with large mass near zero and thick tails, which are automatically adjusted to each parameter based on global and local tuning hyperparameters and shock variances. Posterior distributions of coefficients that are poorly informed by the data (or simply close to zero) are pulled towards zero, resulting in draws that are more likely to be censored by the SAVS estimator, while larger draws are relatively unaffected. I adapt the Makalic and Schmidt (2016) representation that uses auxiliary variables. The priors on the intercepts are maintined to be the same as in Chapter 4, so shrinkage only applies to the K covariate parameters. Coefficient priors are given by[83]

$$
\begin{aligned}
\Gamma_\nu | \lambda_\nu^2, \tau_\nu^2, \sigma_{\zeta c}^2 &\sim \mathcal{N}(m_\Gamma, \Delta_\nu) \\
\Gamma_\mu | \lambda_\mu^2, \tau_\mu^2, \sigma_{\nu c}^2 &\sim \mathcal{N}(m_\Gamma, \Delta_\mu) \\
m_\Gamma &= [0.35, \mathbf{0}_K]' \\
\Delta_\nu &= \begin{bmatrix} 0.5^2 & \mathbf{0}_K \\ \mathbf{0}_K' & \sigma_{\zeta c}^2 \tau_\nu^2 \Lambda_\nu \end{bmatrix} \\
\Lambda_\nu &= diag(\lambda_{\nu,1}^2, \ldots, \lambda_{\nu,K}^2) \\
\Delta_\mu &= \begin{bmatrix} 0.5^2 & \mathbf{0}_K \\ \mathbf{0}_K' & \sigma_{\nu c}^2 \tau_\mu^2 \Lambda_\mu \end{bmatrix} \\
\Lambda_\mu &= diag(\lambda_{\mu,1}^2, \ldots, \lambda_{\mu,K}^2)
\end{aligned}
$$

where $\mathbf{0}_K$ is a row vector of K zeros and $diag(x)$ is a diagonal matrix with vector x on the diagonal. The auxiliary parameters λ, τ and ψ provide the horseshoe prior using the inverse-gamma (IG) distribution

$$
\begin{aligned}
\lambda_{z,j}^2 | \psi_{z,j} &= IG(1/2, 1/\psi_{z,j}); \quad \forall \quad z \in \nu, \mu \quad \text{and} \quad j \in 1, ..., K \\
\tau_z^2 | \xi_z &= IG(1/2, 1/\xi_z); \quad \forall \quad z \in \nu, \mu \\
\psi_{z,1}, ..., \psi_{z,K}, \xi_z &= IG(1/2, 1); \quad \forall \quad z \in \nu, \mu.
\end{aligned}
$$

[83]The same prior structure as Γ_μ is also independently applied to the coefficient vector for the levels control variables Ψ.

5.2.3 Simulation

The estimation algorithm is tested using simulated datasets to gauge the accuracy of inference in selecting covariates and estimating the MPC parameters. Unexplained income and consumption, y and c, are simulated according to the data generating process (DGP) in equations (2.3), (2.4), (4.1) and (4.2), with the Ψ vector set to zero. Simulation results are presented for eight specifications using six different parameterisations of the DGP. The MPC covariates in the design matrix X are all normally distributed random variables with zero mean and variance one, with the exception of the first covariate which is an intercept and modifications in specifications six and seven detailed below. For each specification, the first five covariates are signals with true coefficients

$$\Gamma_{\nu,1:5} = [0.1, 0.05, -0.02, 0.1, -0.05]$$
$$\Gamma_{\mu,1:5} = [0.7, 0.075, -0.025, 0.2, -0.1]$$

and any additional covariates are noise with coefficients of zero. This provides a sample with MPC heterogeneity roughly similar to that seen in the empirical application; an example distribution of true MPCs ($X\Gamma_\nu$ and $X\Gamma_\mu$) is shown in Figure 5.1.

Specifications vary based on whether shrinkage is applied, the number of noise covariates, correlation across covariates and inclusion of binary covariates. The first three specifications have zero correlation between all covariates, while the remaining specifications have all off-diagonal elements of the correlation matrix of X set to 0.5. Specifications six and seven include binary variables that replace covariates four and five. In specification six, they are observation based, taking a value of one if the random variable is greater than 0.5. In specification seven, they are constant within each household, taking a value of one if the household mean of the random variable is greater than zero. Data is simulated for 4,000 households, each with five time observations, and estimation uses 5,000 draws after 3,000 burn in. Reported results are averages over 50 replications of each specification.

I compare results of the dense Γ_ν and Γ_μ vectors under disperse normal priors

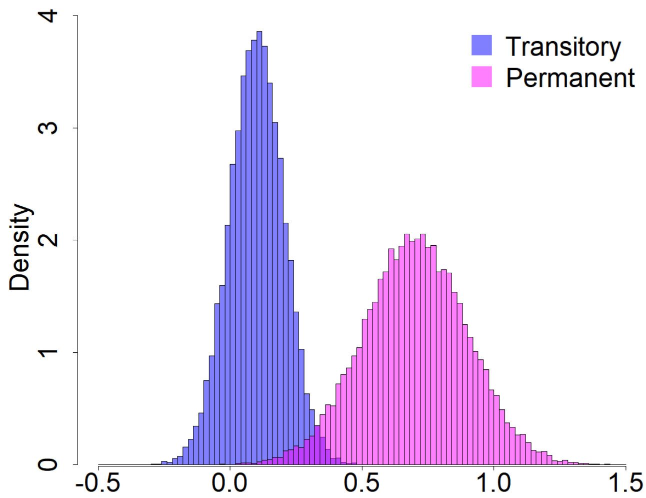

Figure 5.1: Histogram of true transitory and permanent responses, $X\Gamma_\nu$ and $X\Gamma_\mu$, for an example simulation.

with no shrinkage (equivalent to the model in Chapter 4) against results using Horseshoe shrinkage and against sparse SAVS estimates. Accuracy is assessed in terms of parameter posteriors, as well as covariate selection. Although these are strongly related concepts, particularly in terms of the estimation algorithms, they provide different insights into the empirical application. Model selection is achieved through the SAVS routine. A key diagnostic is the PIP, which is the proportion of times in the Markov chain that a covariate is selected as a signal. This can be interpreted as an estimate of the probability that a covariate is part of the true model. Given a particular sparse set of covariates, parameter accuracy is clearly of interest in assessing the effect of each covariate on MPC heterogeneity.

Simulated Model Selection

Table 5.1 reports PIPs from the SAVS estimates, along with true positive and true negative rates. The true positive rate is the share of times a signal is classified correctly in the chain over the total number of true signals. The true negative rate is the equivalent for the noise covariates. Overall, the SAVS routine performs well selecting noise covariates but is worse at selecting signals, reflected in the true positive and true negative rates. Shrinkage improves the selection of noise covariates, but worsens the selection of signals (which is to be expected as it pulls all draws towards zero to some extent). The results also show that covariates driving permanent responses are selected more accurately than transitory ones.

The simulation PIPs indicate that selection depends on the size of the true coefficients and nature of the covariate. The intercepts are classified extremely accurately, but the binary variables tend to be classified more poorly than the standard normal covariates, particularly the household constant binary variables. Correlation across covariates reduces selection accuracy, as in Ray and Bhattacharya (2018). The accuracy of selection is clearly increasing with the magnitude of the coefficient. This is intuitive given the SAVS routine relies on penalised thresholds for categorisation – covariates with small true coefficients appear much like nulls and are subject to a larger penalty.[84] This is an important caveat – even if a coefficient's posterior distribution is extremely tight (accurate inference), the covariate may still be classified as a null if the size of its effect is small.

Simulated Parameter Accuracy

Table 5.2 reports mean absolute errors (MAEs) for sets of parameters. Each error is calculated using posterior means and results are averaged over the parameter vector (and then averaged over replications). Using all covariates, the results show that the Horseshoe prior markedly improves accuracy and that the SAVS estimates, $\mathcal{S}_\nu$ and $\mathcal{S}_\mu$, clearly outperform the dense estimates in the presence of noise. However, MAEs based on all parameters give a somewhat misleading sense of accuracy when there are many noise covariates. That is, parameter accuracy is more concerned with estimating the true value of signal coefficients than estimating noise coefficients close to zero. MAEs based on only the true signal covariates provide contradictory evidence. SAVS estimates perform poorly due

[84]The penalty is set as the reciprocal of the squared parameter draw. Even if the draw is the true value, this disadvantages covariates with small coefficients. For example, the penalty is 2,500 for a parameter draw of 0.02, but only 100 for a parameter draw of 0.1. This suggests that the SAVS routine might be improved for particular applications by tuning the penalisation function to one based on a cut-off value.

Specification	#1	#2	#3	#4	#5	#6	#7	#8
Shrinkage	None	None	HS	None	HS	HS	HS	HS
No. noise variables	0	25	25	25	25	25	25	50
Correlation	0	0	0	0.5	0.5	0.5	0.5	0.5
Binary signals 4 & 5	No	No	No	No	No	Obs	HH	No
Transitory								
$\Gamma_{\nu,1}$ = 0.100	1.00	1.00	1.00	1.00	1.00	1.00	0.98	1.00
$\Gamma_{\nu,2}$ = 0.050	0.81	0.81	0.66	0.78	0.52	0.53	0.50	0.42
$\Gamma_{\nu,3}$ = -0.020	0.18	0.15	0.05	0.29	0.09	0.07	0.08	0.05
$\Gamma_{\nu,4}$ = 0.100	1.00	1.00	1.00	1.00	0.99	0.57	0.54	0.98
$\Gamma_{\nu,5}$ = -0.050	0.81	0.82	0.69	0.75	0.49	0.14	0.18	0.40
$mean(\Gamma_{\nu,i} = 0)$	NA	0.04	0.01	0.13	0.02	0.01	0.01	0.01
$max(\Gamma_{\nu,i} = 0)$	NA	0.63	0.43	0.87	0.51	0.51	0.43	0.75
Permanent								
$\Gamma_{\mu,1}$ = 0.700	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$\Gamma_{\mu,2}$ = 0.075	1.00	1.00	1.00	0.97	0.95	0.95	0.94	0.97
$\Gamma_{\mu,3}$ = -0.025	0.06	0.07	0.05	0.19	0.11	0.12	0.10	0.10
$\Gamma_{\mu,4}$ = 0.200	1.00	1.00	1.00	1.00	1.00	1.00	0.67	1.00
$\Gamma_{\mu,5}$ = -0.100	1.00	1.00	1.00	1.00	1.00	0.84	0.27	1.00
$mean(\Gamma_{\mu,i} = 0)$	NA	0.00	0.00	0.01	0.00	0.00	0.00	0.00
$max(\Gamma_{\mu,i} = 0)$	NA	0.05	0.01	0.64	0.27	0.21	0.22	0.18
All coefficients								
True positive rate	0.79	0.78	0.74	0.80	0.71	0.62	0.53	0.69
True negative rate	NA	0.98	1.00	0.93	0.99	0.99	0.99	0.99

Table 5.1: PIPs, along with true positive and true negative rates; averages over 50 replications.

to both penalisation of the coefficient and misclassification of signal covariates as nulls. Shrinkage improves accuracy for the permanent responses and worsens accuracy for the transitory responses, but is similar to the basic priors overall. The results suggest that, given a sparse set of covariates selected using SAVS, dense estimates should be used when interpreting the economic effects of covariates. Detailed results are available in Appendix D.2.

The simulations provide a guide to interpreting the results in the empirical ap-

Specification	#1	#2	#3	#4	#5	#6	#7	#8
Shrinkage	None	None	HS	None	HS	HS	HS	HS
No. noise variables	0	25	25	25	25	25	25	50
Correlation	0	0	0	0.5	0.5	0.5	0.5	0.5
Binary signals 4 & 5	No	No	No	No	No	Obs	HH	No
All covariates								
Γ_ν	1.29	1.06	0.49	1.43	0.62	0.73	0.81	0.40
Γ_μ	2.40	0.98	0.60	1.16	0.66	0.69	1.07	0.42
$\mathcal{S}_\nu$	1.76	0.33	0.35	0.56	0.43	0.59	0.65	0.25
$\mathcal{S}_\mu$	3.14	0.58	0.53	0.58	0.53	0.63	0.98	0.31
Signals only								
Γ_ν	1.29	1.26	1.41	1.62	1.86	2.77	3.36	2.02
Γ_μ	2.40	2.81	2.55	2.79	2.61	2.85	5.32	2.84
$\mathcal{S}_\nu$	1.76	1.73	2.04	1.86	2.38	3.38	3.78	2.47
$\mathcal{S}_\mu$	3.14	3.51	3.18	3.38	3.17	3.75	5.87	3.42

Table 5.2: Mean absolute error; averages over 50 replications using posterior means.

plication. To select a single 'sparse model' of only robust drivers of MPC heterogeneity, I set a minimum PIP threshold with shrinkage for inclusion. A minimum PIP of 0.5 would suggest that covariates with true coefficient magnitudes greater than 0.05 are typically selected, but any below that may well become false negatives. This would also mean that there is a possibility of a false positive for the transitory responses, but that is very unlikely for permanent responses. The results further indicate that household-constant binary signals (specification seven) are difficult to reliably select at this threshold.

To further investigate whether there are weak, but statistically meaningful, drivers not captured by this sparse model, I select covariates using both PIPs and the highest posterior density (HPD) interval for a 'broad' model. I use a minimum PIP of 0.75 from the estimation *without* shrinkage and also discard covariates where the 99 per cent HPD of the dense estimates includes zero. The broad model could include false positives, given the lower true negative rate of specifications without shrinkage in Table 5.1.

5.2.4 Data

Details of the PSID dataset were provided in Section 2.3, although this chapter uses an expanded set of household covariates. All continuous balance sheet variables (including 'hand-to-mouth' dummies) are lagged one period to minimise potential endogeneity, whereas demographic and other binary variables are contemporaneous. Continuous covariates are standardised to be mean zero and unit variance.[85] This is commonplace in the Bayesian sparsification literature as it rescales the coefficients such that their relative sizes to one another directly reflects their contribution to variation in the prediction. As such, covariates that typically have very large standard deviations and small coefficients, like items on household balance sheets, should not be disadvantaged by sparsification.[86]

A broad suite of demographic and balance sheet covariates are used; several of which have precedence for driving MPC heterogeneity, but many do not. Table 5.3 provides a description of each covariate. Wealth measures are kept as separate components (such as stocks and home equity), rather than being aggregated into measures of net worth or liquid and illiquid wealth as is common in the literature. The functional form of the relationship between balance sheet covariates and MPCs is not obvious, so these covariates are included in levels, log-levels and as a ratio to labour income.[87] Home equity, other real estate and non-labour income are also split into separate covariates for positive and negative values, which is motivated by the results for home equity in Chapter 4 (Figure 4.5).

[85] All categorical variables are separated into binary categories, and all binary covariates are not standardised.

[86] Recall that in a simple linear regression framework, standardisation of variables results in coefficient estimates that are scaled up by the standard deviation of the original variables.

[87] Specifically, this is the case for stocks, checking balances, credit card debt, home equity, other real estate, IRA, other assets, other debt, vehicles net worth, business net worth, mortgage remaining, mortgage payments and vehicle loans. Income and non-labour income are included in levels and log-levels.

Covariate	Description	PSID #	Mean	Std. Dev.
Labour income*	After tax labour income and government transfers (deflated using CPI)		69.71	43.45
Consumption*	Food, utilities, gasoline, car maintenance, public transportation, childcare, health and education expenditures (deflated using CPI)		37.83	21.75
Stocks	Shares of stock in publicly held corporations, mutual funds, or investment trusts	W16	14.39	61.71
Checking account	Checking/savings accounts, money market funds, certificates of deposit, government bonds, or treasury bills	W28	15.16	40.43
Credit card debt	Credit card or store card debt (imputed prior to 2011)	W39a	3.88	7.36
Home equityˆ	Value of home minus mortgage(s) remaining	A20,A24	68.40	111.04
Other real estateˆ	Value of real estate (other than main home) minus debt owed	W2a-b	14.61	71.67
IRA	Private annuities or Individual Retirement Accounts	W22	25.71	83.06
Other assets	Other assets, such as cash value in a life insurance policy, a valuable collection for investment purposes, or rights in a trust or estate	W34	7.69	47.77
Other debt	Student loans, medical bills, legal bills, family loans and all other debts (imputed prior to 2011)	W39b1-7	8.44	22.79
Vehicles net worth	Net value of vehicles (including cars, trucks, motor homes, trailers and boats)	W6	15.18	18.99
Business net worth	Net value of farm or business	W11a-b	28.39	280.62
Mortgage remaining	Remaining principal on first and second mortgages	A24	82.07	109.86
Mortgage payments	Monthly payments on first and second mortgages	A25	0.73	1.17
Months mortgage behind	Total number of months behind on mortgage(s)	A27B	0.07	1.16
Mortgage rate	Mortgage interest rate (weighted average across mortgages)	A25A4	3.36	3.22
Vehicle loans	Principal amount borrowed for vehicle purchases since previous survey	F66	6.74	11.76
Non-labour incomeˆ	After tax asset income (rent, dividends, interest, trusts and asset share of farm or business income)		2.51	11.93
APC	Average propensity to consume (nondurables out of labour income)		0.62	0.38
Has equity loan (bin.)	Whether household has a home equity loan	A23a	0.09	0.28
Has business (bin.)	Whether household owns (part or all of) a farm or business	W10	0.12	0.33
Behind on mortgage (bin.)	Whether household is behind on mortgage payments	A27a	0.01	0.12
Variable rate (bin.)	Whether mortgage is variable rate	A25A3	0.05	0.22
Refinanced (bin.)	Whether mortgage has ever been refinanced	A23B	0.31	0.46
Self-employed (bin.)	Whether head is self-employed	BC22	0.11	0.31
Inherited money (cat.)	Whether family unit received inheritance: in survey year; in previous year	W124a		
Poor HtM (bin.)*	Poor HtM: zero or negative illiquid wealth and either i) liquid assets less than half a month's income or ii) liquid debt more than half a month's income		0.20	0.40
Wealthy HtM (bin.)*	Wealthy HtM: positive illiquid wealth and either i) liquid assets less than half a month's income or ii) liquid debt more than half a month's income		0.24	0.43
Education (cat.)	Education of head: less than high school; high school; $\leq$3 years college; $\geq$4 years college			
Female (bin.)	Head is female		0.14	0.35
Race (cat.)	Race of head: white; black; other person of colour	L40		
Life expectancy	Expected remaining lifespan matched from SSA actuarial life tables		35.66	8.82
Region (cat.)	Region: North East; Midwest; South; West			
Family size (cat.)	Number of people in family unit (1 to 6+)			
Number of children (cat.)	Number of children in household (0 to 4+)			
Cohort group (cat.)	Decade of head birth (1940-1980)			
Other dependents (bin.)	Supports a non-household relation	G103	0.14	0.35
Disabled (bin.)	Head has a disability	H2	0.10	0.30
Veteran (bin.)	Head is a veteran	L42	0.13	0.34

All balance sheet variables in $000; * as defined in Kaplan et al. (2014); ˆ split into positive and negative components

Table 5.3: Covariate data summary; dollar variables are in thousands. Source: PSID (2021); most variables are PSID imputed versions; question numbers relate to the 2017 survey.

5.2.5 Covariate Correlations

The complete set of covariates are checked for perfect multicollinearity and very high correlation, with the resulting design matrix, X, full rank.[88] In standard linear regressions, very high correlation between covariates results in unreliable inference of the individual coefficients. Furthermore, the simulation results in Ray and Bhattacharya (2018) show that the SAVS routine performs worse when the design matrix has high correlation (although SAVS still performs well compared to competitor algorithms). That is, it becomes more difficult to identify signal covariates from noise covariates. This results in a set of 88 covariates, plus an intercept.

Correlation heatmaps are used to assess the likely extent of correlation as a source of misclassification of covariates (see Appendix D.1). Variables with the largest correlation are related to each other by construction. There is surprisingly little negative correlation outside of life-cycle variables (Figure D.2). Family size and the number of children show substantial correlation (as would be expected), but there is enough variation to separate them with an average correlation of 0.70 (Figure D.3). The clearest set of correlations are between the level, log-level and income ratio transformations of the balance sheet variables (Figure D.4). The average correlation is 0.78 between the levels and income ratios, 0.60 between the levels and log-levels, and 0.50 between the income ratios and log-levels. These transformations are included separately in Section 5.5.2 to ensure correlation does not drive the results. Overall, the design matrix does not exhibit a high degree of correlation, with an average absolute off-diagonal correlation of 0.07.

[88] Although not strictly necessary in a Bayesian setting, ensuring the design matrix is full rank is crucial for the interpretation of coefficients individually. Any pairwise correlation greater than 0.95 is eliminated, resulting in the log-level of monthly mortgage payments and the income ratio of negative other real estate being dropped.

5.3 Results

5.3.1 Model Selection

The model is estimated with and without shrinkage using 10,000 draws after 10,000 burn in. The PIPs indicate that demographic variables are the most frequently selected drivers of variation in consumption responses, supporting the finding in Chapter 4 (Figure 5.2). Most of the covariates are classified as nulls the majority of the time, with only a few being regularly selected as signals. As expected, Horseshoe shrinkage priors tighten the posterior distributions around zero for the vast majority of the Γ_ν and Γ_μ parameters (detailed results are provided in Appendix D.3). This makes the covariates much more likely to be identified as nulls by the SAVS routine, as seen by comparing the PIPs with and without shrinkage. Using the sparse model criteria (selected on PIPs alone and using shrinkage), the set of covariates is reduced to only a few drivers of heterogeneity in transitory and permanent MPCs. In contrast, the broad model (selected based on PIPs and HPDs from the model without shrinkage) increases this to over a dozen covariates. Notably, both models select some of the same demographic variables as used in the multiple covariate specifications of Chapter 4, but they do not select the balance sheet variables (home equity, business/non-labour income and the average propensity to consume) or age.

Focussing on the sparse model, the drivers of permanent responses are clearly separated from the rest of the covariates, with few intermediate PIPs. The dummy variables for black household head and those born in the 1980s are selected over 99 per cent of the time, while those with four or more years of college and female household heads are selected 84 and 69 per cent. These variables are (on the most part) constant across time, which is particularly striking given that the simulation results suggest this type of variable would be most difficult to identify. Furthermore, they only reflect demographic characteristics. Selection of the transitory drivers is less clear, consistent with the simulation results. Only the intercept and the ratio of other debt to income are selected more than half the time at 63 and 64 per cent, with the 1980s cohort dummy excluded at 41 per cent.

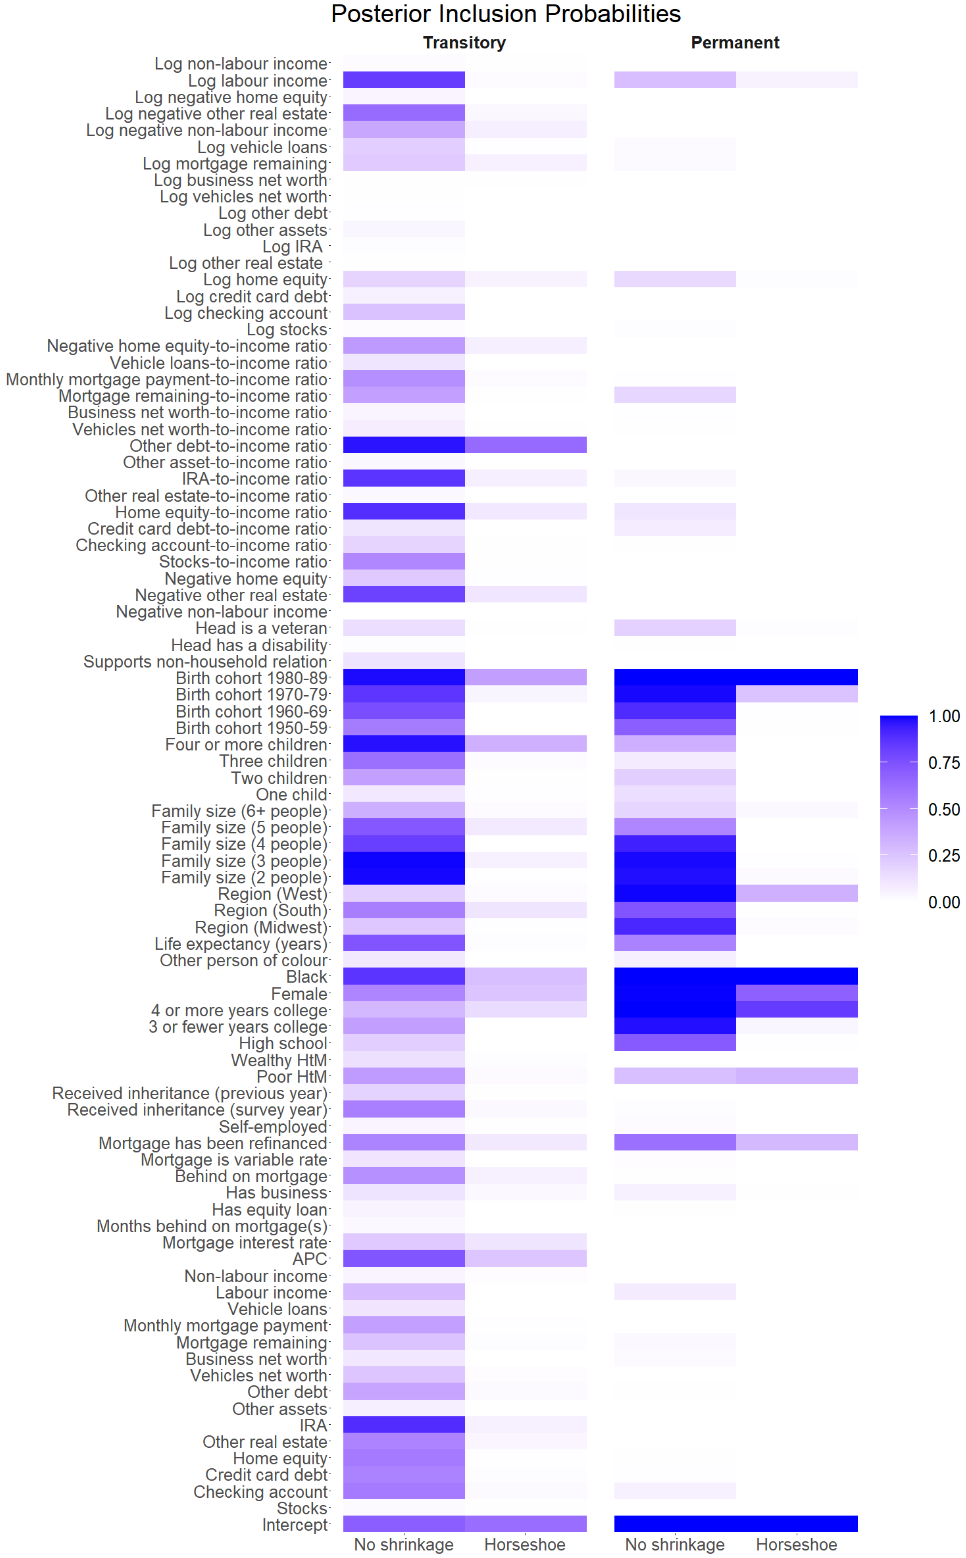

Figure 5.2: PIPs with and without shrinkage.

The specification without shrinkage results in a much smoother distribution of PIPs, making the 0.75 threshold for the broad model somewhat arbitrary (although selection still requires the 99 per cent HPD to exclude zero). Nevertheless, it broadens the selected covariates to include (additional to those identified above): the 1970s cohort, family sizes of 2–4 people, those located in the West region and heads with at least some college for permanent shocks; and the ratios of home equity and individual retirement accounts (IRA) to income, the 1980s cohort, households with four or more children, family sizes of 2–3 people, the IRA balance of households and black household heads for transitory shocks. The transitory intercept is excluded due to the 99 per cent HPD including zero, but is small in any case with a posterior mean of 0.015. Again, the broad model primarily selects demographic covariates.

5.3.2 Driver Effects

The distribution of predicted MPCs varies depending on the selected covariates included, the use of shrinkage and whether dense or SAVS estimates are used for prediction. The simulation results found that the posterior distributions of the dense estimates provide more accurate inference on the signal coefficients than the SAVS estimates. As such, dense estimates from the sparse and broad model are used to gauge the size of effects. Posterior densities of the coefficients can be compared directly, since the continuous covariates have been standardised. However, estimates for the binary variables represent the effect of being in a particular group relative to the reference group, rather than the effect of a one standard deviation change. The posterior distributions of continuous covariates tend to be tight with a small central moment, whereas those of binary covariates are more disperse with large central moments (Figure 5.3 & 5.4). The larger dispersion is in part due to having fewer positive outcomes (i.e. a value of one) in each binary covariate, which limits inference. The share of positive outcomes within covariates ranges between 3 per cent for families with four or more children, to 39 per cent for those with four or more years of college.

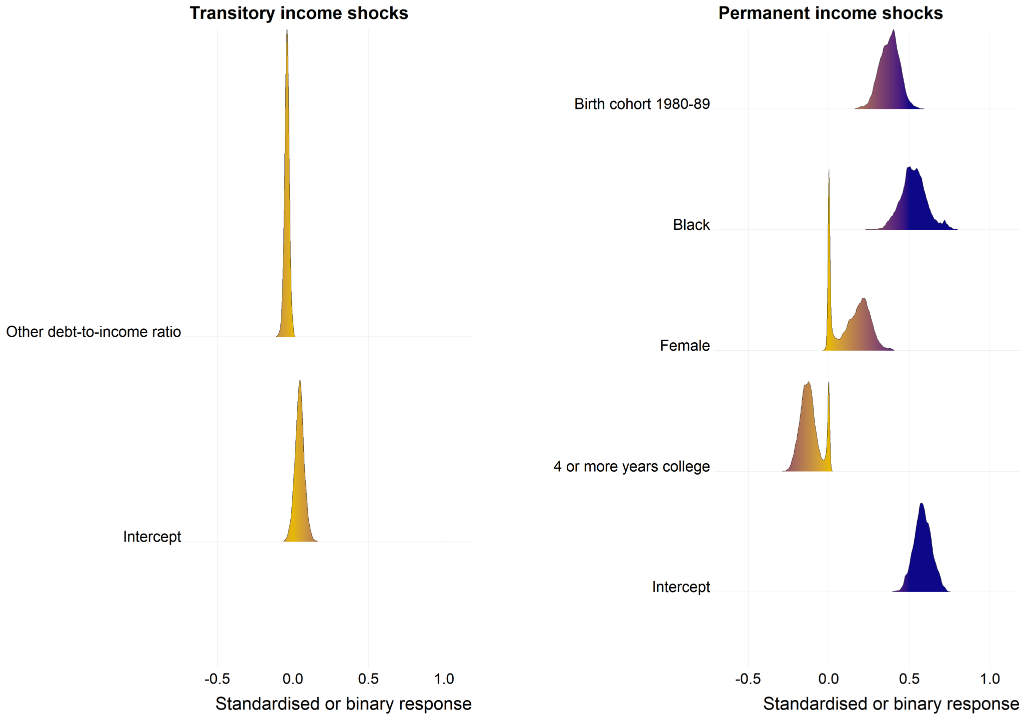

Figure 5.3: Sparse model posterior densities of transitory (left panel) and permanent (right panel) coefficients selected based on a minimum PIP of 0.5; Horseshoe shrinkage.

Sparse Model

The largest effect is found for black household heads, with these households having responses to permanent shocks 0.53 points larger on average than the reference group of white and other heads of colour.[89] The cohort born in the 1980s have permanent responses 0.38 points larger on average than the reference cohort born in the 1940s. Households with female heads have permanent responses 0.14 points larger on average than the male reference group, although this estimate may understate the effect due to shrinkage, as seen in the posterior spike at zero (right panel, Figure 5.3). In contrast, those with four or more years of college have permanent responses 0.12 points lower than households with less than high school education.

