Heterogenous Consumption Responses and Wealth Inequality over the Business Cycle

Citation for published version:
Forshaw, RJ 2024, Heterogenous Consumption Responses and Wealth Inequality over the Business Cycle. in N Tsounis & A Vlachvei (eds), *Applied Economic Research and Trends. ICOAE 2023.* Springer Proceedings in Business and Economics, Springer, pp. 667–700, 2023 International Conference on Applied Economics, Brno, Czech Republic, 29/06/23. https://doi.org/10.1007/978-3-031-49105-4_39

Digital Object Identifier (DOI):
10.1007/978-3-031-49105-4_39

Link:
Link to publication record in Heriot-Watt Research Portal

Document Version:
Peer reviewed version

Published In:
Applied Economic Research and Trends. ICOAE 2023

Heterogenous Consumption Responses and Wealth Inequality over the Business Cycle

Heterogenous Consumption Responses and Wealth Inequality over the Business Cycle

Rachel Forshaw*

Abstract

Recent research has highlighted the importance of heterogeneity in the marginal propensity to consume (MPC) out of transitory income shocks for the efficacy of monetary and fiscal policy. However, work on estimating the distribution of MPCs remains scant, and typically assumes that an individual's MPC remains constant over time. Using the US Panel Study of Income Dynamics (PSID), I calibrate a model of microeconomic wealth and discount rate heterogeneity with aggregate shocks over the Great Recession of 2008, a historical 'upper bound' in terms of the dynamics of the wealth distribution in recessionary times. Using a reduced-form model I estimate the degree of heterogeneity in MPCs, both over the population and over the business cycle and show that a state-invariant MPC distribution is irreconcilable with empirical changes in the wealth distribution.

JEL codes: E2, E3, G5
Keywords: Business cycles, consumption, heterogenous preferences, incomplete markets heterogenous agent models, MPC, wealth inequality

1 Introduction

The standard model of inequality with business cycles, the Krusell and Smith (1998) model (henceforth KS), is becoming an increasingly important tool in the macroeconomist's toolkit. Its popularity is only set to grow given that the barriers to entry

*Heriot-Watt University, Mary Burton Building, Edinburgh, UK; email: r.forshaw@hw.ac.uk.
Thanks to Carlos Carillo-Tudela, Jonathan Thomas, Mike Elsby and Ludo Visschers for their helpful comments. An earlier version of this paper constituted part of my PhD thesis and so received funding from the University of Edinburgh School of Economics.

to this literature are rapidly falling. Historically, a significant barrier has been one of computation. However, recent software developments obviate the need for a computer science degree in order to solve such models[1]. These new applications allow the economist to solve these class of models with only the level of computational savviness needed to operate such widespread applications as Dynare. The second barrier to entry has been the 'so what?' objection. KS showed that, rather than keeping track of the whole distribution of wealth holdings, agents were able to forecast prices with very high precision by only tracking the mean, a finding they term *approximate aggregation*. The corollary to this finding is that macroeconomic aggregates are incredibly similar to those generated by representative agent models. This second barrier is also falling, with a number of recent papers violating the approximate aggregation result and finding that microeconomic heterogeneity can matter significantly for macro, and allowing the study of distributional implications of economic policies across households[2]. Given that these models are becoming increasingly popular, we should care about how well they accord with empirical evidence. Since the original paper's release, a number of additions to the KS model have increased its empirical realism in terms of the degree of wealth inequality seen in the data (see for example Krueger et al. (2016), DeNardi and Fella (2017), Carroll et al. (2017)). However, this literature uses cross-sectional micro data at a point in time to discipline the cross section of the model. The major contribution of this paper is to instead focus on micro-level panel data capturing the distributional dynamics over time.

It is already well-documented that the KS model does a poor job of fitting the cross-sectional wealth distribution in its benchmark form[3]. An increasingly popular method for better-fitting the empirical facts of the cross-sectional wealth distribution in this period is the addition of heterogeneous time preferences. It is backed by growing microeconomic evidence of such heterogeneity in the rate of time discount (Lawrance (1991), Warner and Pleeter (2001)). I calibrate this version of the KS model, which I will call the KS-β-Dist model, to US Panel Survey of Income Dynamics (PSID) data and examine its business cycle characteristics, focusing in particular on its predictions for the marginal propensity to consume (MPC). I focus on MPCs as I establish that variation in the MPC both over the distribution and over the cycle leads to a failure of approximate aggregation. As such, they are the pathway by which microeconomic heterogeneity matters for macroeconomic aggregates.

[1]See for example, Phact by Ahn et al. (2018), Hark by Carroll (2006)

[2]See Ahn et al. (2018) for an overview of this literature.

[3]See Krusell and Smith (1998),Krueger and Perri (2005), De Nardi and Fella (2017).

Using the reduced-form model of Blundell et al. (2008) on the same panel data, I confirm that the KS β-Dist provides a good approximation to the cross-sectional wealth distribution in accordance with previous literature and, new to this paper, the income-poor part of the distribution of cross-sectional MPCs. However, I then repeat this calibration exercise for data during the expansion period leading up to the Great Recession of 2008, and its immediate aftermath. The choice of the 2008 recession is deliberate: it was unprecedented in the postwar period for its severity and duration. It also had a large-reaching influences on consumption and the wealth distribution in the US. Bricker et al. (2012) find that over the period 2007 - 2010, median net worth fell 38.8 percent in real terms, and the Survey of Consumer Finances also documents that net worth decreased considerably relative to income; the median net worth-to-income ratio declined from 8.5 in 2007 to 5.6 in 2010. De Nardi et al. (2011) detail that it took almost 12 quarters for total real personal consumption expenditures to return to the previous peak in 2007 Q4. It therefore represents a historical 'upper bound' in terms of the dynamics of the wealth distribution in recessionary times, and therefore makes the results generalisable to other recessionary periods. In order to fit the dynamics of the wealth distribution in the KS framework, I show that we must modify preference parameters significantly, with large implications for the cross-sectional variation of MPCs. Such variation is ruled out on the basis of my reduced-form evidence.

The paper is organised as follows: section 2 and 3 describe the KS structural model, and the addition of heterogenous discount factors. Section 4 explains the importance of marginal propensities to consume and where the approximate aggregation result fails. 5 describes the dynamics of the KS β-Dist. Section 6 and 7 takes these dynamics to the data, describing the methodology of the reduced-form estimation of MPCs across the distribution. Section 8 discusses the results in the context of the structural model and section 9 concludes.

2 The Krusell Smith Model

KS consider an economy in which there is a continuum of infinitely-lived agents of measure one. Time is discrete, $t = (0, 1, 2, ...)$ and each agent has preferences over flows of consuming the single consumption good, c, that can be described by:

$$\sum_{t=0}^{\infty} \beta^t U(c_t)$$

with constant relative risk aversion (CRRA) utility:

$$U(c) = \lim_{i \to \sigma} \frac{c^{1-i} - 1}{1 - i}$$

Because leisure is not valued, agents spend all of their time - each is endowed with one unit - working when employed. Idiosyncratic risk is introduced through a stochastic shock to labour input:

$$e_i \in E = \begin{cases} 1 & \text{for } i = g, \text{ 'employed'}; \\ 0 & \text{for } i = b, \text{ 'unemployed'}. \end{cases}$$

Aggregate labour L combines with aggregate capital K to make production good Y according to the Cobb-Douglas production function $Y = zK^{\alpha}L^{1-\alpha}$ where $\alpha \in [0, 1]$ and z is the aggregate state of the economy, which feeds through the model as a shock which can take two values:

$$z_i \in Z = \begin{cases} 1 + \delta_z & \text{for } i = g, \text{ 'expansion'}; \\ 1 - \delta_z & \text{for } i = b, \text{ 'recession'}. \end{cases}$$

Where δ_z is a calibration parameter. KS assume that the aggregate state, z and the idiosyncratic shock e follow a first order Markov process with the transition matrices π_z and $\pi_{e,e'|z,z'}$, respectively. Following standard notation a prime signifies the next period's realisation, so that $\pi_{z,z'}$ is the probability that the aggregate state transitions to z' next period from this period's realisation is z. $\pi_{e',z'|e,z}$ is the probability that next period's idiosyncratic shock is e' and that the aggregate shock is z' given that this period's employment realisation is e and the aggregate shock is z. There are no markets for insurance against uncertainty, so agents may only undertake a form of self-insurance by investing in a single asset, capital, which is restricted to take values $k \in \kappa = [0, \infty)$.

Individual Agent's Problem An individual agent's optimisation problem is the following:

$$V(k, e, z, \Gamma) = \max_{c \in \mathbb{R}^+,\, k' \in \kappa} \left\{ U(c) + \beta \sum_{z' \in Z} \sum_{e' \in E} \pi_{e'z'|e,z} V(k', e', z', \Gamma') \right\} \tag{1a}$$

subject to:

$$k(1 + r(K, L, z) - \delta) + [(1 - \tau)\bar{l}e + \mu(1 - e)]w(K, L, z) - c = k' \tag{1b}$$

$$\Gamma' = G(\Gamma, z, z') \tag{1c}$$

$$k' \geq 0. \tag{1d}$$

Where $r(\cdot)$ is the real interest rate and $w(\cdot)$ the wage rate, δ is the constant rate at which capital depreciates, $\bar{l}$ is individual time endowment, μ unemployment insurance as a percentage of the wage rate, τ is a tax on labour income and Γ the measure of agents over wealth and employment status. Equation 1a is a standard Bellman equation. Budget constraint (1b) states that the future individual capital stock is composed of today's capital, compounded by the depreciation-adjusted rental price of capital, $r(K, L, z) - \delta$, the labour income $(1-\tau)\bar{l}w(K, L, z)$ when an agent is employed, or $\mu w(K, L, z)$ if an agent is unemployed, less today's consumption. (1c) is the forecasting rule for the future distribution of Γ and for the aggregate state variable. Agents about the next period's distribution because it determines future prices. Finally, (1d) is the borrowing constraint: which restricts next period's capital choice to be positive.

Government The only role of government in the model is to tax labour income to fund the payment of unemployment insurance, they run a balanced budget each period. This means that:

$$\underbrace{w(K, L, z)\tau\bar{l}L}_{\text{government income}} = \underbrace{w(K, L, z)\mu(1-L)}_{\text{government expenditure}}$$

which implies that the tax rate is:

$$\tau = \frac{\mu(1-L)}{\bar{l}L}$$

where L is total employed labour, and $(1-L)$ the unemployment rate.

