Saving-Constrained Households (Preliminary and Incomplete*)

Jorge Miranda-Pinto*[1], Daniel Murphy†[2], Kieran James Walsh‡[3], and Eric R. Young§[4]

[1]University of Queensland
[2]University of Virginia Darden School of Business
[3]University of Virginia Darden School of Business
[4]University of Virginia and Zhejiang University

February 15, 2019

Abstract

We develop a theory of saving-constrained households to explain the following facts that are difficult to reconcile with existing theories: 1) Consumption is excessively volatile relative to income (established fact), 2) a large fraction of high-debt households exhibit marginal propensities to consume near zero, 3) lagged high expenditure is associated with low contemporaneous spending propensities. Our proposed interpretation of these facts is that household expenditure depends on time-varying minimum consumption thresholds that, if violated, yield substantial utility costs. We demonstrate that such a model can match many features of the joint dynamics of income and consumption. Our theory has implications for the propagation of macroeconomic shocks.

*Email: j.mirandapinto@uq.edu.au
†Email: murphyd@darden.virginia.edu
‡Email: walshk@darden.virginia.edu
§Email: ey2d@virginia.edu

1 Introduction

A large literature, building on Hall (1978), documents that the dynamics of consumption deviate substantially from the predictions of the Permanent Income Hypothesis. Relative to theory, consumption is excessively volatile and insufficiently persistent given the dynamics of income. Theories of consumer behavior featuring credit constraints have some hope of rationalizing this evidence: when consumption is constrained from above by debt limits, consumption moves with income. Yet recent evidence from microdata documents that many high-debt/low-wealth households exhibit behavior that is inconsistent with credit constraints. Rather than exhibiting large marginal propensities to consume (MPCs) out of income shocks, many high-debt/low-wealth households save all additional income (Sahm et al. (2015), Jappelli and Pistaferri (2014), Misra and Surico (2014)). Why do so many households, especially those with low wealth, save rather than spend out of additional income? And how is it possible to reconcile the existence of many low-MPC households with the fact that consumption tends to be excessively volatile?

In this paper we propose a model of consumer behavior to rationalize this evidence. We begin by presenting a new fact that will guide and validate our theory. Using data from the Panel Study of Income and Dynamics (PSID), we show that, conditional on wealth, high lagged expenditure (on broad consumption categories or exclusively nondurables) is associated with lower current MPCs out of income, especially for low-wealth households. The notion that prior nondurable expenditure is inversely related (or related at all) to current MPCs is, at first glance, both counterintuitive and contrary to standard theory. If anything, modifications to standard theory would predict a positive relationship. Theories of habit formation, for example, would predict that high lagged expenditure should be associated with high desired current expenditure and therefore, if income is mean-reverting, a high contemporaneous MPC.

Our proposed interpretation of these facts is that household expenditure depends on time-varying minimum consumption thresholds that, if violated, yield substantial utility costs. These minimum consumption thresholds represent unanticipated shocks such as medical emergencies and auto repairs. When an adverse shock hits a household, it chooses to accumulate debt (reduce wealth) rather than let consumption fall below a threshold level. For example, rather than move out of a house or slash food consumption, households simply accumulate debt when faced with a large, unanticipated expenditure. Maintaining a low net asset position is optimal for these households but it is costly in the sense that, in the event of another adverse shock, households may be forced to consume below the threshold level (e.g., move out of a home or forgo medical care), which is associated with a large utility cost. Therefore, households for which consumption is against the minimum threshold use additional income to pay off debt (increase net assets) as a precautionary measure.

We refer to these households as saving-constrained households to capture the notion that, in the absence of minimum consumption thresholds, households would save rather than reduce their asset position. Minimum consumption thresholds effectively constrain households' saving relative to a frictionless benchmark, just as credit constraints constrain consumption relative to a frictionless benchmark.

We develop a heterogeneous-agent model featuring minimum consumption thresholds and calibrate it to match within-household dynamics of expenditure and income from the PSID. The model with time-varying minimum consumption thresholds easily fits the data, including the large within-household ratio of consumption volatility to income volatility, while standard one-asset heterogeneous agent models do not. Furthermore, the model generates the path-dependence of spending propensities that we observe in the PSID. High lagged expenditure is associated with lower contemporaneous propensities to spend because high expenditure is a proxy for households being saving-constrained by high minimum consumption thresholds.

The model helps to rationalize a number of additional features of the empirical MPC distribution. Households with high (non-mortgage) debt are more prone to being credit-constrained or savings-constrained, and therefore exhibit a higher prevalence of MPCs near unity or near zero. We show that this pattern is consistent with the MPC distribution from Misra and Surico (2014). Our theory is able to rationalize a number of additional features of the MPC distribution, and we expect that future extensions to the theory, such as incorporating additional assets and liquidity-constrained rich households (Kaplan and Violante (2014)), will further improve the match between theory and data.

Our theory has important implications for the propagation of macroeconomic shocks. In Miranda-Pinto et al. (2019), we show that saving constraints are important for understanding the cross-county relationship between fiscal effects on credit markets and inequality. In particular, high inequality is associated with a larger share of saving-constrained households with low MPCs. Fiscal shocks relax credit markets more (increase interest rates less) in countries with high inequality. A tractable extension of our model with minimum consumption constraints shows that the interest rate response to fiscal stimulus depends on income inequality as in the data.