[89]The predicted responses (or elasticities) $X\hat{\Gamma}_\nu$ and $X\hat{\Gamma}_\mu$ are unitless since ν^y, μ^y, μ^c and $ln(C)$ are all in log-real-dollar terms.

For transitory income shocks, an additional one standard deviation in the other debt-to-income ratio (0.56 points) lowers the transitory response by 0.04 points. Notably, the income ratio measure of other debt is selected, not the levels or log-levels version, and the relationship is contrary to expectations. Other debt is driven by student loans, which account for 76 per cent of the value of other debt in the sample for years where disaggregated data are available (2011–19). This may suggest a further role for education.

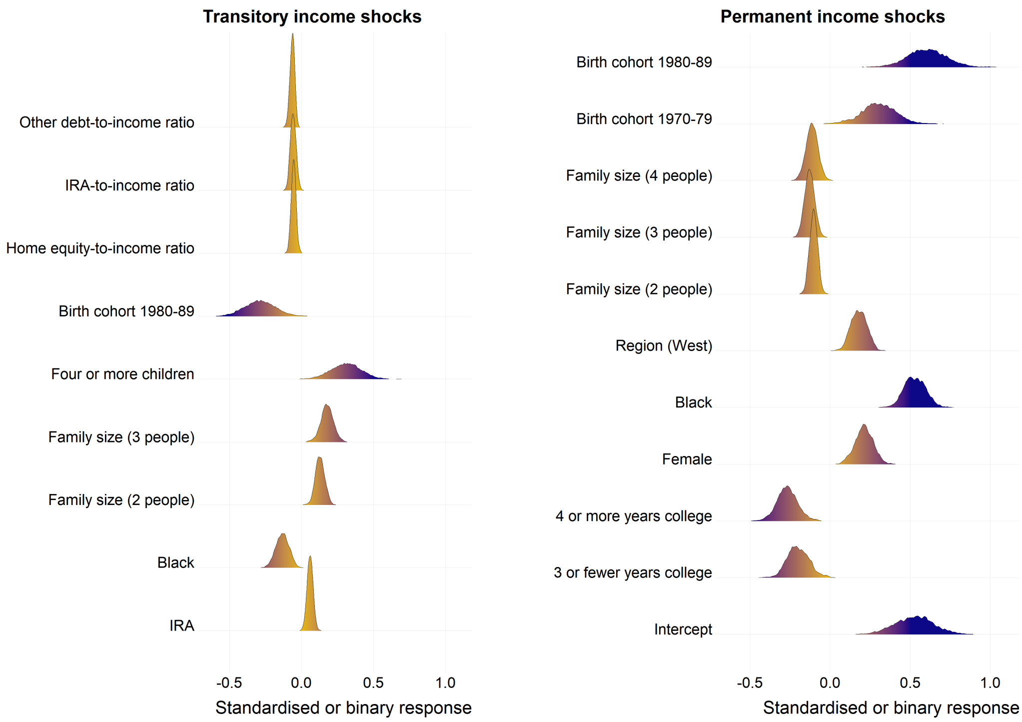

Figure 5.4: Broad model posterior densities of transitory (left panel) and permanent (right panel) coefficients selected based on a minimum PIP of 0.75 and 99 per cent HPD that does not include zero.

Broad Model

The effects of covariates identified in the broad model tend to be larger than those which are also captured in the sparse model, in large part due to the latter's use of shrinkage. In particular, the 1980s cohort now has a permanent elasticity almost twice as large at 0.60 points, with the 1970s cohort also having a large positive response at 0.30 points. However, these covariates have the most disperse posteriors of those selected. The education effect is larger, with all households with at least some college also having lower permanent responses. Family size is found to affect permanent responses. The results are interpreted as a feature of the reference group, single person households, having larger responses of around 0.11 points compared to households with two to four members.

The broad model finds that the home equity-to-income and IRA-to-income ratios have similar effects as that of the other debt-to-income ratio. This could be suggestive of a role for illiquid wealth-to-income, considered in Section 5.5.3; however, given the broad model criteria are more likely to select false positives, and the IRA in levels is selected with the opposite sign, this result is open to contention. The 1980s cohort, single person households and black households have smaller responses to transitory shocks, the opposite direction to their effect on permanent responses.[90]

5.3.3 Variation in MPCs

Along with larger driver coefficients, the broad model exhibits larger overall variation in MPCs, particularly with respect to transitory shocks (Figure 5.5). There is more variation in permanent responses than transitory in both models, consistent with the findings in Chapter 4. Both models feature predicted transitory MPC distributions with substantial mass below zero, suggesting some households have a negative MPC out of transitory shocks (left panel, Figure 5.5).[91] The broad model also features more centered distributions of both MPCs, predicting some households have relatively large transitory responses and some have relatively small permanent responses. This is mostly driven by family composition and education (right panels of Figures 5.7 & 5.8). A smaller response to permanent shocks for families with two or three members is consistent with additional smoothing through labour supply decisions within the household (Blundell, Pistaferri, and Saporta-Eksten (2016)).

[90]The simulations indicate that estimates of transitory and permanent responses to the same covariate are pulled toward each other, if anything, rather than driven apart, so this does not appear to be driven by estimation issues.

[91]This feature is even more pronounced in the joint model in Appendix D.5.

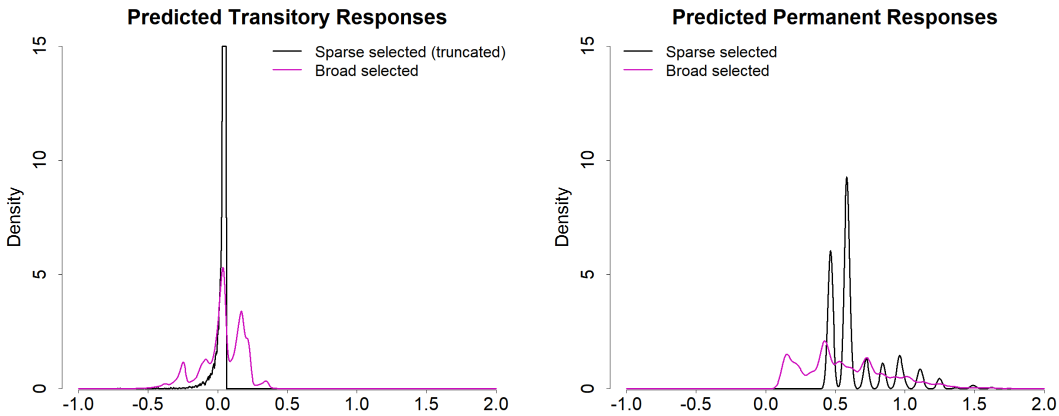

Figure 5.5: Kernel densities of predicted transitory (left panel) and permanent (right panel) responses of the sparse and broad models; selected covariates only, using dense posterior means; the sparse transitory density is truncated as the only selected covariate is other debt-to-income ratio, which is zero for most households.

Comparing the coefficients directly does not provide a complete picture of what drives variation in the predicted responses since the binary variables have not been standardised. That is, the coefficients do not account for the share of the population that fall into each category. Figure 5.6 decomposes the variance of predicted permanent responses into variance and covariance terms that are attributed to each selected covariate. For a given vector of predictions $X\hat{\Gamma}$, the variance is

$$var(X\hat{\Gamma}) = \hat{\Gamma}'\Sigma_X\hat{\Gamma} = \sum_i^M \omega_i \quad \text{for selected covariates } 1, ..., M$$

$$\omega_i = \hat{\Gamma}_i^2\sigma_i^2 + \hat{\Gamma}_i \sum_{j\neq i}^M \hat{\Gamma}_j\sigma_{ij}$$

with σ_{ij} being the ijth element of the covariate unconditional covariance matrix Σ_X.[92] Each ω_i is the contribution of covariate i to the overall variance of the predictions. Within each ω_i, the first term is the contribution from the covariate variance, and the second term that of the covariance.

[92]Note that this is a deterministic decomposition of the prediction since the true MPCs are unobserved. In addition, the coefficients are estimated conditional on the covariance of the full set of covariates, whereas here only the selected covariates are accounted for.

Both sparse and broad models only select demographic drivers of permanent MPCs, supporting the finding of Chapter 4 that demographics drive most of the variation in MPCs. The decade birth cohorts and the black household head dummy drive most of the variation, with sex and education also contributing (Figure 5.6). The positive covariance contributions reflect the tendency for black households to have comparatively more female heads and the negative covariance contributions in the broad model are largely by construction of the dummy variables. Family composition and location play a minor role in MPC heterogeneity in the broad model.

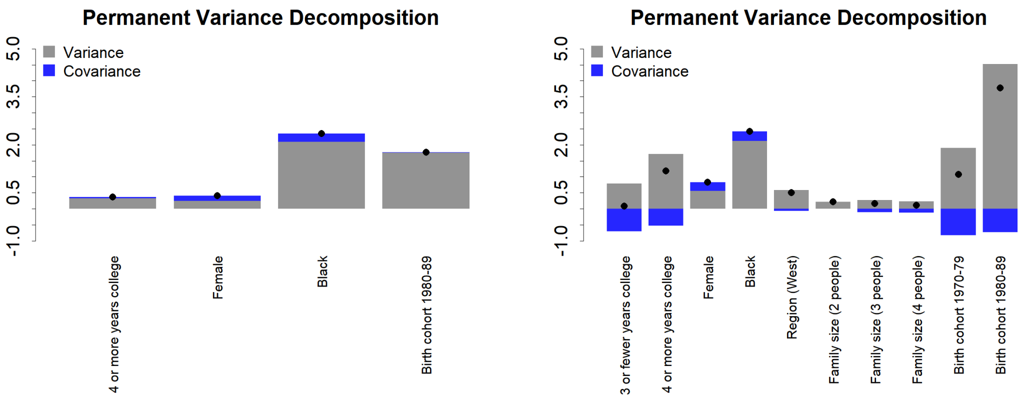

Figure 5.6: Variance decomposition of predicted permanent responses for the sparse model (left panel) and broad model (right panel); bars have been multiplied by 100; predictions use posterior means of dense estimates.

Figure 5.7 decomposes the predicted permanent responses for each observation into the contributions from each selected covariate (excluding the intercept). This provides insight into how household characteristics vary across the distribution of MPCs. As would be expected, the 1980s birth cohort and black households drive much of the positive side of the distributions, with the long positive tails reflecting households with compounding characteristics. In contrast, the 1970s birth cohort is more evenly distributed in the broad model.

Variation in transitory MPCs is shown only for the broad model (Figure 5.8), since the sparse model only selects a single covariate. The 1980s cohort again drives

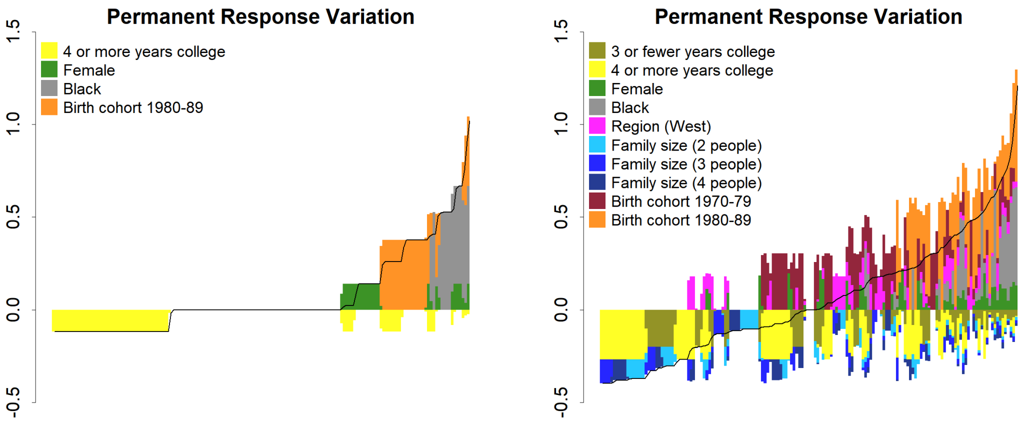

Figure 5.7: Variation in predicted permanent responses for the sparse model (left panel) and broad model (right panel) driven by each selected covariate (excluding intercepts); bars are average contributions over 150 observations, ordered by the total response size; predictions use posterior means of dense estimates.

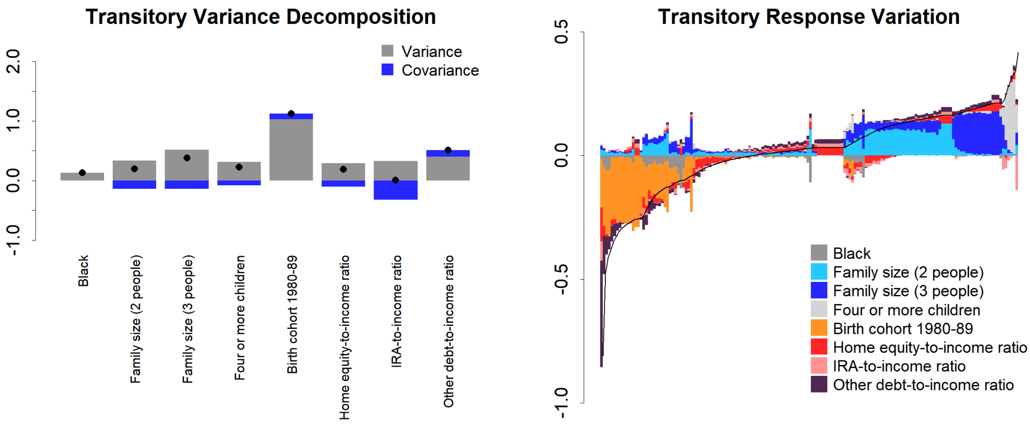

Figure 5.8: Transitory response variance decomposition (left panel) and ordered predicted variation (right panel) for the broad model; left panel: bars have been multiplied by 100; right panel: bars are average contributions over 150 observations, ordered by the total response size; predictions use posterior means of dense estimates.

most of the variation, but reduces responsiveness to transitory shocks. The long negative tail of the distribution is driven by the other debt-to-income ratio and may reflect the presence of outliers rather than genuine heterogeneity. Family composition drives most of the positive side of the distribution. Large families tend to have smaller permanent MPCs – thus more consumption smoothing – but larger transitory MPCs. The smaller response to permanent shocks is consistent with Blundell et al. (2016), who find that family labour supply plays a large role in smoothing permanent wage shocks.

5.4 Demographics and Consumption Smoothing

Following the Blundell et al. (2008) interpretation of one minus the permanent response coefficient as a measure of partial insurance, the results indicate that black, female and the 1980s demographic groups have much less partial insurance against permanent shocks. That is, they are much less able or willing to smooth consumption relative to their reference groups. These characteristics are (typically) fixed across time; differences could be attributable to innate preferences or persistent socio-economic divisions. There is a substantial literature documenting income, wealth and consumption inequality across racial, gender and education characteristics (for example, Oliver et al. (2006), Hardy et al. (2018), Kuhn et al. (2020) and Bhutta et al. (2020) on race; Ruel and Hauser (2013), Kassenboehmer and Sinning (2014) and Blau and Kahn (2017) on gender; and Gregorio and Lee (2002), Boshara et al. (2015), Lindley and Machin (2016) and Bartscher et al. (2020) on education). However, there is comparatively little on households' consumption smoothing.

Ganong et al. (2020) and De Giorgi et al. (2020) are notable exceptions that address racial differences in consumption smoothing in the US. Both papers find that black households are substantially more responsive than white households to income shocks, but that differences in wealth explain almost all of the gap. They rationalise their results in consumption models by varying parameters external to households – interest rates, life expectancy and borrowing constraints (i.e. not preferences) – across black and white households. This approach is consistent with studies that find differing access to credit markets by race (e.g. Harkness (2016), Bayer, Ferreira, and Ross (2018), Bartlett, Morse, Stanton, and Wallace (2019)). Thus, race becomes a proxy for financial constraints in the model, which provide the underlying mechanism. In contrast, the result in this paper finds a larger response for black households *conditional on a large set of financial covariates*. That is, measures of wealth and financial constraints do not explain the gap.[93]

A core strand of the consumption smoothing literature posits that preferences are crucial to understanding the empirical evidence (Carroll et al. (2017), Aguiar et al. (2020), Gelman (2021)). The results in this paper *could* be consistent with such an interpretation; however, the available evidence for gender differences in preferences tends to point in the opposite direction. That is, women tend to be more patient (higher discount rate) and have greater risk aversion (translating to higher consumption smoothing under CRRA utility), indicating they should be *less* responsive to income shocks (Croson and Gneezy (2009), Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2011), Castillo, Ferraro, Jordan, and Petrie (2011), Gränsmark (2012)). Nevertheless, Aguiar et al. (2020) find that differences in the intertemporal elasticity of substitution (distinct from risk aversion under Epstein-Zin utility) are crucial determinants of MPCs. There does not appear to be sufficient evidence on whether the intertemporal elasticity of substitution varies by gender or race to support or refute this explanation.

The results are consistent with financial literacy being an underlying driver of households' partial insurance (see Lusardi and Mitchell (2014) for a survey of the literature). College education can be interpreted directly through this lens as education correlates strongly with financial literacy (Duca and Kumar (2014) use education as a proxy for financial literacy in the PSID). There is a body of evidence that black and female survey respondents have lower levels of financial literacy than the white/male reference group (e.g. Lusardi and Mitchell (2011a), Cupák, Fessler, Schneebaum, and Silgoner (2018), Angrisani, Barrera, Blanco, and Contreras (2020)). Financial literacy also increases with age towards retirement, although may decline thereafter (Agarwal, Driscoll, Gabaix, and Laibson (2009), Lusardi and Mitchell (2011b), van Rooij, Lusardi, and Alessie (2011), Xiao, Chen, and Sun (2015)).[94]

Low financial literacy has been linked to various undesirable economic outcomes

[93] The set of covariates specifically includes mortgage interest rates reported by each household, although life expectancy is not matched on race.

[94] Although not directly equivalent to cohort groups, age will play a similar role. The 1980s cohort are not yet middle aged over the PSID sample.

and behaviours, such as more costly debt, lower participation in financial markets, lower wealth accumulation and less planning for retirement (see Lusardi and Mitchell (2014) and Stolper and Walter (2017); Behrman, Mitchell, Soo, and Bravo (2012) find a plausibly causal connection with wealth accumulation).[95] Lusardi et al. (2017) present a stochastic life-cycle model where endogenous financial knowledge plays a key role in generating wealth inequality. In the model, more highly educated households – with larger incomes and less access to social security – are incentivised to invest more in financial education to smooth lifetime marginal utility (because they face potentially larger fluctuations in income). This provides a potential mechanism to rationalise the empirical findings of differing demographic responses to permanent income shocks, which is considered further in the next section.

5.4.1 Testing Financial Literacy

Financial literacy offers a testable mechanism to explain the results. Furthermore, the financial-literature life-cycle (FLLC) model of Lusardi et al. (2017) offers a theoretical benchmark that can be used to validate the estimation strategy. In the FLLC model, financial literacy is costly to acquire but facilitates access to a sophisticated investment technology with higher expected returns. Thus financial literacy affects household consumption and savings decisions primarily through financial markets, leaving it unclear whether household financial characteristics could provide sufficient statistics. In this section, I test whether the empirical model selects financial literacy as a driver of MPCs in i) data simulated from the FLLC model and ii) PSID data incorporating a measure of financial literacy from the 2016 PSID supplemental Wellbeing and Daily Life survey. Financial literacy is found to affect consumption smoothing, but it does not eliminate the effects of demographics.

[95] Related to consumption smoothing, Lührmann, Serra-Garcia, and Winter (2018) conduct a randomised field experiment on adolescents to test whether a financial education program affects intertemporal choices. They find that financial education increases participants' sophistication and understanding of intertemporal choice, resulting in them behaving in a manner that is more consistent with standard consumption theory. However, the results are mixed in terms of the implications for consumption smoothing under uninsurable income risk; the treatment decreases both present bias and intertemporal smoothing of the experimental payments.

Validation on Simulated Data

The FLLC model incorporates a range of features including stochastic income, asset returns and out-of-pocket medical expenditures, as well as social insurance guaranteeing a consumption floor and self-funded retirement. Crucially, the income process is similar to that used in the empirical model; however, it varies across three education groups (less than high school, high school and at least some college). Log household income is given by

$$ln(Y_{e,i,t}) = g_e(t) + \mu^y_{i,t} + \nu^y_{i,t}$$
$$\mu^y_{i,t} = \rho_e \mu^y_{i,t-1} + \zeta^y_{i,t}$$

where e is the education group, $g_e(t)$ is a deterministic quadratic age profile, ρ_e is the autoregressive coefficient on the persistent component of income, and $\nu^y_{i,t}$ and $\zeta^y_{i,t}$ are transitory and permanent income shocks as previously given. I use the same FLLC model simulations as in the baseline results of Lusardi et al. (2017), which set all transitory income shocks to zero and further allow the variance of the permanent income shock to vary with the education group

$$\zeta^y_{i,t} \sim \mathcal{N}(0, \sigma^2_{e,\zeta y}).$$

The process differs to the empirical model in terms of the persistence of shocks and the education-specific variance of shocks. Persistence is set to $\rho_e = 0.955$ for less than high school and college households and $\rho_e = 0.946$ for high school households, which is not a large deviation from the assumption of $\rho = 1$ in the empirical model. The shock standard deviations are set to $\sigma_{e,\zeta y} = 0.182$ for less than high school, $\sigma_{e,\zeta y} = 0.158$ for high school, and $\sigma_{e,\zeta y} = 0.126$ for college households. The lack of transitory income shocks means that only permanent MPCs can be tested, which have the strongest results in the empirical model in any case.

The dataset contains simulations of 5,000 households over their life cycle on an annual basis, but due to the numerical burden of the MPC estimation routine I randomly drop 60 per cent of households leaving a sample of 2,008 households.

I include households aged 25 to 63 and drop every second year of data to replicate the biennial nature of the PSID, resulting in 35,609 observations. Income and consumption are stripped of predictable variation using year dummies interacted with education dummies and a social insurance dummy indicating whether households received a transfer.[96] This is equivalent to the first-stage life-cycle regressions in the empirical model.

The FLLC model uses a continuous indicator of financial knowledge that households acquire. I split this into four categories – very low, low, typical and high (with high being the reference group) – according to the shares observed in the proxy measure from the 2016 PSID supplemental survey (8.9, 20.4, 40.8 and 29.9 per cent; see below). Age, education dummies, lagged labour income, lagged APC (consumption over labour income) and beginning of period household wealth are also included as covariates. Estimation assumes all transitory income shocks (and transitory MPCs) are zero, but still allows for transitory consumption shocks, and does not use shrinkage.

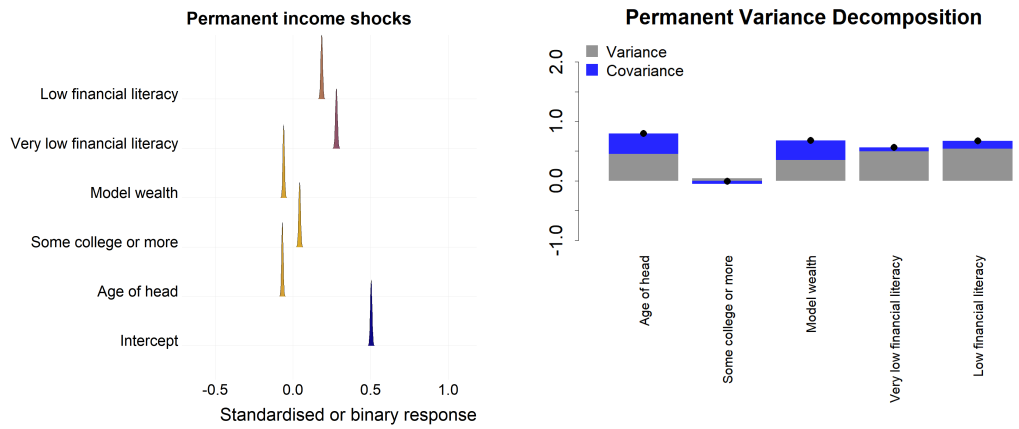

Figure 5.9: Kernel densities of predicted permanent responses (left panel) and permanent response variance decomposition driven by each selected covariate (right panel) for the simulated dataset of Lusardi et al. (2017).

The posterior densities are extremely tight for all parameters, with none of the 90 per cent HPDs including zero (Table D.3). Nonetheless, the PIPs clearly identify a selection of main drivers. Very low and low financial literacy dummies,

[96]Years are equivalent to age in the simulation data, so this removes the age profile of income.

age and wealth are all selected 100 per cent of the time, with the dummy for at least some college selected 80 per cent of the time. The financial literacy dummies have the largest coefficients and together drive much of the variation in MPCs (Figure 5.9). The signs of all covariate coefficients are as expected, with the exception of those with some college, who have larger MPCs than the reference group of households with less than high school.[97] The exercise validates the ability of the empirical framework to identify key theoretical drivers of consumption smoothing, and tightens the connection between financial literacy and MPCs.

Financial Literacy in the PSID Data

The 2016 PSID supplemental Wellbeing and Daily Life survey collected information from PSID households regarding wellbeing, personality traits and every day skills. I follow Bialowolski et al. (2021) and use three of these questions to construct a proxy measure of financial literacy for the PSID. The three questions target understanding of probability, division and compound interest, and require respondents to provide numerical answers. The questions are different to those typically used in the financial literacy literature, which target understanding of compound interest, inflation and diversification (e.g. Lusardi and Mitchell (2011a)). Bialowolski et al. (2021) interpret their measure as an 'instrument' for financial literature related to the numeracal foundations of financial literacy.[98] The questions are:

1. If the chance of getting a disease is 10 per cent, how many people out of 1,000 would be expected to get the disease?
2. If 5 people all have the winning numbers in the lottery and the prize is $2 million, how much will each of them get?
3. Suppose you have $200 in a savings account. The account earns 10 per cent interest each year. How much would you have in the account at the end of two years?

---

[97] This drives little of the overall variation. The variance of unexplained consumption over income is largest for this group in the simulated data, despite having smaller variances of each component individually. This may be due to the larger variance of asset returns for college educated households, which could be resulting in a larger MPC. These features are not found in the PSID data.

[98] I refer to the measure as a proxy as it is not used as an instrument in the econometric sense.

The proxy measure is the sum of correct answers, ranging from zero to three. Of all household heads, 8.9 per cent recorded zero, 20.4 per cent recorded one, 40.8 per cent recorded two and 29.9 per cent recorded three correct answers. I only use responses provided by the household head and match these to households in the 2015 survey year of the main sample used in this thesis, which results in 1,748 matched households. All unmatched households are dropped. As there is only one observation of the proxy per household, it is assumed to be constant across time within the household. The financial literacy proxy is split into dummy variables (with the reference category being those who answered all questions correctly) and included in the set of household covariates.

	Transitory		Permanent	
	None	Horseshoe	None	Horseshoe
Very low financial literacy	0.41	0.01	0.98	0.15
Low financial literacy	0.14	0.04	0.90	0.02
Typical financial literacy	0.22	0.04	0.54	0.00
4 or more years college	0.47	0.11	0.96	0.22
Female	0.49	0.05	0.97	0.67
Black	0.80	0.07	1.00	1.00
Birth cohort 1980-89	0.97	0.46	1.00	1.00

Table 5.4: PIPs of key covariates using the financial literacy subsample.

Estimation using an intercept only model results in posterior statistics similar to those of the full sample (Table D.4). Using the expanded set of covariates, PIPs show that the financial literacy proxies are frequently selected as signals when no shrinkage is applied, but this is severely reduced when using Horseshoe shrinkage (Table 5.4). Due to the reduced number of observations, the selection criteria for the broad model is loosened to 95 per cent of the HPD not including zero (down from 99 per cent), but retaining the 0.75 PIP threshold.[99] Under both sparse and broad model criteria, the 1980s cohort, sex and black covariates continue to be selected for permanent shocks. Out of the financial literacy dummies, none are

[99]The reduction of the HPD criterion results in the inclusion of families with three and four members, the Midwest region, those with four or more years of college, and the low financial literacy group for permanent shocks.

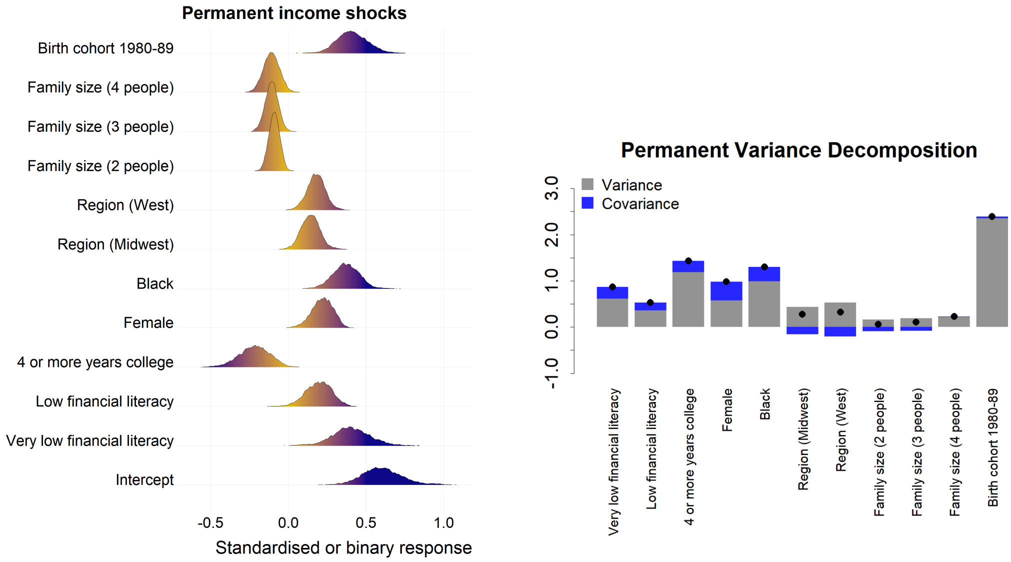

Figure 5.10: Kernel densities of predicted permanent responses (left panel) and permanent response variance decomposition (right panel) for the broad model; decomposition bars have been multiplied by 100; financial literacy subsample.

selected under the sparse criteria, but the low and very low categories (one and zero questions correct) are selected under the broad model criteria for permanent shocks. The posterior densities show a pattern that is consistent with financial literacy decreasing households' MPCs out of permanent income shocks, although the densities are quite disperse (left panel, Figure 5.10). In the broad model, most of the variation in permanent MPCs is still driven by the 1980s cohort, black and female households, but the financial literacy proxies make a substantial contribution, accounting for 16 per cent of the variation (right panel, Figure 5.10).

Overall, these results suggest financial literacy plays a role in household consumption smoothing, and one which appears to be distinct from broader measures of education. However, financial literacy does not eliminate the results for the 1980s cohort, black, and female headed households. That is, the larger MPCs out of permanent income shocks of these households is not primarily driven by financial literacy. This would be consistent with Darity and Hamilton (2017), who argue that financial outcomes like the racial wealth gap are due to structural socioeconomic and political barriers rather than underinvestment in education or financial literacy.