Firm's Problem Factor prices follow from the competitive firm's optimisation problem, as in the standard representative agent model:

MPL $= \frac{\partial Y}{\partial L} = (1-\alpha)z(K/L)^{\alpha} = w(K, L, z)$;
MPK $= \frac{\partial Y}{\partial K} = \alpha z(K/L)^{\alpha-1} = r(K, L, z)$, where the last equalities hold due to competitive markets, and where:

$$K = \int_A k' \quad d\Gamma \tag{2}$$

$$L = \int_A e \quad d\Gamma. \tag{3}$$

where $A = \kappa \times E$ is the type space of agents over capital holdings and employment status. The associated measurable space is $M = (A, \mathcal{B}(A))$ where $\mathcal{B}(A) = \mathcal{B}(\kappa) \times \mathcal{P}(E)$ is the Borel σ-algebra of A and $\mathcal{P}(E)$ is the power set of E. The set of all measures on M is $\mathcal{M}$, and we shall require that Γ is an element of $\mathcal{M}$.

2.1 Recursive Competitive Rational Expectations Equilibrium

The KS recursive competitive equilibrium can be defined in the following way:

Definition 1. Recursive competitive rational expectations equilibrium
A recursive competitive rational expectations equilibrium consists of:
(i) A value function: $V^*(k, e, z, \Gamma) : A \times Z \times \mathcal{M} \to \mathbb{R}$ which solves the individual's optimisation problem (1a), with the associated optimal decision rules for capital and consumption:
$g_{k'}(k, e, z, \Gamma) : A \times Z \times \mathcal{M} \to \mathbb{R},\ g_{k'}(k, e, z, \Gamma) = k'^*$
$g_{c'}(k, e, z, \Gamma) : A \times Z \times \mathcal{M} \to \mathbb{R},\ g_{c'}(k, e, z, \Gamma) = c'^*$
(ii) Pricing functions:
$w^*(K, L, z) : \mathcal{M} \times Z \to \mathbb{R},\ w^*(K, L, z) = (1-\alpha)z(K/L)^{\alpha}$
$r^*(K, L, z) : \mathcal{M} \times Z \to \mathbb{R},\ r^*(K, L, z) = \alpha z(K/L)^{\alpha-1}$
which solve the firm's optimisation problem.
(iii) An equilibrium transition function:
$G^*(\Gamma, z, z') : Z \times Z \times \mathcal{M} \to \mathcal{M},\ G^*(\Gamma, z, z') = \Gamma'$
that is consistent with the law of motion for Γ implied by individual decision rules, $g_{k'}(\cdot)$, $g_{c'}(\cdot)$, and the Markov process $\pi_{e',z'|e,z}$.

3 Improving the distributional fit (β-Dist)

The addition of heterogenous discount factors to improve the fit of the wealth distribution goes back to Krusell and Smith (1998) in which they added the assumption of a stochastic discount factor, $\tilde{\beta}$, which can take values $\{0.9858, 0.9894, 0.9930\}$. $\tilde{\beta}$ follows a three-state Markov chain which generates an invariant distribution for discount factors that is symmetric around its mean. KS gave the intuition that this was a way of modelling overlapping generations with different levels of patience - it in fact creates three types of agents (impatient, baseline, patient). I use the specification from Castañeda et al. (2003) because it is much quicker to solve, and so is more often used in the literature. The major difference between this and the stochastic discount factor model is that agents' discount parameters do not change over time, as such the model generates much less mobility in the wealth distribution. Specifically, they choose time preference parameters to be distributed uniformly in the population between $\grave{\beta} \pm \nabla$ to fit the proportion of wealth w held by richest 20,40, 60 and 80%, ie:

$$\{\grave{\beta}, \nabla\} = \underset{\beta,\nabla}{\operatorname{argmin}} \left(\sum_{i=20,40,60,80} (w_i(\beta, \nabla) - w_i)^2 \right)^{1/2} \tag{4}$$

Using a distribution of discount factors helps to better fit the skewness of the empirical distribution as it attenuates the precautionary saving motive for agents with smaller discount factors to generate a larger mass at the lower end of the distribution. In a sense, they become myopic to likelihood of hitting the borrowing constraint. It also fits the upper parts of the distribution by heightening the desire to save for agents with higher discount factors. It is not only a useful device for better fitting the empirical distribution, but also backed by empirical evidence of heterogeneity in discount factors. In a study of military drawdown payments, Warner and Pleeter (2001) find discount factors between 0.76 and 1.0 when military personnel were given the choice between a lump-sum payment or annuity. In order to best fit the minimisation problem of equation 4, I find that a much smaller variation in $\grave{\beta}$ of $\nabla = 0.01$ is sufficient to fit the empirical distribution. This much smaller variance is likely due to coarse fitting points, in accordance with the literature I use only 4 datapoints and do not fit the maximum or minimum (0 % or 100 %) quantiles.[4]

Figure 1 shows Lorenz Curves generated from solving for the $\grave{\beta}$s and ∇s which solve equation 4 for PSID waves 2007 (i.e., pre-recession) and 2009 (recession data). The dashed line labelled KS-JEDC is the benchmark KS model. The solid line labelled β-Dist is the KS model with the addition of heterogeneous discount factors. Clearly, a relatively small variation in discount factors can dramatically improve the fit to the empirical wealth distribution. The baseline KS model exhibits very little wealth inequality while the β-Dist models capture the pre-recession and recession data closely.

[4] To fit the distribution more closely is a small modification to the algorithm. The contradiction in business cycle dynamics still exists in this case.

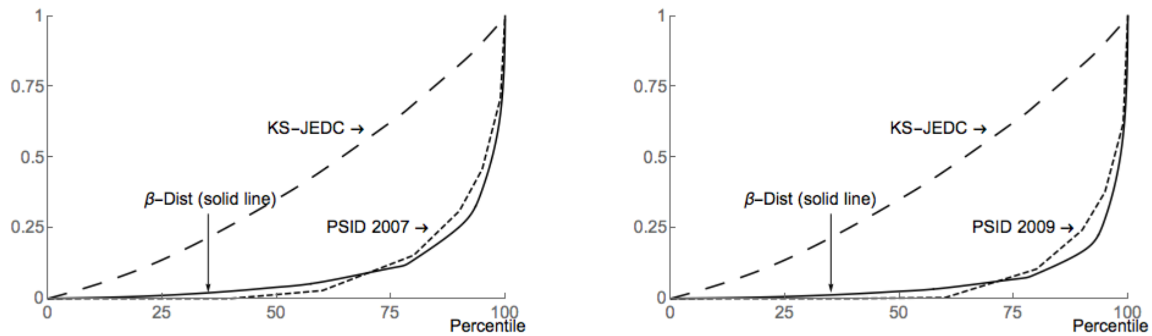

Figure 1: Distribution of Wealth (Lorenz Curves)

Wealth distributions for Krusell-Smith (KS-JEDC) vs. β distribution, calibrated to 2007 (left) and 2009 (right) PSID data.

4 Why are MPCs So Important?

Before exploring the KS model's predictions for the cross-sectional distribution of MPCs, and the dynamics over the cycle, it is worthwhile emphasising why this paper puts such emphasis on them. For this, it is important to consider when approximate aggregation holds, and when it does not. Consider the following simple example. Without individual or aggregate uncertainty in the KS model, if agents had linear consumption policies, each individual i having the same intercept ϕ_0 and constant marginal propensity to save ϕ_1, the individual policy function would take the form:

$$k_i' = \phi_0 + \phi_1 k_i$$

Then aggregation is easy:

$$K' = \phi_0 + \phi_1 K.$$

Because we are working with probability space and normalise total labour supply to equal one, the first moment is equal to total asset holding, K. We can *exactly* aggregate all of the individual policy functions and the first moment of the wealth distribution is a the only statistic we need to perfectly forecast capital stock tomorrow.

Approximate aggregation holds in the standard KS model because the vast majority of agents act approximately in this manner, having near-linear savings functions and constant marginal propensities to save. Only a few that are constrained have very low zero marginal propensity to save out of income, but they are too small in number to affect the aggregate result dramatically. Figure 2 plots the individual capital policy functions for the baseline KS economy. Though functions exhibits non-linearity for the poorest agents, there are very few agents in this region, and low amounts of redistribution to this area.

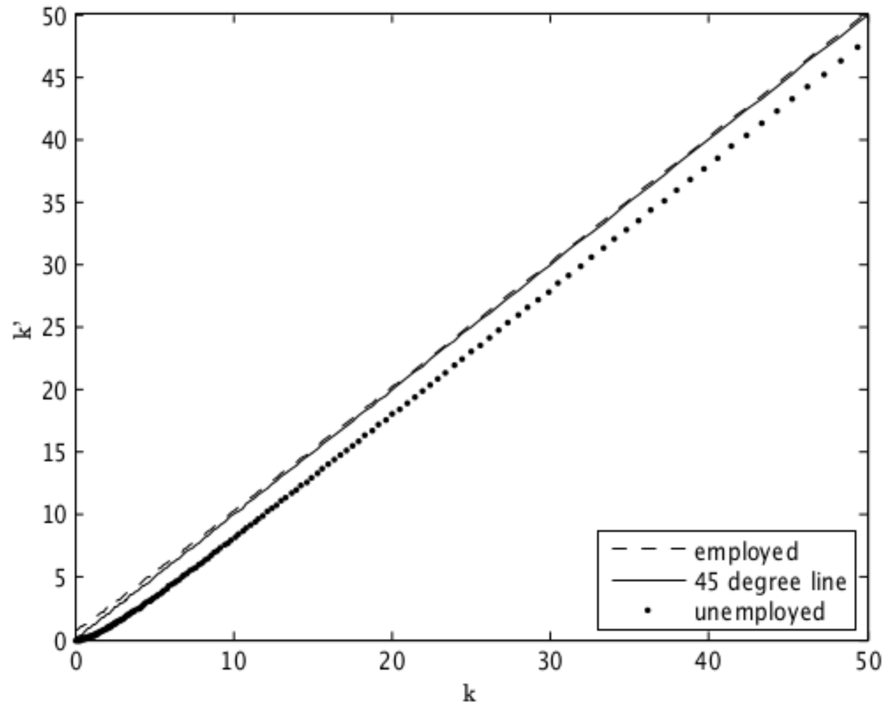

Figure 2: Individual Capital Policy Functions in the KS Model

This period's individual capital (k) plotted against next period's capital choice (k') for aggregate capital=39 and good aggregate state.