2 The Joint Dynamics of Household-Level Income and Consumption

Here we highlight evidence on household-level income and consumption that, when taken together, are difficult to reconcile with existing theories of heterogeneous households with uninsurable indisyncratic income risk. The relevant facts are the following: (1) within-household consumption is excessively volatile relative to predictions of standard Bewley models. (2) We present new evidence that lagged high expenditure is associated with low current spending propensities. (3) We extend the analysis in Misra and Surico to document that a large share of high-debt households have spending propensities of effectively zero. Furthermore, the fraction of households with zero-MPCs is increasing with debt.

Our analysis with respect to facts (1) and (2) relies on data from the 1999-2015 Panel Study of Income Dynamice (PSID), which is a biennial panel study of over 3,000 households that are representative of the US population. The PSID is the most comprehensive dataset that tracks household-level expenditure and income over an extended period of time. Starting in 1999 the PSID began collecting detailed data on a range of consumption categories, including expenditures on health, housing, food, transportation, and education. In more recent waves the PSID added information on clothing and travel expenditures.

2.1 Fact 1: Excess Volatility of Consumption

It is well-documented that, on average, within-household consumption (C) is nearly as volatile as income (I). Here we demonstrate this fact for households in the PSID. For each household, we compute the coefficient of variation of income, which is defined as the within-household standard deviation of income divided by the within-household mean of income. We also estimate the persistence of income by running an AR(1) model with household and time fixed effects on the panel data. We then calibrate a standard Bewley model to match these moments of the income distribution.

Table 1 compares the volatility and persistence of income and consumption from the data with the statistics from the Bewley model. Model statistics are from quarterly simulations aggregated to a 2-year frequency. While the income processes are nearly identical, the consumption process in the data is much more volatile and less persistent. The average coefficient of variation of consumption in the data is much higher than in the model, and consumption in the data is nearly as volatile as income.

Table 1
Volatility and persistence of consumption and income: PSID and Bewley model

	PSID	Bewley Model
Avg. coeff. of variation (I_t)	0.47	0.46
AR coeff. (I_t)	0.09	0.09
Avg. coeff. of variation (C_t)	0.40	0.14
AR coeff. (C_t)	0.21	0.77

Note: The coefficient of variation for consumption (income) is defined as the within-household ratio between the mean and standard deviation of consumption (income) ($\frac{\sigma}{\mu}$). The PSID AR coefficients are estimated after removing household and time fixed effects. The model is a standard Bewley model with uninsured idiosyncratic shocks, one asset, and borrowing constraints.

2.2 Fact 2: Lagged High Expenditure is associated with low contemporaneous spending propensities

In heterogeneous agent models with uninsurable idiosyncratic income risk, current assets and income are sufficient information to infer households' optimal consumption decisions. Therefore, lagged expenditure contains no additional relevant information for determining agents' consumption. Here we show that lagged expenditure indeed contains additional relevant information, and in particular that lagged high expenditure is associated with low contemporaneous propensities to spend out of additional income.

Our first step is to estimate spending propensities out of additional income.[1] We then identify episodes of prior high consumption in the past to determine whether prior high expenditure is associated with differential contemporaneous spending propensities. We identify a household as experiencing high expenditure when its expenditure exceeds its within-household average by a standard deviation. Our main specification is

$$\log C_{it} = \beta_0 \log I_{it} + \beta_1 HighC_{it-1} + \beta_2 \log I_{it} \times HighC_{it-1} + \gamma X_{it} + \epsilon_{it}, \tag{1}$$

where $HighC_{it-1}$ is a dummy variable that equals one when the expenditure of household i exceeds the within-household average by a standard deviation at period $t-1$. X_{it} includes log(wealth), the interaction of the wealth term with high income, and a quadratic in age of the head of the household. We also control for household and time fixed effects.

Table 2 shows that the elasticity of expenditure with respect to income is 0.18 (Column 1). Column (2) shows our main specification, which includes lagged high expenditure and its interaction with income. Having high lagged expenditure reduces the spending propensity

[1] One limitation of the PSID is that it is impossible to distinguish between anticipated and unanticipated changes in income. The CEX, on the other hand, can be merged with information on tax rebates to estimate marginal spending propensities. However, the CEX does not have the unique panel structure that the PSID contains and hence is uninformative for the relationship between lagged expenditure and current spending propensities. The implication of potential anticipation effects in the PSID is that estimated contemporaneous spending propensities out of observed income changes should be lower than spending propensities out of unanticipated income changes.

by 0.041, or 22.1% of the average effect (0.18) of income, indicating a large state-dependence of spending propensities. In a standard theoretical framework, beginning-of-period wealth subsumes any information conveyed by prior consumption, so it is possible that the heterogeneous effect associated with lagged expenditure simply reflects spending propensities that vary by wealth. To address this possibility, in column (3) we include then interaction of (log) wealth and current (log) income. The coefficient on the interaction with lagged high expenditure is similar (slightly lower at -0.031, or 17.1% of the average elasticity), indicating that lagged expenditure affects spending propensities beyond the effect of household wealth.

In heterogeneous agent models high consumption is caused by high income and is not associated with lower future spending propensities. In the model we develop below high consumption can also arise for reasons unrelated to contemporaneous income, and it is these idiosyncratic consumption episodes that cause lower future spending propensities. Therefore, to isolate the mechanism that we will propose, it is important to isolate episodes of high consumption that are not associated with high contemporaneous income. To do so, we identify "high expenditure episodes" as periods in which households have high expenditure but *not* high income. In particular a high expenditure episode is a dummy variable that equals unity when household consumption exceeds its within-household mean by a standard deviation and household income does not exceed its within-household mean by a standard deviation. These episodes capture periods of high consumption that are not driven by high contemporaneous income (see Table 3 for a summary of the indicator variables and their precise definitions). Column (5) of Table 2 replaces the indicator for high expenditure with the indicator for a high expenditure episode. Here, the interaction term is even more negative: a high expenditure episode is associated with a spending elasticty that is lower by 0.40, or 22.1% of the average elasticity of spending with respect to income. High expenditure that is not associated with high contemporaneous income is associated with substantially lower future spending propensities. The negative association of high consumption is unique to episodes with low income. When we replace the indicator for a high consumption episode with an indicator for high consumption *and* high income, the effect on future spending propensities is no longer negative (column 6). Therefore, the negative relationship between high expenditure and future low spending propensities is driven by episodes in which high expenditure is not associated with high income.