5.5 Robustness

5.5.1 Shock Variance Heterogeneity

The model efficiently captures heterogeneity in MPCs under the assumption of common shock variances across the sample; however, if households face different shocks, the common variance restriction can lead to misleading inference on the MPCs. Estimating the model on separate subsamples allows the variances of shocks to differ across household groups. Using a simple intercept only model with no shrinkage, there is evidence that this is the case for the 1980s cohort and female households with respect to permanent MPCs (left column, Figure 5.11). That is, the posterior densities of the permanent MPCs for these two groups overlap and have similar means as those of the remainder of the sample. Instead, the difference is found in the larger variance of permanent income shocks $\sigma^2_{\zeta y}$ and transitory consumption shocks $\sigma^2_{\nu c}$. A larger variance of permanent income shocks allows for a smaller MPC to generate the same amount of variation in consumption. There is some evidence that black households retain larger, and college graduates smaller, permanent MPCs; although, the posterior densities show non-negligible overlap. There are clear differences in transitory MPCs for female households and college graduates. Subsampling suggests that a model which also allows for group-specific shock variances would be preferred.

The estimated variance of transitory consumption shocks show the starkest difference for all four demographic groups. The inclusion of this shock is motivated by measurement error in Blundell et al. (2008), which does not offer an economic rationale as to why the variance would differ across groups. However, Commault (2021) provides evidence that unexplained consumption deviates from a random walk, implying that transitory consumption shocks have meaningful variation.[100] Her 'robust' estimator permits consumption to respond to past income shocks – a feature also found by Kaplan and Violante (2010) in simulated data from a life-cycle model when households are near their credit constraint. The model and estimation technique used in this thesis does not capture consumption responses

[100] She finds $cov(\Delta c_{i,t}, \Delta y_{i,t+2}) \neq 0$, which would violate moment conditions of a random walk consumption model when the MA(1) parameter lies between zero and one (as it does).

to past income shocks, but could be adapted to do so (indeed, this is the approach taken by Cho et al. (2021)). This may provide some insight into differences in transitory consumption shocks across groups.

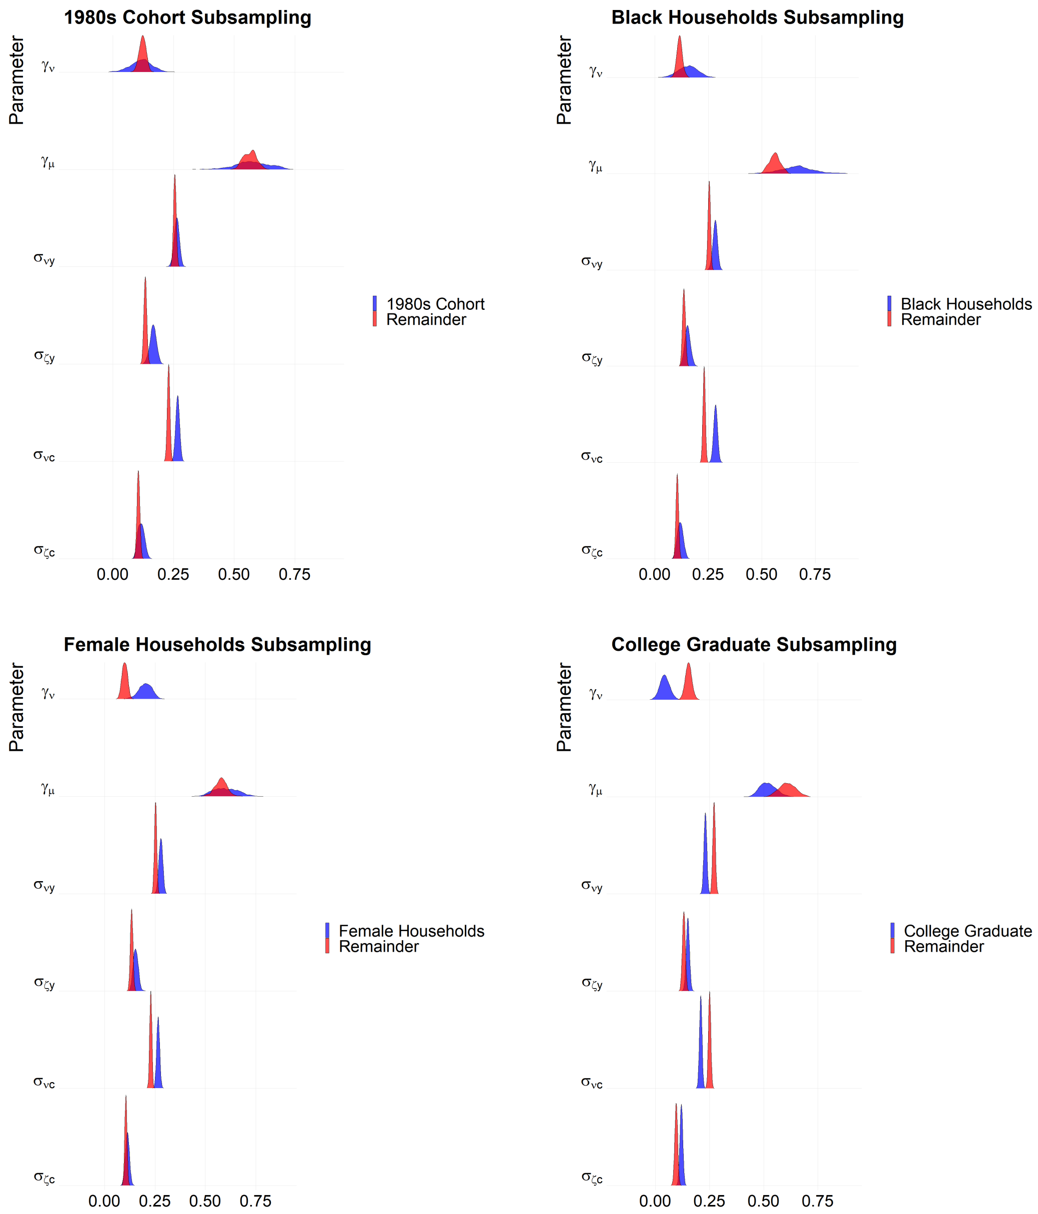

Figure 5.11: Kernel densities of intercept only model parameters using subsampling.

5.5.2 Balance Sheet Transformations

To test whether the inclusion of all three transformations of balance sheet variables (level, log-levels and income ratios) may be impeding the selection of these covariates due to correlation, estimation is rerun using three subsamples of each transformation separately (and retaining all other covariates). PIPs estimates (under shrinkage) for permanent drivers remain very stable across the three specifications, with the 1980s cohort, sex, black and college covariates always selected. The other debt-to-income ratio becomes a marginal null for transitory responses, with a PIP of 0.43. The APC and the poor hand-to-mouth (HtM) dummy are now selected at the 0.5 PIP threshold for transitory and permanent responses, respectively, in some subsamples (Table 5.5). Both variables correlate with the transformation of home equity present in the subsamples where they are not selected, perhaps suggesting a common underlying driver.[101] Nevertheless, the posterior mean estimates indicate that the effect on variation in MPCs is small. The results provide further evidence that balance sheet characteristics are not the main drivers of MPC heterogeneity.

	APC (transitory)		Poor HtM (permanent)	
	PIP	Mean	PIP	Mean
Baseline (all)	0.24	-0.025	0.32	0.047
Levels only	0.62	-0.038	0.61	0.063
Log levels only	0.65	-0.038	0.26	0.046
Income ratios only	0.25	-0.027	0.51	0.059

Table 5.5: PIPs and posterior means using balance sheet subsamples.

5.5.3 Liquid and Illiquid Wealth

The lack of a role for balance sheet variables might be driven by other restrictions of the specification, in particular the linear relationship assumed between covariates and MPCs and the disaggregation of balance sheet covariates. Some

[101] That is, the APC correlates with the home equity-to-income ratio and the poor HtM dummy (negatively) correlates with log home equity (Table D.4).

of the empirical literature finds that larger holdings of liquid assets reduce household MPCs (Jappelli and Pistaferri (2014), Bunn et al. (2018), Fagereng et al. (2018)). The two-asset model of Kaplan and Violante (2014) and Kaplan et al. (2018) emphasises the distinction between liquid and illiquid assets in affecting consumption decisions. The model predicts nonlinear relationships between these variables and MPCs, which have been proxied by categorical definitions such as the poor hand-to-mouth and wealthy hand-to-mouth covariates used in the baseline specification in this Chapter.

To test the robustness of demographic drivers being more important than these financial frictions, I estimate a targeted specification based on the sparse model (with shrinkage) that includes liquid and illiquid wealth as ratios of labour income.[102] The targeted model includes covariates for those with four or more years of college, female and black household heads, the cohort born in the 1980s and the other debt-to-income ratio. I use three specifications that include liquid and illiquid wealth as ratios of labour income under different functional forms; linear, linear spline and cubic spline models. The spline models include nodes at the 0.5th, 20th, 50th, 70th, 90th and 99th percentiles (with the outer nodes motivated by removing the effects of outliers). These provide flexible specifications to assess liquid and illiquid assets as drivers of MPCs, controlling for key demographic features.

The PIPs indicate that liquid and illiquid wealth do not play a statistically meaningful role in MPC heterogeneity (Table 5.6). The reported PIPs average over all liquid and illiquid wealth terms, but there is also no individual coefficient PIP for the liquid and illiquid wealth terms larger than 0.1. This is compelling evidence that liquid and illiquid wealth do not drive a large amount of variation in MPCs. This can be seen in Figure 5.12, which plots the (posterior mean) predicted responses using only the liquid and illiquid wealth terms. The predicted effect is small across the full distribution of wealth positions.

[102] The results are very similar using the level of liquid and illiquid wealth.

	Transitory			Permanent		
	Linear	Linear spline	Cubic spline	Linear	Linear spline	Cubic spline
Liquid wealth-to-income*	0.00	0.02	0.00	0.00	0.00	0.00
Illiquid wealth-to-income*	0.00	0.01	0.00	0.00	0.00	0.00
Intercept	1.00	1.00	1.00	1.00	1.00	1.00
4 or more years college	0.24	0.22	0.04	0.92	0.76	0.83
Female	0.33	0.28	0.05	0.92	0.65	0.89
Black	0.50	0.18	0.04	1.00	1.00	1.00
Birth cohort 1980-89	0.43	0.28	0.07	1.00	1.00	1.00
Other debt-to-income ratio	0.48	0.50	0.25	0.00	0.00	0.00

* Average over all terms

Table 5.6: PIPs for sparse model with liquid and illiquid wealth as ratios of labour income; Horseshoe shrinkage.

5.6 Discussion

Overall, the results present a puzzle – what is the mechanism through which race, sex, education and birth cohort affect consumption smoothing? Although there is no smoking gun, there are a few potential explanations that can guide future research. In addition, the results could have important implications for policy, which deserves some attention.

The ability of households to smooth permanent income shocks may be influenced by various factors external to household balance sheets. Familial support through in kind giving could help to insure against adverse shocks. Implicit financial support could also play a role – working age households may factor their parents' wealth into their lifetime consumption profile. The pattern of inheritance provides some circumstantial evidence in favour of this. Black and female household heads are less likely to receive inheritance in the sample, whereas those with four or more years of college are more likely (with the differences statistically significant). Dummies for the year in which households receive inheritance are included in the sparsification analysis, but this only captures a contemporaneous effect at the time of transfer, rather than inclusion in expectations of lifetime income.

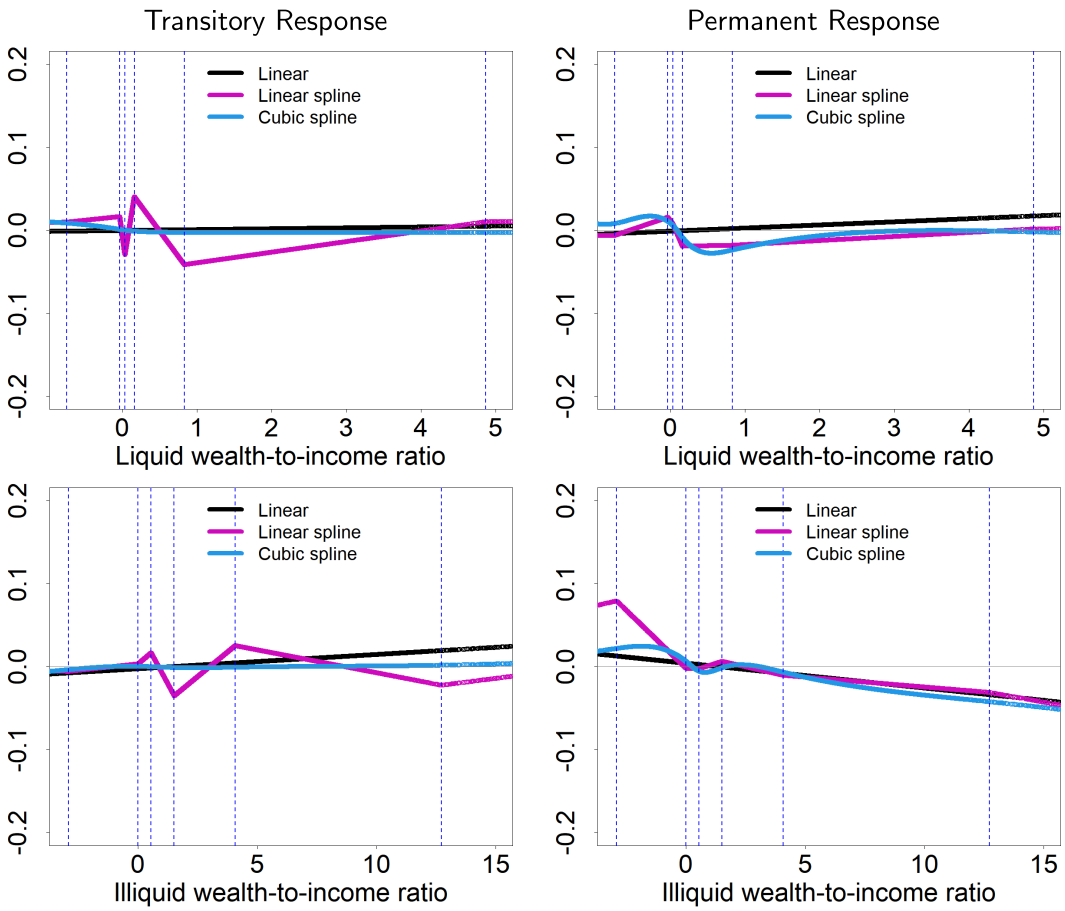

Figure 5.12: Predicted variation in transitory (left column) and permanent (right column) responses driven by liquid wealth (top row) and illiquid wealth (bottom row) using posterior means of dense estimates with shrinkage; vertical dashed lines are spline nodes at the 0.5th, 20th, 50th, 70th, 90th and 99th percentiles.

Familial support fails to explain the results for the 1980s cohort. The most likely suspect for this result is a change in the macroeconomic conditions these households face, given their age. The first-stage regressions remove various time-varying components from the level of income and consumption; however, they do not remove cohort-specific time-varying components (means or volatility). The Great Recession is an obvious candidate. Not only did the Great Recession weigh more heavily on the 1980s cohort at the time (e.g. Bell and Blanchflower (2011)), but youth face persistent earnings reductions when entering the workforce in a weak labour market (Schwandt and von Wachter (2019)). The 1980s result could reflect a downgrade to, or increased uncertainty about, lifetime earnings profiles resulting from the Great Recession. Cho et al. (2021) test for structural breaks in MPCs and find that transitory consumption responses to transitory income shocks increased after the Great Recession. They find no break for MPCs out of permanent income shocks, but do not consider splits of the data by age.

As noted previously, the pattern of results aligns with studies highlighting gaps in income, wealth, consumption and financial literacy across race, gender and education. It is also consistent with labour market disadvantage (e.g. Altonji and Blank (1999), Aaronson, Daly, Wascher, and Wilcox (2019)), which has the potential to drive unequal financial outcomes irrespective of how financially savvy a household may be. Labour market disadvantage could affect consumption smoothing via increased uncertainty over lifetime income. The pervasiveness of similar findings across various economic and financial outcomes is indicative of deeper causes, particularly for the results presented in this chapter which control for a large range of potential confounders. I find the narrative of systemic socioeconomic barriers being the root cause of the findings for race and sex to be compelling. That is, discrimination satisfies Occam's razor. A large body of evidence supports labour market discrimination in particular (e.g. Darity, Dietrich, and Guilkey (2001), Bertrand and Mullainathan (2004), Lang and Manove (2011), Neumark (2018)).

From a statistical point of view, the MPC parameters reflect the transmission of income volatility to consumption. The framework applied in this thesis is limited by the assumption of common variances across the population. In some respects, the results are to be expected given the unconditional moments of the unexpected income and consumption data – the volatility of consumption is systematically larger than the volatility of income for groups with larger MPCs (left panel, Figure 5.13). The results could in part reflect the fact that the MPC is the only degree of freedom to explain this cross-sectional variation. However, the model is surprisingly flexible as the implied *estimates* of the transitory income shocks follow a similar pattern, reducing this concern (right panel, Figure 5.13). Combined with the robustness results by subsample, this suggests that heterogeneity in transitory shock variance of both income and consumption could play an important role.

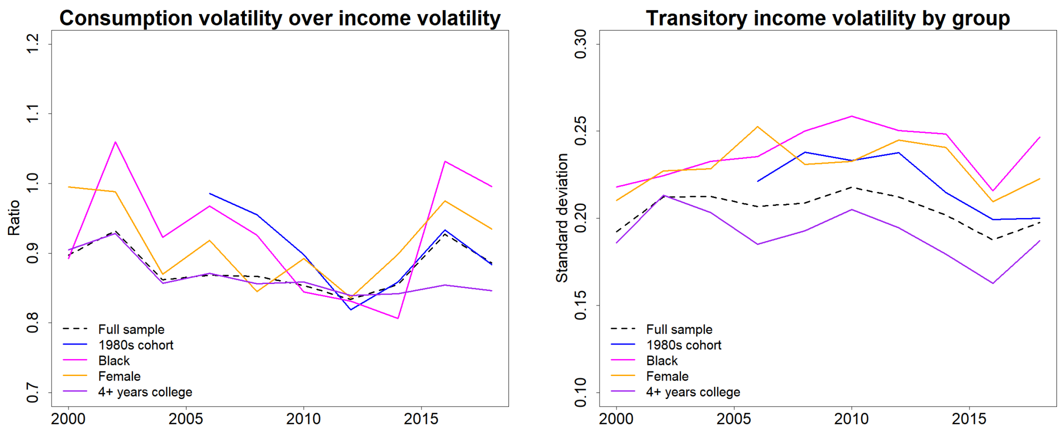

Figure 5.13: The ratio of the standard deviation of unexplained consumption over the standard deviation of unexplained income by year (left panel) and the standard deviation of estimated transitory income shocks backed out using posterior means (right panel).

Ultimately, the strength of policy implications is contingent on identifying the underlying drivers of the results. A crucial finding is that addressing gaps in financial literacy might reduce, but will not eliminate, greater sensitivity to income shocks. Of course, measures to promote financial literacy may well be beneficial for a multitude of other reasons beyond consumption smoothing. A similar takeaway applies to liquidity; however, the benefits of promoting household liquidity beyond improving consumption smoothing are less clear. The results also provide direction for the tax and transfer system. Permanent adjustments to tax incidence or household transfers will have a larger effect on consumption when they fall on black, female, lower educated and younger households.

5.7 Conclusion

I identify race, sex, education and birth cohorts as the key drivers of heterogeneity in household consumption responses to permanent and transitory income shocks. I apply Bayesian shrinkage and sparsification techniques to the empirical model in Chapter 4, which allows MPCs out of transitory and permanent income to vary with observable household characteristics. A rich set of demographic and financial covariates is used, with no *a priori* information regarding which covariates affect MPCs. The Signal Adaptive Variable Selector estimator of Ray and Bhattacharya (2018) is used to select a sparse set of covariates that drive variation in MPCs, with and without a Horseshoe shrinkage prior. The approach is guided by simulation exercises before being applied to a dataset taken from the PSID over 1999–2019.

Demographic factors are found to be the main drivers of heterogeneity in MPCs. Households with heads who are black, female or born in the 1980s have larger responses to permanent income shocks, while those with four or more years of college have lower responses. These four covariates drive the overwhelming majority of heterogeneity in permanent MPCs. Flexible specifications that allow for different functional forms of the relationship between financial variables and MPCs find little role for these covariates. This stands in contrast to much of the focus of the existing literature.

The pattern of results is consistent with financial literacy being an underlying driver of heterogeneity in MPCs. Including a proxy for financial literacy from the 2016 PSID supplemental Wellbeing and Daily Life survey finds that it does contribute to variation in MPCs, but the demographic covariates still drive most of the variation. Allowing shocks to vary across household groups weakens the findings. Overall, the results suggest a future research agenda further interrogating the mechanisms through which demographic groups are faced with unique consumption smoothing dynamics.

Chapter 6

Conclusion

Building on the 'partial insurance' framework of Blundell et al. (2008), I estimate households' marginal propensity to consume (MPC) out of transitory and permanent income shocks using data from the US Panel Study of Income Dynamics over 1999-2019. I use Bayesian methods for estimation, which provide a natural setting for estimating high dimensional models with unobserved components. I extend the model in different directions by allowing for: i) asymmetry in MPCs with respect to the sign of transitory income shocks; ii) observable and unobservable heterogeneity in MPCs; iii) robust aggregation of population statistics; and iv) selection of the main drivers of MPC variation from a rich set of covariates.

I find clear evidence of *positive* asymmetry – larger responses to positive transitory income shocks – for households with substantial home equity. This contrasts the existing survey evidence on 'reported preferences' and is suggestive of a behavioural interpretation. However, evidence of *negative* asymmetry for so called 'poor hand-to-mouth' households (who have low liquidity and low wealth) is consistent with the Kaplan-Violante two-asset model. Households with heads aged 55–64 are found to exhibit negative asymmetry, irrespective of their financial position, which is interpreted as a savings preference before retirement.

Under the assumption of symmetry, I look for evidence of variation in MPCs that would accord with theoretical predictions. Distributions of MPCs along balance sheet positions show some evidence of financial frictions, but demographic characteristics account for most of the variation in MPCs. I estimate *unobserved* MPC heterogeneity and find that it is large for transitory MPCs, but that observed variation explains much of the heterogeneity for permanent shocks. US population-representative dollar-for-dollar MPCs are carefully aggregated from microdata elasticities using an extended dataset that includes retired households. Households increase permanent nondurables consumption by 7 cents per dollar for transitory income shocks, and 38 cents for permanent shocks.

Using a rich set of demographic and financial covariates, Bayesian shrinkage and sparsification techniques are applied to identify the key drivers of MPC heterogeneity. Demographic traits are found to be the main drivers of heterogeneity in MPCs. Households with heads who are black, female or born in the 1980s have larger responses to permanent income shocks, while those with four or more years of college have lower responses. The pattern of results is consistent with financial literacy being an underlying mechanism; however, incorporating a proxy for financial literacy does not remove the demographic effects.

This thesis also suggests direction for further research. Life-cycle models that capture behavioural consumption and savings traits may better fit the data and help to generate meaningful variation in transitory consumption. In addition, more comprehensive modelling of demographic differences across the life cycle will help to determine whether financial characteristics drive outcomes, or whether social and structural barriers continue to impede consumption smoothing for certain groups. Finally, the results in this thesis will likely contain some universal aspects, along with features unique to the US; applying the methods developed here to household panel surveys from other countries would help to paint a more colourful picture of household behaviour.

Overall, this thesis presents a nuanced view of how different households respond to income shocks and how well the life-cycle literature can explain the results. Through detailed modelling of the distribution of consumption decisions, policy institutions can design more targeted interventions and develop a deeper understanding of their effects. The estimates from this thesis could be used to guide such endeavours.

Bibliography

Aaronson, S., M. C. Daly, W. L. Wascher, and D. W. Wilcox (2019). Okun Revisited: Who Benefits Most from a Strong Economy. Finance and Economics Discussion Series 2019-072, Board of Governors of the Federal Reserve System (U.S.).

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2017). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. NBER Working Papers 23732, National Bureau of Economic Research, Inc.

Agarwal, S., J. C. Driscoll, X. Gabaix, and D. Laibson (2009). The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation. *Brookings Papers on Economic Activity* 40(2 (Fall)), 51–117.

Agarwal, S. and W. Qian (2017). Access to Home Equity and Consumption: Evidence from a Policy Experiment. *The Review of Economics and Statistics* 99(1), 40–52.

Aguiar, M. A., M. Bils, and C. Boar (2020). Who Are the Hand-to-Mouth? NBER Working Papers 26643, National Bureau of Economic Research, Inc.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics* 109(3), 659–684.

Aizenman, J. (1998). Buffer stocks and precautionary savings with loss aversion. *Journal of International Money and Finance* 17(6), 931 – 947.

Altonji, J. G. and R. M. Blank (1999). Chapter 48 race and gender in the labor market. *Handbook of Labor Economics*, volume 3, 3143–3259. Elsevier.

Angeletos, G.-M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001). The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. *Journal of Economic Perspectives* 15(3), 47–68.

Angrisani, M., S. Barrera, L. R. Blanco, and S. Contreras (2020). The racial/ethnic gap in financial literacy in the population and by income. *Contemporary Economic Policy* 1–13.

Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica* 85(3), 693–734.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review* 109(6).

Baker, S. R. (2018). Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data. *Journal of Political Economy* 126(4), 1504–1557.

Ballantyne, A. (2020). Household Consumption: MPC Asymmetry and Financial Frictions. Working paper.

Banks, J., R. Blundell, and T. Tanner (1995). Is there a retirement-savings puzzle? IFS Working Papers W95/04, Institute for Fiscal Studies.

Bartlett, R., A. Morse, R. Stanton, and N. Wallace (2019). Consumer-Lending Discrimination in the FinTech Era. NBER Working Papers 25943, National Bureau of Economic Research, Inc.

Bartscher, A. K., M. Kuhn, and M. Schularick (2020). The College Wealth Divide: Education and Inequality in America, 1956-2016. *Review* 102(1), 19–49.

Baugh, B., I. Ben-David, H. Park, and J. A. Parker (2021). Asymmetric consumption smoothing. *American Economic Review* 111(1), 192–230.

Bauwens, L., M. Lubrano, and J.-F. Richard (2000). *Bayesian Inference in Dynamic Econometric Models*. Number 9780198773139 in OUP Catalogue. Oxford University Press.

Bayer, P., F. Ferreira, and S. L. Ross (2018). What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders. *Review of Financial Studies* 31(1), 175–205.

Behrman, J. R., O. S. Mitchell, C. K. Soo, and D. Bravo (2012). How financial literacy affects household wealth accumulation. *American Economic Review* 102(3), 300–304.

Bell, D. N. F. and D. G. Blanchflower (2011). Young people and the Great Recession. *Oxford Review of Economic Policy* 27(2), 241–267.

Benito, A. (2009). Who Withdraws Housing Equity and Why? *Economica* 76(301), 51–70.

Bertrand, M. and S. Mullainathan (2001). Do people mean what they say? implications for subjective survey data. *American Economic Review* 91(2), 67–72.

Bertrand, M. and S. Mullainathan (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review* 94(4), 991–1013.

Bhutta, N., A. C. Chang, L. J. Dettling, and J. W. Hsu (2020). Disparities in Wealth by Race and Ethnicity in the 2019 Survey of Consumer Finances. FEDS Notes 2020-09-28-2, Board of Governors of the Federal Reserve System (U.S.).

Bialowolski, P., A. Cwynar, J. J. Xiao, and D. Weziak-Bialowolska (2021). Consumer financial literacy and the efficiency of mortgage-related decisions: New evidence from the panel study of income dynamics. *International Journal of Consumer Studies* (n/a), 1–14.

Blau, D. M. (2008). Retirement and Consumption in a Life Cycle Model. *Journal of Labor Economics* 26(1), 35–71.

Blau, F. D. and L. M. Kahn (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature* 55(3), 789–865.

Blundell, R., H. Low, and I. Preston (2004). Income risk and consumption inequality: a simulation study. IFS Working Papers W04/26, Institute for Fiscal Studies.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.

Blundell, R. and I. Preston (1998). Consumption Inequality and Income Uncertainty. *The Quarterly Journal of Economics* 113(2), 603–640.

Boar, C., D. Gorea, and V. Midrigan (2017). Liquidity Constraints in the U.S. Housing Market. NBER Working Papers 23345, National Bureau of Economic Research, Inc.

Boshara, R., W. R. Emmons, and B. J. Noeth (2015). The Demographics of Wealth - How Age, Education and Race Separate Thrivers from Strugglers in Today's Economy. Essay No. 2: The Role of Education. *Demographics of Wealth* (2), 1–28.

Browning, M. and T. F. Crossley (2001). The Life-Cycle Model of Consumption and Saving. *Journal of Economic Perspectives* 15(3), 3–22.

Browning, M., T. F. Crossley, and G. Weber (2003). Asking Consumption Questions in General Purpose Surveys. *The Economic Journal* 113(491), F540–F567.

Bunn, P., J. Le Roux, K. Reinold, and P. Surico (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics* 96(C), 1–15.

Cagetti, M. (2003). Wealth Accumulation over the Life Cycle and Precautionary Savings. *Journal of Business & Economic Statistics* 21(3), 339–353.

Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics* 112(1), 1–55.

Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives* 15(3), 23–45.

Carroll, C. D., J. Slacalek, K. Tokuoka, and M. N. White (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics* 8(3), 977–1020.

Carvalho, C. M., N. G. Polson, and J. G. Scott (2010). The horseshoe estimator for sparse signals. *Biometrika* 97(2), 465–480.

Castillo, M., P. J. Ferraro, J. L. Jordan, and R. Petrie (2011). The today and tomorrow of kids: Time preferences and educational outcomes of children. *Journal of Public Economics* 95(11), 1377–1385. Special Issue: International Seminar for Public Economics on Normative Tax Theory.

Chan, J. C. and I. Jeliazkov (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation* 1(1-2), 101–120.

Chatterjee, A., J. Morley, and A. Singh (2021). Estimating household consumption insurance. *Journal of Applied Econometrics* .

Chen, A. and H. H. Jensen (1985). Home equity use and the life cycle hypothesis. *Journal of Consumer Affairs* 19(1), 37–56.

Chib, S. and E. Greenberg (1994). Bayes inference in regression models with arma (p, q) errors. *Journal of Econometrics* 64(1), 183 – 206.

Cho, Y., J. Morley, and A. Singh (2021). Marginal propensities to consume before and after the great recession. Working Paper 2019-11, University of Sydney.

Christelis, D., D. Georgarakos, T. Jappelli, L. Pistaferri, and M. van Rooij (2017). Asymmetric consumption effects of transitory income shocks. DNB Working Papers 551, Netherlands Central Bank, Research Department.

Clarida, R. H. (1991). Aggregate Stochastic Implications of the Life Cycle Hypothesis. *The Quarterly Journal of Economics* 106(3), 851–867.

Commault, J. (2021). Does consumption respond to transitory shocks? *American Economic Journal: Macroeconomics forthcoming.*

Crawley, E. and A. Kuchler (2020). Consumption Heterogeneity: Micro Drivers and Macro Implications. Finance and economics discussion series 2020-005, Board of Governors of the Federal Reserve System.

Croson, R. and U. Gneezy (2009). Gender differences in preferences. *Journal of Economic Literature* 47(2), 448–74.

Crossley, T. F. and J. K. Winter (2013). Asking households about expenditures: What have we learned? Working Paper 19543, National Bureau of Economic Research.

Cupák, A., P. Fessler, A. Schneebaum, and M. Silgoner (2018). Decomposing gender gaps in financial literacy: New international evidence. *Economics Letters* 168, 102–106.

Darity, W., J. Dietrich, and D. K. Guilkey (2001). Persistent advantage or disadvantage?: Evidence in support of the intergenerational drag hypothesis. *American Journal of Economics and Sociology* 60(2), 435–470.

Darity, W. A. and D. Hamilton (2017). The Political Economy of Education, Financial Literacy, and the Racial Wealth Gap. *Review* 99(1), 59–76.