With the addition of heterogeneous βs, approximate aggregation still holds. To see why, it is useful to examine the consumption function c', which can be recovered using the savings policy and the budget constraint. Defining individual income as $m_i = w[(1-\tau))\bar{l}e_i + \mu(1-e_i)] + (1+r-\delta)k_i$, if the consumption function at the individual level were:

$$c_i' = \tilde{\phi}_0(z,\Gamma)E(m_i) + \tilde{\phi}_1(z,\Gamma)E(m_i\ln(m_i)) + \tilde{\phi}_2(z,\Gamma)E(m_i\beta_j)$$

Then running a regression on the macro data generated by this economy would result in estimating:

$$c_i' = \tilde{\phi}_0(z,\Gamma)E(m_i) + \tilde{\phi}_1\pi_1(z,\Gamma)E(m_i)\ln E(m_i) + \tilde{\phi}_2\pi_2(z,\Gamma)E(m_i)E(\beta_j)$$

where

$$\pi_1 = \frac{E(m_i \ln(m_i))}{E(m_i)\ln E(m_i)}$$

and

$$\pi_2 = \frac{E(m_i\beta_j)}{E(m_i)E(\beta_j)}$$

The regression coefficients recovered from the micro-level regression will only differ from the coefficients obtained by a regression on the corresponding macro aggregates will only differ substantially from each other if π_1 and π_2 vary highly over the cycle. They both capture how the ϕ coefficients are adjusted if individual consumption is evaluated at its average determinants. π_1 measures the degree of inequality in the economy, since $E(m_i \ln m_i)$ captures entropy. π_2 measures the fraction of income held by heterogenous sub-population $E(m_i\beta_j)$ relative to the size of that sub-population $E(\beta_j)$. Since neither varies significantly over the cycle in the KS β-Dist, the approximate aggregation result holds.

5 Dynamics of the KS β-Dist MPCs

Table 1 shows how the addition of heterogenous discount factors to the benchmark KS model accomplishes a better fit to the wealth distribution: by creating much greater variation in the marginal propensity to consume over the cross section. In the 2007 calibration, the poorest agents[5] consume over 30% of a shock to income, this figure decreases as income levels increase, with the richest consuming 15%. High MPCs are concentrated in unemployed agents, who consume nearly 60% of an income shock in the 2007 calibration, compared to just over 20% for employed agents. In the 2009 calibration, the dispersion in marginal propensity to consume is much greater. The poorest agents consume over 50% of a shock to income, with the richest increasing their share to 25%. Both employed and unemployed agents increase their MPC compared to the 2007 calibration by around 10 percentage points, with unemployed agents consuming over 70% of an income shock.

[5] I report the MPCs for annual income quintile, a very similar pattern of dispersion exists when looking in terms of wealth quintiles.

	Krusell–Smith (KS): β-Dist							
Model	2007 Calibration				2009 Calibration			
Scenario	Baseline	Recession	Expansion		Baseline	Recession	Expansion	
Overall average	0.25	0.27	0.24		0.35	0.37	0.33	
By Income Quintile								
Q1	0.33	0.39	0.26		0.51	0.58	0.43	
Q2	0.2	0.2	0.2		0.34	0.34	0.33	
Q3	0.2	0.21	0.2		0.32	0.33	0.32	
Q4	0.19	0.19	0.18		0.3	0.3	0.29	
Q5	0.15	0.16	0.15		0.25	0.26	0.24	
By employment status								
Employed	0.22	0.23	0.22		0.31	0.31	0.31	
Unemployed	0.58	0.6	0.56		0.73	0.74	0.72	
Time preference parameters[‡]								
$\grave{\beta}$		0.9837				0.9787		
∇		0.00108				0.0172		
PSID 2007 % of wealth held by the richest:	20%	40%	60%	80%				
	80.6	95.1	99.9	100.8				
PSID 2009 % of wealth held by the richest:					20%	40%	60%	80%
					85.5	97.3	100.8	101.3

Table 1: Marginal Propensity to Consume over the Business Cycle - 2006 and 2008 Calibrations Compared

Notes: Annual MPC is calculated by $1-(1-\text{quarterly MPC})^4$. *The scenarios are calculated for the β-Dist models calibrated to the net worth distributions described. For the KS aggregate shocks, the results are obtained by running the simulation over 1,000 periods, and the scenarios are defined as 'Recessions/Expansions': bad/good realization of the aggregate state.* ‡*: Discount factors are uniformly distributed over the interval* $[\grave{\beta}-\nabla, \grave{\beta}+\nabla]$.

Table 1 also shows a contradiction of the model including heterogeneous discount factors. When looking the dispersion of the MPC across the distribution's internal business cycle dynamics (comparing columns 'Recession' and 'Expansion' within a calibration), it is clear to see that in both the 2007 and 2009 calibrations, the MPCs vary very little over the cycle. In the 2007 calibration, the aggregate marginal propensity to consume increases from 0.25 to 0.27 in a recession, and falls to 0.24 in an expansion. Similarly, the 2009 calibration increases from an MPC of 0.35 to 0.37 in a recession and falls to 0.33 in an expansion. However, the implied difference in MPC given the wealth distributions *across* the calibrations are 10 percentage points higher in the recession sample than the pre-recession sample. To make this point more clearly, figure 3 plots these MPC values within and across calibrations. While a recession within a calibration increases the MPC for the very poorest quintile significantly, the rest of the distribution barely changes. In contrast, to fit the wealth distributions across the calibrations requires a shift in the MPCs, and therefore individual policy functions, for the whole of the distribution. The within-calibration MPC changes are consistent with a greater mass of agents becoming closer to the borrowing constraint. The across-calibration MPCs suggests that preferences across the distribution have fundamentally changed, or something else has occurred in the Great Recession to shift individual policy functions that the model is not capturing.[6]

Is it plausible that consumption functions could have shifted during the Great Recession? According to the logic of the KS-β-Dist model, the determinants that could plausibly generate a shift in cross-sectional MPCs are 1) an increase in the variance of income shocks, manifesting as a shift in the consumption function; 2) borrowing constraints becoming tighter during recessions. Point 1) is described extensively in Carroll et al. (2014) in a calibrated incomplete markets heterogeneous agents model. They consider two types of income shocks: permanent (highly persistent shocks) and transitory ones. Decomposing income into these two components is a standard way of thinking about income shocks in the literature, and goes back to Friedman (1957). Carroll et al. (2014) show that increases in the variance of permanent shocks do not have a significant effect on the consumption function - this stands to reason, since recessions are temporary, there is little correlation between permanent shocks and aggregate uncertainty. However, an increase in the variance in

[6]Note that these results are obtained without making any special assumptions about the nature of the Great Recession in terms of its severity or duration. Krueger et al. (2016) explores calibrating the KS model aggregate shocks to the frequency of observed severe recessions - defined as historical periods where unemployment exceeded 9% for at least one quarter and remains above 7% thereafter. Employing such a calibration does not change the qualitative result.

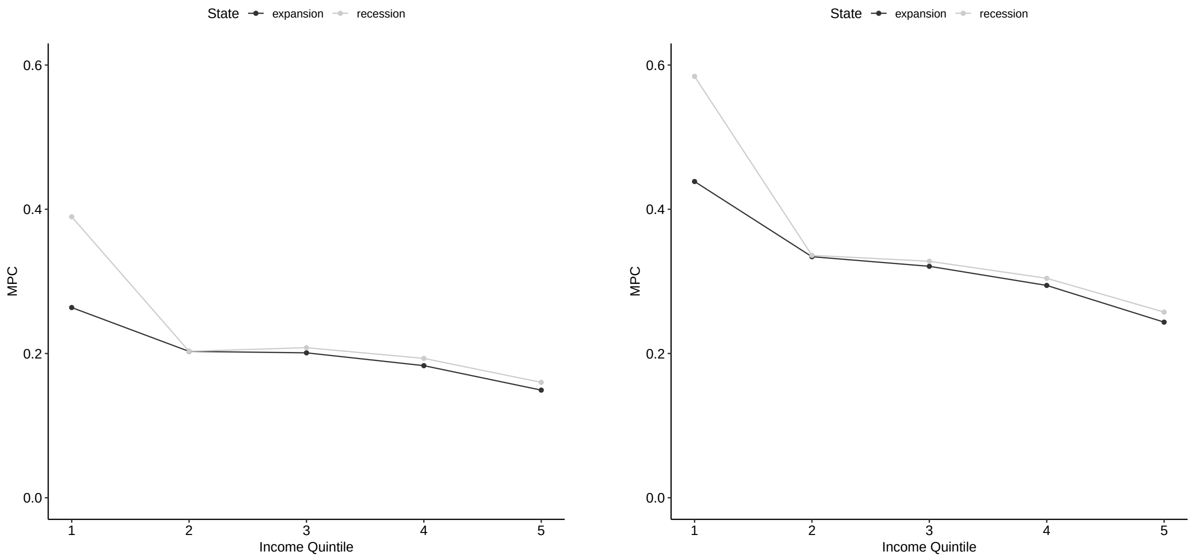

Figure 3: KS β-Dist Simulated MPC at the Income Quintiles

MPC at the income quintiles generated from 2007 calibration in recession (left) and 2009 calibration in expansion (right). Annual MPC is calculated by $1 - (1-\textit{quarterly MPC})^4$. *Calculated for the* β*-Dist models calibrated to the net worth distribution for a given year. The results are obtained by running the simulation over 1,000 periods. Recession/Expansion MPCs are defined as averaging over bad/good realizations of the aggregate state.*

transitory shock affects the consumption function significantly, by shifting it upwards and steepening the function for the poorest households, and also manifesting as a larger aggregate MPC. There is certainly a large literature on the deleterious effects of recessions on labour markets, see for example Elsby et al. (2010) on the Great Recession. Moreover, Guvenen et al. (2014) find that the left-skewness of income shocks is countercyclical.

For point 2), there is ample evidence that borrowing constraints are tighter during periods of falling economic activity. Particularly in financial crises, the sharp drop in lending can worsen and prolong economic downturns (Bernanke and Gertler (1989); Kiyotaki and Moore (1997)). Ludvigson (1999) finds that predictable growth in consumer credit is significantly related to consumption growth in the macroeconomic time series. Gross (1994) estimate using bankruptcy data that MPC out of liquid funds is 20-30% higher during the Great Recession. However, as I show in the appendix, in order to fit the proportion of households with negative wealth, the standard incomplete markets model actually implies a *decrease* (ie loosening) of the borrowing constraint over the Great Recession. While this mechanism certainly deserves to be explored, because the current study is concerned with reconciling the incomplete market heterogenous agent model with empirical evidence, I leave it to future research.

The focus of the following section is to focus on point 1). I do this by estimating the marginal propensity to consume out of transitory income over the cross-section of income and the business cycle while controlling for point 2), borrowing constraint changes.