Decomposing High Expenditure Episodes. Which categories of expenditure drive high expenditure episodes (henceforth referred to as 'episodes')? Are episodes primarily driven by subset of expenditure, or do all components of expenditure contribute to these episodes? To address these questions, we first examine how much consumers spend on different categories during episodes relative to average spending on each category. Column 1 of Table 4 shows average (across households) expenditure shares for each category of expenditure, where the categories correspond to PSID classification schemes. Column (2) shows the expenditure

Table 2
Lagged Expenditure and Spending Propensities in the PSID.

Dependent variable: $\log(C_t)$	(1)	(2)	(3)	(4)	(5)
$\log(I_t)$	0.181***	0.216***	0.241***	0.241***	0.224***
	(0.011)	(0.013)	(0.019)	(0.019)	(0.019)
High C_{t-1}		0.483***	0.373**		
		(0.165)	(0.150)		
$\log(I_t) \cdot$ High C_{t-1}		-0.041***	-0.031**		
		(0.015)	(0.014)		
High C_{t-1} & not high I_t				0.477***	
				(0.156)	
$\log(I_t) \cdot$ (High C_{t-1} & not high I_t)				-0.040***	
				(0.014)	
High C_{t-1} & high I_{t-1}					0.096***
					(0.014)
$\log(I_t) \cdot$ (High C_{t-1} & high I_{t-1})					0.048***
					(0.001)
Control for wealth and income	NO	NO	YES	YES	YES
N	53648	46942	46942	46942	46942
R^2	0.67	0.67	0.69	0.68	0.70

Note: A household has high income, High I_t, (expenditure, High C_t) in periods in which income (expenditure) is over a standard deviation above average income (expenditure) for the household. A high expenditure episode (High C_{t-1} not high I_t) is a dummy that equals unity when a household experiences high expenditure but not high income. All regressions control for log(wealth) and a quadratic in age of the head of household. Robust standard errors in parentheses. ***p<0.01, ** p<0.05, *, p<0.1.

shares during episodes.[2] Expenditure shares during episodes (column 2) is generally similar to average expenditure shares (column 1). Two categories are noticeably more prevalent during episodes: education and transportation. Food and housing are less prevalent during episodes.

Columns (3) and (4) offer an alternative approach to examining the relevance of different expenditure categories in driving episodes. Here, we identify episodes for each category and examine the extent to which category-specific episodes predict aggregate expenditure episodes. For example, the dummy for a transportation episode is set to unity whenever a

[2]To compute the statistics in column 2, we first compute category-specific expenditure for each household relative to that household's mean. We then average over households to obtain average demeaned expenditure for a category. We then do the same for total expenditure and take the ratio of the two.

Table 3
Indicator Variable Definitions and Incidence

Indicator Variable	Description	Incidence
High Expenditure	$Expenditure_t >$ Mean(($Expenditure_t$)+SD($Expenditure_t$)	0.16
High Income	$Income_t >$ Mean($Income_t$)+SD($Income_t$)	0.26
High Expenditure Episode	High Expenditure and Not High Income	0.10
High Expenditure & High Income	High Expenditure and High Income	0.07

household's transportation expenditure exceeds its within-household average by a standard deviation (and income is not high). We then regress the dummy for an episode on indicators for the category-specific episodes. The pattern that emerges from both OLS (column 3) and probit (column 4) models is that high sub-category expenditure is associated with a high-expenditure episode for each category of expenditure. Furthermore, the likelihood that any given category-specific episode is associated with a total expenditure episode is broadly proportional to that category's share of total expenditure and follows a similar pattern of relevance that is depicted in Table 2.

Table 4
Decomposing High Expenditure Episodes

	Share of Total Expenditure	Ratio of Category Expenditure Relative to Total Expenditure during High Expenditure Episodes	Coefficient from Linear Probability Model	Coefficient from Probit Model
	(1)	(2)	(3)	(4)
Food	0.18	0.07	0.16	0.95
Housing	0.43	0.37	0.33	1.67
Transportation	0.26	0.39	0.38	1.86
Education	0.04	0.11	0.13	0.89
Child Care	0.01	0.01	0.05	0.33
Health	0.07	0.06	0.11	0.77

Note: This table presents statistics for the broad categories that make up the measure of total expenditure. Expenditure on Clothing, Trips, and Other Recreation are not included in the measure of total expenditure since they were only recorded beginning in 2005. In column (2), expenditure during episodes is relative to within-household averages. In columns (3) and (4), the depicted statistics are the coefficients from a regression of an indicator for a high expenditure episode on indicator variables for high sub-category expenditure. All regression coefficients are significant at the $p<0.01$ level.