Darity, W. A. and P. L. Mason (1998). Evidence on Discrimination in Employment: Codes of Color, Codes of Gender. *Journal of Economic Perspectives* 12(2), 63–90.

Davies, J. B. (1981). Uncertain lifetime, consumption, and dissaving in retirement. *Journal of Political Economy* 89(3), 561–577.

De Giorgi, G., L. Gambetti, and C. Naguib (2020). Life-cycle inequality: Blacks and whites differentials in life expectancy, savings, income, and consumption. CEPR Discussion Papers 15182, C.E.P.R. Discussion Papers.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica* 59(5), 1221–1248.

Deaton, A. (2016). Measuring and understanding behavior, welfare, and poverty. *American Economic Review* 106(6), 1221–43.

Disney, R., J. Gathergood, and A. Henley (2010). House Price Shocks, Negative Equity, and Household Consumption in the United Kingdom. *Journal of the European Economic Association* 8(6), 1179–1207.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner (2011). Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences. *Journal of the European Economic Association* 9(3), 522–550.

Domeij, D. and M. Flodén (2010). Inequality trends in sweden 1978–2004. *Review of Economic Dynamics* 13(1), 179 – 208. Special issue: Cross-Sectional Facts for Macroeconomists.

Duca, J. V. and A. Kumar (2014). Financial literacy and mortgage equity withdrawals. *Journal of Urban Economics* 80, 62–75.

Durbin, J. (2000). The foreman lecture: the state space approach to time series analysis and its potential for official statistics (with discussion). *Australian & New Zealand Journal of Statistics* 42(1), 1–23.

Fagereng, A., M. B. Holm, and G. J. Natvik (2018). MPC heterogeneity and household balance sheets. Working Papers 4, Department of the Treasury, Ministry of the Economy and of Finance.

Fernández-Villaverde, J. and J. F. Rubio-Ramírez (2007). Estimating Macroeconomic Models: A Likelihood Approach. *The Review of Economic Studies* 74(4), 1059–1087.

Ferrari, S. and F. Cribari-Neto (2004). Beta regression for modelling rates and proportions. *Journal of Applied Statistics* 31(7), 799–815.

Fuster, A., G. Kaplan, and B. Zafar (2018). What would you do with $500? Spending responses to gains, losses, news, and loans. Staff Reports 843, Federal Reserve Bank of New York.

Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, and C. Wheat (2020). Wealth, race, and consumption smoothing of typical income shocks. Working Paper 27552, National Bureau of Economic Research.

Gelman, M. (2021). What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics. *Journal of Monetary Economics* 117(C), 521–542.

Gordon, N., D. Salmond, and A. Smith (1993). Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)* 140, 107–113(6).

Gourinchas, P.-O. and J. A. Parker (2002). Consumption Over the Life Cycle. *Econometrica* 70(1), 47–89.

Gregorio, J. D. and J. Lee (2002). Education and Income Inequality: New Evidence From Cross-Country Data. *Review of Income and Wealth* 48(3), 395–416.

Griffiths, W. E., R. G. Drynan, and S. Prakash (1979). Bayesian estimation of a random coefficient model. *Journal of Econometrics* 10(2), 201–220.

Gränsmark, P. (2012). Masters of our time: Impatience and self-control in high-level chess games. *Journal of Economic Behavior & Organization* 82(1), 179–191.

Gross, I. (2019). Anticipated House Purchases and Household Consumption. Working paper, Monash University.

Gunawan, D., A. Panagiotelis, W. Griffiths, and D. Chotikapanich (2020). Bayesian weighted inference from surveys. *Australian & New Zealand Journal of Statistics* 62(1), 71–94.

Hall, R. E. (2009). By How Much Does GDP Rise If the Government Buys More Output? *Brookings Papers on Economic Activity* 40(2 (Fall)), 183–249.

Hall, R. E. and F. S. Mishkin (1982). The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica* 50(2), 461–481.

Hardy, B., J. Morduch, W. Darity, and D. Hamilton (2018). Wealth Inequality, Income Volatility, and Race. Working paper.

Harkness, S. K. (2016). Discrimination in lending markets: Status and the intersections of gender and race. *Social Psychology Quarterly* 79(1), 81–93.

Harvey, A. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Heathcote, J., K. Storesletten, and G. L. Violante (2004). The Cross-Sectional Implications of Rising Wage Inequality in the United States. CEPR Discussion Papers 4296, C.E.P.R. Discussion Papers.

Heathcote, J., K. Storesletten, and G. L. Violante (2009). Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics* 1(1), 319–354.

Heckman, J. J. (2001). Micro data, heterogeneity, and the evaluation of public policy: Nobel lecture. *Journal of Political Economy* 109(4), 673–748.

Herring, C. and L. Henderson (2016). Wealth inequality in black and white: Cultural and structural sources of the racial wealth gap. *Race and Social Problems* 8.

Hsiao, C. and M. H. Pesaran (2004). Random Coefficient Panel Data Models. CESifo Working Paper Series 1233, CESifo.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary Saving and Social Insurance. *Journal of Political Economy* 103(2), 360–399.

Huber, F., G. Koop, and L. Onorante (2020). Inducing sparsity and shrinkage in time-varying parameter models. *Journal of Business & Economic Statistics* 0(0), 1–15.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control* 17(5-6), 953–969.

Hurst, E. (2003). Grasshoppers, ants, and pre-retirement wealth: A test of permanent income. Working Paper 10098, National Bureau of Economic Research.

Hurst, E. (2008). The Retirement of a Consumption Puzzle. NBER Working Papers 13789, National Bureau of Economic Research, Inc.

Hurst, E. and F. Stafford (2004). Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption. *Journal of Money, Credit and Banking* 36(6), 985–1014.

Jappelli, T. and F. Modigliani (1998). The Age-Saving Profile and the Life-Cycle Hypothesis. CSEF Working Papers 09, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics* 2(1), 479–506.

Jappelli, T. and L. Pistaferri (2014). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics* 6(4), 107–136.

Jappelli, T. and L. Pistaferri (2020). Reported mpc and unobserved heterogeneity. *American Economic Journal: Economic Policy* 12(4), 275–97.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review* 96(5), 1589–1610.

Kahneman, D. and A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47(2), 263–291.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to hank. *American Economic Review* 108(3), 697–743.

Kaplan, G. and G. Violante (2016). Wealthy 'hand-to-mouth' households: key to understanding the impacts of fiscal stimulus. *Microeconomic Insights* URL.

Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2(4), 53–87.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica* 82(4), 1199–1239.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity* 45(1 (Spring)), 77–153.

Kassenboehmer, S. C. and M. G. Sinning (2014). Distributional changes in the gender wage gap. *ILR Review* 67(2), 335–361.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *The Review of Economic Studies* 65(3), 361–393.

Kocherlakota, N. R. (2000). Creating business cycles through credit constraints. *Quarterly Review* (Sum), 2–10.

Kovacs, A. and P. Moran (2021). Breaking the Commitment Device: The Effect of Home Equity Withdrawal on Consumption, Saving, and Welfare. Working paper.

Kronmal, R. A. (1993). Spurious correlation and the fallacy of the ratio standard revisited. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 156(3), 379–392.

Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and Wealth Inequality in America, 1949–2016. *Journal of Political Economy* 128(9), 3469–3519.

La Cava, G., H. Hughson, and G. Kaplan (2016). The Household Cash Flow Channel of Monetary Policy. RBA Research Discussion Papers rdp2016-12, Reserve Bank of Australia.

Lafayette, L., G. Sauter, L. Vu, and B. Meade (2016). Spartan Performance and Flexibility: An HPC-Cloud Chimera. URL.

Lang, K. and M. Manove (2011). Education and labor market discrimination. *American Economic Review* 101(4), 1467–96.

Lewis, D. J., D. Melcangi, and L. Pilossoph (2020). Latent Heterogeneity in the Marginal Propensity to Consume. Staff Reports 902, Federal Reserve Bank of New York.

Lindley, J. and S. Machin (2016). The rising postgraduate wage premium. *Economica* 83(330), 281–306.

Love, D. A., M. G. Palumbo, and P. A. Smith (2009). The trajectory of wealth in retirement. *Journal of Public Economics* 93(1-2), 191–208.

Lührmann, M., M. Serra-Garcia, and J. Winter (2018). The impact of financial education on adolescents' intertemporal choices. *American Economic Journal: Economic Policy* 10(3), 309–32.

Lusardi, A., P.-C. Michaud, and O. S. Mitchell (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy* 125(2), 431–477.

Lusardi, A. and O. S. Mitchell (2007). Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics* 54(1), 205–224.

Lusardi, A. and O. S. Mitchell (2011a). Financial Literacy and Planning: Implications for Retirement Wellbeing. NBER Working Papers 17078, National Bureau of Economic Research, Inc.

Lusardi, A. and O. S. Mitchell (2011b). Financial Literacy around the World: An Overview. NBER Working Papers 17107, National Bureau of Economic Research, Inc.

Lusardi, A. and O. S. Mitchell (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature* 52(1), 5–44.

Makalic, E. and D. F. Schmidt (2016). A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters* 23(1), 179–182.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica* 72(1), 1–32.

Mian, A., K. Rao, and A. Sufi (2013). Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics* 128(4), 1687–1726.

Miranda-Pinto, J., D. Murphy, K. Walsh, and E. R. Young (2020). A Model of Expenditure Shocks. Working Papers 202004, Federal Reserve Bank of Cleveland.

Misra, K. and P. Surico (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics* 6(4), 84–106.

Morley, J. C., C. R. Nelson, and E. Zivot (2003). Why are the beveridge-nelson and unobserved-components decompositions of gdp so different? *The Review of Economics and Statistics* 85(2), 235–243.

Nakajima, J. (2018). The role of household debt heterogeneity on consumption: Evidence from Japanese household data. BIS Working Papers 736, Bank for International Settlements.

Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature* 56(3), 799–866.

Oliver, M., T. Shapiro, and T. Shapiro (2006). *Black Wealth, White Wealth: A New Perspective on Racial Inequality*. Routledge library editions: History and philosophy of science. Routledge.

Pagan, A. and T. Robinson (2020). Too many shocks spoil the interpretation. CAMA Working Papers 2020-28, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

Parker, J. A. (2017). Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment. *American Economic Journal: Macroeconomics* 9(4), 153–183.

Parker, J. A. and B. Preston (2005). Precautionary saving and consumption fluctuations. *American Economic Review* 95(4).

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review* 103(6), 2530–53.

Pfeffer, F. T., R. F. Schoeni, A. Kennickell, and P. Andreski (2016). Measuring wealth and wealth inequality: Comparing two U.S. surveys. *Journal of Economic and Social Measurement* (2), 103–120.

Price, F., B. Beckers, and G. La Cava (2019). The Effect of Mortgage Debt on Consumer Spending: Evidence from Household-level Data. RBA Research Discussion Papers rdp2019-06, Reserve Bank of Australia.

Primiceri, G. and T. van Rens (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of Monetary Economics* 56(1), 20–39.

PSID (2021). Panel Study of Income Dynamics public use dataset. Technical report, Produced and distributed by the Institute for Social Research, University of Michigan, Ann Arbor, MI.

Ramey, V. A. (2019). Ten years after the financial crisis: What have we learned from the renaissance in fiscal research? *Journal of Economic Perspectives* 33(2), 89–114.

Ray, P. and A. Bhattacharya (2018). Signal adaptive variable selector for the horseshoe prior.

Ritter, C. and M. A. Tanner (1992). Facilitating the gibbs sampler: The gibbs stopper and the griddy-gibbs sampler. *Journal of the American Statistical Association* 87(419), 861–868.

Ruel, E. and R. M. Hauser (2013). Explaining the gender wealth gap. *Demography* 50(4), 1155–1176.

Schwandt, H. and T. von Wachter (2019). Unlucky cohorts: Estimating the long-term effects of entering the labor market in a recession in large cross-sectional data sets. *Journal of Labor Economics* 37(S1), S161–S198.

Skinner, J. (2007). Are you sure you're saving enough for retirement? NBER Working Papers 12981, National Bureau of Economic Research, Inc.

Sodini, P., S. Van Nieuwerburgh, R. Vestman, and U. von Lilienfeld-Toal (2016). Identifying the benefits from home ownership: A swedish experiment. Working Paper 22882, National Bureau of Economic Research.

Solon, G., S. Haider, and J. Wooldridge (2013). What are we weighting for? NBER Working Papers 18859, National Bureau of Economic Research, Inc.

StataCorp (2019). *Stata 16 Base Reference Manual.* Stata Press.

Stolper, O. A. and A. Walter (2017). Financial literacy, financial advice, and financial behavior. *Journal of Business Economics* 87(5), 581–643.

Swamy, P. A. V. B. and J. S. Mehta (1975). Bayesian and non-bayesian analysis of switching regressions and of random coefficient regression models. *Journal of the American Statistical Association* 70(351a), 593–602.

Telyukova, I. A. (2013). Household Need for Liquidity and the Credit Card Debt Puzzle. *Review of Economic Studies* 80(3), 1148–1177.

van den Heuvel, M., B. Vandermarliere, and K. Schoors (2019). The Asymmetric Response Of Consumption To Income Changes And The Effect Of Liquid Wealth. Working Paper 19/958, Ghent University, Faculty of Economics and Business Administration.

van Rooij, M. C., A. Lusardi, and R. J. Alessie (2011). Financial literacy and retirement planning in the netherlands. *Journal of Economic Psychology* 32(4), 593–608. Financial Capability.

Wong, A. (2018). Transmission of Monetary Policy to Consumption and Population Aging. Working paper.

Xiao, J. J., C. Chen, and L. Sun (2015). Age differences in consumer financial capability. *International Journal of Consumer Studies* 39(4), 387–395.

Zeger, S. L. and M. R. Karim (1991). Generalized linear models with random effects; a gibbs sampling approach. *Journal of the American Statistical Association* 86(413), 79–86.

Appendix A

Appendices to Chapter 2

A.1 Reconciling the BPP Literature

Several similar empirical specifications have been used since the original BPP paper; however, there are substantial differences across them. Key differences relate to i) estimation methods, ii) structure of income and consumption shocks, iii) assumptions of constant or stochastic variance, and iv) sample definitions. This section draws out these differences.

Paper	ν^y model	c model	Variances	γ_ν	γ_μ
Blundell et al. (2008)	MA(1)	ARIMA(0,1,1)	Time-varying	0.053	0.642
Kaplan et al. (2014)*	MA(0)	ARIMA(0,1,0)	Constant	0.064ˆ	0.320†
Chatterjee et al. (2021)	MA(1)	UC model	Time-varying	0.030	0.450
Ballantyne (2021)*	MA(1)	UC model	Constant	0.055	0.486

* Estimation method does not account for missing values; reduced dataset has 10,814 observations from 1,617 households.

ˆ Estimate generated using method in replication files.

† Estimate generated using IV regression of $\Delta c_{i,t}$ on $\Delta y_{i,t}$, instrumented by $\Delta y_{i,t+1} + \Delta y_{i,t} + \Delta y_{i,t-1}$.

Table A.1: Benchmark estimates from the literature using BPP dataset.

A.1.1 Estimation Methods

The literature employs a variety of estimation methods, which can contribute to quite different estimates of MPCs. BPP use minimum-distance generalised method of moments (GMM) estimation, Kaplan et al. (2014) use an instrumental variables (IV) two-stage least-squares method and CMS utilise QMLE estimation via the Kalman filter (under the assumption of Gaussian shocks). Kaplan and Violante (2010) use the same moment conditions as BPP but, working with a very large simulated dataset, they use the moment conditions directly as parameter estimates. This thesis uses Bayesian methods, which share similarities with QMLE. The IV and direct moment methods are partial information, whereas the other techniques estimate all parameters jointly (taking unexplained income and consumption as given).[103] CMS find that QMLE results in more accurate estimates than GMM. Their results indicate that the BPP GMM estimates of the permanent MPC γ_μ are upwardly biased, consistent with the finding in Kaplan and Violante (2010). However, the IV estimates, which use the same moment condition as GMM to identify γ_μ, are substantially lower (Table A.1).

Bayesian estimation is comparable to QMLE when used on the same dataset, but more precise (Table A.2). The additional precision is likely due to the assumption of constant volatility of shocks, whereas the QMLE estimation uses time-varying volatility. The Bayesian estimate of the transitory MPC is a little larger than under QMLE, whereas the permanent MPC is a little smaller, but the parameter estimate intervals (HPD and CI) overlap.

Differences in estimation method can be difficult to disentangle from differences in the model structure, particularly when using IV and moment-based estimation. For both of these methods, identification and estimation rely directly on the model structure, which can result in misleading estimates. For example, the bias found in Kaplan and Violante (2010) is attributed to a violation of the 'short memory' assumption near the borrowing limit; that is, consumption is assumed to *not* respond to lags of income shocks, but it *does* respond in the simulated theoretical model. This is equivalent to (empirical) model misspecification that results in bias in the moment conditions. Specifically, they find that $cov(\Delta c_{i,t}, \nu^y_{i,t-2}) < 0$,

[103] That is, all of the previous literature estimates equations (2.1)–(2.2) separately, even if equations (2.3)–(2.6) are estimated jointly.

	Bayesian (HPD)			QMLE (CI)		
	Mean	5%	95%	Mean	5%	95%
γ_ν	0.111	0.093	0.130	0.085	0.068	0.102
γ_μ	0.533	0.497	0.569	0.554	0.504	0.604
θ	0.087	0.070	0.105	0.088	0.061	0.114
$\sigma_{\nu y}$	0.260	0.257	0.264	0.258	0.243	0.273
$\sigma_{\zeta y}$	0.153	0.148	0.159	0.153	0.131	0.175
$\sigma_{\nu c}$	0.243	0.240	0.245	0.243	0.232	0.254
$\sigma_{\zeta c}$	0.104	0.099	0.108	0.106	0.100	0.112
Observations	29645					
Households	4528					

Table A.2: Benchmark model using Bayesian and QMLE estimation on the main dataset; Bayesian statistics reflect the HPD and QMLE statistics represent the confidence interval (CI); Bayesian statistics reflect 10,000 draws after 10,000 burn-in; QMLE estimation uses the replication code of CMS, with statistics for variances given by the averages over time-varying estimates.

implying a correctly specified model would be

$$\Delta c_{i,t} \simeq \gamma_\mu \zeta^y_{i,t} + \gamma^0_\nu \nu^y_{i,t} + \gamma^1_\nu \nu^y_{i,t-1} + \gamma^2_\nu \nu^y_{i,t-2} + \zeta^c_{i,t} + \Delta \nu^c_{i,t}. \tag{A.1}$$

Identification of γ_μ in BPP and Kaplan and Violante (2010) relies on the moment condition

$$\gamma_\mu = \frac{cov(\Delta c_{i,t}, \Delta y_{i,t+1} + \Delta y_{i,t} + \Delta y_{i,t-1})}{cov(\Delta y_{i,t}, \Delta y_{i,t+1} + \Delta y_{i,t} + \Delta y_{i,t-1})} \;;$$

however, this condition is violated when $cov(\Delta c_{i,t}, \nu^y_{i,t-2}) < 0$. In IV terminology, the instrument $\Delta y_{i,t+1} + \Delta y_{i,t} + \Delta y_{i,t-1}$ fails the exclusion restrictriction. This is a similar point to the focus of Commault (2021), who derives a GMM estimator for γ_ν that is robust to past income shocks affecting consumption.

The same model misspecification under QMLE or Bayesian methods would result in biased estimates of γ_ν or γ_μ to the extent that the implied shocks from the estimation are correlated. That is, suppose the true data generating process were

$$c_{i,t} = \gamma_\mu \mu^y_{i,t} + \mu^c_{i,t} + \nu^c_{i,t} \tag{2.5}$$

$$\mu^c_{i,t} = \mu^c_{i,t-1} + \gamma^0_\nu \nu^y_{i,t} + \gamma^1_\nu \nu^y_{i,t-1} + \gamma^2_\nu \nu^y_{i,t-2} + \zeta^c_{i,t}, \tag{A.2}$$

which is equivalent to equation (A.1), but estimation uses the now misspecified model in equations (2.5)–(2.6). The estimates would likely be biased if $cov(\hat{\zeta}^y_{i,t}, \hat{\nu}^y_{i,t-s}) \neq 0$ or $cov(\hat{\nu}^y_{i,t}, \hat{\nu}^y_{i,t-s}) \neq 0$, where the hat variables are shocks backed out using parameter estimates. Pagan and Robinson (2020) show that this is likely to be the case.

A.1.2 Shock Structure

The specification of income shocks directly affects estimates of the MPC coefficients, which regulate the share of the variance of income shocks that drives consumption (Commault (2021)). This is particularly clear for GMM and IV estimation, as the specification of shocks changes the moment conditions that identify the MPC parameters. Equations (2.3)–(2.6) follow the CMS representation, but they are empirically equivalent (in terms of autocovariance structure) to the reduced-form specifications (equations (2.10)–(2.11)) used to generate results in Tables (6)–(8) of BPP. (See Durbin (2000) or Morley, Nelson, and Zivot (2003) for discussion of such equivalence.) In contrast, Kaplan et al. (2014) make simplifying assumptions. Using the reduced form representation, they omit the transitory consumption shock, so the unexplained consumption model is given by

$$\Delta c_{i,t} \simeq \gamma_\mu \zeta^y_{i,t} + \gamma_\nu \nu^y_{i,t} + \zeta^c_{i,t}.$$

BPP note that the moment conditions that identify the MPC parameters are the same with and without the transitory consumption shock, so in theory this should not affect the point estimates.[104] However, it does change the overall autocovariance structure of the model. The reduced form model becomes an autoregressive integrated moving average ARIMA(0,1,0) model instead of an ARIMA(0,1,1) and so is no longer equivalent to the UC representation. Kaplan et al. (2014) also assume the transitory income shock is given by an MA(0) process, so the unex-

[104] In practice, minimum distance estimation with the full set of moment conditions (as used in BPP) could provide different results to the partial information estimation used in Kaplan et al. (2014).

plained income model is

$$\Delta y_{i,t} \simeq \zeta^y_{i,t} + \Delta \nu^y_{i,t}.$$

Overall, the estimation method and shock structure used in Kaplan et al. (2014) is quite different to the rest of the empirical literature, providing several potential explanations for the notable differences in estimates relative to other methods (Table A.1).

	MA(0)		MA(1)		
Sample	γ_ν	γ_μ	γ_ν	γ_μ	θ
KVW(ext)	0.131	0.448	0.135	0.478	0.087
KVW	0.141	0.431	0.137	0.488	0.088

Table A.3: Comparison of MA(1) assumption using Bayesian estimation.

I estimate specifications with and without the MA(1) term to quantify what bearing this issue has for Bayesian estimation (Table A.3). The MA assumption affects the estimates of MPCs out of permanent income γ_μ much more than those for transitory shocks γ_ν. The results are consistent with Domeij and Flodén (2010), who find that failing to account for the serial correlation results in inflated estimates of the variance of permanent income. This in turn reduces the permanent MPC to attenuate the transmission of additional volatility into consumption. The MA assumption does not appear to have a strong bearing on the transitory MPCs. Simulations also suggest that misspecification due to incorrectly including an MA(1) term when the true data is MA(0) is preferable over incorrectly omitting an MA(1) term when the true data is MA(1). I retain the MA(1) term in my baseline specification, consistent with BPP and CMS.

A.1.3 Constant and Stochastic Variance

Just as the structure of shocks can affect MPC estimates, so too can assumptions regarding the variances of shocks. As the data has both time and cross-section dimensions, there are various choices that can be made in this respect. BPP and

CMS allow the variances of transitory and permanent income shocks, and the variance of transitory consumption shocks, to vary over time. Studies concerned with inequality often take this approach, or go further by allowing group-specific variances (e.g. Blundell and Preston (1998), Heathcote, Storesletten, and Violante (2004), Domeij and Flodén (2010)). In contrast, Kaplan et al. (2014) and this paper take the variances as common over time and households. This choice may have implications for inference on the MPC estimates. I conduct Wald tests using QMLE estimates (of the Huber-White covariance matrix) to test for the constant variance restriction on shocks. Tests are conducted separately for transitory income shocks, permanent income shocks and transitory consumption shocks, using both the BPP dataset and the main dataset used in this thesis. The restriction tested is that the variance at time t equals the variance at time $t+1$ for all $t = 1 : (T-1)$. For the BPP data, estimates of the transitory income shock variance show a clear change in 1985, so Wald tests are also conducted on the sample before and after.[105]

Sample*	Transitory Income	Permanent Income	Transitory Consumption
BPP (1979-1992)	0.003	0.754	0.001
1979-1984	0.940	–	–
1985-1992	0.376	–	–
Ballantyne (2001-2019)	0.000	0.592	0.000

* Note that variances for the first period in both datasets are not estimated.

Table A.4: P-values of Wald tests on time-varying shock volatility; joint restriction that variance at time t equals variance at time $t+1$ for all $t = 1 : (T-1)$.

The results in Table A.4 are mixed. For the BPP data, the constant variance restriction is rejected for transitory consumption shocks, but not for both income shocks (taking into account the series break in 1985). For the new dataset, the restriction cannot be rejected for permanent income shocks only. However, preliminary versions of the new dataset over 1999–2017 (prior to inclusion of the 2019 survey) failed to reject the restricted model at the 10 per cent level for all

[105] 1985 was the first year in which interviews with spouses were conducted, through which spousal income was recorded. Prior to this, spousal income was provided by the reference person (household 'head').

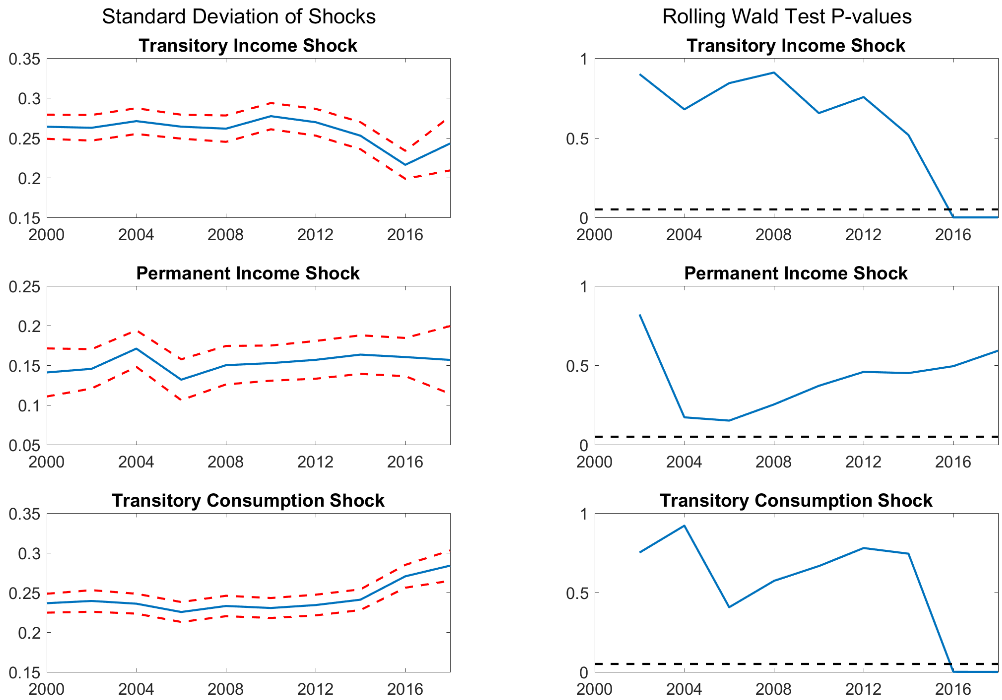

Figure A.1: Standard deviations of shocks over time (left panel) and P-values of rolling Wald tests (right panel); dashed red lines are 2-standard error bands; dashed black line is 5 per cent significance level.

shocks. Rolling Wald tests – testing only the first period restrictions, then adding subsequent periods to the set of restrictions – show that differences in variances are driven by the inclusion of the 2016 and 2018 income years (corresponding to the 2017 and 2019 survey years; right panel, Figure A.1). Overall, the assumption of constant variance shocks used in this thesis is not perfectly supported by the data; however, it is argued that the findings are not driven by time-varying shock variances as MPC estimates for the 2016 and 2018 income years show no clear difference to those of other years (e.g. Figure B.1).

A.1.4 PSID Samples

All of the studies mentioned use data from the PSID, but the different samples used provide a further source of variation across the literature. BPP introduce a novel dataset that spans 1979-1992 by using imputed consumption from the Con-

Sample	γ_ν	γ_μ
BPP	0.055	0.486
Kaplan et al. (2014)	0.134	0.495
Ballantyne	0.111	0.533

Table A.5: Posterior mean estimates of MPC coefficients across different samples using Bayesian estimation.

sumer Expenditure Survey (CEX). This is also used in CMS. In contrast, Kaplan et al. (2014) use PSID data over 1999-2011 as this contains direct measures of consumption. This paper extends the Kaplan et al. (2014) dataset to cover 1999-2019 and contains some differences in filtering the raw data. Differences in estimates across samples, using the same model and estimation technique, can be seen in Table A.5. A comparison of the filtering criteria for the various datasets is provided in Appendix A.4. The BPP dataset clearly suggests smaller responses to transitory shocks, whereas the dataset introduced in this thesis (Ballantyne) features larger responses to permanent shocks. This thesis focusses on the *relative* size of MPCs estimated within a single model, so the overall size of estimates is of less importance.