6 Reduced-form Estimates of the MPC

In the first stage regressions, detailed in section 7, I estimate the predictable elements of consumption and income changes in order to leave only permanent and transitory shocks in the residual for the second stage regressions which recover the MPC. For this first stage, I use data from the Panel Study of Income Dynamics public use dataset, a nationally representative longitudinal survey of US households which is completed biennially. I combine it with PSID wealth supplements which provides detailed data on wealth and assets. For the pre- and recession periods (2003-2007 and 2009-2013) I joined each adjacent year by finding those households that had the same household head. In other words, in order to increase the sample size, I only require the households to be the same within the periods, but they can differ across the periods. The appendix repeats the analysis with a smaller, balanced panel over 2003-2013 and finds very similar results. The additional criteria for inclusion in the sample include having a household head between the 22-65 and non-blank data not only for consumption and income, but also all demographic and other control variables. In this section, I detail the properties of the dependent and independent variables for the first state regression.

It is important to note that the PSID is known to undersample the most wealthy 1-2% (Pfeffer et al. (2016)), so I use the weights and strata information provided. The design of the PSID is a complex survey, therefore working with unweighted PSID data would violate the assumption that observations are i.i.d., since the complex survey design creates data with correlations between observations and unequal sampling probabilities. It is also a top-coded survey for purposes of anonymity, I drop these observations since the true values are unknown. This is innocuous since the very top of the income and wealth distributions have so little mass, they do not matter for aggregate MPC.

6.1 Consumption data in the PSID

Since its release, consumption data in the PSID is starting to become much more widely used in research, notably Blundell et al. (2016). However its introduction has been gradual, with extra expenditure categories added over time. Figure 4 shows the amounts calculated including food, transport, childcard, healthcare, education and housing (consumption). In the 2004 survey, vacations, recreations and clothing were added to the survey (consumption plus). There is some difference in the consumption amounts including and excluding the extra categories, both in levels and in co-movement with the cycle. Unfortunately, because the method of estimating marginal propensities requires a minimum of three time periods (see section 7 below), I use consumption rather than consumption plus in the estimations. I drop observations with negative consumption values, around 0.01% of the sample per year.[7]

Table 2 compares the mean values of consumption and consumption plus to the corresponding values reported from the Consumer Expenditure Survey (CEX). Clearly, the PSID means are smaller relative to the CEX means, even when including the additional expenditure categories. It is likely that, even with extra consumption categories, because the PSID coverage is not as extensive as the CEX, it is underestimating consumption. It is also possible that the mean values are lower because the PSID is know to undersample the richest households. However, the dynamics of consumption over the Great Recession are very similar in both the CEX and the PSID.

[7]Negative consumption values come from imputed consumption values from the PSID, which uses a linear regression to predict missing consumption values.

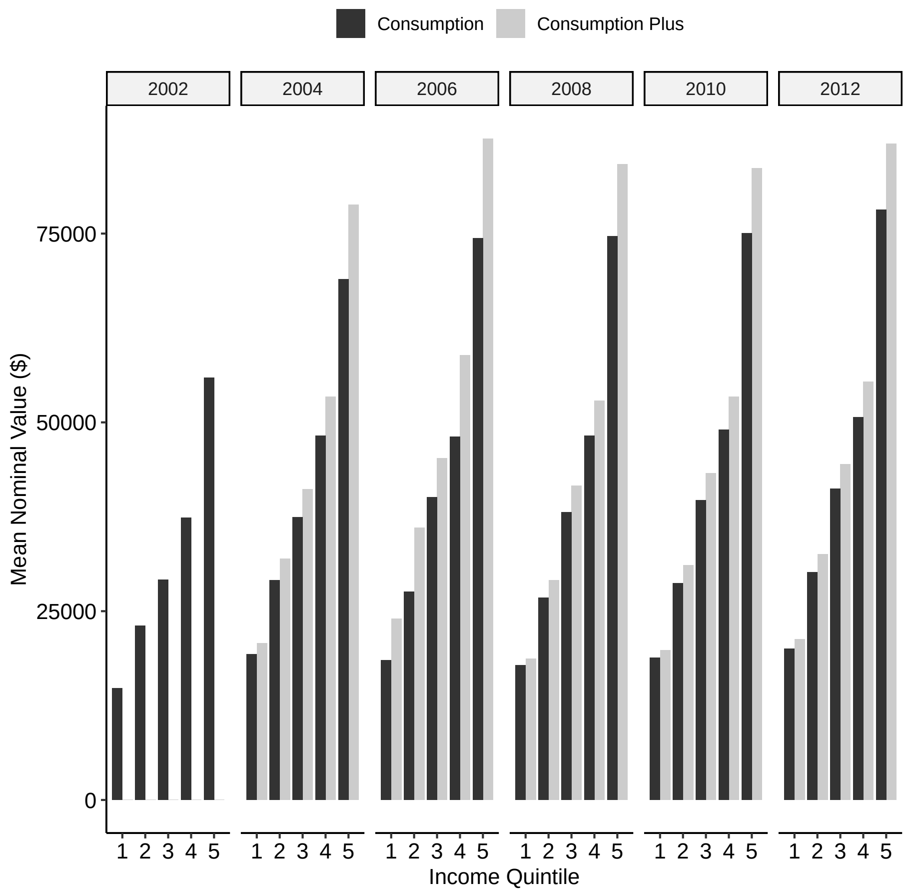

Figure 4: Mean Consumption in the PSID

Consumption = food + transport + childcare + healthcare + education + housing; Consumption Plus = consumption + vacations + recreation + clothing

Table 2: Mean Consumption Comparison

Year	$\bar{c}_{i,t}^{CEX}$	$\bar{c}_{i,t}$	$\bar{c}_{i,t}^{\text{PLUS}}$
2002	40677	28573	NA
2004	43395	35612	39559
2006	48400	38870	43112
2008	50486	37816	41837
2010	48109	37890	41502
2012	51442	38213	41878

$\bar{c}^{CEX}$ is mean consumption available from the CEX
$\bar{c}_t$ is the corresponding estimate in the PSID data
$\bar{c}_t^{\text{PLUS}}$ is mean PSID consumption plus extra expenditure categories, author's calculations.

6.2 Measuring After-Tax Income

The PSID reports total taxable income of the household in the preceding year to the survey, ie., $y_{it} + T_t$ where T is a lump sum tax. However, for the purposes of understanding consumption and saving responses to income changes, I calculate after-tax income. It is important to use after-tax income because changes in taxation could change consumption responses. To do this, I use the TAXSIM program from NBER, which estimates the tax lump-sum given details of income, family composition and deductions. I use the method outlined in Butrica and Burkhauser (1997) by adapting code from Kimberlin et al. (2014) to include 2013 data. Clearly, it is very important that after-tax income is well measured, so table 3 compares average (mean) after-tax income estimated from the CEX versus the corresponding quantity estimated from the PSID using TAXSIM.

Table 3: After-tax Income Comparison

Year	$\bar{y}_t^{CEX}$	$\bar{y}_t^{alt}$
2002	46934	53858
2004	52287	61565
2006	58101	65033
2008	61774	62466
2010	60712	62766
2012	63370	67629

$\bar{y}^{CEX}$ is mean after-tax income available from the CEX
$\bar{y}_t$ is the corresponding estimate in the PSID data using TAXSIM from the NBER, author's calculations

The PSID values, although larger, are closer to CEX estimates in this case, likely because the PSID covers more income categories, though it could also be due to an underestimate of taxes in the TAXSIM programme. However, the dynamics over the sample are again similar in both surveys.

6.3 Explanatory variables for Predictable Consumption/Income

6.3.1 Demographic and Economic variables

Demographic controls are all taken to be the values of the household head (which in the PSID are overwhelmingly male). Controls include a dummy for year of birth, dummies for education which takes three levels: up to high school education (low), college educated (medium), or some postgraduate (high). I control for employment status which can take values employed, unemployed, retired or inactive. Also included are race dummies (taking values white, black, and other), and continuous variables for family size and number of kids in the family unit. I include dummies for whether the family has extra income coming from those outside the household, extra dependents outside the family unit and categorical dummies for region: North East, Midwest, South, West. Finally, I include controls for total net wealth level in thousands of dollars.

6.3.2 Controls for the borrowing constraint

Following Kaplan and Violante (2014) I distinguish two types of household: poor hand-to-mouth (P-HtM) and wealthy hand-to-mouth (W-HtM) which, due to a lack of liquid assets on hand, can come up against binding borrowing constraints. Poor hand-to-mouth households are defined as being at the credit limit when their illiquid wealth holdings and liquid wealth holdings are not positive, and their cash-on-hand and available credit is less than half their yearly income; i.e.:

$$a_{it} \leq 0, \quad m_{it} \leq 0 \quad \text{and} \quad m_{it} \leq y_{it}/2 - \underline{m}_{it}$$

where a_{it} is holdings of illiquid wealth by household i in period t, m_{it} is average balances of liquid wealth over period t, y_{it} is total household income and $\underline{m}_{it}$ is the household's credit limit.

Wealthy hand-to-mouth households are defined similarly, with the key difference being that they own positive illiquid assets, i.e.:

$$a_{it} > 0, \quad m_{it} \leq 0 \quad \text{and} \quad m_{it} \leq y_{it}/2 - \underline{m}_{it}$$

In the PSID, illiquid wealth is calculated as the sum of home equity, other real estate equity, private annuities and other assets. Liquid wealth is the sum of checking and saving accounts, money market funds, certificates of deposits, savings bonds, treasury bills, stocks net of liquid debt which includes all debt other than mortgage debt. Figures 5 and 6 show the fraction of Poor HtM and Wealthy HtM households in the PSID, the fractions are consistent with those reported by Kaplan and Violante (2014). The fraction of wealthy hand-to-mouth households falls over the Great Recession while the fraction of poor rises. Interestingly, calculated for the quintiles of income, the fraction of P-HtM exhibits a strong negative relationship, while the fraction of W-HtM is an inverted U-shape in each year. The fractions within a given quintile for the wealthy hand-to-mouth are relatively stable over the Great Recession, while the fraction of P-HtM in the poorest income quintiles rises markedly. This is suggestive evidence that borrowing constraints became more likely to bind for the very poorest hand-to-mouth households during the recession, with less of a noticeable effect for the wealthy hand-to-mouth. To control for binding borrowing constraints in the consumption and income regressions, I use a dummy variable for Poor and Wealthy HtM households, equal to 1 when a household falls into the respective categories and 0 otherwise.