Alternative specifications. Here we extend the analysis in Table 2 to examine different measures of income and expenditure. First, to help isolate the unanticipated component of income, we first replace the continuous measure of income in the regressions with an indicator for high income (defined above). Under the assumption that abnormal realizations of income are less likely to be anticipated, the indicator variable is more likely to isolate unanticipated changes in income.[3] We also examine an alternative measure of expenditure that excludes purchases of durables. Specifically, we identify high expenditure episodes based on extreme realizations of a expenditure net of purchases of automobiles and furniture. This will help guide the theory that we develop to explain the joint dynamics of expenditure and income.

According to the results in Table 5, replacing the income measure with an indicator for high income produces a similar pattern: the propensity to spend in response to a high income realization is lower in the presence of a lagged high expenditure episode (columns 1 and 3). If anything, the magnitude of the negative interaction term is larger as a fraction of the average effect of high income. Furthermore, the negative effect of lagged high expenditure on spending propensities is just as strong when limiting the expenditure measure used to identify episodes to nondurables (columns 2 and 4). This indicates that modeling durable goods may not be crucial for understanding the effect of lagged expenditure on spending propensities.

Table 5
Specification with Alternative Expenditure and Income Measures

Dependent variable: log(expenditure)				
Income measure:	log(income)		High Income	
Expenditure measure used to identify high-expenditure episodes:	All	Nondurables	All	Nondurables
	(1)	(2)	(1)	(2)
Coefficient on Interaction term	-0.043*** (0.013)	-0.043*** (0.013)	-0.073*** (0.019)	-0.067*** (0.019)
Fraction of average effect of income measure	-0.24	-0.24	-0.32	-0.30

Note: Nondurable expenditure is defined as total expenditure net of purchases of automobiles and furniture. An episode is defined as a period in which a household has high expenditure but not high income. A household has high expenditure (income) in periods in which expenditure (income) is over a standard deviation above average expenditure (income) for the household. All regressions include the income measure, the episode indicator, a quadratic in age, log(wealth) and its interaction with the income measure, and time and household fixed effects. The table reports only the coefficient on the interaction between lagged episode and the income measure. Robust Standard Errors in parentheses. ***p<0.01, ** p<0.05, *, p<0.1.

3 A Theory of Saving-Constrained Households

Here we present a theory to rationalize the large volatility and persistence of consumption. We introduce random minimum consumption thresholds into a standard Bewley model with

[3]Note that anticipation effects should, if anything, reduce our coefficient estimates toward zero. The fact that spending propensities are positive suggests that, in the absence of credit constraints, households behave as if changes in income are to some extent unanticipated.

capital and calibrate it to match the consumption and income dynamics from the PSID. We then demonstrate that the model replicates the dependence of spending propensities on lagged high expenditure (Section 2). Finally, we examine the model's predictions relative to recent evidence that many high-debt households have MPCs near zero.

3.1 Model

The economy consists of a measure one of infinitely–lived households (indexed by subscript $i \in [0,1]$ when necessary for clarity of exposition) that are ex ante identical and a representative firm that hires capital and labor to produce the single tradable consumption good. The households and firm participate in a global capital market with exogenous interest/rental rate r. The labor market, in which the firm hires household labor at wage rate w is, however, purely domestic. The recursive problem of a household is

$$V(k,z,x,\underline{c}) = \max_{c \geq 0, k' \geq b} \left\{ \log(c) - \lambda \max\{\underline{c} - c, 0\} + \beta E_{z',x',\underline{c}'}\left[V\left(k', z', x', \underline{c}'\right) | z, x, \underline{c} \right] \right\}$$

subject to the budget constraint

$$c + k' \leq (1 + r - \delta + \phi \mathbf{1}(k \leq 0))\, k + w \exp(z + x)\, \overline{h},$$

where V is the value function, c is consumption (the numeraire), k is capital wealth (which exogenously depreciates at rate $\delta \geq 0$), x is persistent idiosyncratic household productivity, z is (nearly) permanent idiosyncratic productivity, and $\underline{c}$ is a persistent necessary expenditure shock. For any variable q, q' represents its value in the subsequent period.

If the household consumes less than $\underline{c}$, it must pay utility cost $\lambda(\underline{c} - c)$ for some $\lambda \geq 0$. Given a particular realization of $\underline{c}$, a wealthy household voluntarily consumes above it, extremely poor borrowers pay the utility penalty, and households with intermediate wealth borrow/dis-save just enough to consume the minimum. In MPC terms, for a given $\underline{c}$, the extremely poor are effectively hand–to–mouth, the rich have normal MPCs, and many intermediate wealth households have MPCs of exactly zero. Indeed, for an intermediate range of wealth, households facing a minimum consumption shock borrow/dis-save simply to avoid paying the utility penalty. In the absence of the shock, they would have saved more and are thus *saving*–constrained: with an unexpected splash of income, all they do is save and pay down debts, that is, their MPCs are zero.

Additionally, borrowing ($k < 0$) entails a cost $\phi(-k)$, for some $\phi \geq 0$, so the household interest rate on borrowing is higher than the interest rate on saving. In each period, the household inelastically supplies effective labor $\exp(z + x)\overline{h}$ at wage w. We assume that household productivity and necessary expenditure levels evolve according to:

$$\begin{aligned} z' &= \rho_z z + \epsilon_z' \\ x' &= \rho_x x + \epsilon_x' \\ \underline{c}' &= (1-\rho_c)\mu_c + \rho_c \underline{c} + \epsilon_c', \end{aligned}$$

where ϵ_j, $j \in \{z, x, c\}$, is an idiosyncratic mean–zero shock with standard deviation σ_j. Define the stationary aggregate labor supply to be $\overline{H}$.