A.2 Estimation Details

Estimation Setup

The derivation proceeds using the benchmark model in equations (2.3)–(2.6). Stacking the data vertically along both the time and household dimensions of the panel data, the state space model can be written as[106]

$$
\begin{aligned}
\tilde{y} &= \tilde{\mu}^y + F\tilde{\nu}^y + \theta\tilde{\nu}_0^y \\
H\tilde{\mu}^y &= \tilde{\mu}_0^y + \tilde{\zeta}^y \\
\tilde{c} &= \gamma_\mu\tilde{\mu}^y + \tilde{\mu}^c + \tilde{\nu}^c
\end{aligned}
$$

$$H\tilde{\mu}^c = \tilde{\mu}_0^c + \gamma_\nu \tilde{\nu}^y + \tilde{\zeta}^c.$$

where: $\tilde{z} = (z_{1,1}, ..., z_{1,T}, z_{2,1}, ...z_{N,T})'$ for $z = y, \mu^y, \nu^y, \zeta^y, c, \mu^c, \nu^c, \zeta^c$; $H = (I_N \otimes H_T)$, where H_T is the $(T \times T)$ first difference matrix; $\tilde{z}_0 = E_N \ddot{z}_0$ for $\tilde{z}_0 = \tilde{\nu}_0^y, \tilde{\mu}_0^y, \tilde{\mu}_0^c$, where $\ddot{z}_0 = (z_{1,0}, ..., z_{N,0})'$ and E_N is an $NT \times N$ selection matrix constructed by placing $(T-1) \times N$ zeros below each row of I_N; and $F = (I_N \otimes F_T)$ with

$$F_T = \begin{bmatrix} 1 & 0 & 0 & 0 \\ \theta & 1 & 0 & 0 \\ 0 & \theta & 1 & \ddots \\ 0 & 0 & \ddots & \ddots \end{bmatrix}.$$

The data are distributed as conditionally normal

$$\begin{aligned} \tilde{y}|\tilde{\mu}^y, \theta, \tilde{\nu}_0^y, \sigma_{\nu y}^2 &\sim \mathcal{N}(\tilde{\mu}^y - \theta\tilde{\nu}_0^y, \sigma_{\nu y}^2 \Omega) \\ \tilde{c}|\tilde{\mu}^c, \tilde{\mu}^y, \gamma_\mu, \sigma_{\nu c}^2 &\sim \mathcal{N}(\gamma_\mu \tilde{\mu}^y + \tilde{\mu}^c, \sigma_{\nu c}^2 I_{NT}) \end{aligned}$$

where $\Omega = (I_N \otimes \Omega_T)$ with

$$\Omega_T = \begin{bmatrix} 1 & \theta & 0 & 0 \\ \theta & 1+\theta^2 & \theta & 0 \\ 0 & \theta & 1+\theta^2 & \ddots \\ 0 & 0 & \ddots & \ddots \end{bmatrix}$$

which has the property $|\Omega_T| = 1$. The priors for $\tilde{\mu}^y$ and $\tilde{\mu}^c$ are given by

$$\begin{aligned} H\tilde{\mu}^y|\tilde{\mu}_0^y, \sigma_{\zeta y}^2 &\sim \mathcal{N}(\tilde{\mu}_0^y, \sigma_{\zeta y}^2 I_{NT}). \\ H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \tilde{\nu}_0^y, \tilde{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2 &\sim \mathcal{N}(\tilde{\mu}_0^c + \gamma_\nu \tilde{\nu}^y, \sigma_{\zeta c}^2 I_{NT}) \end{aligned}$$

[106]The notation here assumes a balanced panel but is easily adjusted for an unbalanced panel.

where $\tilde{\nu}^y = F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}_0^y]$. The priors for the variance of the shocks are given by inverse-gamma 2 (IG2) functions (see Appendix A in Bauwens et al. (2000))

$$\sigma_{\nu y}^2, \sigma_{\nu c}^2, \sigma_{\zeta y}^2, \sigma_{\zeta c}^2 \sim IG2(a, b)$$

and the priors for $\tilde{\mu}_0^y$, $\tilde{\mu}_0^c$ and $\tilde{\nu}_0^y$ are given in terms of $\ddot{\mu}_0^y$, $\ddot{\mu}_0^c$ and $\ddot{\nu}_0^y$ to eliminate zero elements

$$\begin{aligned} \ddot{\mu}_0^y &\sim \mathcal{N}(0, v_\mu I_N) \\ \ddot{\mu}_0^c &\sim \mathcal{N}(0, v_\mu I_N) \\ \ddot{\nu}_0^y|\sigma_{\nu y}^2 &\sim \mathcal{N}(0, \sigma_{\nu y}^2 I_N) \end{aligned}$$

with variance hyperparameter v_μ. I set the scale $a = 0.15$, shape $b = 0.5$, and variance $v_\mu = 1.4$. The priors on the coefficients γ_μ, γ_ν and θ are given by

$$\gamma_\mu, \gamma_\nu, \theta \sim \mathcal{N}(m_p, v_p)$$

where $m_p = 0.35$ and $v_p = 0.5^2$. Further discussion of priors and hyperparameters is in the main text. The general form of the joint posterior distribution is given by

$$\begin{aligned} p(\Theta|\tilde{y}, \tilde{c}) \propto\ & p(\tilde{y}|\tilde{\mu}^y, \theta, \ddot{\nu}_0^y, \sigma_{\nu y}^2)p(\tilde{c}|\tilde{y}, \tilde{\mu}^c, \tilde{\mu}^y, \gamma_\mu, \sigma_{\nu c}^2) \\ & \times p(H\tilde{\mu}^y|\ddot{\mu}_0^y, \sigma_{\zeta y}^2)p(H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \ddot{\nu}_0^y, \ddot{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2) \\ & \times p(\ddot{\mu}_0^y)p(\ddot{\mu}_0^c)p(\ddot{\nu}_0^y|\sigma_{\nu y}^2) \\ & \times p(\gamma_\mu)p(\gamma_\nu)p(\theta) \\ & \times p(\sigma_{\nu y}^2)p(\sigma_{\nu c}^2)p(\sigma_{\zeta y}^2)p(\sigma_{\zeta c}^2) \end{aligned}$$

Conditional Posteriors – Variances

The conditional posterior for $\sigma^2_{\nu y}$ is given by

$$\begin{aligned} p(\sigma^2_{\nu y}|\ \cdot\ ) &\propto p(\tilde{y}|\tilde{\mu}^y,\theta,\tilde{\nu}^y_0,\sigma^2_{\nu y})p(\sigma^2_{\nu y})p(\ddot{\nu}^y_0|\sigma^2_{\nu y}) \\ &\propto |\sigma^2_{\nu y}\Omega|^{-\frac{1}{2}}exp\left(-\frac{1}{2}(\tilde{y}-\tilde{\mu}^y-\theta\tilde{\nu}^y_0)'(\sigma^2_{\nu y}\Omega)^{-1}(\tilde{y}-\tilde{\mu}^y-\theta\tilde{\nu}^y_0)\right) \\ &\times |\sigma^2_{\nu y}I_N|^{-\frac{1}{2}}exp\left(-\frac{1}{2\sigma^2_{\nu y}}(\ddot{\nu}^y_0)'(\ddot{\nu}^y_0)\right)(\sigma^2_{\nu y})^{-\frac{1}{2}(b+2)}exp\left(-\frac{a}{2\sigma^2_{\nu y}}\right) \end{aligned}$$

which is proportional to the kernel of an IG2 distribution $IG2(a^y_\nu, b_{NNT})$ with $a^y_\nu = (\tilde{y}-\tilde{\mu}^y-\theta\tilde{\nu}^y_0)'(\Omega)^{-1}(\tilde{y}-\tilde{\mu}^y-\theta\tilde{\nu}^y_0)+(\ddot{\nu}^y_0)'(\ddot{\nu}^y_0)+a$ and $b_{NNT}=NT+N+b$.

The conditional posterior for $\sigma^2_{\nu c}$ is given by

$$\begin{aligned} p(\sigma^2_{\nu c}|\ \cdot\ ) &\propto p(\tilde{c}|\tilde{y},\tilde{\mu}^c,\tilde{\mu}^y,\gamma_\mu,\sigma^2_{\nu c})p(\sigma^2_{\nu c}) \\ &\propto (\sigma^2_{\nu c})^{-\frac{NT}{2}}exp\left(-\frac{1}{2\sigma^2_{\nu c}}(\tilde{c}-U^c)'(\tilde{c}-U^c)\right)(\sigma^2_{\nu c})^{-\frac{1}{2}(b+2)}exp\left(-\frac{a}{2\sigma^2_{\nu c}}\right) \end{aligned}$$

which is proportional to the kernel of an IG2 distribution $IG2(a^c_\nu, b_{NT})$ with $a^c_\nu = (\tilde{c}-U^c)'(\tilde{c}-U^c)+a$ and $b_{NT}=NT+b$, where $U^c=\gamma_\mu\tilde{\mu}^y+\tilde{\mu}^c$

The conditional posterior for $\sigma^2_{\zeta y}$ is given by

$$\begin{aligned} p(\sigma^2_{\zeta y}|\ \cdot\ ) &\propto p(H\tilde{\mu}^y|\tilde{\mu}^y_0,\sigma^2_{\zeta y})p(\sigma^2_{\zeta y}) \\ &\propto (\sigma^2_{\zeta y})^{-\frac{NT}{2}}exp\left(-\frac{1}{2\sigma^2_{\zeta y}}(H\tilde{\mu}^y-\tilde{\mu}^y_0)'(H\tilde{\mu}^y-\tilde{\mu}^y_0)\right)(\sigma^2_{\zeta y})^{-\frac{1}{2}(b+2)}exp\left(-\frac{a}{2\sigma^2_{\zeta y}}\right) \end{aligned}$$

which is proportional to the kernel of an IG2 distribution $IG2(a_\zeta^y, b_{NT})$ with $a_\zeta^y = (H\tilde{\mu}^y - \tilde{\mu}_0^y)'(H\tilde{\mu}^y - \tilde{\mu}_0^y) + a$.

The conditional posterior for $\sigma_{\zeta c}^2$ is given by

$$
\begin{aligned}
p(\sigma_{\zeta c}^2 | \;\cdot\;) &\propto p(H\tilde{\mu}^c | \tilde{y}, \tilde{\mu}^y, \tilde{\nu}_0^y, \tilde{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2) \\
&\propto (\sigma_{\zeta c}^2)^{-\frac{NT}{2}} exp\left(-\frac{1}{2\sigma_{\zeta c}^2}(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)'(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)\right) \\
&\times (\sigma_{\zeta c}^2)^{-\frac{1}{2}(b+2)} exp\left(-\frac{a}{2\sigma_{\zeta c}^2}\right)
\end{aligned}
$$

which is proportional to the kernel of an IG2 distribution $IG2(a_\zeta^c, b_{NT})$ with $a_\zeta^c = (H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)'(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y) + a$ and $\tilde{\nu}^y$ is given by $F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta \tilde{\nu}_0^y]$.

Conditional Posteriors – Initial Values

The conditional posterior for $\tilde{\mu}_0^y$ is written in terms of $\ddot{\mu}_0^y$, given by

$$
\begin{aligned}
p(\ddot{\mu}_0^y | \;\cdot\;) &\propto p(H\tilde{\mu}^y | E_N \ddot{\mu}_0^y, \sigma_{\zeta y}^2) p(\ddot{\mu}_0^y) \\
&\propto exp\left(-\frac{1}{2\sigma_{\zeta y}^2}(H\tilde{\mu}^y - E_N \ddot{\mu}_0^y)'(H\tilde{\mu}^y - E_N \ddot{\mu}_0^y)\right) \\
&\times exp\left(-\frac{1}{2v_\mu}(\ddot{\mu}_0^y - \ddot{m}_y)'(\ddot{\mu}_0^y - \ddot{m}_y)\right)
\end{aligned}
$$

which is proportional to the kernel of a normal distribution $\mathcal{N}(p_0^y, \Sigma_0^y)$ with

$$
\Sigma_0^y = \left(\sigma_{\zeta y}^{-2} E_N' E_N + v_\mu^{-1} I_N\right)^{-1}
$$

$$p_0^y = \Sigma_0^y \left(\sigma_{\zeta y}^{-2} E_N' H\tilde{\mu}^y + v_\mu^{-1}\ddot{m}_y\right).$$

The conditional posterior for $\tilde{\mu}_0^c$ is written in terms of $\ddot{\mu}_0^c$, given by

$$\begin{aligned}
p(\ddot{\mu}_0^c| \ \cdot \ ) &\propto p(H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \tilde{\nu}_0^y, E_N\ddot{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2)p(\ddot{\mu}_0^c) \\
&\propto exp\left(-\frac{1}{2\sigma_{\zeta c}^2}(H\tilde{\mu}^c - E_N\ddot{\mu}_0^c - \gamma_\nu\tilde{\nu}^y)'(H\tilde{\mu}^c - E_N\ddot{\mu}_0^c - \gamma_\nu\tilde{\nu}^y)\right) \\
&\times exp\left(-\frac{1}{2v_\mu}(\ddot{\mu}_0^c - \ddot{m}_c)'(\ddot{\mu}_0^c - \ddot{m}_c)\right)
\end{aligned}$$

which is proportional to the kernel of a normal distribution $\mathcal{N}(p_0^c, \Sigma_0^c)$ with

$$\begin{aligned}
\Sigma_0^c &= \left(\sigma_{\zeta c}^{-2} E_N' E_N + v_\mu^{-1} I_N\right)^{-1} \\
p_0^c &= \Sigma_0^c \left(\sigma_{\zeta c}^{-2}[E_N' H\tilde{\mu}^c - \gamma_\nu E_N'\tilde{\nu}^y] + v_\mu^{-1}\ddot{m}_c\right)
\end{aligned}$$

with $\tilde{\nu}^y$ is given by $F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}_0^y]$.

The conditional posterior for $\tilde{\nu}_0^y$ is written in terms of $\ddot{\nu}_0^y$, given by

$$\begin{aligned}
p(\ddot{\nu}_0^y| \ \cdot \ ) &\propto p(\tilde{y}|\tilde{\mu}^y, \theta, E_N\ddot{\nu}_0^y, \sigma_{\nu y}^2)p(H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, E_N\ddot{\nu}_0^y, \tilde{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2)p(\ddot{\nu}_0^y|\sigma_{\nu y}^2) \\
&\propto exp\left(-\frac{1}{2}(\tilde{y} - \tilde{\mu}^y - \theta E_N\ddot{\nu}_0^y)'(\sigma_{\nu y}^2\Omega)^{-1}(\tilde{y} - \tilde{\mu}^y - \theta E_N\ddot{\nu}_0^y)\right) \\
&\times exp\left(-\frac{1}{2\sigma_{\zeta c}^2}(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta E_N\ddot{\nu}_0^y])'\right. \\
&\qquad \left.(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta E_N\ddot{\nu}_0^y])\right) \\
&\times exp\left(-\frac{1}{2\sigma_{\nu y}^2}(\ddot{\nu}_0^y)'(\ddot{\nu}_0^y)\right)
\end{aligned}$$

which is proportional to the kernel of a normal distribution $\mathcal{N}(p_0^\nu, \Sigma_0^\nu)$, and along with the property $(F^{-1})'F^{-1} = (FF')^{-1} = \Omega^{-1}$, gives

$$\Sigma_0^\nu = \left(\theta^2 E_N'\Omega^{-1}E_N[\sigma_{\nu y}^{-2} + \sigma_{\zeta c}^{-2}\gamma_\nu^2] + \sigma_{\nu y}^{-2}I_N\right)^{-1}$$
$$p_0^\nu = \Sigma_0^\nu\left(\theta E_N'\Omega^{-1}[\tilde{y} - \tilde{\mu}^y](\sigma_{\nu y}^{-2} + \sigma_{\zeta c}^{-2}\gamma_\nu^2) + \sigma_{\zeta c}^{-2}\gamma_\nu\theta E_N'(F^{-1})'[\tilde{\mu}_0^c - H\tilde{\mu}^c]\right).$$

Note that the posterior distribution of this parameter vector under the asymmetry model in Chapter 3 is of unknown form. Given it is not a crucial part of estimation – it is often set to zero in the literature – I make a simplifying assumption that the draw of any $\nu_{i,0}^y$ cannot affect the sign of $\tilde{\nu}_{i,1}^y$. That is, the posterior is conditional on $\mathbb{1}^{(\tilde{\nu}^y)\geq 0}$, which is reasonable for small values of θ (otherwise a Metropolis-Hastings algorithm must be used).

Conditional Posterior – Coefficients

The conditional posterior for γ_ν is given by

$$\begin{aligned}
p(\gamma_\nu|\ \cdot\ ) &\propto p(H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \tilde{\mu}_0^c, \tilde{\nu}_0^y, \theta, \gamma_\nu, \sigma_{\zeta c}^2)p(\gamma_\nu) \\
&\propto exp\left(-\frac{1}{2\sigma_{\zeta c}^2}(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu\tilde{\nu}^y)'(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu\tilde{\nu}^y)\right) \\
&\times exp\left(-\frac{1}{2v_p}(\gamma_\nu - m_p)'(\gamma_\nu - m_p)\right)
\end{aligned}$$

which is proportional to the kernel of a normal distribution $\mathcal{N}(p_\gamma^\nu, \Sigma_\gamma^\nu)$ with

$$\Sigma_\gamma^\nu = \left(\sigma_{\zeta c}^{-2}(\tilde{\nu}^y)'\tilde{\nu}^y + v_p^{-1}\right)^{-1}$$
$$p_\gamma^\nu = \Sigma_\gamma^\nu\left(\sigma_{\zeta c}^{-2}(\tilde{\nu}^y)'[H\tilde{\mu}^c - \tilde{\mu}_0^c] + v_p^{-1}m_p\right)$$

with $\tilde{\nu}^y$ is given by $F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}_0^y]$.

The conditional posterior for γ_μ is given by

$$\begin{aligned} p(\gamma_\mu| \;\cdot\;) &\propto p(\tilde{c}|\tilde{y}, \tilde{\mu}^c, \tilde{\mu}^y, \gamma_\mu, \sigma^2_{\nu c})p(\gamma_\mu) \\ &\propto exp\left(-\frac{1}{2\sigma^2_{\nu c}}(\tilde{c} - \gamma_\mu\tilde{\mu}^y - \tilde{\mu}^c)'(\tilde{c} - \gamma_\mu\tilde{\mu}^y - \tilde{\mu}^c)\right) \\ &\times exp\left(-\frac{1}{2v_p}(\gamma_\mu - m_p)'(\gamma_\mu - m_p)\right) \end{aligned}$$

which is proportional to the kernel of a normal distribution $\mathcal{N}(p^\mu_\gamma, \Sigma^\mu_\gamma)$ with

$$\begin{aligned} \Sigma^\mu_\gamma &= \left(\sigma^{-2}_{\nu c}\tilde{\mu}^{y\prime}\tilde{\mu}^y + v_p^{-1}\right)^{-1} \\ p^\mu_\gamma &= \Sigma^\mu_\gamma\left(\sigma^{-2}_{\nu c}[\tilde{\mu}^{y\prime}\tilde{c} - \tilde{\mu}^{y\prime}\tilde{\mu}^c] + v_p^{-1}m_p\right) \end{aligned}$$

The conditional posterior for θ is given by

$$\begin{aligned} p(\theta| \;\cdot\;) &\propto p(\tilde{y}|\tilde{\mu}^y, \theta, \tilde{\nu}^y_0, \sigma^2_{\nu y})p(H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \tilde{\nu}^y_0, \tilde{\mu}^c_0, \theta, \gamma_\nu, \sigma^2_{\zeta c})p(\theta) \\ &\propto |\sigma^2_{\nu y}\Omega|^{-\frac{1}{2}} exp\left(-\frac{1}{2}(\tilde{y} - \tilde{\mu}^y - \theta E_N\ddot{\nu}^y_0)'(\sigma^2_{\nu y}\Omega)^{-1}(\tilde{y} - \tilde{\mu}^y - \theta E_N\ddot{\nu}^y_0)\right) \\ &\times exp\left(-\frac{1}{2\sigma^2_{\zeta c}}(H\tilde{\mu}^c - \tilde{\mu}^c_0 - \gamma_\nu F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}^y_0])'(H\tilde{\mu}^c - \tilde{\mu}^c_0 - \gamma_\nu F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}^y_0])\right) \\ &\times exp\left(-\frac{1}{2v_p}(\theta - m_p)'(\theta - m_p)\right) \end{aligned}$$

which has a distribution of unknown form. The density can be evaluated easily for a given θ so I use a Griddy-Gibbs sampler on the interval (0,1) to ensure invertibility of the variance-covariance matrix.

Conditional Posteriors – Permanent Components

The conditional posterior for $\tilde{\mu}^y$ is given by

$$
\begin{aligned}
p(\tilde{\mu}^y|\ \cdot\ ) &\propto p(\tilde{y}|\tilde{\mu}^y, \theta, \tilde{\nu}_0^y, \sigma_{\nu y}^2)p(\tilde{c}|\tilde{y}, \tilde{\mu}^c, \tilde{\mu}^y, \gamma_\mu, \sigma_{\nu c}^2) \\
&\times p(H\tilde{\mu}^y|\tilde{\mu}_0^y, \sigma_{\zeta y}^2)p(H\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \tilde{\nu}_0^y, \tilde{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2) \\
&\propto exp\left(-\frac{1}{2}(\tilde{y} - \tilde{\mu}^y - \theta E_N \ddot{\nu}_0^y)'(\sigma_{\nu y}^2 \Omega)^{-1}(\tilde{y} - \tilde{\mu}^y - \theta E_N \ddot{\nu}_0^y)\right) \\
&\times exp\left(-\frac{1}{2\sigma_{\nu c}^2}(\tilde{c} - \gamma_\mu \tilde{\mu}^y - \tilde{\mu}^c)'(\tilde{c} - \gamma_\mu \tilde{\mu}^y - \tilde{\mu}^c)\right) \\
&\times exp\left(-\frac{1}{2\sigma_{\zeta y}^2}(H\tilde{\mu}^y - E_N \ddot{\mu}_0^y)'(H\tilde{\mu}^y - E_N \ddot{\mu}_0^y)\right) \\
&\times exp\left(-\frac{1}{2\sigma_{\zeta c}^2}(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)'(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)\right)
\end{aligned}
$$

part of which is the kernel of a normal distribution $\mathcal{N}(p_\mu^y, \Sigma_\mu^y)$ with

$$
\begin{aligned}
\Sigma_\mu^y &= \left((\sigma_{\nu y}^2 \Omega)^{-1} + \sigma_{\nu c}^{-2}\gamma_\mu^2 I_{NT} + \sigma_{\zeta y}^{-2} H'H\right)^{-1} \\
p_\mu^y &= \Sigma_\mu^y \left(\sigma_{\nu y}^{-2}(\Omega^{-1})'[\tilde{y} - \theta E_N \ddot{\nu}_0^y] + \sigma_{\nu c}^{-2}\gamma_\mu[\tilde{c} - \tilde{\mu}^c] + \sigma_{\zeta y}^{-2} H' E_N \ddot{\mu}_0^y\right).
\end{aligned}
$$

So the full conditional posterior can be written as

$$
\begin{aligned}
p(\tilde{\mu}^y|\ \cdot\ ) &\propto exp\left(-\frac{1}{2}(\tilde{\mu}^y - p_\mu^y)'(\Sigma_\mu^y)^{-1}(\tilde{\mu}^y - p_\mu^y)\right) \\
&\times exp\left(-\frac{1}{2\sigma_{\zeta c}^2}(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)'(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)\right)
\end{aligned}
$$

This model is linear and so posterior draws can be constructed using the precision sampler by further deriving the kernel of the posterior. However, the asymmetry model in Chapter 3 replaces $\gamma_\nu \tilde{\nu}^y = \gamma_\nu^+ \tilde{\nu}^{y,+} + \gamma_\nu^- \tilde{\nu}^{y,-}$, where $\tilde{\nu}^{y,+} = \mathbb{1}^{(\tilde{\nu}^y)\geq 0}(F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}_0^y])$ and so the posterior distribution is of an unknown form. I use a Metropolis-Hastings routine for each $\mu_{i,t}^y$ to generate posterior draws.

The conditional posterior for $\tilde{\mu}^c$ is given by

$$
\begin{aligned}
p(\tilde{\mu}^c | \ \cdot \ ) &\propto p(\tilde{c}|\tilde{y}, \tilde{\mu}^c, \tilde{\mu}^y, \gamma_\mu, \sigma_{\nu c}^2) p(H\tilde{\mu}^c | \tilde{y}, \tilde{\mu}^y, \tilde{\nu}_0^y, \tilde{\mu}_0^c, \theta, \gamma_\nu, \sigma_{\zeta c}^2) \\
&\propto exp\left(-\frac{1}{2\sigma_{\nu c}^2} (\tilde{c} - \gamma_\mu \tilde{\mu}^y - \tilde{\mu}^c)'(\tilde{c} - \gamma_\mu \tilde{\mu}^y - \tilde{\mu}^c) \right) \\
&\times exp\left(-\frac{1}{2\sigma_{\zeta c}^2} (H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y)'(H\tilde{\mu}^c - \tilde{\mu}_0^c - \gamma_\nu \tilde{\nu}^y) \right)
\end{aligned}
$$

which is proportional to the kernel of a normal distribution $\mathcal{N}(p_\mu^c, \Sigma_\mu^c)$ with

$$
\begin{aligned}
\Sigma_\mu^c &= \left(\sigma_{\nu c}^{-2} I_{NT} + \sigma_{\zeta c}^{-2} H'H \right)^{-1} \\
p_\mu^c &= \Sigma_\mu^c \left(\sigma_{\nu c}^{-2} [\tilde{c} - \gamma_\mu \tilde{\mu}^y] + \sigma_{\zeta c}^{-2} [H'\tilde{\mu}_0^c + \gamma_\nu H'\tilde{\nu}^y] \right)
\end{aligned}
$$

where $\tilde{\nu}^y$ is given by $F^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\tilde{\nu}_0^y]$.

A.3 Imputation of Credit Card Debt

The broad definition of 'other debt' in the PSID prior to 2011 includes credit card debt, student loans, medical bills, legal bills, family loans, and any other debt (excluding mortgages on the main home and vehicle loans). I use the periods for which the separate components are available to estimate a predictive model of the share of credit card debt in the broader definition of other debt. The distribution of this variable is unusual (left panel Figure A.2), so I use a two stage approach.

The first stage sorts households into two groups; one has extreme values of the credit card share (either zero or one) and the other has credit card debt comprising a fraction of other debt (for example, a large student loan and small credit card balance). This uses a probit model, with the predicted probabilities split between the two groups based on a threshold that matches the prediction average to the sample average. Define the binary variable $D_{i,t}^{E}$ as taking a value of one if the credit card share $CCS_{i,t}$ is either one or zero. The first stage probit model is given by

$$Pr(D_{i,t}^{E} = 1) = \Phi\left(X_{i,t}\beta\right)$$

where $X_{i,t}$ is a vector of covariates, β a corresponding vector of parameters, and $\Phi()$ is the CDF of the standard normal distribution. Covariates comprise log income, employment status, age and its square, family size, race, year (linear term), log other debt, checking and savings balance, dummies for having a vehicle loan and for having other real estate, and the inverse hyperbolic sine of the net value of other real estate.

The second stage is different for each group. The same probit model is used as above for the extreme value group, but with a dependent variable $D_{i,t}^{CC}$, which reflects the extreme value credit card share (i.e. zero or one). A simple 50 per cent probability threshold is used to create binary predictions. For the fractional group, the credit card share $CCS_{i,t}$ is modelled by a beta regression where the conditional mean is a logistic function of the above covariates, and the scale parameter is given by a logarithmic constant.[107] The final imputed values for credit card debt are calculated by multiplying other debt by the predicted credit share for surveys prior to 2011.

[107]For details on the beta regression, see Ferrari and Cribari-Neto (2004) or StataCorp (2019).

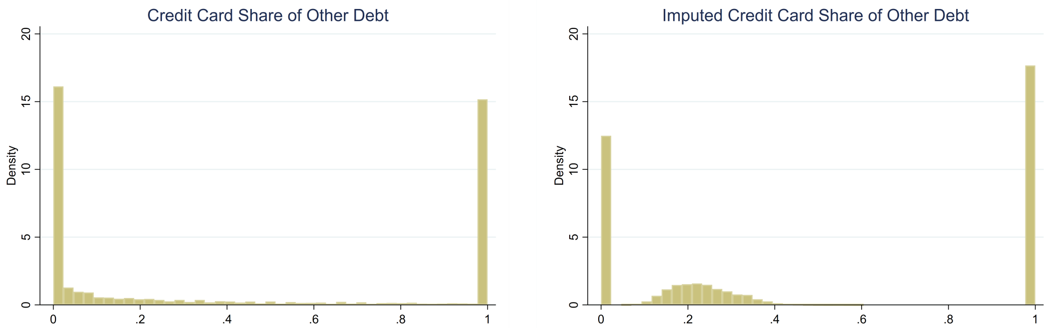

Figure A.2: Histograms of the credit card share of other debt (left panel) and the imputed version (right panel); PSID surveys 2011–17.

A.4 Comparison of Datasets Filtering

	Blundell et al. (2008)	Kaplan et al. (2014)	Cho et al. (2021)	Ballantyne
Years	1979–1992	1999–2011	1999–2017	1999–2019
Minimum T per household	2 (consecutive)	3 (consecutive)	2 (consecutive)	3 (consecutive)
Drop SEO	Yes	Yes	Yes	Yes
Drop unstable composition	Yes	No	Yes	No
Age	30–65	25–55	25–64	25–64
Drop if income level	$< \$0$	$< \$100$ OR growth $> 500\%$ OR growth $< -80\%$	$< \$100$ OR growth $> 500\%$ OR growth $< -80\%$	$< \$100$ OR growth $> 500\%$ OR growth $< -80\%$
Drop zero food expenditure	Yes	No	No	No
Consumption level deflation	None*	None*	CPI components	Headline CPI
Income level deflation	Headline CPI	Headline CPI	Headline CPI	Headline CPI
Drop if unexplained (log) income or consumption	No	No	No	< -2 OR > 2
Drop if liquid wealth	N/A	No	No	$< -\$250,000$ OR $> \$1$ million
Drop if illiquid wealth	N/A	No	No	$< -\$250,000$ OR $> \$2$ million
Tax Adjustment	Survey/Taxsim9	Taxsim9	Taxsim9	Taxsim32

* Unexplained consumption is nonetheless equivalent to deflation with Headline CPI due to the inclusion of time dummies when removing explained components.

Table A.6: Filtering conditions used in various datasets.

Appendix B

Appendices to Chapter 3

B.1 Simulation Analysis Discussion

The simulation results in Tables B.1 and B.2 provide further evidence that the posterior distributions correctly represent the true parameters when groups with different parameters are present in the DGP. Every observation is randomly assigned to one of two groups that differ in γ_ν; that is, MPCs vary across households, and across time within households. The estimation assumes perfect knowledge of the group. Changing this assumption would introduce misclassification error, which would pull the posterior distributions of group-specific MPCs toward each other.

The simulation results in Tables 3.2 and 3.3 show that for some parameters the true value does not align with the middle of the 90 per cent highest posterior density (HPD) interval. With increasing N, the HPD intervals shrink, but estimates do not appear to converge around the true values (Table B.3). From a frequentist perspective, this would suggest the estimators are biased. However, there is some tentative evidence that the mean of the HPD converges to the true value with increasing T; and that it may be driven by sampling volatility since the sign of the 'bias' switches in different simulations (Table B.4).

B.2 Additional Asymmetry Results

Results split by high and low income and MIR terciles do not present compelling evidence with respect to asymmetry. The largest tercile of MIR households shows

	True	HPD Mean	5%	95%
Group 1				
γ_ν	0.300	0.330	0.189	0.467
γ_μ	0.600	0.554	0.428	0.686
Group 2				
γ_ν	-0.100	-0.049	-0.197	0.097
γ_μ	0.600	0.542	0.414	0.675
Common				
θ	0.200	0.157	0.072	0.234
$\sigma^2_{\nu y}$	0.060	0.058	0.050	0.066
$\sigma^2_{\zeta y}$	0.040	0.041	0.033	0.050
$\sigma^2_{\nu c}$	0.120	0.124	0.114	0.135
$\sigma^2_{\zeta c}$	0.040	0.034	0.025	0.044

Table B.1: Benchmark symmetry model with heterogeneous MPCs. Averages over 100 replications. $N = 500, T = 5$.

	True	HPD Mean	5%	95%
Group 1				
γ_ν^+	0.500	0.527	0.363	0.694
γ_ν^-	0.100	0.144	-0.021	0.308
γ_μ	0.600	0.544	0.419	0.672
Group 2				
γ_ν^+	-0.100	-0.028	-0.198	0.139
γ_ν^-	-0.100	-0.043	-0.214	0.126
γ_μ	0.600	0.532	0.407	0.661
Common				
θ	0.200	0.158	0.073	0.235
$\sigma^2_{\nu y}$	0.060	0.058	0.051	0.066
$\sigma^2_{\zeta y}$	0.040	0.041	0.033	0.050
$\sigma^2_{\nu c}$	0.120	0.124	0.114	0.135
$\sigma^2_{\zeta c}$	0.040	0.034	0.024	0.044

Table B.2: Baseline asymmetry model with heterogeneous MPCs. Averages over 100 replications. $N = 500, T = 5$.

	True	Posterior mean 'bias' N=250	N=500	N=1000
γ_ν	0.300	0.028	0.030	0.030
γ_μ	0.600	-0.024	-0.017	-0.017
θ	0.200	-0.044	-0.025	-0.026
$\sigma^2_{\nu y}$	0.060	-0.002	-0.001	0.000
$\sigma^2_{\zeta y}$	0.040	0.003	0.001	0.000
$\sigma^2_{\nu c}$	0.120	0.002	0.004	0.005
$\sigma^2_{\zeta c}$	0.040	-0.005	-0.007	-0.008

Table B.3: Averages over 100 replications. $T = 5$.