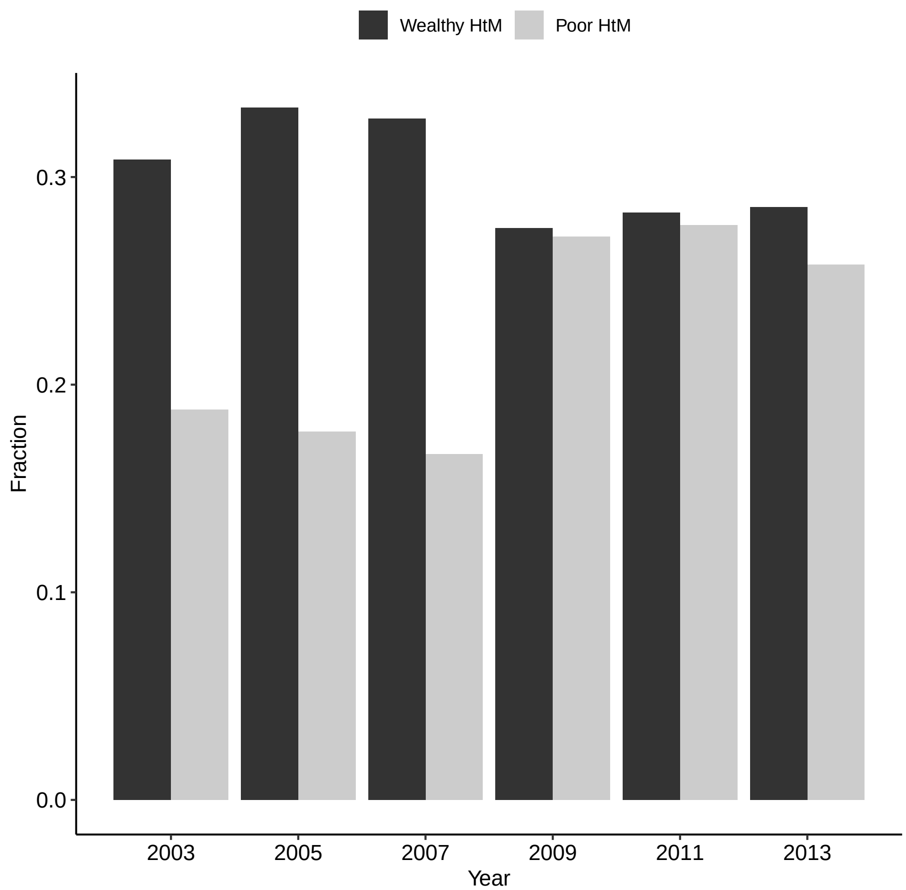

Figure 5: Hand-to-mouth households in the PSID

Fraction of Poor hand-to-mouth (P-HtM) and Wealthy hand-to-mouth (W-HtM) consumers in the PSID by year, author's calculations

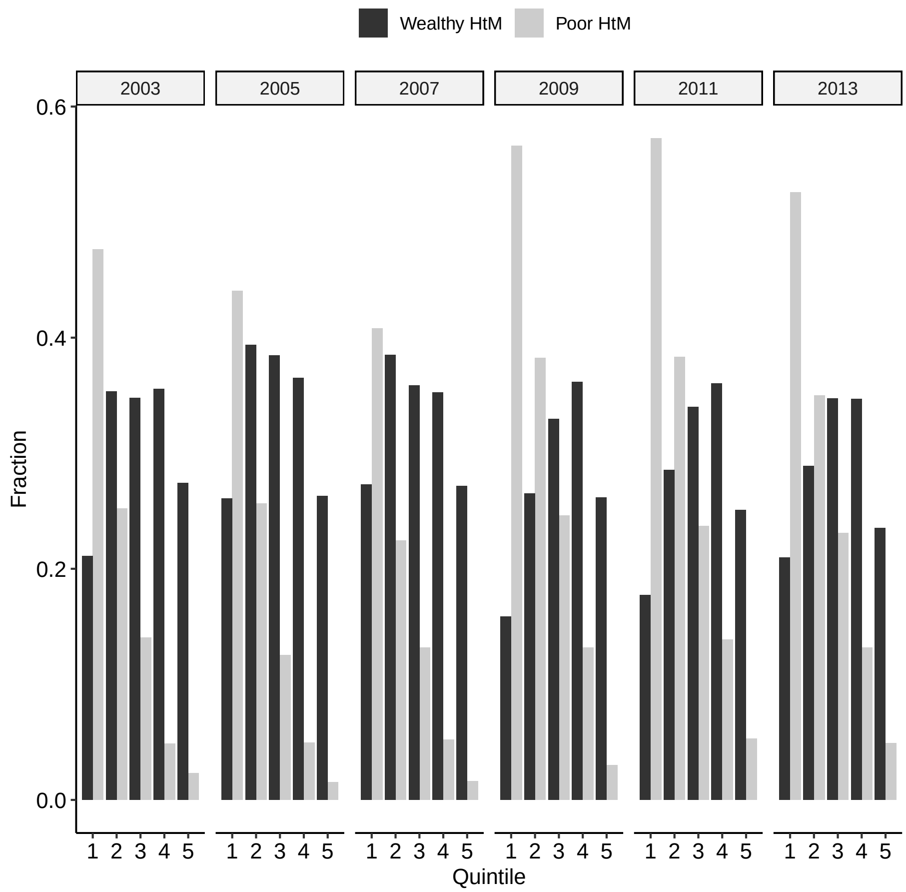

Figure 6: Hand-to-mouth households by income quintile in the PSID

Fraction of Poor hand-to-mouth (P-HtM) and Wealthy hand-to-mouth (W-HtM) consumers in the PSID by year and quintile, author's calculations.

7 Method

7.1 Estimating the Marginal Propensity to Consume and Save

Following Blundell et al. (2008) and Kaplan et al. (2014), assume income follows the process:

$$\log Y_{i,t} = \mathbf{Z}'_{it}\mathbf{\Phi}_t + P_{i,t} + \epsilon_{i,t} \tag{5}$$

where i is an individual at time t, Y is income, $\mathbf{Z}$ is a set of observable income characteristics, $P_{i,t} = P_{i,t-1} + \xi_{i,t}$ is a martingale permanent income process with i.i.d. shock ξ; and ϵ is an i.i.d. transitory income shock. Blundell et al. (2008) show that such a specification provides a good approximation to the solution of a life cycle optimization problem where agents have CRRA utility.

By estimating equation 5 and recovering the first-differenced residuals, I obtain unexplained income growth:

$$\Delta\widehat{y_{i,t}} = \xi_{i,t} + \Delta\epsilon_{i,t}$$

where $\widehat{y_{i,t}} = \log Y_{i,t} - \mathbf{Z}'_{it}\hat{\mathbf{\Phi}}_t$.

Consumption is assumed to be subject to the same processes but with loading factors (marginal propensities) on permanent and transitory income shocks $\psi^P_{i,t}$ and $\psi^T_{i,t}$, giving unexplained growth as:

$$\Delta\widehat{c_{i,t}} = \psi^P_{i,t}\xi_{i,t} + \psi^T_{i,t}\Delta\epsilon_{i,t}$$

where $\Delta\widehat{c_{i,t}}$ is estimated first-differenced consumption residuals $\widehat{c_{i,t}} = \log Y_{i,t} - \mathbf{Z}'_{it}\hat{\mathbf{\Psi}}_t$.

The covariance restriction necessary for identification of the marginal propensity to save from transitory income shocks are that individuals have no foresight about future shocks, i.e. $\text{cov}(\Delta c_{i,t}, \epsilon_{i,t+1}) = \text{cov}(\Delta c_{i,t}, \xi_{i,t+1}) = 0$.

The true marginal propensity to consume out of a transitory shock is given by:

$$\text{MPC}_t = \frac{\text{cov}(\Delta c_{i,t}, \epsilon_{i,t})}{\text{var}(\epsilon_{i,t})}$$

Which, using the covariance restrictions, can be estimated consistently via an instrumental variable regression of $\Delta\widehat{c_{it}}$ on $\Delta\widehat{y_{i,t}}$ instrumented by $\Delta\widehat{y_{i,t+1}}$:

$$\widehat{\text{MPC}_t} = \frac{\text{cov}(\Delta\widehat{c_{i,t}}, \Delta\widehat{y_{i,t+1}})}{\text{cov}(\Delta\widehat{y_{i,t}}, \Delta\widehat{y_{i,t+1}})} \tag{6}$$

By using this methodology on a panel simulated from an incomplete insurance model in which they can compare estimated and true values, Kaplan and Violante (2010) show that this method works well for estimating transitory shocks and is not biased in the presence of binding borrowing constraints. The estimation requires 3 periods of data for consumption: $t-1, t, t+1$ so I drop households with fewer than 3 consecutive years of observations in each of the two time periods. To get an estimate of the marginal propensity to consume across the distribution, I estimate equation 6 for the income quintiles at the beginning of the pre-recession and recession sample period (2003 and 2009).

8 Results

In this section I begin with a summary of the first-stage regression output, the result of estimating the income process in equation 5 and the corresponding consumption equation. I then discuss the second-stage regression results which use the residuals from the first-stage estimation to estimate the marginal propensity to consume. I first discuss the results over the entire distribution for the pre-recession and recession periods and then show the estimates over the income quintiles.

Table 4 details the results of the first-stage regressions to extract the predictable parts of consumption and income. Though a means to estimating the marginal propensity to consume, these results merit inspection in their own right. Year dummies show time growth in consumption and income which is greater prior to the Great Recession. As we might expect, income levels are greater by approximately 6 and 17 percent for the medium and high educated, respectively, relative to the low educated. This greater income level does not transfer fully into consumption, which are greater by approximately 3 and 10 percent for the same groups. Being a race other than white is associated with lower income and consumption levels, while family size increases both. More kids in the household is associated with lower consumption and income levels, as is being unemployed, retired or inactive. Comparing the pre-recession and recession periods by employment status, the penalty to consumption and income increases for all non-employment in the recession period. Being in a region that is not the North East is associated with a consumption and income penalty of between 2 and 5 percent, but this does not seem to change dramatically with the business cycle. Both Poor Hand-To-Mouth households (those with cash on hand and available credit at less than half their yearly income and a negative illiquid net worth position), and Wealthy Hand-To-Mouth households (those with cash on hand and available credit at less than half their yearly income but a positive illiquid net worth position) see a penalty to consumption and income. Those that are P-HtM have lower consumption levels of approximately 7 percent relative to non-HtM households, and lower income levels of 12 percent. Meanwhile the associated reduction for W-HtM households is not significant from zero for consumption, but associated with 5 percent lower income levels. This is suggestive evidence that borrowing constraints are more likely to be binding for P-HtM, since it is associated with lower consumption.