In each period, the representative firm chooses capital K and effective labor L to solve

$$\max_{K,L} \{K^\alpha L^{1-\alpha} - rK - wL\}.$$

We examine *stationary equilibria*, which consist of constant firm capital K^* and labor L^*, a constant wage w^*, a constant household wealth distribution Ω^*, and household value and policy functions V^*, c^*, and k'^* such that (1) the value and policy functions solve the household problem given prices, (2) K^* and L^* solve the firm problem:

$$r = \alpha (L^*/K^*)^{1-\alpha}, \; w^* = (1-\alpha)(K^*/L^*)^\alpha = (1-\alpha)\,(r/\alpha)^{\frac{\alpha}{\alpha-1}},$$

(3) the labor market clears: $L^* = \overline{H}$, and (4) Ω^* is generated by k'^*. Let $\overline{K}$ denote aggregate household capital, define $y = (1+r-\delta)k + w\exp(z+x)\overline{h}$ to be household income, and let Y be aggregate household income.

3.2 Calibration

We assume a period is one quarter and calibrate the model with a two–step procedure. First, we choose the productivity process parameters to approximate household income from the PSID and set the borrowing cost, capital share, depreciation rate, global interest rate, household labor endowment, and borrowing constraint to reasonable values: $\phi = .03$, $\alpha = .36$, $\delta = .0125$, $r = .0225$, $\overline{h} = .33$, $b = -2$, $\rho_z = .5$, $\sigma_z = .494$, $\rho_x = .99$, and $\sigma_x = .174$. In the second step, we choose the remaining parameters (discount rate, utility cost, and $\underline{c}$ process) to target the following quarterly frequency moments:[4]

$$\begin{aligned}
\frac{\overline{K}}{Y} &= 12 \\
Corr\,(y,c) &= 0.2 \\
\frac{Std\,(c)}{mean\,(c)} &= 0.79 \\
Fraction\,(k<0) &= 0.1 \\
Corr\,(c_t, c_{t-1}) &= 0.975 \\
Fraction\,(|c-\underline{c}|<0.01) &= 0.4 \\
mean\,(MPC) &= 0.2
\end{aligned}$$

Via global optimization, the best–fit parameter values are $\beta = 1/1.025$, $\lambda = 5.014$, $\rho_c = .826$, $\sigma_c = 1.179$, and $\mu_c = .024$.

3.3 Results

The zero-MPC households are evident from the consumption functions of households with different realizations of $\underline{c}$. Figure 1 shows that households with higher values of $\underline{c}$ tend to have consumption that is flat with respect to wealth (MPCs of zero) until wealth is sufficiently high that the minimum consumption threshold is no longer binding. These constrained households use all additional wealth/income to save. However, not all low-wealth households with high realizations of $\underline{c}$ have MPCs of zero. The poorest households are credit-constrained and cannot even achieve the minimum consumption level. They experience severe dis-utility and consume all additional income, up until they are able to consume at the threshold.

What does this imply for the MPC distribution across households? Figure 2 shows, at different levels of productivity, consumption functions and MPCs averaged across realizations of and $\underline{c}$. Comparing with the wealth distributions in Figure 1, we see that in the ergodic distribution many agents have an MPC of zero and the higher MPCs come from the very rich and poor by wealth. Especially for low and middle productivity households, MPCs are U–shaped in wealth. The lowest MPCs households are not rich. They are poor enough for the minimum consumption level to matter but not so poor that they violate it.

In Figure 3 we explicitly report the fraction of zero MPC households across wealth quintiles, for households with low wealth. We observe that most of saving-constrained households (binding at $\underline{c}$) are in the second and third quantile of wealth. The very poor households in

[4]Given parameters, we use Rouwenhorst's method to discretize the productivity and $\underline{c}$ processes as Markov processes, three states for each productivity process and seven states for the expenditure shock. Given exogenous r and aggregate labor supply $\overline{H}$, firm capital and the equilibrium wage follow trivially from firm optimality. We then solve the household problem with standard global methods, yielding the household policy functions and the stationary wealth distribution.