	True	Posterior mean 'bias' T=5	T=10	T=20
γ_ν	0.300	0.028	-0.005	-0.009
γ_μ	0.600	-0.024	-0.018	-0.031
θ	0.200	-0.044	-0.017	-0.011
$\sigma^2_{\nu y}$	0.060	-0.002	-0.001	-0.001
$\sigma^2_{\zeta y}$	0.040	0.003	0.003	0.006
$\sigma^2_{\nu c}$	0.120	0.002	-0.001	0.000
$\sigma^2_{\zeta c}$	0.040	-0.005	0.000	0.001

Table B.4: Averages over 100 replications. $N = 250$.

marked positive asymmetry, similar to households with substantial equity. However, this tercile is far from a homogeneous group. Origination of mortgages with MIRs of under four is commonplace and not considered exceptional leverage; this represents over 85 per cent of households in the third tercile of MIRs.[108] However, there are also a number of households in this group with exceptionally high MIRs, who would be expected to behave like credit constrained households. Households with no mortgage display evidence of negative asymmetry. This is similar to households with no home equity, but again the MIR split is a mixed group as it includes households who own their home outright.

Price et al. (2019) and Cho et al. (2021) find that households with high levels of debt are more sensitive to shocks, which is not corroborated by the results here for transitory shocks, but is for permanent shocks. The permanent MPC for the third tercile of MIR households is larger than those in the first tercile. This also marks a difference between high MIR households and non-HtM households; despite similar asymmetry, high MIR households have less partial insurance against permanent income shocks. It suggests that mortgages hinder the ability of households to self insure. Households with low MIR show distinctly larger tendency to self-insure against permanent shocks (smaller MPC), which accords with their much higher average wealth.

Asymmetry is fairly consistent over time for the sample as a whole (Figure B.1). Income year 2002 is the only period in which the 90 per cent highest posterior density (HPD) interval of the difference distribution does not contain zero. The magnitude of asymmetry in 2002 is small compared with the asymmetry driven by balance sheet factors. Overall, this suggests that time-varying shock variance is unlikely to be driving the results.

[108] A common rule of thumb (cited in popular press) for baseline maximum mortgage expenses as a proportion of gross income is 28 per cent. For the 2017 subsample, the average annual income is around \$67,000. The average new mortgage size was \$310,000 in January 2017, with an interest rate of 4.1 per cent for a 30-year term. Such a mortgage results in repayments of around \$18,000 per year, which is about 27 per cent of income for an annual income of \$67,000. This results in an MIR of 4.6.

	Mean	5%	95%
Lower income households			
γ_ν^+	0.111	0.080	0.142
γ_ν^-	0.102	0.075	0.126
γ_μ	0.578	0.529	0.635
(Group obs.	17787	)	
Higher income households			
γ_ν^+	0.122	0.099	0.145
γ_ν^-	0.102	0.066	0.138
γ_μ	0.486	0.442	0.539
(Group obs.	11858	)	

Table B.5: Asymmetry model, split by higher income (top two quintiles) and lower income (bottom three quintiles).

	Mean	5%	95%
No mortgage			
γ_ν^+	0.057	0.029	0.087
γ_ν^-	0.112	0.082	0.145
γ_μ	0.583	0.530	0.625
(Group obs.	12760	)	
First tercile MIR (low: 0–1)			
γ_ν^+	0.119	0.084	0.153
γ_ν^-	0.147	0.101	0.190
γ_μ	0.470	0.427	0.510
(Group obs.	5629	)	
Second tercile MIR (medium: 1–1.75)			
γ_ν^+	0.135	0.095	0.173
γ_ν^-	0.094	0.050	0.137
γ_μ	0.517	0.473	0.556
(Group obs.	5629	)	
Third tercile MIR (high: 1.75+)			
γ_ν^+	0.237	0.190	0.285
γ_ν^-	0.049	0.016	0.085
γ_μ	0.560	0.510	0.600
(Group obs.	5627	)	

Table B.6: Asymmetry model, split by mortgage-to-income ratio terciles.

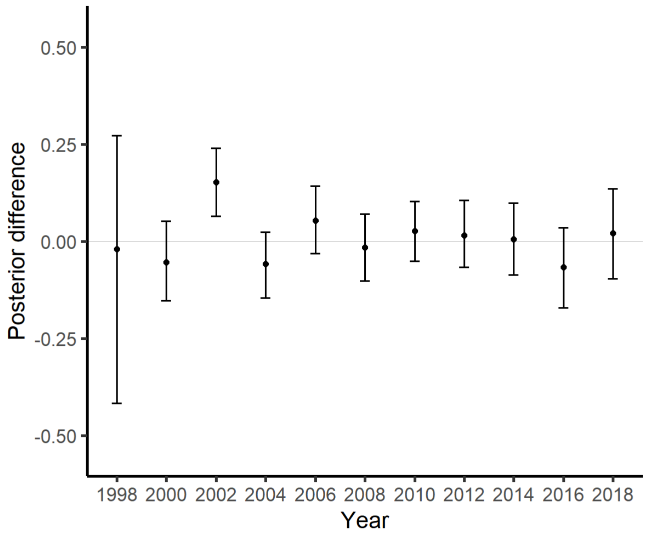

Figure B.1: MPC asymmetry by income year; posterior means of $\gamma_\nu^+ - \gamma_\nu^-$ with 90 per cent HPD interval.

B.3 Diagnostics of Baseline Asymmetry Estimation

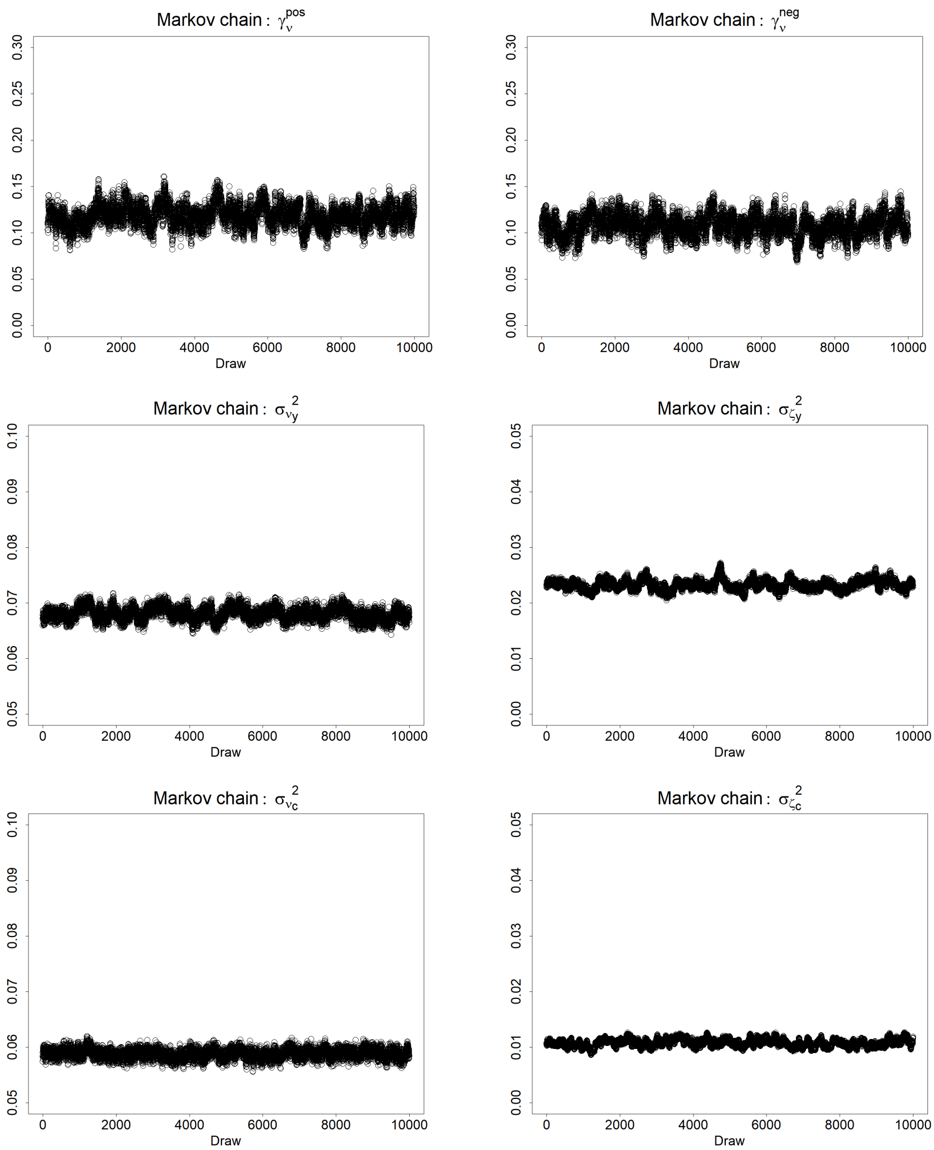

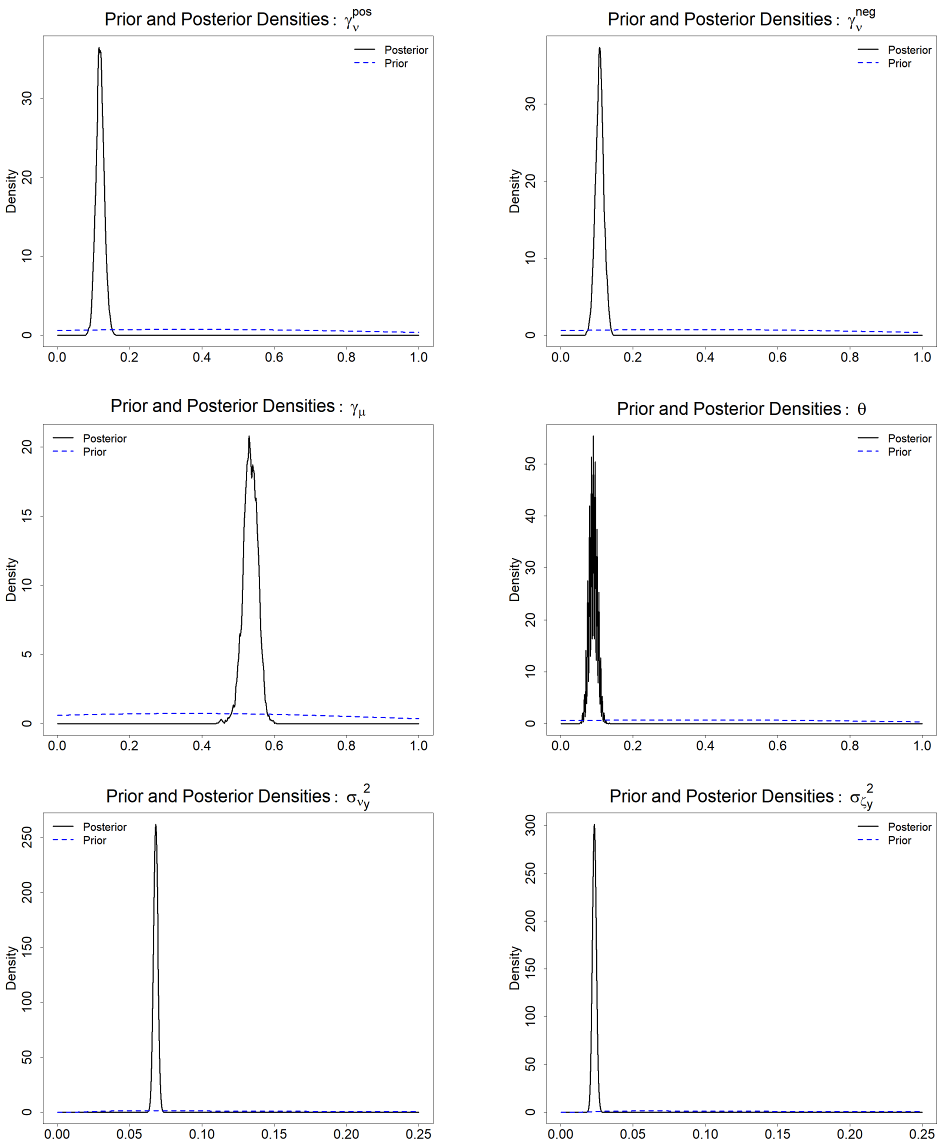
Prior and Posterior Densities : γ_v^{pos}
Prior and Posterior Densities : γ_v^{neg}
Prior and Posterior Densities : γ_μ
Prior and Posterior Densities : θ
Prior and Posterior Densities : $\sigma_{v_y}^2$
Prior and Posterior Densities : $\sigma_{\zeta_y}^2$
Posterior
Prior
Density

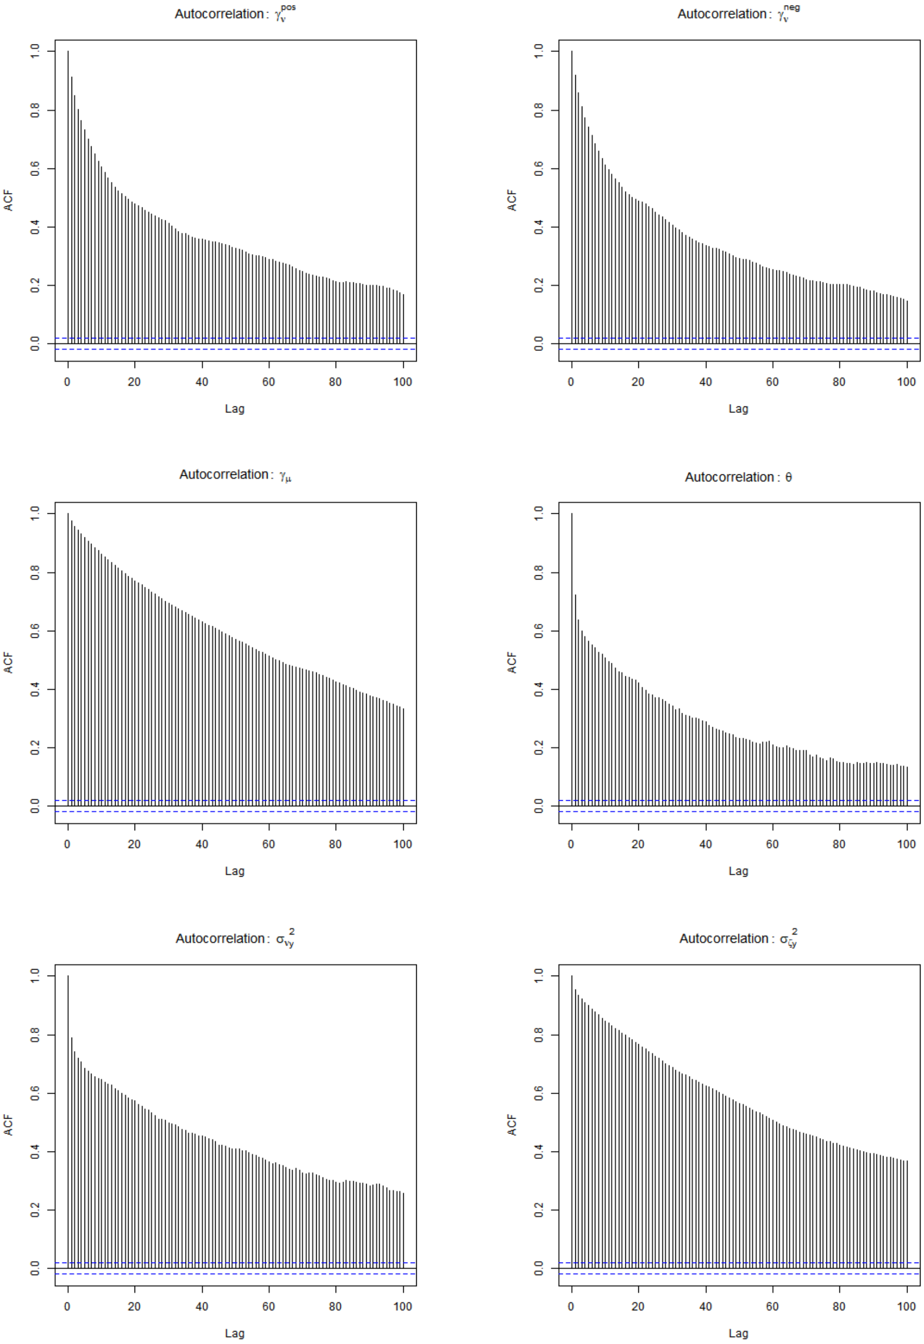
Autocorrelation : γ_v^{pos}
Autocorrelation : γ_v^{neg}
Autocorrelation : γ_μ
Autocorrelation : θ
Autocorrelation : σ_{vy}^2
Autocorrelation : $\sigma_{\xi y}^2$
ACF
Lag
0.0
0.2
0.4
0.6
0.8
1.0
0
20
40
60
80
100

Appendix C

Appendices to Chapter 4

C.1 Derivation of Joint Posterior Distribution – Interaction Model

The derivation follows that in Appendix A.2, with the following modifications. The consumption model of equations (4.1) and (4.2) can be written as

$$\begin{aligned} \tilde{c} &= I_\mu^{X\Gamma}\tilde{\mu}^y + \tilde{X}\Psi + \tilde{\mu}^c + \tilde{\nu}^c \\ H\tilde{\mu}^c &= \tilde{\mu}_0^c + I_\nu^{X\Gamma}\tilde{\nu}^y + \tilde{\zeta}^c \end{aligned}$$

where $I_\mu^{X\Gamma}$ is a $NT \times NT$ diagonal matrix with the vector $\tilde{X}\Gamma_\mu$ as the diagonal, and similarly for $I_\nu^{X\Gamma}$. The consumption data are distributed

$$\tilde{c}|\tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \Psi, \sigma_{\nu c}^2 \sim \mathcal{N}(I_\mu^{X\Gamma}\tilde{\mu}^y + \tilde{X}\Psi + \tilde{\mu}^c, \sigma_{\nu c}^2 I_{NT})$$

The priors for Γ_μ and Γ_ν are taken as $(k+1) \times 1$ vectors of independent normal distributions

$$\Gamma_\mu, \Gamma_\nu \sim \mathcal{N}(m_\Gamma, V_\Gamma)$$

where the first element of the mean vector m_Γ and the diagonal covariance matrix V_Γ (i.e. the prior on the constant term of the heterogeneous MPC) are the hyperparameters from the homogeneous model, $m_p = 0.35$ and $v_p = 0.5^2$. The remaining elements provide a zero-mean prior with a standard deviation of two (left panel Figure C.1). The prior for Ψ is a vector of independent normal distributions with zero mean and a standard deviation of two. The prior for $\tilde{\mu}^c$ is given by

$$H\tilde{\mu}^c|\tilde{y}, \tilde{X}, \tilde{\mu}^y, \ddot{\mu}_0^c, \ddot{\nu}_0^y, \theta, \Gamma_\nu, \sigma^2_{\zeta c} \sim \mathcal{N}(\ddot{\mu}_0^c + I_\nu^{X\Gamma}\tilde{\nu}^y, \sigma^2_{\zeta c} I_{NT})$$

where $\tilde{\nu}^y$ is given by $\tilde{F}^{-1}[\tilde{y} - \tilde{\mu}^y - \theta\ddot{\nu}_0^y]$. Other priors are as in Appendix A.2.

Defining the set of parameters as $\Theta = \{\Gamma_\mu, \Gamma_\nu, \Psi, \theta, \tilde{\mu}^y, \tilde{\mu}^c, \ddot{\mu}_0^y, \ddot{\mu}_0^c, \ddot{\nu}_0^y, \sigma^2_{\nu y}, \sigma^2_{\nu c}, \sigma^2_{\zeta y}, \sigma^2_{\zeta c}\}$, the general form of the joint posterior distribution can be written as

$$\begin{aligned} p(\Theta|\tilde{y}, \tilde{c}, \tilde{X}) \propto\ & p(\tilde{y}|\tilde{\mu}^y, \theta, \ddot{\nu}_0^y, \sigma^2_{\nu y}) p(\tilde{c}|\tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \Psi, \sigma^2_{\nu c}) \\ & \times p(H\tilde{\mu}^y|\ddot{\mu}_0^y, \sigma^2_{\zeta y}) p(H\tilde{\mu}^c|\tilde{y}, \tilde{X}, \tilde{\mu}^y, \ddot{\mu}_0^c, \ddot{\nu}_0^y, \theta, \Gamma_\nu, \sigma^2_{\zeta c}) \\ & \times p(\ddot{\mu}_0^y) p(\ddot{\mu}_0^c) p(\ddot{\nu}_0^y|\sigma^2_{\nu y}) \\ & \times p(\Gamma_\mu) p(\Gamma_\nu) p(\Psi) p(\theta) \\ & \times p(\sigma^2_{\nu y}) p(\sigma^2_{\nu c}) p(\sigma^2_{\zeta y}) p(\sigma^2_{\zeta c}) \end{aligned}$$

C.2 Relationship between the Empirical Model and Two-asset Model

The MPC concept used in Kaplan et al. (2018) (KMV), and developed in Achdou, Han, Lasry, Lions, and Moll (2017), measures the change in steady-state expected cumulative consumption over a given period when all households have an additional \$500 of liquid assets. That is, all households are shifted \$500 along their liquid asset state variable, then steady-state cumulative consumption at the original state space is subtracted from that at the new state space, and the result is

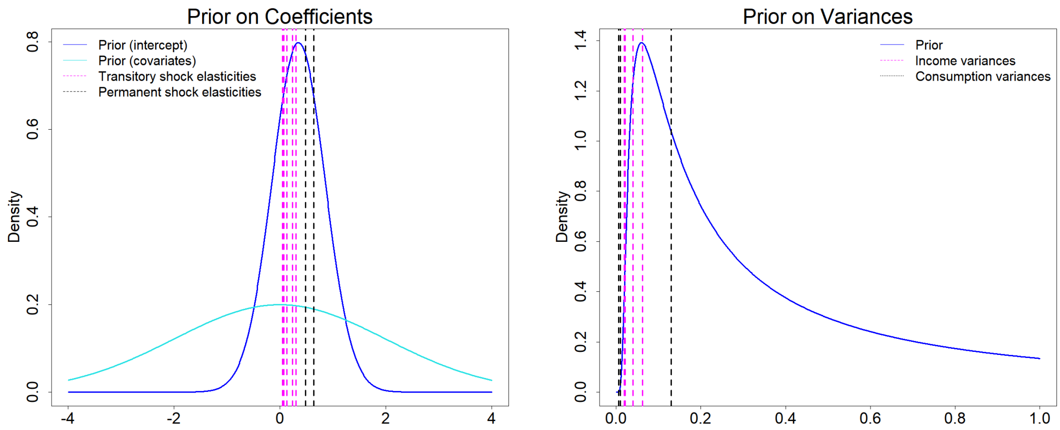

Figure C.1: Priors on coefficients (left panel) and variance parameters (right panel) along with estimates from the literature (dashed vertical lines).

divided by \$500 (Figure C.2, left panel). A shock to liquid assets is conceptually similar to what occurs when the government delivers an (unexpected) direct payment to households. However, the MPC concept identified in the empirical model is different. Holdings of liquid assets in the PSID data are point-in-time stock measures, whereas income and consumption are annualised flow measures. Changes to liquid assets will reflect the *result* of numerous income and consumption decisions between observations and hence a shock to liquid assets would be a less meaningful concept.

The empirical framework instead identifies bienniel shocks to annual consumption driven by bienniel shocks to annual income. The income process in KMV has some close similarities with the empirical model, being separated into transitory and persistent components. This allows me to construct an MPC measure in the two-asset model that is more comparable to the empirical MPC by shifting all households \$500 along their transitory or persistent income state (Figure C.2, middle and right panels). However, there are also important differences between the income processes. The continuous time formulation used in KMV means that households experience transitory shocks every three years and permanent shocks every 38 years, instead of experiencing both shocks every two years as in the PSID. Also, the KMV estimated variances are similar for transitory and persistent components, and are much larger than those in the PSID. These features likely affect how households response to income shocks.

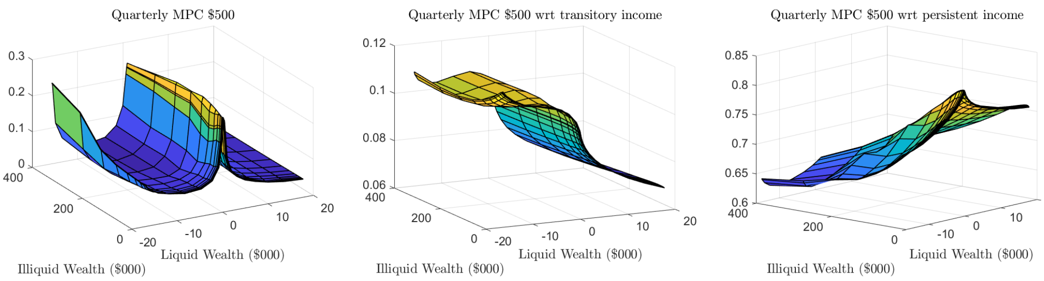

Figure C.2: Kaplan et al. (2018) MPC distributions across liquid and illiquid assets generated using the replication code; MPC out of liquid assets (left panel), out of transitory income (middle panel) and out of persistent income (right panel).

Intuitively, *pure* transitory income shocks and one-off liquid wealth shocks should be equivalent in the model since households pool liquid assets and income when optimising. However, the consumption policy function is conditional on labour income and liquid asset states, as they have different implications for expected future income streams.[109] Furthermore, the transitory income shocks in the model are not pure, but rather have some persistence. Nevertheless, it is surprising that constructing MPCs based on shifts to the income state practically eliminates the effect of the budget kink at zero liquid wealth; further work analysing the model MPC concepts may be warranted. The excericse does confirm the high level relationship posited in the two-asset model – liquid wealth is used to buffer against transitory shocks and illiquid wealth against persistent shocks.

[109]Asset income is based on the interest rate and labour income is based on an exogenous productivity shock process.

C.3 Derivation of Joint Posterior Distribution – Stochastic Interaction Model

The derivation proceeds as for the deterministic interactions model, with some further changes. Define the $NT \times 1$ vectors of $\gamma^{\nu}_{i,t}$ and $\gamma^{\mu}_{i,t}$ as $\tilde{\gamma}_{\nu}$ and $\tilde{\gamma}_{\mu}$. The consumption model in equations (4.3)-(4.8) imply

$$
\begin{aligned}
\tilde{c}|\tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \tilde{\gamma}_{\mu}, \Psi, \sigma^2_{\nu c} &\sim \mathcal{N}(I^{\gamma}_{\mu}\tilde{\mu}^y + \tilde{X}\Psi + \tilde{\mu}^c, \sigma^2_{\nu c} I_{NT}) \\
\tilde{H}\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \ddot{\mu}^c_0, \ddot{\nu}^y_0, \theta, \tilde{\gamma}_{\nu}, \sigma^2_{\zeta c} &\sim \mathcal{N}(\tilde{\mu}^c_0 + I^{\gamma}_{\nu}\tilde{\nu}^y, \sigma^2_{\zeta c} I_{NT}) \\
\tilde{\gamma}_{\mu}|\tilde{X}, \Gamma_{\mu}, \sigma^2_{\varepsilon\mu} &\sim \mathcal{N}(\tilde{X}\Gamma_{\mu}, \sigma^2_{\varepsilon\mu} I_{NT}) \\
\tilde{\gamma}_{\nu}|\tilde{X}, \Gamma_{\nu}, \sigma^2_{\varepsilon\nu} &\sim \mathcal{N}(\tilde{X}\Gamma_{\nu}, \sigma^2_{\varepsilon\nu} I_{NT})
\end{aligned}
$$

where I^{γ}_{ν} is a diagonal $NT \times NT$ matrix with elements given by $\tilde{\gamma}_{\nu}$, and similar for I^{γ}_{μ}. The covariance structure of the $\tilde{\gamma}_{\mu}$ and $\tilde{\gamma}_{\nu}$ parameters could be further examined in future work, but here they are taken as independent for simplicity. Additional variance priors are given by

$$
\sigma^2_{\varepsilon\mu}, \sigma^2_{\varepsilon\nu} \sim IG2(a, b).
$$

The general form of the joint posterior distribution can be written as

$$
\begin{aligned}
p(\Theta|\tilde{y}, \tilde{c}) \propto\ & p(\tilde{y}|\tilde{\mu}^y, \theta, \ddot{\nu}^y_0, \sigma^2_{\nu y}) p(\tilde{c}|\tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \tilde{\gamma}_{\mu}, \Psi, \sigma^2_{\nu c}) \\
& \times p(\tilde{H}\tilde{\mu}^y|\ddot{\mu}^y_0, \sigma^2_{\zeta y}) p(\tilde{H}\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \ddot{\mu}^c_0, \ddot{\nu}^y_0, \theta, \tilde{\gamma}_{\nu}, \sigma^2_{\zeta c}) \\
& \times p(\ddot{\mu}^y_0) p(\ddot{\mu}^c_0) p(\ddot{\nu}^y_0|\sigma^2_{\nu y}) \\
& \times p(\tilde{\gamma}_{\mu}|\tilde{X}, \Gamma_{\mu}, \sigma^2_{\varepsilon\mu}) p(\tilde{\gamma}_{\nu}|\tilde{X}, \Gamma_{\nu}, \sigma^2_{\varepsilon\nu}) p(\Psi) p(\theta) \\
& \times p(\sigma^2_{\nu y}) p(\sigma^2_{\nu c}) p(\sigma^2_{\zeta y}) p(\sigma^2_{\zeta c}) \\
& \times p(\Gamma_{\mu}) p(\Gamma_{\nu}) p(\sigma^2_{\varepsilon\mu}) p(\sigma^2_{\varepsilon\nu})
\end{aligned}
$$

C.4 Random Intercept Model

Model Extension

A further extension to the empirical framework takes an agnostic approach to the drivers of heterogeneity and estimates household-specific responses from the unexplained income and consumption data alone. The random intercept model posits that MPCs are household-specific, constant over time and drawn from a common distribution. This has some similarities to a panel 'fixed-effects' model where each household-specific MPC is independent; however, as the panel dimension of the data varies between three and nine observations, a pure fixed effects approach has insufficient information for inference.[110] (That is, the posterior distributions chiefly reflect the prior.) By specifying a common distribution, the random intercept model incorporates information from all households when estimating each household-specific MPC. The model is a particular case of the more general 'random coefficients' model (see Griffiths et al. (1979), Zeger and Karim (1991), Hsiao and Pesaran (2004)). Equations (2.5) and (2.6) are replaced with

$$c_{i,t} = \gamma_i^{\mu} \mu_{i,t}^{y} + \mu_{i,t}^{c} + \nu_{i,t}^{c} \tag{C.1}$$

$$\mu_{i,t}^{c} = \mu_{i,t-1}^{c} + \gamma_i^{\nu} \nu_{i,t}^{y} + \zeta_{i,t}^{c} \tag{C.2}$$

$$\gamma_i^{\nu} \sim \mathcal{N}(\bar{\gamma}_{\nu}, \sigma_{\gamma\nu}^2) \tag{C.3}$$

$$\gamma_i^{\mu} \sim \mathcal{N}(\bar{\gamma}_{\mu}, \sigma_{\gamma\mu}^2). \tag{C.4}$$

Posterior distributions of γ_i^{μ} and γ_i^{ν} are generated for all N households, along with the common mean and variance parameters $\bar{\gamma}_{\nu}$, $\bar{\gamma}_{\mu}$, $\sigma_{\gamma\nu}^2$ and $\sigma_{\gamma\mu}^2$. Estimation details are provided below.

Model Derivation

[110] Jappelli and Pistaferri (2020) use a frequentist fixed effects specification to *control* for unobserved heterogeneity in a short panel. This is distinct to inference on unobserved heterogeneity.