Table 4: First Stage Regressions

	$\log(\widehat{c_{it}})$		$\log(\widehat{y_{it}})$	
	2002-2006	2008-2012	2002-2006	2008-2012
	(1)	(2)	(3)	(4)
Year=2004	0.053***		0.029***	
	(0.002)		(0.003)	
Year=2006	0.077***		0.051***	
	(0.004)		(0.003)	
Year=2010		0.007**		0.006
		(0.002)		(0.003)
Year=2012		0.017***		0.025***
		(0.002)		(0.004)
Education=Medium	0.034***	0.038***	0.063***	0.063***
	(0.006)	(0.006)	(0.008)	(0.009)
Education=High	0.105***	0.109***	0.167***	0.175***
	(0.006)	(0.007)	(0.010)	(0.010)
Race=Black	−0.045***	−0.037***	−0.069***	−0.068***
	(0.007)	(0.007)	(0.008)	(0.007)
Race=Other	−0.009	−0.004	−0.047**	−0.053**
	(0.014)	(0.009)	(0.014)	(0.015)
Family Size	0.069***	0.084***	0.091***	0.108***
	(0.004)	(0.004)	(0.005)	(0.005)
Number of Kids	−0.047***	−0.058***	−0.071***	−0.081***
	(0.004)	(0.004)	(0.005)	(0.006)
Status=Unemployed	−0.028*	−0.058***	−0.082**	−0.107***
	(0.011)	(0.009)	(0.019)	(0.017)
Status=Retired	−0.055***	−0.075***	−0.089***	−0.106***
	(0.006)	(0.005)	(0.010)	(0.008)
Status=Inactive	−0.061***	−0.097***	−0.121***	−0.152***
	(0.007)	(0.007)	(0.010)	(0.011)
Extra Family Income	0.012*	0.003	0.027**	0.027**
	(0.005)	(0.007)	(0.007)	(0.008)
Region=Midwest	−0.043***	−0.053***	−0.036**	−0.053**
	(0.008)	(0.010)	(0.011)	(0.012)
Region=South	−0.027**	−0.033**	−0.032*	−0.038**
	(0.009)	(0.010)	(0.013)	(0.012)
Region=West	−0.022*	−0.042***	−0.037**	−0.050***
	(0.008)	(0.009)	(0.009)	(0.010)
Kids outside Family Unit	0.025**	0.023***	0.059***	0.062***
	(0.006)	(0.005)	(0.012)	(0.008)
Poor-HtM	−0.078***	−0.061***	−0.128***	−0.122***
	(0.006)	(0.006)	(0.008)	(0.009)
Rich-HtM	−0.005	−0.012	−0.053***	−0.056***
	(0.005)	(0.006)	(0.007)	(0.008)
Total Wealth ($1000s)	0.0002**	0.0004*	0.001**	0.001**
	(0.0001)	(0.0001)	(0.0001)	(0.0002)
Constant	11.383***	11.668***	11.731***	11.693***
	(0.015)	(0.022)	(0.022)	(0.052)
Year of Birth	Yes	Yes	Yes	Yes
R^2	0.4	0.45	0.43	0.51
N	13,281	14,820	13,281	14,820

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Survey-weighted generalised least squares regression for the years 2003, 2005, 2007 (columns 1 and 3) and 2009, 2011, 2013 (columns 2 and 4) in the PSID.

Table 5 details the results of the instrumental variable regression estimating the marginal propensities to consume across the whole distribution over the pre-recession and recession periods - giving an estimate of the aggregate MPC. The findings for the marginal propensity to consume are in line with other empirical evidence, suggesting that the MPC out of transitory income is non-zero, a result which is significant at the 1% level. The results suggest that on average, households consumed 10% of a transitory income shock over the pre-recession period 2003-07, and this increased to 16% in the recession period, 2009-13. The difference in the estimated MPCs is also significant at 1% over the pre-recession and recession periods. Note that the R^2 statistics in these regressions are much lower than those of the first-stage regressions in table 4. We should expect this by construction since we are running a regression on residuals. Indeed, classical theory would predict an R^2 of zero, although we have known since at least Hall (1978) that there is a predictive relationship, albeit much lower than the relationship to changes in permanent income.

	MPC	
	03-07	09-13
	0.1^{***}	0.162^{***}
	(0.027)	(0.029)
R^2	0.052	0.075
N	4535	5160

Table 5: Whole distribution estimates of MPC

Estimates of the marginal propensity to consume MPC estimated on the whole sample of the income distribution on PSID data. Pre-recession: 2003-2007, recession: 2009-2013. Standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Turning to the estimates over the income distribution, tables 6 and 7 reports the regression output over the pre-recession and recession periods which figure 7 shows graphically. On the left of figure 7, the MPC is downward sloping, implying an upward-sloping consumption function. The very poorest quintile consumes between 20-40% of a transitory income shock, this reduces down to a value that is not significantly different to 0% for the highest income quintiles. Note that we see downward sloping relationship between MPC and income despite controlling for borrowing constrained households, a factor which is known to generate such a relationship. R^2 values are surprisingly high for the first, second and third quintile, but drop to near zero - as would be predicted by the permanent income hypothesis - for the richest 2 quintiles. This is suggests that there is an certainly an important role for hetergeneous preferences in matching the MPC over the distribution of income - with poorer households acting as hand-to-mouth consumers and richer households as consumption smoothers. Indeed, the point estimates for some of the distribution in the pre-recession sample are remarkably similar to the KS β-Dist generated 2007 MPCs. Comparing table 1 with the reduced-form estimates, we see that the poorest quintiles both consume approximately 30% of their transitory income, Q2 and 3 consume 20% in the KS β-Dist model and 28 and 13% respectively in the estimates. There is a stark difference however, in the upper end of the distribution - KS β-Dist calibration generates MPCs of 19 and 15%, whereas the estimation does not find coefficients significantly different from zero.

Lastly, we see that the estimates of the MPCs are noisily estimated and, although point estimates are close together, we cannot completely rule in or out the possibility that the consumption function shifted over the Great Recession. However, we certainly can rule out the 2009 calibration values of the KS β-Dist for MPCs - which generated MPCs for the poorest of over 50%, declining to 25% for the income-rich. At all points, the generated MPCs exceed the estimated MPCs beyond the confidence intervals of the estimation.

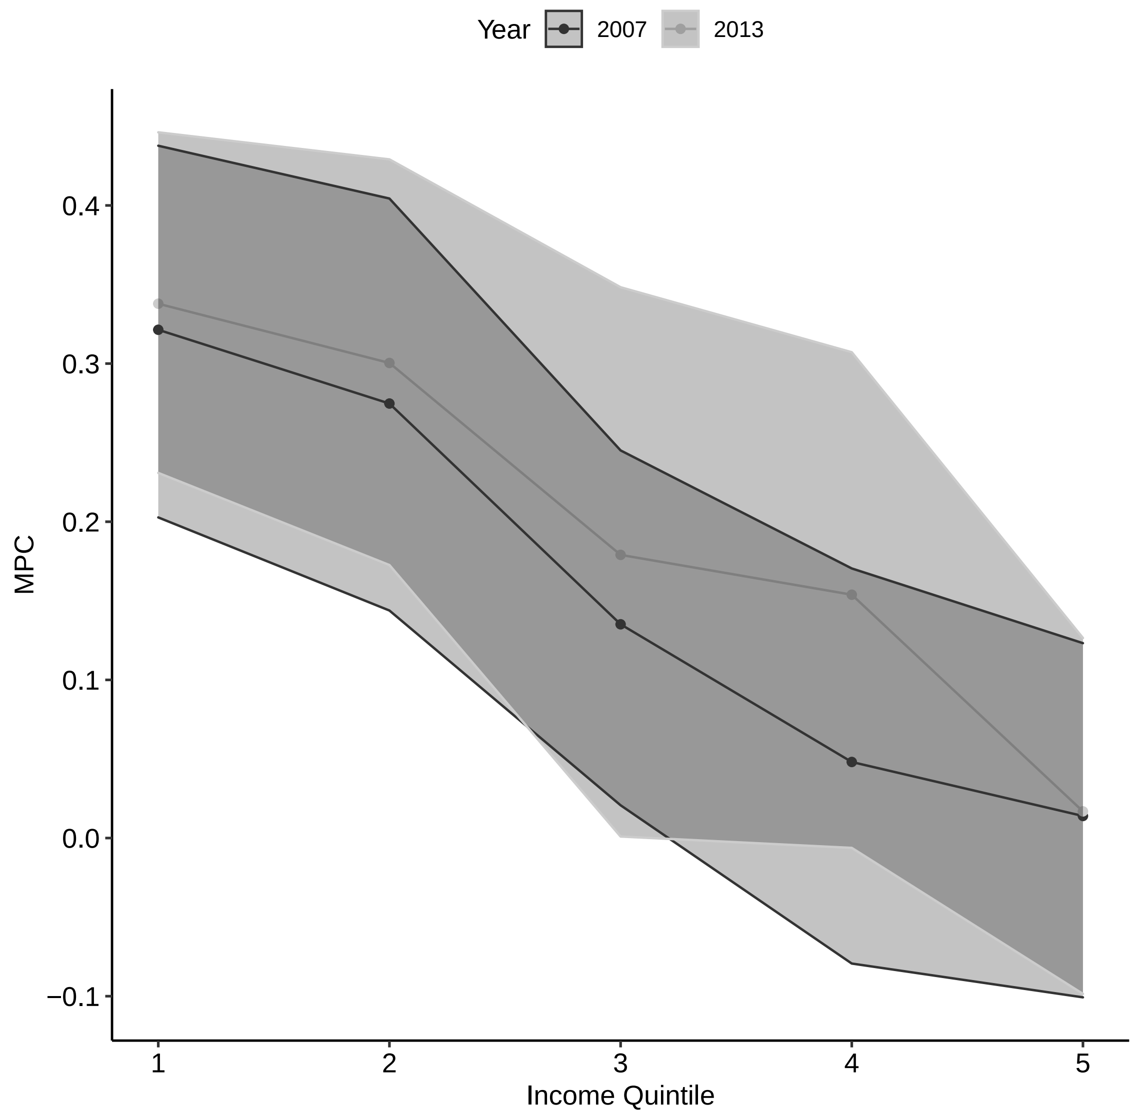

Figure 7: Empirical Distribution of MPC by year and quintile

Estimated marginal propensity to consume by income quintile and time: pre-recession period (2003-2007), recession (2009-2013).