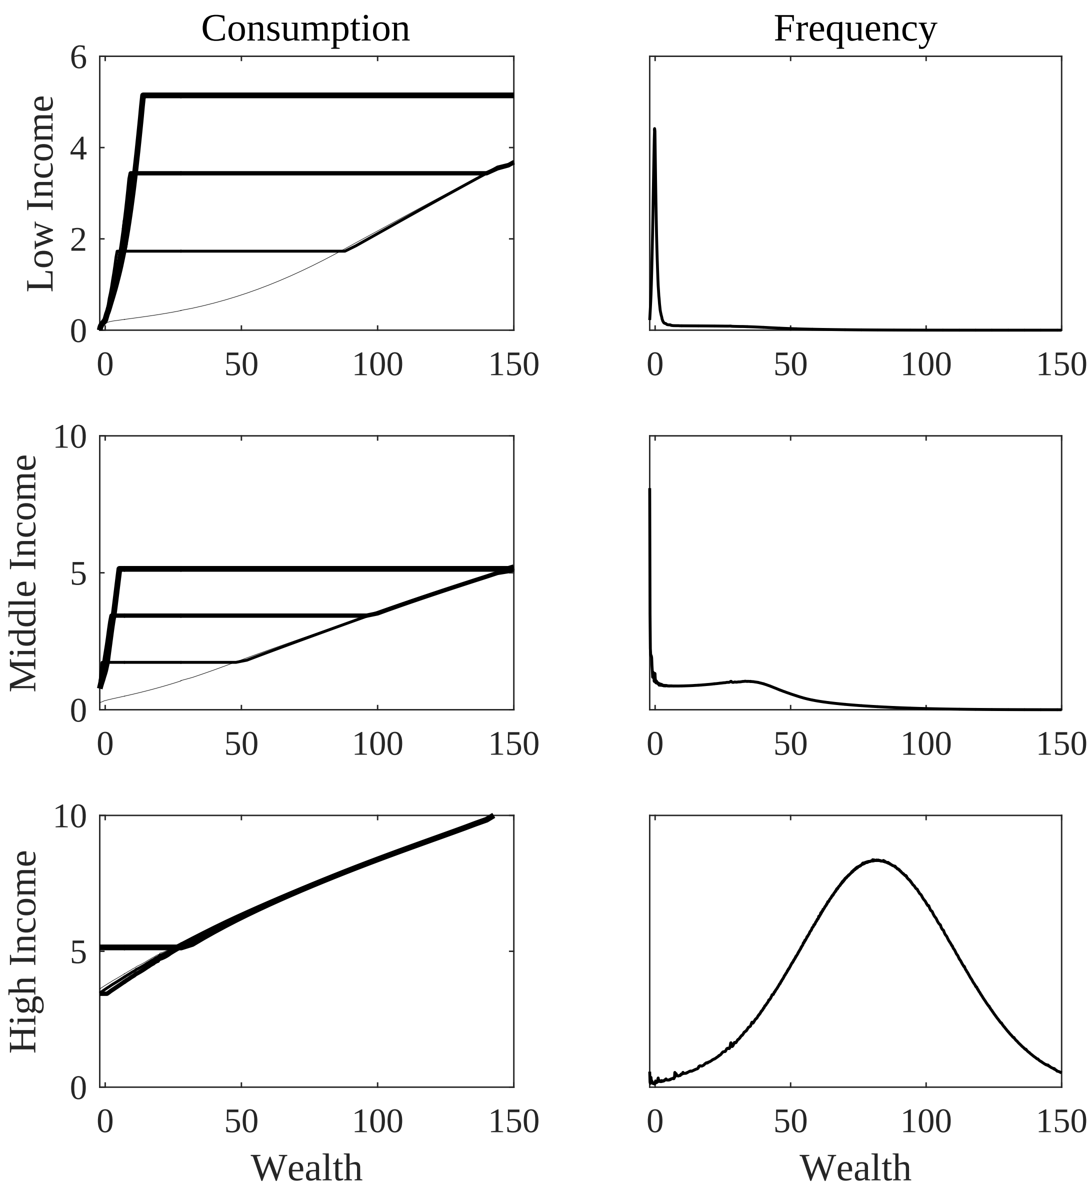

Figure 1

Note: The left column shows the model consumption policy functions at different income levels, and the right column shows the steady–state wealth distribution conditional on these income levels. Line thickness corresponds to the value of $\underline{c}$. Low (Middle, High) income means both the permanent and persistent components of productivity are at their low (middle, high) discretized values.

quintile one pay the utility cost of consuming below $\underline{c}$, while the richer households in quintile 5 are able to move away from the constraint.