The derivation proceeds as for the deterministic interactions model, with only minor changes. Define the $N \times 1$ vectors of γ_i^ν and γ_i^μ as Γ_ν and Γ_μ. The consumption model in equations (C.1) and (C.2) imply the data are distributed as

$$\begin{aligned}\tilde{c}|\tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \sigma^2_{\nu c} &\sim \mathcal{N}(I^\gamma_\mu \tilde{\mu}^y + \tilde{\mu}^c, \sigma^2_{\nu c} I_{NT}) \\ \tilde{H}\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \ddot{\mu}^c_0, \ddot{\nu}^y_0, \theta, \Gamma_\nu, \sigma^2_{\zeta c} &\sim \mathcal{N}(\ddot{\mu}^c_0 + I^\gamma_\nu \tilde{\nu}^y, \sigma^2_{\zeta c} I_{NT}).\end{aligned}$$

where I^γ_ν is a diagonal $NT \times NT$ matrix with elements given by a vector that stacks γ_i^ν for each household observation, and similar for I^γ_μ. The priors for Γ_ν and Γ_μ are given by equations (C.3) and (C.4). The priors for $\bar{\gamma}_\nu$, $\bar{\gamma}_\mu$, $\sigma^2_{\gamma\nu}$ and $\sigma^2_{\gamma\mu}$ are given by

$$\begin{aligned}\sigma^2_{\gamma\mu}, \sigma^2_{\gamma\nu} &\sim IG2(a, b) \\ \bar{\gamma}_\mu, \bar{\gamma}_\nu &\sim \mathcal{N}(m_p, v_p).\end{aligned}$$

The general form of the joint posterior distribution can be written as

$$\begin{aligned}p(\Theta|\tilde{y}, \tilde{c}) \propto\ & p(\tilde{y}|\tilde{\mu}^y, \theta, \ddot{\nu}^y_0, \sigma^2_{\nu y}) p(\tilde{c}|\tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \sigma^2_{\nu c}) \\ & \times p(\tilde{H}\tilde{\mu}^y|\ddot{\mu}^y_0, \sigma^2_{\zeta y}) p(\tilde{H}\tilde{\mu}^c|\tilde{y}, \tilde{\mu}^y, \ddot{\mu}^c_0, \ddot{\nu}^y_0, \theta, \Gamma_\nu, \sigma^2_{\zeta c}) \\ & \times p(\ddot{\mu}^y_0) p(\ddot{\mu}^c_0) p(\ddot{\nu}^y_0|\sigma^2_{\nu y}) \\ & \times p(\Gamma_\mu|\bar{\gamma}_\mu, \sigma^2_{\gamma\mu}) p(\Gamma_\nu|\bar{\gamma}_\nu, \sigma^2_{\gamma\nu}) p(\theta) \\ & \times p(\sigma^2_{\nu y}) p(\sigma^2_{\nu c}) p(\sigma^2_{\zeta y}) p(\sigma^2_{\zeta c}) \\ & \times p(\bar{\gamma}_\mu) p(\bar{\gamma}_\nu) p(\sigma^2_{\gamma\mu}) p(\sigma^2_{\gamma\nu})\end{aligned}$$

Model Results

Diagnostics of the model are similar to the simple intercept only model and indicate reasonable inference on the response mean and variance parameters $\bar{\gamma}_\nu$, $\bar{\gamma}_\mu$, $\sigma^2_{\varepsilon\nu}$ and $\sigma^2_{\varepsilon\nu}$ (Table C.1).[111] The posterior distribution of $\bar{\gamma}_\mu$ is very similar to that of γ_μ in the intercept only model; however the posterior of $\bar{\gamma}_\nu$ is clearly lower than

[111] The same sample as the previous section is used for comparison; however, a larger sample is available as the model does not require lagged covariates. Estimation on the larger sample produces similar results, but with narrower posterior distributions of parameters.

that of γ_ν. Given the purely statistical nature of this model, the interpretation of this result is unclear.

	Mean	5%	95%
$\bar{\gamma}_\nu$	0.090	0.066	0.113
$\sigma_{\gamma\nu}$	0.284	0.258	0.313
$\bar{\gamma}_\mu$	0.550	0.510	0.590
$\sigma_{\gamma\mu}$	0.611	0.558	0.659
θ	0.059	0.035	0.085
$\sigma_{\nu y}$	0.239	0.235	0.243
$\sigma_{\zeta y}$	0.174	0.169	0.179
$\sigma_{\nu c}$	0.206	0.202	0.211
$\sigma_{\zeta c}$	0.073	0.067	0.079
Obs.	20286		
Households	3479		

Table C.1: Random intercept model posterior statistics; labour-force households.

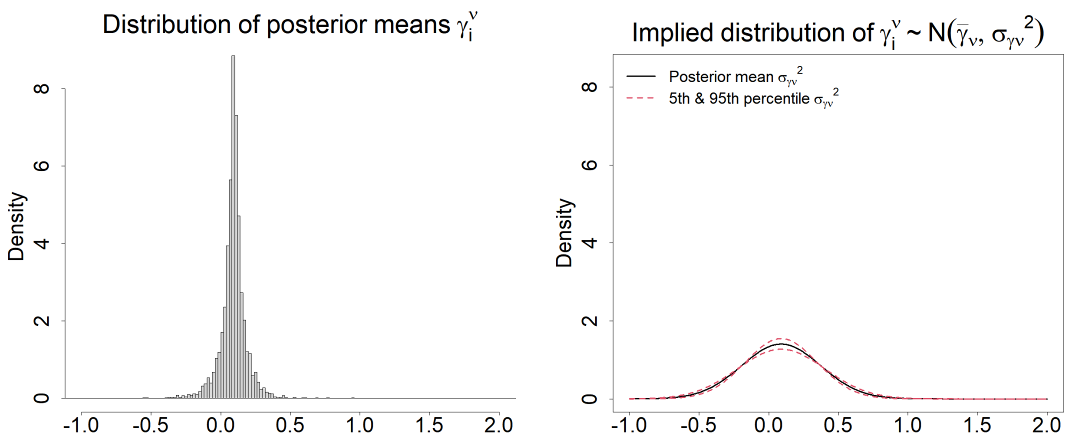

Figure C.3: Random intercept model: distribution of household-specific transitory elasticities given by the posterior means of γ_i^ν (left panel) and the implied normal distribution of γ_i^ν (right panel); the implied distribution uses the posterior mean of $\bar{\gamma}_\nu$ and the posterior mean, 5th and 95th percentiles of $\sigma^2_{\gamma\nu}$; 3,479 labour-force households.

Estimates of the response variance parameters $\sigma^2_{\gamma\nu}$ and $\sigma^2_{\gamma\mu}$ are quite large, particularly for permanent shocks, implying very disperse distributions of household MPCs (right panels Figures C.3-C.4). The distribution of predicted MPCs (using posterior means of each γ_i^ν and γ_i^μ) is similarly disperse, and has substantial ex-

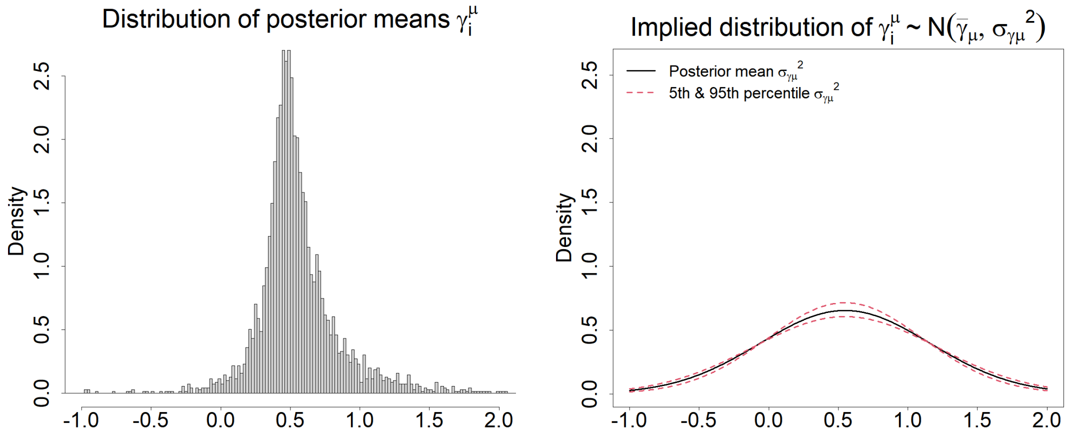

Figure C.4: Random intercept model: distribution of household-specific permanent elasticities given by the posterior means of γ_i^μ (left panel) and the implied normal distribution of γ_i^μ (right panel); the implied distribution uses the posterior mean of $\bar{\gamma}_\mu$ and the posterior mean, 5th and 95th percentiles of $\sigma^2_{\gamma\mu}$; 3,479 labour-force households.

cess kurtosis (left panels Figures C.3-C.4).[112] Overall, the results suggest there is a large amount of unobserved MPC heterogeneity across households, particularly for permanent shocks. However, this agnostic measure of heterogeneity does not provide insight into whether observable characteristics drive any of the variation. The results contrast those of the stochastic multiple covariate model, which finds more unobserved variation in transitory MPCs and less in permanent, due to observable heterogeneity accounting for most of the variation in permanent responses.

[112] Excess kurtosis means the tails of the distribution have larger probability mass than in a normal distribution. This is consistent with the large estimates of response variances.

C.5 Univariate Deterministic Interactions – Linear Spline Model

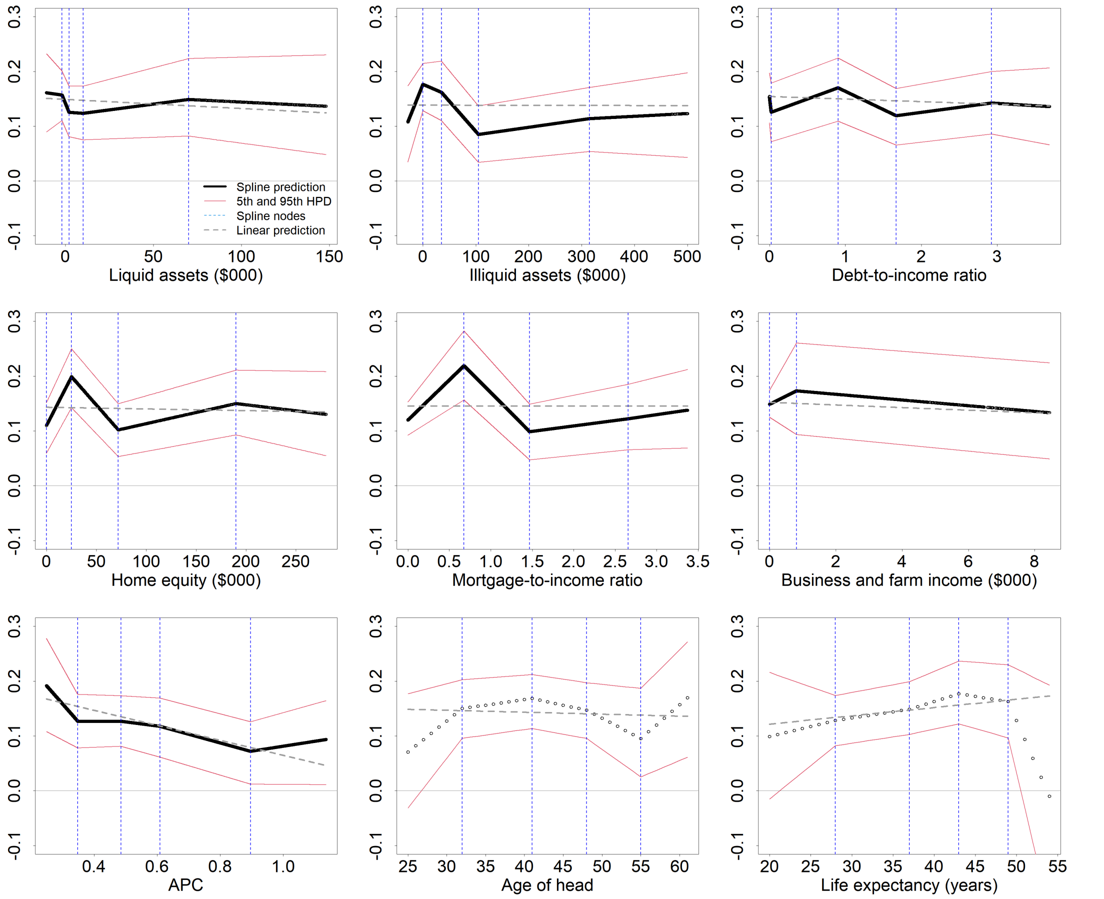

Figure C.5: Predicted responses to transitory shocks using the posterior mean values of Γ_{ν}; labour-force households; covariates plotted over 5–95th percentiles, vertical dashed lines are spline nodes at the 20th, 50th, 70th and 90th percentiles.

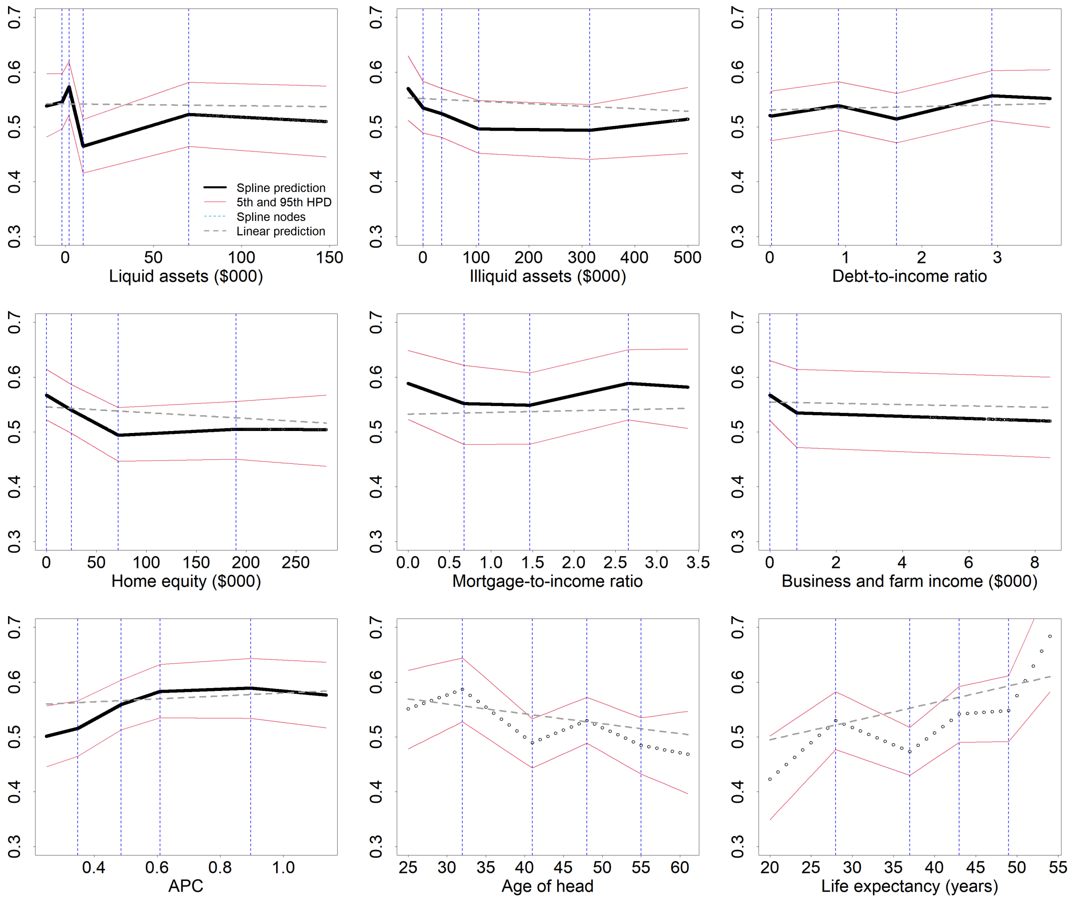

Figure C.6: Predicted responses to permanent shocks using the posterior mean values of Γ_μ; labour-force households; covariates plotted over 5–95th percentiles, vertical dashed lines are spline nodes at the 20th, 50th, 70th and 90th percentiles.

C.6 Detailed Results for Multiple Covariate Models

	Γ_ν			Γ_μ		
	Mean	5%	95%	Mean	5%	95%
Intercept	0.385	0.256	0.500	0.747	0.559	0.951
Home equity	0.020	0.002	0.038	-0.011	-0.023	0.001
Business & farm income	-0.016	-0.029	-0.004	-0.011	-0.020	-0.002
APC	-0.057	-0.076	-0.038	0.007	-0.002	0.016
Age of head	-0.015	-0.036	0.007	-0.024	-0.045	-0.003
Black	-0.220	-0.293	-0.146	0.432	0.326	0.560
Female	0.028	-0.040	0.097	0.311	0.176	0.419
High school	-0.096	-0.178	-0.010	-0.072	-0.236	0.085
3 or fewer years college	-0.110	-0.199	-0.028	-0.120	-0.289	0.034
4 or more years college	-0.149	-0.233	-0.069	-0.218	-0.368	-0.078
Obs.	20286					
Households	3479					

Table C.2: Deterministic model posterior statistics for covariates; statistics for home equity, business income, APC and age are scaled up by their standard deviation; labour-force households.

	Γ_ν			Γ_μ		
	Mean	5%	95%	Mean	5%	95%
Intercept	0.262	0.109	0.411	0.731	0.595	0.864
Home equity	0.017	-0.007	0.039	-0.010	-0.023	0.004
Business and farm income	-0.022	-0.042	-0.004	-0.009	-0.019	-0.000
APC	-0.055	-0.082	-0.028	-0.004	-0.015	0.005
Age of head	0.004	-0.022	0.029	-0.039	-0.058	-0.021
Black	-0.241	-0.334	-0.139	0.340	0.236	0.443
Female	0.051	-0.025	0.137	0.184	0.073	0.287
High school	-0.098	-0.188	-0.006	-0.069	-0.180	0.049
3 or fewer years college	-0.064	-0.150	0.038	-0.193	-0.292	-0.084
4 or more years college	-0.111	-0.202	-0.011	-0.194	-0.296	-0.076
Obs.	20286					
Households	3479					

Table C.3: Stochastic model posterior statistics for covariates; statistics for home equity, business income, APC and age are scaled up by their standard deviation; labour-force households.

C.7 Distribution of Weights

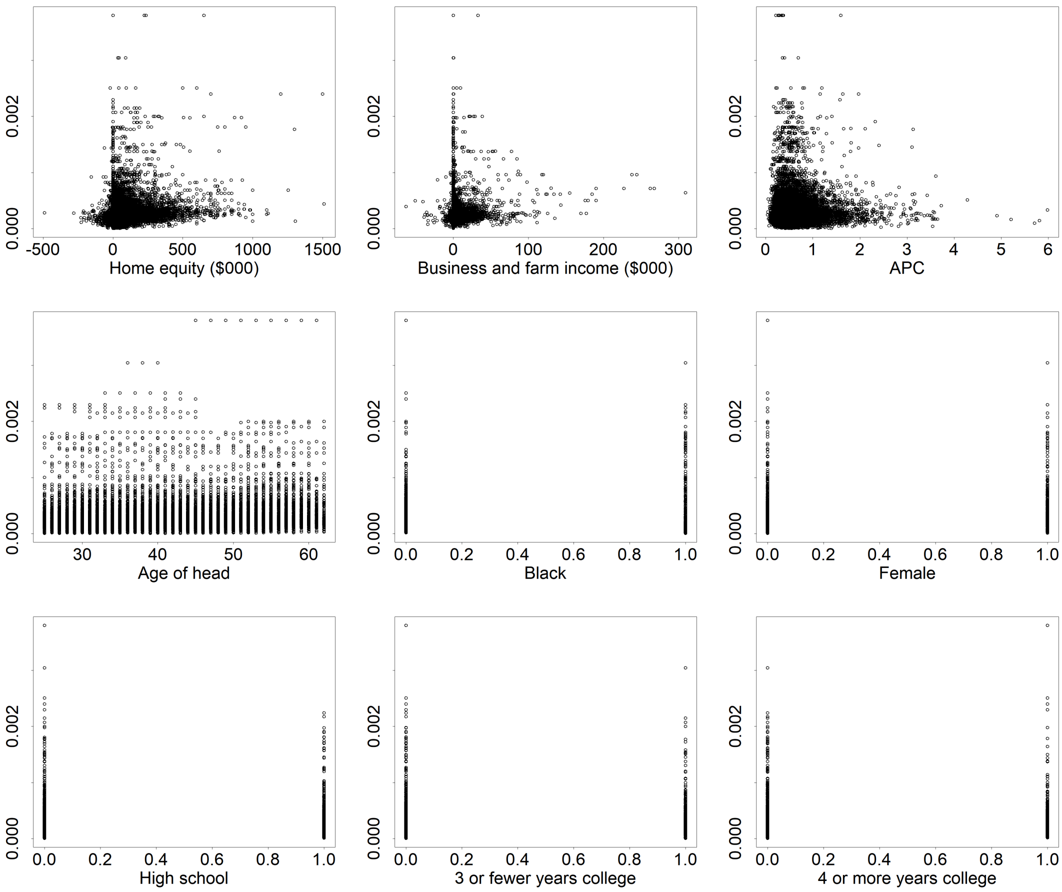

Figure C.7: Normalised probability weights by covariate; labour-force sample.

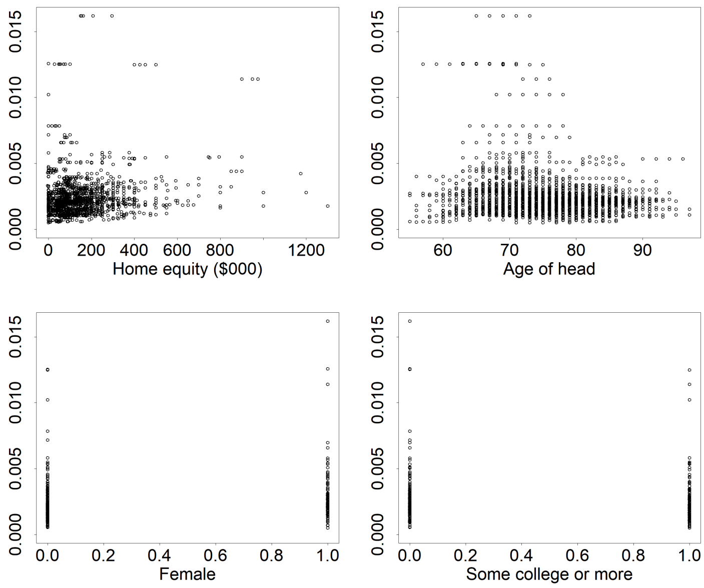

Figure C.8: Normalised probability weights by covariate; retired sample.

Appendix D

Appendices to Chapter 5

D.1 Detailed Covariate Correlations

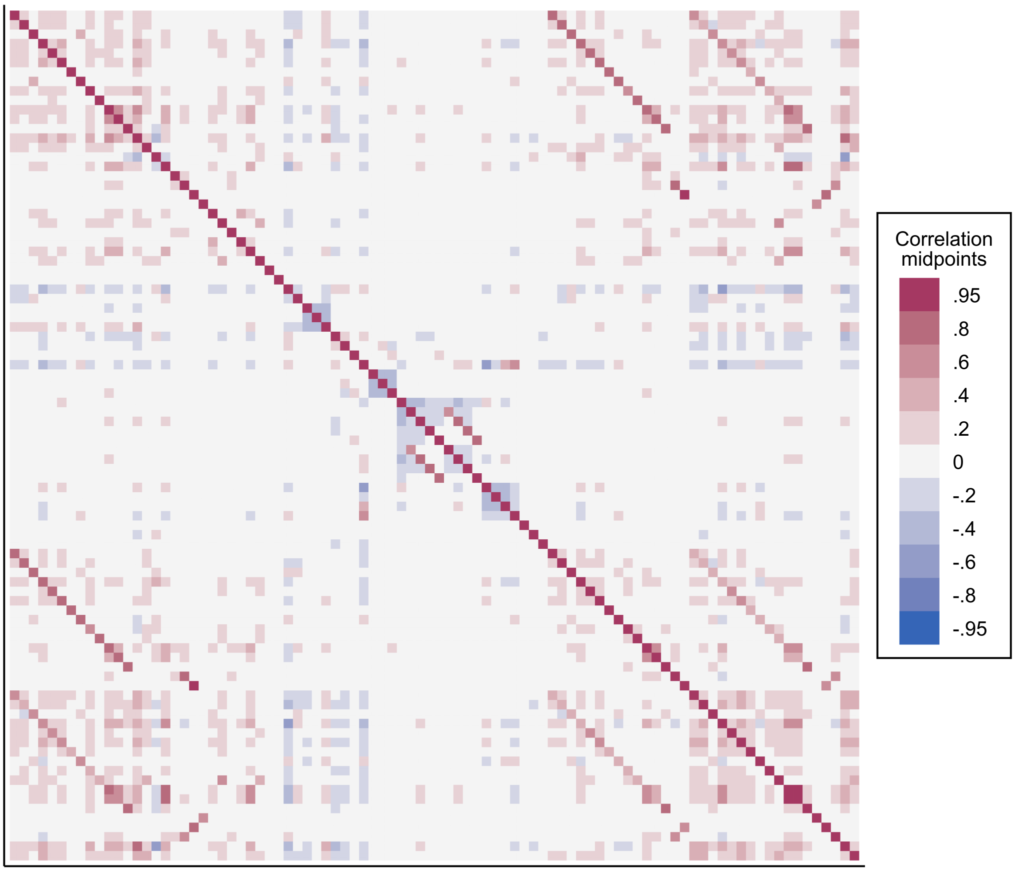

Figure D.1: Correlation heatmap of all 90 covariates (including two that are subsequently dropped); see figures below for expanded sections with labels.

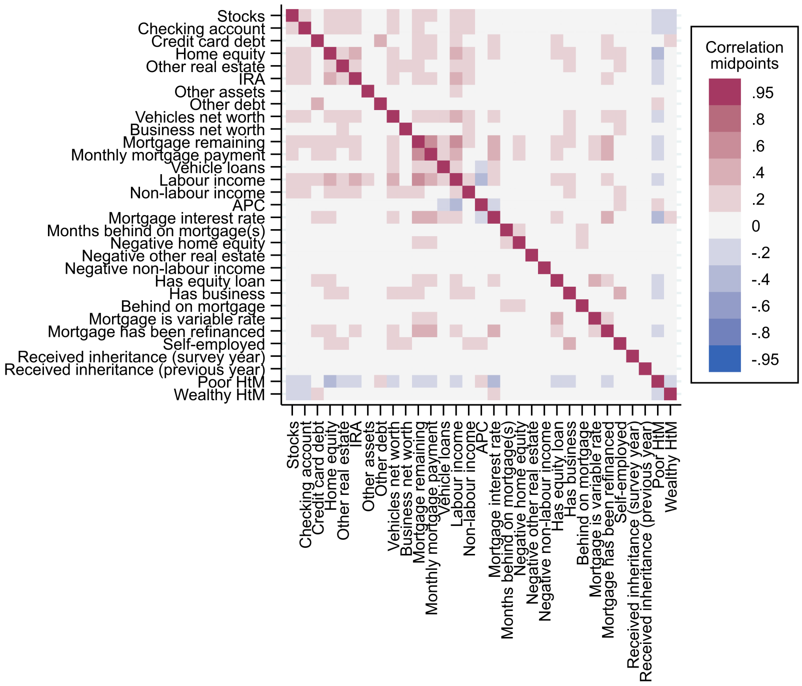

Figure D.2: Correlation heatmap of balance sheet covariates in levels.

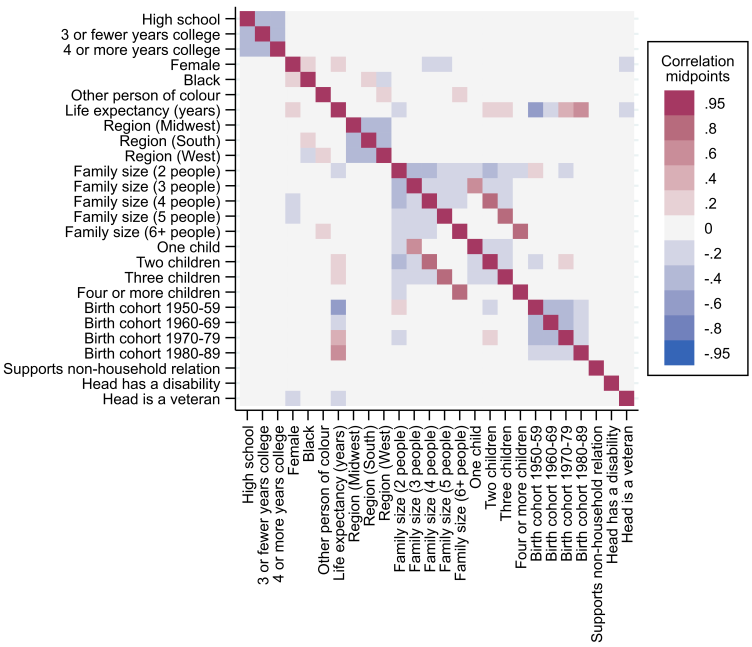

Figure D.3: Correlation heatmap of demographic covariates.

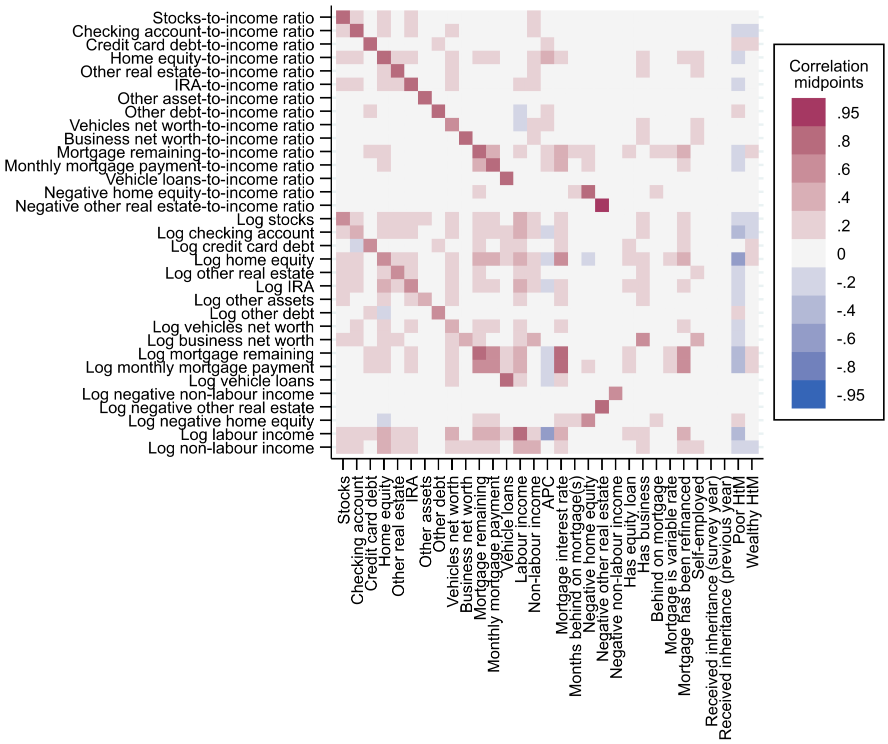

Figure D.4: Correlation heatmap of balance sheet covariates in levels against versions in logarithms and as income ratios.