Table 6: IV regression: $\Delta c_{i,t}$ 2003-2007 for Income Quintiles

	$\Delta\widehat{c_{i,t}}$				
	(1)	(2)	(3)	(4)	(5)
$\Delta\widehat{y_{i,t}}$	0.321***	0.275***	0.135**	0.048	0.014
	(0.059)	(0.065)	(0.057)	(0.063)	(0.056)
N	879	882	878	881	856
R^2	0.258	0.154	0.073	0.019	0.006
Residual Std. Error	0.078 (df = 877)	0.091 (df = 880)	0.103 (df = 876)	0.116 (df = 879)	0.140 (df = 854)

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Instrumental variable regression of $\Delta\widehat{c_t}$ on $\Delta\widehat{y_t}$ instrumented by $\Delta\widehat{y_{t+1}}$ for the income quintiles, 2003-2007

Table 7: IV regression: $\Delta c_{i,t}$ 2008-2012 for Income Quintiles

			$\Delta\widehat{c_{i,t}}$		
	(1)	(2)	(3)	(4)	(5)
$\Delta\widehat{y_{i,t}}$	0.338***	0.300***	0.179**	0.154*	0.017
	(0.054)	(0.065)	(0.087)	(0.079)	(0.057)
N	984	984	981	978	960
R^2	0.175	0.127	0.091	0.072	0.008
Residual Std. Error	0.088 (df = 982)	0.098 (df = 982)	0.114 (df = 979)	0.118 (df = 976)	0.136 (df = 958)

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Instrumental variable regression of $\Delta\widehat{c_t}$ on $\Delta\widehat{y_t}$ instrumented by $\Delta\widehat{y_{t+1}}$ for the income quintiles, 2009-2013

9 Conclusion

In this paper, I explored the distributional dynamics of the standard incomplete-markets, heterogenous agent model known as the Krusell-Smith model. Since software developments in the field allow these types of models to be applied in policy questions much more readily, and particularly in questions where the distribution matters, it is incumbent on economists to make sure that these dynamics are realistic. I estimated the marginal propensity to consume out of transitory income over the pre-recession and recession periods and over the quintiles of the distribution. I provided evidence that the MPC varies greatly over the distribution of income, with income-poor housholds consuming around 20-40% of a transitory income shock and income-rich housholds consuming none. As Krusell and Smith themselves point out, while it fits the income-poor end of the distribution reasonably well at a point in time, the Krusell-Smith model does not capture the richer end of the distribution. Moreover, changing the dispersion of heterogeneous preferences to fit the dynamics of the wealth distribution implies a shift in the consumption function, and distribution of MPCs which is at odds with the data. The magnitude of the shift in the consumption implied by a standard incomplete markets heterogeneous agent model that is calibrated to the wealth distribution in the pre-recession and recession periods is far too high to be consistent with the evidence presented in this paper. Therefore, analysis of the distributional implications of economic policies with such models could result in predictions far at odds from true values. Future research should focus on bridging this gap.

References

Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2018). When inequality matters for macro and macro matters for inequality. *NBER macroeconomics annual.*, 32(1):1–75.

Bernanke, B. and Gertler, M. (1989). Agency costs, net worth, and business fluctuations. *The American Economic Review*, 79(1):14–31.

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

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

Bricker, J., Kennickell, A. B., Moore, K. B., and Sabelhaus, J. (2012). Changes in U.S. family finances from 2007 to 2010: evidence from the Survey of Consumer Finances. *Federal Reserve Bulletin*, (June).

Butrica, B. and Burkhauser, R. (1997). Estimating federal income tax burdens for panel study of income dynamics (PSID) families using the national bureau of economic research TAXSIM model. Working Paper 12, Syracuse University.

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

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.

Carroll, C. D., Slacalek, J., and Tokuoka, K. (2014). The distribution of wealth and the MPC: Implications of new European data. *The American Economic Review*, 104(5):107–111.

Castañeda, A., Giménez, J. D., and Rull, J. V. R. (2003). Accounting for the U.S. earnings and wealth inequality. *Journal of Political Economy*, 111(4):818–857.

De Nardi, M. and Fella, G. (2017). Saving and Wealth Inequality. *Review of Economic Dynamics*, 26:280–300.

De Nardi, M., French, E., and Benson, D. (2011). Consumption and the great recession. Working Paper 17688, National Bureau of Economic Research.

Elsby, M., Hobijn, B., and Sahin, A. (2010). The labor market in the great recession. *Brookings Papers on Economic Activity*, 41(1 (Spring)):1–69.

Friedman, M. (1957). *A Theory of the Consumption Function*. National Bureau of Economic Research, Inc.

Gross, D. B. (1994). The investment and financing decisions of liquidity constained firms.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Hall, R. E. (1978). Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy*, 86(6):971–87.

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

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

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. Working Paper 20073, National Bureau of Economic Research.

Kimberlin, S., Kim, J., and Shaefer, H. L. (2014). An updated method for calculating income and payroll taxes from psid data using the nber's taxsim, for psid survey years 1999 through 2011. Manuscript, University of Michigan.

Kiyotaki, N. and Moore, J. (1997). Credit Cycles. *Journal of Political Economy*, 105(2):211–248.

Krueger, D., Mitman, K., and Perri, F. (2016). *Macroeconomics and Household Heterogeneity*, volume 2 of *Handbook of Macroeconomics*, pages 843–921. Elsevier.

Krueger, D. and Perri, F. (2005). Understanding Consumption Smoothing: Evidence from the U.S. Consumer Expenditure Data. *Journal of the European Economic Association*, 3(2-3):340–349.

Krusell, P. and Smith, Jr., A. A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):pp. 867–896.

Lawrance, E. C. (1991). Poverty and the rate of time preference: Evidence from panel data. *Journal of Political Economy*, 99(1):54–77.

Ludvigson, S. (1999). Consumption and credit: A model of time-varying liquidity constraints. *The Review of Economics and Statistics*, 81(3):434–447.

Maliar, L., Maliar, S., and Valli, F. (2010). Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm. *Journal of Economic Dynamics and Control*, 34(1):42–49.

Pfeffer, F., Schoeni, R., Kennickell, A., and Andreski, P. (2016). Measuring wealth and wealth inequality: Comparing two US surveys. *Journal of Economic and Social Measurement*, 41(2):103–120.

Warner, J. T. and Pleeter, S. (2001). The personal discount rate: Evidence from military downsizing programs. *American Economic Review*, 91(1):33–53.

Appendix - Borrowing Constraint

I consider a second modification to the benchmark KS model which implies a contradiction when compared to PSID data. This modification to the benchmark model is the lowering of the borrowing constraint. To my knowledge, no other paper has considered this modification - probably because it is trivial - in terms of the solution it simply shifts the distribution by a constant and has no effects on the dynamics.[8] However, I consider it firstly because a significant proportion of individuals in the PSID hold zero or negative net wealth: 12.7% in 2006, 16.1% in 2008. Secondly, I will show that changing the borrowing constraint implies a contradiction with empirical evidence that finds that credit constraints are higher in recessions than in normal economic times.[9] To fit the evidence of significant holdings of zero or negative wealth, I add an outer loop to the algorithm which solves the model, which uses the 2006 β-Dist calibration and lowers the budget constraint by 0.1 until the gap between the model-generated percentage of agents holding negative wealth is within $\epsilon = 0.01$ of the 2006 empirical figure. Because the algorithm fits the percentage of people with negative wealth, rather than the cumulative wealth held, it doesn't do a great job of improving the fit of the empirical Lorenz Curves - compare figures 8, the empirical Lorenz curve estimated from the PSID and the KS β-Dist model generated 9. In the data, both the fraction of individuals holding negative wealth and the amounts of negative wealth increase substantially in the recession relative to the pre-recession estimation. Figure 9 shows the Lorenz curve in the good and bad aggregate state within the 2007 calibration; neither the fraction of agents holding negative wealth nor the amounts of negative wealth held change significantly in the dynamics over the cycle (which should not be surprising, given that this modification does not change the dynamics of the β-Dist model).

[8] Not to mention, it takes a long time to solve - one iteration of the Carroll et al. (2017) Mathematica code takes around 8 hours on my MacBook Pro 2013 laptop with 2.6 GHz processor and 8 GB RAM; I rewrote the algorithm based on Maliar et al. (2010) code in Matlab R2016b to get a total running time of 24 hours for all iterations of the borrowing constraint.

[9] Chapter 3 provides an overview of this evidence.

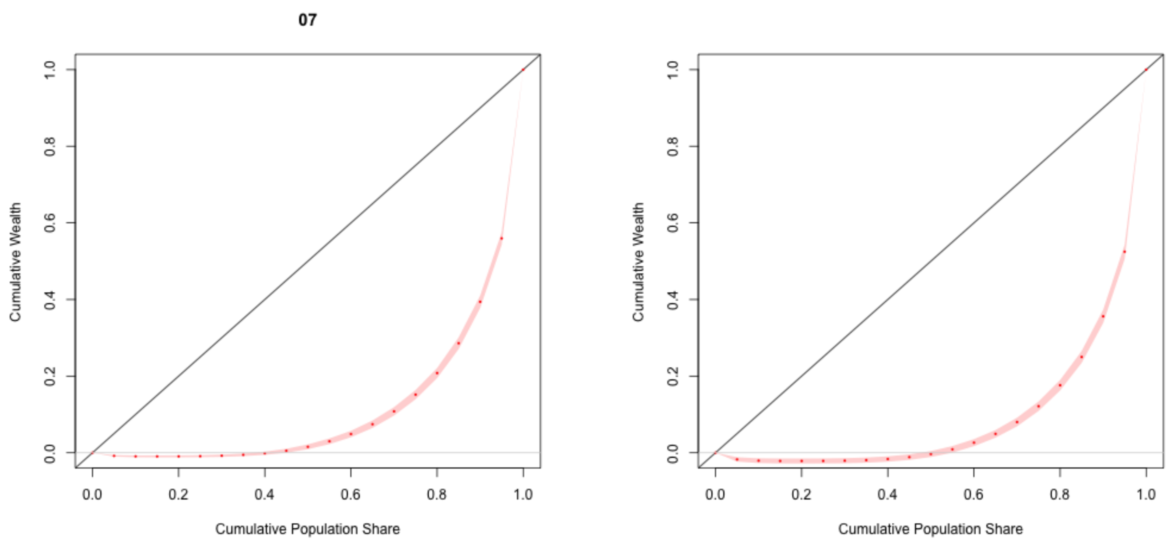

Figure 8: Empirical distribution of Wealth (Lorenz Curves)

Lorenz curves estimated from the PSID 2007 (left) and 2009 (right). Red bands are survey-weighted confidence intervals.