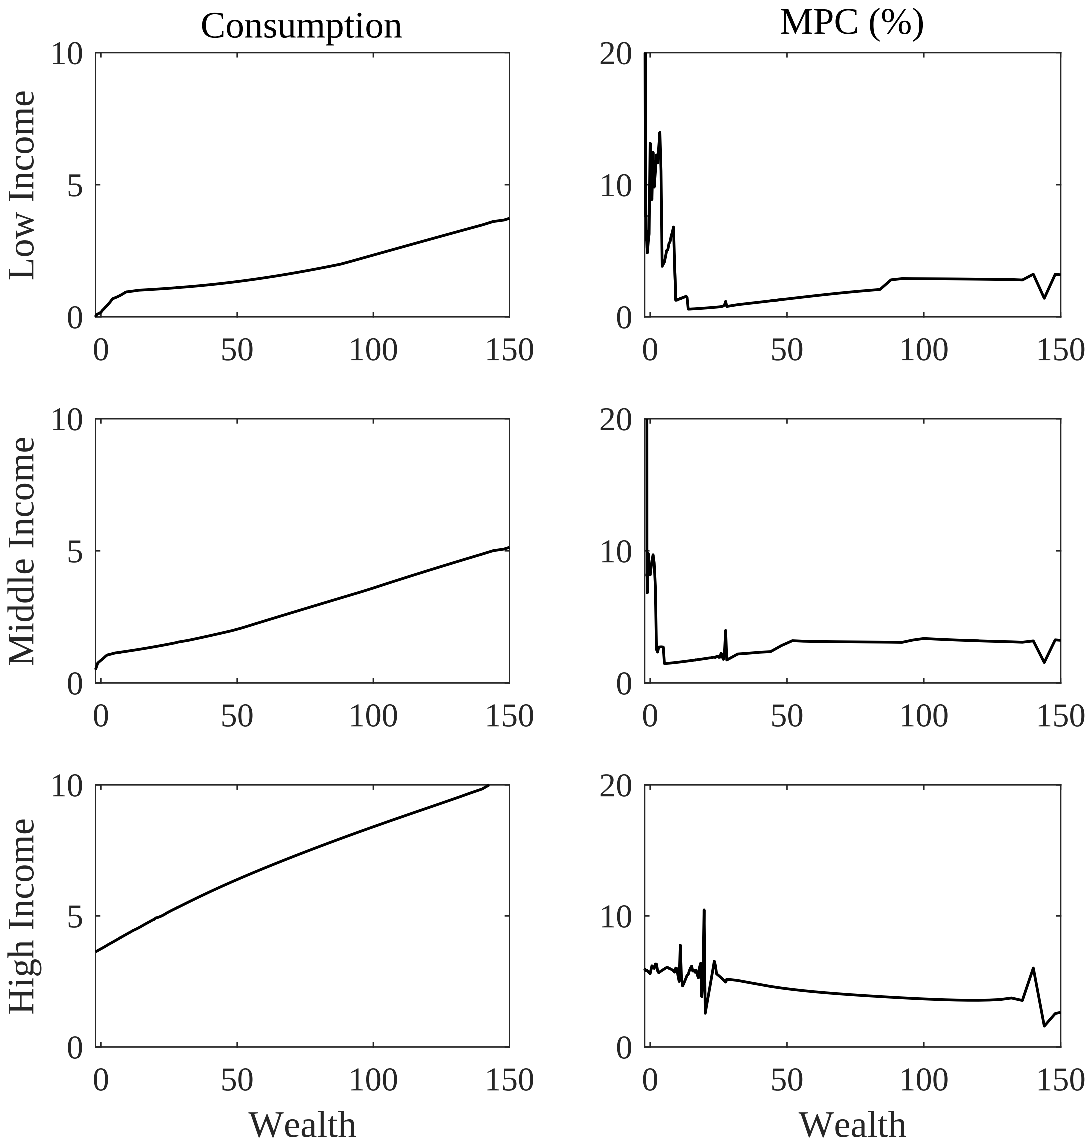

Figure 2
Note: The left column shows the model consumption policy functions at different income levels (averaged across $\underline{c}$ levels), and the right column shows the percent marginal propensity to consume conditional on these income levels (averaged across $\underline{c}$ levels). Low (Middle, High) income means both the permanent and persistent components of productivity are at their low (middle, high) discretized values.

3.3.1 Model Fit: Comparison to PSID

Here we examine the model's ability to improve the fit to the PSID along two dimensions: (1) the high volatility of consumption relative to income, and (2) the effect of lagged high expenditure on current spending propensities. To make the results comparable with the PSID, we aggregate eight-quarter episodes into a period and define income and consumption

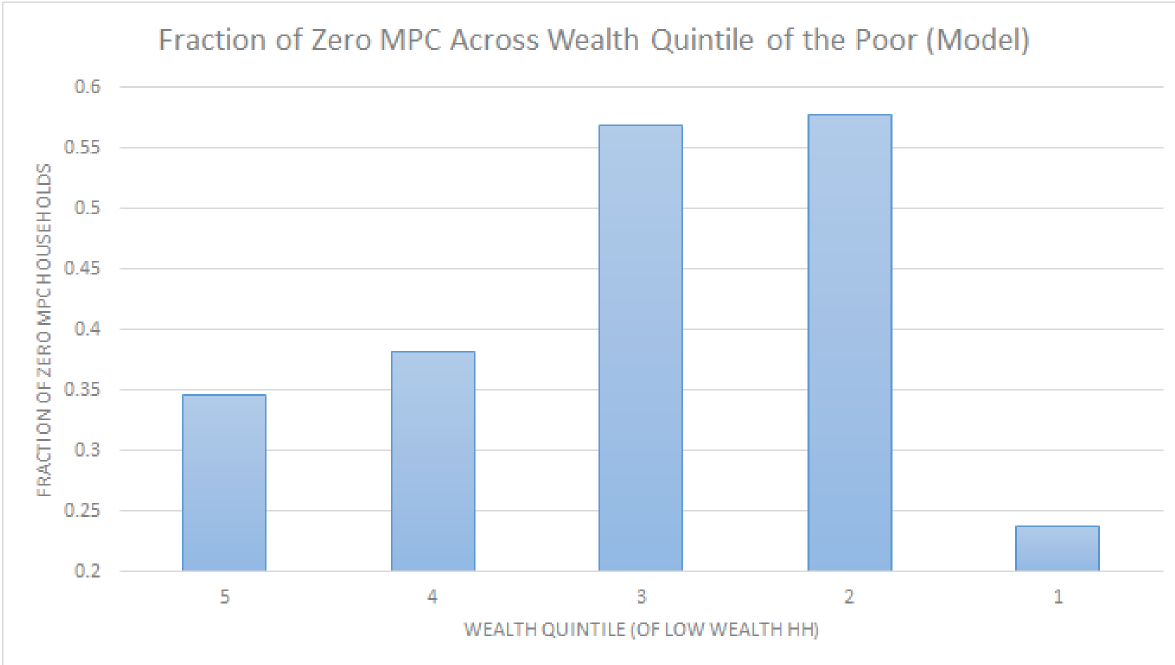

Figure 3

Note: The figure shows MPC across wealth quintiles households with wealth below median.

over that period as their eight-quarter average.

With respect to (1), Table 6 shows that including the minimum consumption shock improves the fit of the simple model to the PSID. Without this new feature, consumption is far too persistent and has an unreasonably low coefficient of variation. Introducing $\underline{c}$, the coefficient of variation rises to and autocorrelation falls to levels much closer to the PSID numbers of .40 and .20.

With respect to (2), we first demonstrate that the theory fits the empirical relationship between lagged expenditure and spending propensities. We then discuss how the empirical specification captures the behavior of saving constrained households.

Table 7 shows that spending propensities out of income are lower when households experienced high expenditure in the previous period (column 2). This relationship is even stronger when isolating high expenditure episodes that don't coincide with high income (column 3), a similar patter to what what is observed in the PSID. Columns (4) through (6) show regressions from a Bewley model without minimum consumption threshold shocks. Here, the dependence of spending propensities is no longer negative. If anything, lagged expenditure is associated with higher spending propensities.