D.2 Detailed Simulation Results

	Spec1	Spec2	Spec3	Spec4	Spec5	Spec6	Spec7	Spec8
Shrinkage	None	None	HS	None	HS	HS	HS	HS
No. noise variables	0	25	25	25	25	25	25	50
Correlation	0	0	0	0.5	0.5	0.5	0.5	0.5
Binary signals 4 & 5	No	No	No	No	No	Obs	HH	No
$\Gamma_{\nu,1}$ = 0.100	0.120	0.122	0.120	0.122	0.122	0.125	0.125	0.120
	(0.092, 0.148)	(0.094, 0.149)	(0.092, 0.148)	(0.094, 0.151)	(0.094, 0.149)	(0.090, 0.159)	(0.081, 0.167)	(0.093, 0.148)
$\Gamma_{\nu,2}$ = 0.050	0.051	0.052	0.044	0.056	0.038	0.039	0.036	0.031
	(0.031, 0.071)	(0.032, 0.073)	(0.022, 0.067)	(0.028, 0.084)	(0.010, 0.069)	(0.011, 0.070)	(0.010, 0.067)	(0.005, 0.062)
$\Gamma_{\nu,3}$ = -0.020	-0.020	-0.020	-0.009	-0.022	-0.010	-0.008	-0.008	-0.006
	(-0.039, 0.000)	(-0.040, 0.001)	(-0.028, 0.003)	(-0.050, 0.006)	(-0.033, 0.005)	(-0.031, 0.005)	(-0.029, 0.005)	(-0.024, 0.005)
$\Gamma_{\nu,4}$ = 0.100	0.104	0.106	0.103	0.105	0.097	0.067	0.071	0.102
	(0.084, 0.124)	(0.085, 0.126)	(0.083, 0.123)	(0.077, 0.134)	(0.068, 0.125)	(0.015, 0.128)	(0.021, 0.141)	(0.074, 0.130)
$\Gamma_{\nu,5}$ = -0.050	-0.053	-0.054	-0.046	-0.054	-0.036	-0.018	-0.021	-0.029
	(-0.074, -0.033)	(-0.074, -0.033)	(-0.068, -0.024)	(-0.082, -0.025)	(-0.067, -0.010)	(-0.063, 0.006)	(-0.065, 0.004)	(-0.061, -0.004)
$\Gamma_{\mu,1}$ = 0.700	0.630	0.613	0.631	0.618	0.635	0.638	0.634	0.631
	(0.575, 0.687)	(0.559, 0.668)	(0.577, 0.687)	(0.564, 0.673)	(0.580, 0.692)	(0.578, 0.701)	(0.556, 0.714)	(0.577, 0.687)
$\Gamma_{\mu,2}$ = 0.075	0.068	0.068	0.067	0.065	0.062	0.062	0.062	0.064
	(0.056, 0.081)	(0.056, 0.080)	(0.054, 0.079)	(0.048, 0.082)	(0.045, 0.079)	(0.045, 0.079)	(0.045, 0.079)	(0.046, 0.081)
$\Gamma_{\mu,3}$ = -0.025	-0.022	-0.021	-0.015	-0.022	-0.014	-0.014	-0.013	-0.013
	(-0.034, -0.010)	(-0.033, -0.009)	(-0.029, -0.004)	(-0.039, -0.005)	(-0.030, -0.001)	(-0.031, -0.002)	(-0.029, -0.001)	(-0.030, -0.001)
$\Gamma_{\mu,4}$ = 0.200	0.176	0.175	0.175	0.176	0.176	0.174	0.135	0.174
	(0.164, 0.188)	(0.163, 0.187)	(0.163, 0.187)	(0.159, 0.193)	(0.159, 0.193)	(0.143, 0.206)	(0.047, 0.240)	(0.158, 0.191)
$\Gamma_{\mu,5}$ = -0.100	-0.088	-0.089	-0.089	-0.091	-0.091	-0.083	-0.040	-0.085
	(-0.100, -0.076)	(-0.101, -0.077)	(-0.101, -0.076)	(-0.108, -0.074)	(-0.108, -0.074)	(-0.116, -0.050)	(-0.111, 0.001)	(-0.102, -0.068)

Table D.1: Dense posterior mean estimates from simulations; 5th and 95th percentile HPD interval in parentheses; averages over 50 replications.

	Spec1	Spec2	Spec3	Spec4	Spec5	Spec6	Spec7	Spec8
Shrinkage	None	None	HS	None	HS	HS	HS	HS
No. noise variables	0	25	25	25	25	25	25	50
Correlation	0	0	0	0.5	0.5	0.5	0.5	0.5
Binary signals 4 & 5	No	No	No	No	No	Obs	HH	No
$\Gamma_{\nu,1}$ = 0.100	0.116	0.118	0.116	0.118	0.118	0.121	0.120	0.117
	(0.086, 0.145)	(0.087, 0.147)	(0.085, 0.145)	(0.088, 0.148)	(0.088, 0.147)	(0.083, 0.157)	(0.072, 0.165)	(0.086, 0.145)
$\Gamma_{\nu,2}$ = 0.050	0.031	0.033	0.024	0.039	0.024	0.024	0.022	0.018
	(0.004, 0.060)	(0.007, 0.062)	(0.003, 0.054)	(0.008, 0.076)	(0.002, 0.056)	(0.002, 0.056)	(0.002, 0.054)	(0.000, 0.047)
$\Gamma_{\nu,3}$ = -0.020	-0.004	-0.003	-0.001	-0.008	-0.003	-0.002	-0.002	-0.002
	(-0.016, 0.000)	(-0.014, 0.000)	(-0.004, 0.000)	(-0.031, 0.001)	(-0.012, 0.000)	(-0.010, 0.000)	(-0.010, 0.000)	(-0.005, 0.000)
$\Gamma_{\nu,4}$ = 0.100	0.099	0.101	0.098	0.100	0.090	0.050	0.062	0.096
	(0.076, 0.121)	(0.078, 0.123)	(0.075, 0.120)	(0.067, 0.131)	(0.055, 0.122)	(0.005, 0.114)	(0.018, 0.133)	(0.063, 0.126)
$\Gamma_{\nu,5}$ = -0.050	-0.035	-0.035	-0.027	-0.036	-0.021	-0.008	-0.015	-0.016
	(-0.063, -0.009)	(-0.064, -0.009)	(-0.056, -0.005)	(-0.074, -0.006)	(-0.052, -0.002)	(-0.035, 0.000)	(-0.045, -0.001)	(-0.046, 0.000)
$\Gamma_{\mu,1}$ = 0.700	0.630	0.613	0.631	0.618	0.635	0.638	0.634	0.631
	(0.575, 0.686)	(0.558, 0.668)	(0.577, 0.687)	(0.564, 0.673)	(0.580, 0.692)	(0.578, 0.701)	(0.556, 0.714)	(0.577, 0.687)
$\Gamma_{\mu,2}$ = 0.075	0.057	0.056	0.054	0.051	0.047	0.047	0.046	0.049
	(0.039, 0.073)	(0.038, 0.072)	(0.036, 0.071)	(0.025, 0.074)	(0.022, 0.071)	(0.021, 0.070)	(0.021, 0.070)	(0.023, 0.073)
$\Gamma_{\mu,3}$ = -0.025	-0.001	-0.001	-0.001	-0.004	-0.002	-0.002	-0.002	-0.002
	(-0.004, 0.000)	(-0.004, 0.000)	(-0.003, 0.000)	(-0.013, 0.000)	(-0.009, 0.000)	(-0.009, 0.000)	(-0.008, 0.000)	(-0.007, 0.000)
$\Gamma_{\mu,4}$ = 0.200	0.174	0.173	0.173	0.175	0.174	0.169	0.130	0.173
	(0.162, 0.187)	(0.161, 0.186)	(0.161, 0.186)	(0.157, 0.192)	(0.157, 0.191)	(0.135, 0.202)	(0.044, 0.234)	(0.156, 0.190)
$\Gamma_{\mu,5}$ = -0.100	-0.081	-0.083	-0.082	-0.085	-0.084	-0.059	-0.035	-0.077
	(-0.095, -0.067)	(-0.096, -0.068)	(-0.096, -0.068)	(-0.104, -0.065)	(-0.103, -0.064)	(-0.103, -0.017)	(-0.096, -0.005)	(-0.097, -0.055)

Table D.2: SAVS posterior mean estimates from simulations; 5th and 95th percentile HPD interval in parentheses; averages over 50 replications.

D.3 Detailed Posterior Results

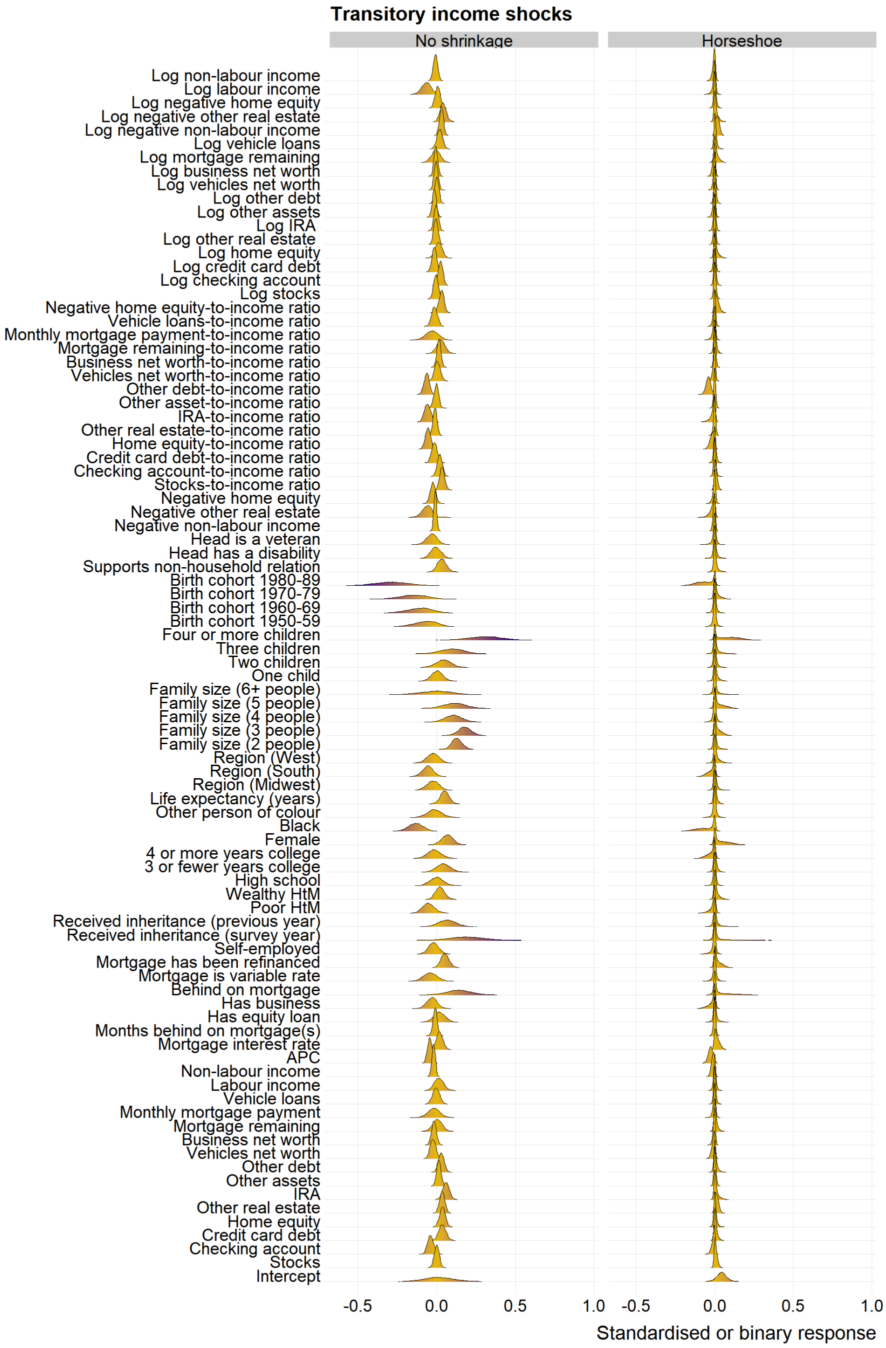

Figure D.5: Posterior densities of transitory coefficients with and without shrinkage.

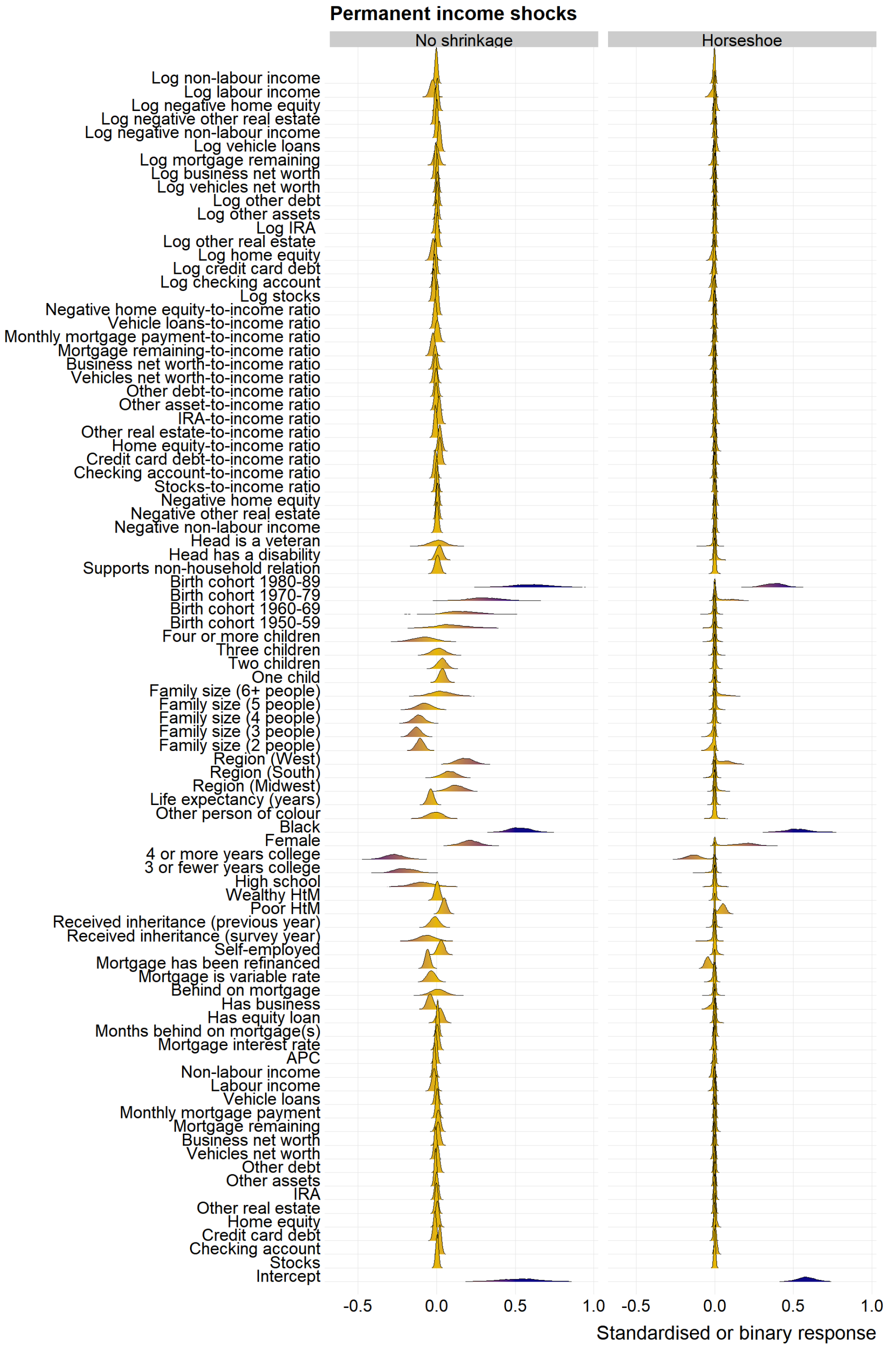

Figure D.6: Posterior densities of permanent coefficients with and without shrinkage.

D.4 Detailed Financial Literacy Results

	Mean	5%	95%
Γ_μ: Intercept	0.505	0.500	0.509
Γ_μ: Age of head	-0.067	-0.069	-0.066
Γ_μ: High school	-0.009	-0.013	-0.005
Γ_μ: Some college or more	0.043	0.038	0.049
Γ_μ: Labour income	-0.023	-0.024	-0.021
Γ_μ: APC	-0.009	-0.011	-0.008
Γ_μ: Model wealth	-0.059	-0.062	-0.057
Γ_μ: Very low financial literacy	0.280	0.274	0.287
Γ_μ: Low financial literacy	0.186	0.180	0.191
Γ_μ: Typical financial literacy	0.028	0.024	0.032
$\sigma_{\zeta y}$	0.243	0.241	0.245
$\sigma_{\nu c}$	0.016	0.015	0.016
$\sigma_{\zeta c}$	0.040	0.040	0.041
Observations	35609		
Households	2008		

Table D.3: Posterior statistics of estimation using the simulated dataset of Lusardi et al. (2017).

	Financial Literacy			Full Sample		
	Mean	5%	95%	Mean	5%	95%
γ_ν	0.105	0.077	0.132	0.118	0.098	0.138
γ_μ	0.544	0.491	0.599	0.579	0.525	0.629
θ	0.068	0.035	0.095	0.098	0.075	0.120
$\sigma_{\nu y}$	0.237	0.231	0.242	0.257	0.253	0.261
$\sigma_{\zeta y}$	0.141	0.134	0.149	0.137	0.131	0.144
$\sigma_{\nu c}$	0.228	0.224	0.232	0.236	0.233	0.239
$\sigma_{\zeta c}$	0.112	0.106	0.118	0.105	0.100	0.110
Observations	12981			23642		
Households	1748			3781		

Table D.4: Posterior statistics of an intercept only model on the financial literacy subsample and the full sample.

D.5 Joint Estimation

The previous literature takes unexplained income and consumption, y and c, as data and estimates equations equivalent to (2.3), (2.4), (4.1) and (4.2) (Blundell et al. (2008), Kaplan et al. (2014), Chatterjee et al. (2021)). That is, the dependent variables are the residuals from the first stage regressions in equations (2.1) and (2.2). Blundell et al. (2008) posit that the life-cycle control variables represent deterministic changes to the gradient of the consumption path known to households at time t. By estimating equations (2.1) and (2.2) as a separate first stage, previous work applies a partial information approach – estimates of Υ_C are not affected by income shock realisations. Partial information appears to be consistent with the Blundell et al. (2008) interpretation. I relax this assumption by jointly estimating the full model in equations (2.1)-(2.4), (4.1) and (4.2), which implies that the deterministic changes to the gradient of the consumption path are now conditional on realised income shocks at time t. The joint model can be written as

$$ln(Y_{i,t}) = Z_{i,t}\Upsilon_Y + \mu^y_{i,t} + \nu^y_{i,t} + \theta\nu^y_{i,t-1} \tag{D.1}$$

$$\mu^y_{i,t} = \mu^y_{i,t-1} + \zeta^y_{i,t}. \tag{2.4}$$

$$ln(C_{i,t}) = \begin{bmatrix} Z_{i,t} & X_{i,t} \end{bmatrix} \begin{bmatrix} \Upsilon_C \\ \Psi \end{bmatrix} + X_{i,t}\Gamma_\mu\mu^y_{i,t} + \mu^c_{i,t} + \nu^c_{i,t} \tag{D.2}$$

$$\mu^c_{i,t} = \mu^c_{i,t-1} + X_{i,t}\Gamma_\nu\nu^y_{i,t} + \zeta^c_{i,t}. \tag{4.2}$$

In equation (D.2), the life-cycle control variables Z and covariates X are augmented so that a single parameter vector can be drawn that contains both Υ_C and Ψ. This ensures that any correlation between Z and X is accounted for in estimation.[113] The joint model utilises the full information set in estimating all parameters. In specifications where shrinkage is used, it is also applied to the Υ_Y, Υ_C and Ψ vectors.

Table D.5 reports posterior estimates using the simple model used in the main text, whereas Table D.6 reports those of the joint model (with Horseshoe shrink-

[113] In addition, the augmented matrix is checked for multicollinearity and covariates are dropped to ensure full rank. This does not affect the covariates interacted with ν^y and μ^y, however.

age priors). Both specifications use an intercept as the only X covariate, with estimates generated from 10,000 draws after 10,000 burn in. The posterior distributions of most parameters have substantial overlap, or when they clearly differ it is not by a substantive amount. However, the posteriors for the transitory MPC γ_ν are both statistically and economically different. The joint model finds there is no consumption response to transitory shocks, with or without shrinkage being applied to the life-cycle control variables. The transitory coefficient γ_ν is only selected by the SAVS routine 1.9 per cent of the time in the joint model.

	Mean	5%	95%
γ_ν	0.118	0.098	0.138
γ_μ	0.579	0.525	0.629
θ	0.098	0.075	0.120
$\sigma_{\nu y}$	0.257	0.253	0.261
$\sigma_{\zeta y}$	0.137	0.131	0.144
$\sigma_{\nu c}$	0.236	0.233	0.239
$\sigma_{\zeta c}$	0.105	0.100	0.110
Observations	23642		
Households	3781		

Table D.5: Simple intercept only model posterior statistics.

	Mean	5%	95%
γ_ν	-0.009	-0.032	0.014
γ_μ	0.537	0.493	0.578
θ	0.117	0.095	0.140
$\sigma_{\nu y}$	0.230	0.224	0.234
$\sigma_{\zeta y}$	0.144	0.138	0.151
$\sigma_{\nu c}$	0.228	0.225	0.231
$\sigma_{\zeta c}$	0.113	0.108	0.118
Observations	23642		
Households	3781		

Table D.6: Joint intercept only model posterior statistics; Horseshoe shrinkage.

The distributions of permanent responses from the joint and simple models using all covariates are similar, but the transitory responses from the joint model are shifted lower compared to the simple model (Figure D.7). This shift results in the majority of predicted responses lying below zero. This would suggest that households decrease their consumption in response to transitory increases in income. Although possible – for example due to anticipated house purchases as in Gross (2019) – it is unlikely that incentives or constraints that drive such behaviour affect the majority of the sample. The results for the joint model suggest a substantial amount of dispersion in transitory responses, similar to Chapter 4 but shifted lower.

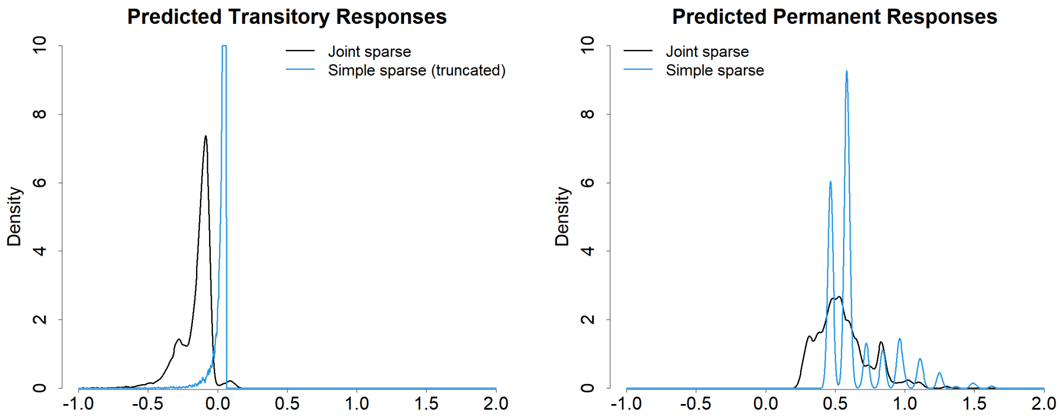

Figure D.7: Kernel densities of sparse predicted transitory (left panel) and permanent (right panel) responses using posterior means of joint and simple models with shrinkage; the simple sparse transitory distribution is truncated as the only selected covariate is other debt-to-income ratio, which is zero for most households.

Investigation of the results regarding transitory responses finds that it is primarily driven by differences in estimates of the latent consumption process $\hat{\mu}^c$, which reflects permanent shifts in consumption unrelated to permanent income. Average household correlations between the simple and joint model for a variety of model statistics are shown in Table D.7. The first difference of the latent consumption process $\Delta\hat{\mu}^c$ exhibits particularly weak correlation. This is further confirmed by simple linear regressions of this variable on the transitory income shock $\hat{\nu}^y$ (mimicking equation (4.2)), while varying the model across independent and dependent variables. That is, regressing $\Delta\hat{\mu}^c$ from the joint model on $\hat{\nu}^y$ from the simple model yields a coefficient estimate of effectively zero. However, regressing $\Delta\hat{\mu}^c$ from the simple model on $\hat{\nu}^y$ from the joint model yields a coefficient estimate of 0.10, which is close to $\hat{\gamma}_\nu = 0.12$ in the simple model. This confirms that differences in the estimates of $\Delta\hat{\mu}^c$ drive the result, rather than estimates of transitory shocks themselves. This implies that shifts in permanent consumption that were previously attributed to transitory income shocks are in fact driven by the life-cycle control variables when households have information about their correlations with income shocks.

The joint result for transitory shocks contrasts previous literature that finds a statistically significant positive response (using the simple model), such as that

	$\hat{y}$	$\hat{c}$	$\hat{\nu}^y$	$\hat{\mu}^y$	$\Delta\hat{\mu}^y$	$\Delta\hat{\mu}^c$
Correlation	0.88	0.95	0.89	0.82	0.80	0.59

Table D.7: Average household correlations between the simple and joint models.

emphasised in Commault (2021). The results of the joint model suggest that consumers who are given full information regarding the correlation of income shocks and deterministic life-cycle factors do not respond to transitory income shocks. This also implies that estimates relying on household panel data for identification of income and consumption shocks may have fundamental limitations compared with natural experiments that can more plausibly identify an income shock that is unrelated to other factors.

D.6 Derivation of Joint Model

The derivation proceeds using the definitions and notation in Appendix C.1, and with $\tilde{Y}$ and $\tilde{C}$ being the logarithm of income and consumption stacked by time and individual. Equations (D.1),(2.4), (D.2) and (4.2) imply the data are distributed

$$\begin{aligned}
\tilde{Y}|\tilde{Z}, \Upsilon_Y, \tilde{\mu}^y, \theta, \ddot{\nu}_0^y, \sigma_{\nu y}^2 &\sim \mathcal{N}(\tilde{Z}\Upsilon_Y + \tilde{\mu}^y + \theta\tilde{\nu}_0^y, \sigma_{\nu y}^2\Omega) \\
\tilde{H}\tilde{\mu}^y|\ddot{\mu}_0^y, \sigma_{\zeta y}^2 &\sim \mathcal{N}(\tilde{\mu}_0^y, \sigma_{\zeta y}^2 I_{NT}) \\
\tilde{C}|\tilde{Z}, \Upsilon_C, \tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \Psi, \sigma_{\nu c}^2 &\sim \mathcal{N}(\begin{bmatrix}\tilde{Z} & \tilde{X}\end{bmatrix}\begin{bmatrix}\Upsilon_C \\ \Psi\end{bmatrix} + I_\mu^{X\Gamma}\tilde{\mu}^y + \tilde{\mu}^c, \sigma_{\nu c}^2 I_{NT}) \\
\tilde{H}\tilde{\mu}^c|\tilde{Y}, \tilde{X}, \tilde{\mu}^y, \ddot{\mu}_0^c, \ddot{\nu}_0^y, \theta, \Gamma_\nu, \sigma_{\zeta c}^2 &\sim \mathcal{N}(\tilde{\mu}_0^c + I_\nu^{X\Gamma}\tilde{\nu}^y, \sigma_{\zeta c}^2 I_{NT}).
\end{aligned}$$

Shrinkage is included as per Section 5.2.2, extended to also cover Υ_Y and Υ_C.

The full set of parameters is $\Theta = \big\{\Upsilon_Y, \Upsilon_C, \Gamma_\mu, \Gamma_\nu, \Psi, \theta, \tilde{\mu}^y, \tilde{\mu}^c, \ddot{\mu}_0^y, \ddot{\mu}_0^c, \ddot{\nu}_0^y, \sigma^2_{\nu y}, \sigma^2_{\nu c}, \sigma^2_{\zeta y}, \sigma^2_{\zeta c}, \lambda^2_\nu, \lambda^2_\mu, \lambda^2_\Psi, \lambda^2_Y, \lambda^2_C, \tau^2_\nu, \tau^2_\mu, \tau^2_\Psi, \tau^2_Y, \tau^2_C, \psi^2_\nu, \psi^2_\mu, \psi^2_\Psi, \psi^2_Y, \psi^2_C, \xi_\nu, \xi_\mu, \xi_\Psi, \xi_Y, \xi_C\big\}$. The likelihood is given by

$$\begin{aligned}
p(\tilde{Y}, \tilde{C}, \tilde{X}, \tilde{Z}|\Theta) &= p(\tilde{Y}|\tilde{Z}, \Theta)p(\tilde{C}|\tilde{Y}, \tilde{Z}, \tilde{X}, \Theta) \\
&= p(\tilde{Y}|\tilde{Z}, \Upsilon_Y, \tilde{\mu}^y, \theta, \ddot{\nu}_0^y, \sigma^2_{\nu y})p(\tilde{C}|\tilde{Z}, \Upsilon_C, \tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \Psi, \sigma^2_{\nu c})
\end{aligned}$$

The general form of the joint posterior distribution can be written as

$$\begin{aligned}
p(\Theta|\tilde{Y}, \tilde{C}, \tilde{X}, \tilde{Z}) \propto\ & p(\tilde{Y}|\tilde{Z}, \Upsilon_Y, \tilde{\mu}^y, \theta, \ddot{\nu}_0^y, \sigma^2_{\nu y})p(\tilde{C}|\tilde{Z}, \Upsilon_C, \tilde{X}, \tilde{\mu}^c, \tilde{\mu}^y, \Gamma_\mu, \Psi, \sigma^2_{\nu c}) \\
&\times p(\tilde{H}\tilde{\mu}^y|\ddot{\mu}_0^y, \sigma^2_{\zeta y})p(\tilde{H}\tilde{\mu}^c|\tilde{Y}, \tilde{X}, \tilde{\mu}^y, \ddot{\mu}_0^c, \ddot{\nu}_0^y, \theta, \Gamma_\nu, \sigma^2_{\zeta c}) \\
&\times p(\ddot{\mu}_0^y)p(\ddot{\mu}_0^c)p(\ddot{\nu}_0^y|\sigma^2_{\nu y}) \\
&\times p(\Upsilon_Y|\lambda^2_Y, \tau^2_Y, \sigma^2_{\nu y})p(\Upsilon_C|\lambda^2_C, \tau^2_C, \sigma^2_{\nu c}) \\
&\times p(\Gamma_\mu|\lambda^2_\mu, \tau^2_\mu, \sigma^2_{\nu c})p(\Gamma_\nu|\lambda^2_\nu, \tau^2_\nu, \sigma^2_{\zeta c})p(\Psi|\lambda^2_\Psi, \tau^2_\Psi, \sigma^2_{\nu c})p(\theta) \\
&\times p(\sigma^2_{\nu y})p(\sigma^2_{\nu c})p(\sigma^2_{\zeta y})p(\sigma^2_{\zeta c}) \\
&\times p(\lambda^2_\nu|\psi_\nu)p(\lambda^2_\mu|\psi_\mu)p(\lambda^2_\Psi|\psi_\Psi)p(\lambda^2_Y|\psi_Y)p(\lambda^2_C|\psi_C) \\
&\times p(\tau^2_\nu|\xi_\nu)p(\tau^2_\mu|\xi_\mu)p(\tau^2_\Psi|\xi_\Psi)p(\tau^2_Y|\xi_Y)p(\tau^2_C|\xi_C) \\
&\times p(\psi_\nu)p(\psi_\mu)p(\psi_\Psi)p(\psi_Y)p(\psi_C)p(\xi_\nu)p(\xi_\mu)p(\xi_\Psi)p(\xi_Y)p(\xi_C)
\end{aligned}$$