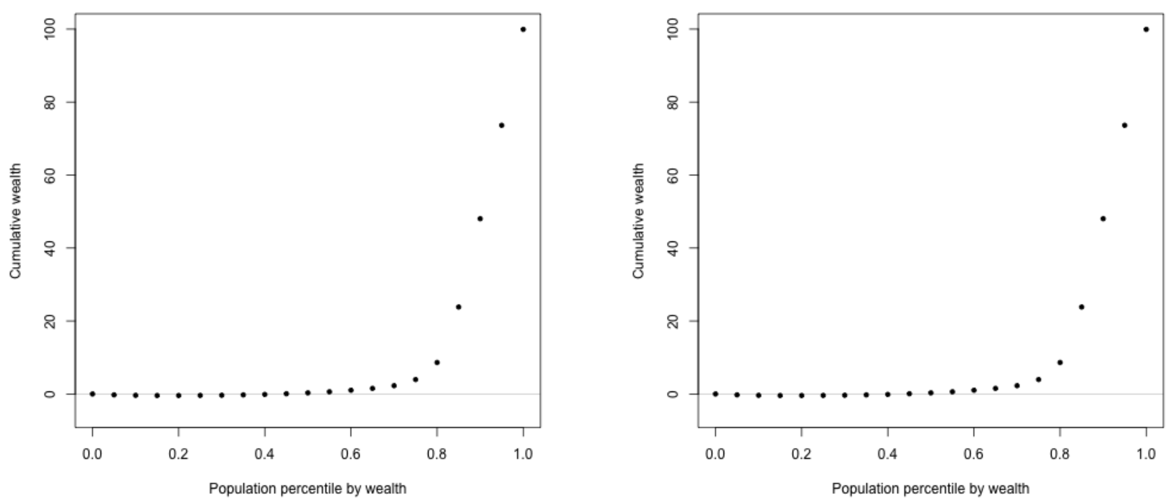

Figure 9: KS β-Dist Simulated Lorenz Curves

Lorenz curves generated from 2006 calibration with BC=-3.3 in good aggregate state (left) and bad aggregate state (right).

What is surprising is the implications for the borrowing constraint over the cycle. I find that the borrowing constraint that fits the percentage of negative wealth holdings in 2006 is -3.3. Figure 10 plots the percentage of the population holding negative wealth for the β-Dist model solved with progressively looser borrowing constraints, lowered by 0.1 on each iteration. The relationship between the fraction holding negative wealth and the borrowing constraint is positive and convex. This implies that, in order to better fit the greater mass of agents holding zero or negative wealth in the 2008 recession, it implies that the borrowing constraint has to be *looser*, i.e. that agents can borrow more, not less, in a recession.

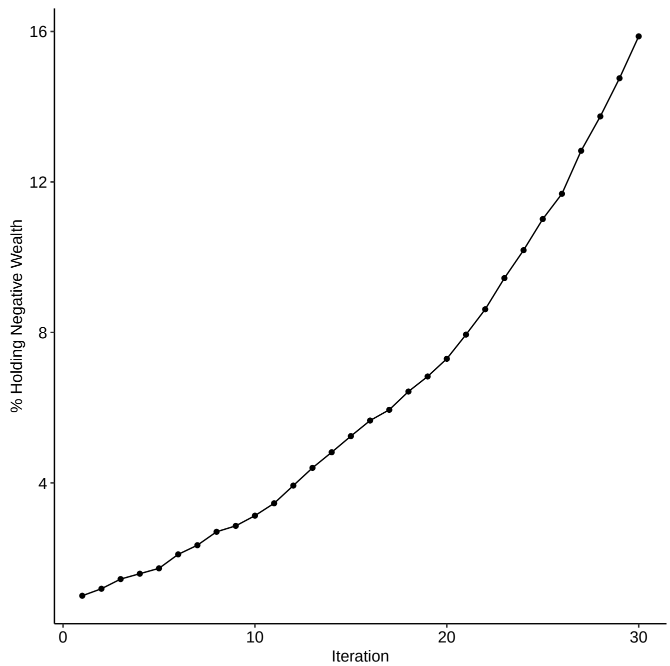

Figure 10: Percent holding negative wealth by iteration through borrowing constraints

KS β-Dist model iterated through negative borrowing constraints. Each iteration is the solution to the model which, starting at zero, lowers the borrowing constraint by 0.1.

Appendix - Balanced Panel

This appendix contains the first- and second-stage regressions in the main text repeated on a balanced panel. In other words, while the main text allowed different households over the pre-recession and recession periods, this appendix keeps only households that have non-missing data for the full 2003-2013 period. The estimated equations are the same, and the results are qualitatively identical and quantitatively very similar. The only major difference is in the sample size, and as a result, the precision of the estimates.

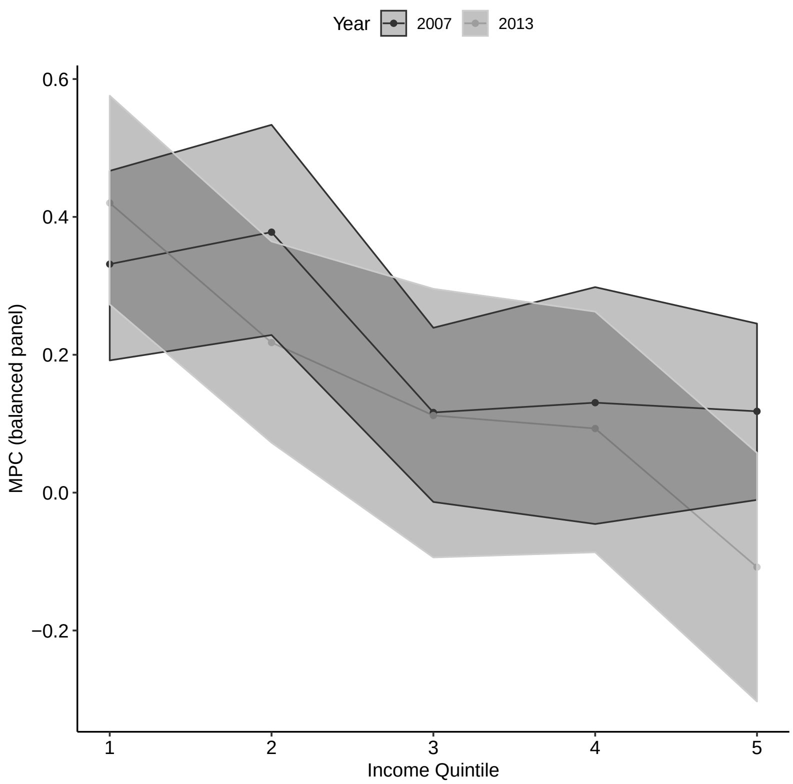

Figure 11: Estimated Distribution of MPC year and quintile, balanced panel

Estimated marginal propensity to consume by income quintile and time, balanced panel: 2003 - 2013

Table 8: IV regression: $\Delta c_{i,t}$ 2003-2007 for Income Quintiles (balanced panel)

	$\Delta\widehat{c_{i,t}}$				
	(1)	(2)	(3)	(4)	(5)
$\Delta\widehat{y_{i,t}}$	0.255^{***}	0.350^{***}	0.099	0.099	0.052
	(0.074)	(0.087)	(0.069)	(0.085)	(0.073)
N	553	556	555	555	539
R^2	0.223	0.171	0.059	0.042	0.021
Residual Std. Error	0.082 (df = 551)	0.093 (df = 554)	0.101 (df = 553)	0.111 (df = 553)	0.132 (df = 537)

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Instrumental variable regression of $\Delta\widehat{c_t}$ on $\Delta\widehat{y_t}$ instrumented by $\Delta\widehat{y_{t+1}}$ for the income quintiles, 2003-2007

Table 9: IV regression: $\Delta c_{i,t}$ 2008-2012 for Income Quintiles (balanced panel)

	$\Delta\widehat{c_{i,t}}$				
	(1)	(2)	(3)	(4)	(5)
$\Delta\widehat{y_{i,t}}$	0.450***	0.302***	0.157	0.105	−0.077
	(0.078)	(0.075)	(0.100)	(0.084)	(0.084)
N	556	554	554	553	543
R^2	0.196	0.077	0.105	0.054	-0.064
Residual Std. Error	0.091 (df = 554)	0.100 (df = 552)	0.118 (df = 552)	0.124 (df = 551)	0.144 (df = 541)

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Instrumental variable regression of $\Delta\widehat{c_t}$ on $\Delta\widehat{y_t}$ instrumented by $\Delta\widehat{y_{t+1}}$ for the income quintiles, 2008-2012

Table 10: First Stage Regressions

	$\log(\widehat{c_{it}})$		$\log(\widehat{y_{it}})$	
	2003-2007	2009-2013	2003-2007	2009-2013
	(1)	(2)	(3)	(4)
Year=2004	0.056***		0.035***	
	(0.003)		(0.006)	
Year=2006	0.082***		0.059***	
	(0.005)		(0.005)	
Year=2010		0.006*		0.002
		(0.003)		(0.004)
Year=2012		0.015**		0.022**
		(0.003)		(0.006)
Education=Medium	0.038***	0.037**	0.070***	0.068***
	(0.007)	(0.009)	(0.011)	(0.013)
Education=High	0.111***	0.116***	0.168***	0.191***
	(0.007)	(0.010)	(0.012)	(0.014)
Race=Black	−0.047***	−0.038**	−0.073***	−0.067***
	(0.010)	(0.010)	(0.011)	(0.011)
Race=Other	−0.009	−0.015	−0.053*	−0.076**
	(0.022)	(0.012)	(0.024)	(0.018)
Family Size	0.068***	0.087***	0.091***	0.108***
	(0.006)	(0.005)	(0.006)	(0.006)
Number of Kids	−0.045***	−0.062***	−0.073***	−0.080***
	(0.006)	(0.007)	(0.008)	(0.008)
Status=Unemployed	−0.074**	−0.056***	−0.129**	−0.108***
	(0.020)	(0.010)	(0.031)	(0.019)
Status=Retired	−0.059***	−0.081***	−0.102***	−0.118***
	(0.009)	(0.007)	(0.011)	(0.011)
Status=Inactive	−0.069***	−0.098***	−0.130***	−0.153***
	(0.011)	(0.009)	(0.011)	(0.015)
Extra Family Income	0.017*	0.003	0.031**	0.028*
	(0.007)	(0.009)	(0.008)	(0.010)
Region=Midwest	−0.043**	−0.055***	−0.041*	−0.059**
	(0.013)	(0.012)	(0.015)	(0.014)
Region=South	−0.025	−0.037**	−0.031	−0.043*
	(0.015)	(0.013)	(0.020)	(0.017)
Region=West	−0.018	−0.048**	−0.034*	−0.059**
	(0.014)	(0.012)	(0.013)	(0.013)
Kids outside Family Unit	0.022*	0.034***	0.058**	0.070***
	(0.009)	(0.005)	(0.015)	(0.010)
Poor-HtM	−0.081***	−0.064***	−0.136***	−0.130***
	(0.007)	(0.009)	(0.011)	(0.014)
Rich-HtM	−0.012	−0.013	−0.066***	−0.064***
	(0.006)	(0.008)	(0.009)	(0.012)
Total Wealth ($1000s)	0.0002*	0.0003*	0.0005*	0.001**
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Constant	11.615***	11.659***	11.723***	11.666***
	(0.044)	(0.035)	(0.070)	(0.079)
Year of Birth	Yes	Yes	Yes	Yes
N	8,373	8,373	8,373	8,373

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.