What explains the negative dependence of spending propensities on lagged expenditure? The theory implies that saving constraints are associated with lower spending propensities. Saving constraints are persistent (due to the persistence of the minimum consumption threshold), which implies that lagged saving constraints are also associated with lower

Table 6
Consumption and Income Volatility and Persistence in the Model

	PSID	Bewley model (cbar=0)	Saving Constraint Model (cbar>0)	Bewley Model, same income process (cbar=0)
	(1)	(2)	(3)	(4)
Avg Coeff. of Variation (I_t)	0.47	0.46	0.31	0.31
AR coefficient (I_t)	0.09	0.09	0.09	0.09
Avg Coeff. of Variation (C_t)	0.40	0.14	0.38	0.19
AR coefficient (C_t)	0.21	0.77	0.46	0.74

Note: the PSID AR coefficients are estimated after removing household and time fixed effects. The model in Column 2 is a standard Bewley model with uninsured idiosyncratic shocks, one asset, and borrowing constraints calibrated to match the income process in the PSID. The model in column 4 is the same Bewley model but with an income process that matches the calibrated model with saving constraints (column 3). All models are simulated at the quarterly frequency and then aggregated to two-year periods

spending propensities. High expenditure in the presence of high income is a proxy for saving constraints: when households receive a high realization of $\underline{c}$, their expenditure increases. Therefore, the theory implies that high expenditure, which is a proxy for saving constraints, tends to be associated with lower future spending propensities. Below we examine these relationships in more detail using quarterly data generated by the model.

3.3.2 Analysis based on quarterly data

Table 8 shows summary statistics and regression coefficients based on the quarterly data produced by the model. The advantage of analyzing quarterly data is that we can observe whether households are saving constrained at any point in time. When aggregating to the two-year frequency, households may move in and out of being constrained over 8 quarters, and hence their 8-quarter consumption will not match their 8-quarter consumption threshold even if they experience episodes of saving constraints and those episodes affect future spending behavior.

Rows (1) and (2) demonstrate that replacing a high expenditure episode with a direct indicator of being saving constrained (specifically, $c_t = \underline{c_t}$) leads to a similar pattern in the regressions: lagged saving constraints are associated with lower spending propensities. Approximately 15% of households are saving constrained in any given quarter (row 3), which implies that a substantial share of households have MPCs near zero. Three percent of households experience saving constraints for eight consecutive quarters (column 4). This persistence helps explain how the effects of savings constraints can be detected even in time-aggregated data.

Saving constraint episodes do not perfectly correspond to high expenditure episodes. While the incidence of being saving constrained is 15%, the incidence of a high expenditure

Table 7
Lagged Expenditure and Spending Propensities in the Model.

Dependent variable: $\log(C_t)$	Saving Constraint Model			Bewley Model		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(I_t)$	0.105***	0.167***	0.031***	0.125***	0.120***	0.091***
High C_{t-1}		0.234***			0.054***	
$\log(I_t) \cdot$ High C_{t-1}		-0.061***			0.023***	
High C_{t-1} & not high I_t			0.235***			0.034***
$\log(I_t) \cdot$ (High C_{t-1} & not high I_t)			-0.073***			0.017***
Control for wealth and income	NO	YES	YES	NO	YES	YES
R^2	0.92	0.93	0.93	1.00	1.00	1.00

Note: A household has high income (consumption) in periods in which income (consumption) is over a standard deviation above average income (consumption) for the household. A high expenditure episode is a dummy that equals unity when a household experiences high expenditure but not high income. All regressions control for log(wealth). Standard errors not shown since regression estimates are highly precise. Models are simulated at the quarterly frequency and then averaged to two-year periods before estimating regressions. The Bewley model is simulated using the same income process as the one derived from the calibrated Savings Constraint model.

episode is only 6% (column 4) because it is possible to be saving constrained even with low levels of consumption (if, for example, income falls but $c_t = \underline{c_t}$ is positive). How well does a high expenditure episode proxy for saving constraints? Approximately half of households that have a high-expenditure episode but are also saving constrained (column 9), suggesting that the proxy used in the regressions does indeed capture a large share of households that have MPCs of zero.

3.4 Dynamic Consumption Responses to Temporary Income Shocks

[To Be Completed]

4 Evidence of Saving-Constrained Households from the Consumer Expenditure Survey (CEX)

Here we examine the theory's predictions of low-wealth, zero-MPC households using recent estimates of MPC heterogeneity from the CEX. The evidence here complements recent work that documents that large shares of high-debt households plan to save rather than spend

Table 8
Summary statistics from quarterly data

(1)	Coefficient on Interaction between log(income) and saving constrained	-0.09
(2)	Interaction Coefficient as fraction of average Spending Elasticity	-1.30
	Share of agents who	
(3)	Are on saving constraint in a given quarter	0.15
(4)	Are on saving constraint for eight consecutive quarters	0.03
(5)	Have high consumption	0.11
(6)	Have high consumption but not high income (episode)	0.06
(7)	Pay the utility cost (consumption$<$ minimum threshold)	0.11
	Correspondence between high consumption episode and saving constraint:	
(8)	Incidence of episode if saving constrained	0.21
(9)	Incidence of saving constraint if episode	0.49

Note: In rows (1) and (2), coefficients are based on quarterly regressions with log(consumption) as the dependent variable. Regressions control for log (wealth) and household fixed effects.

additional disposable income (e.g., Sahm et al. (2015)).

We follow Misra and Surico (2014) and estimate households' consumption response to the 2001 tax rebates in the U.S. We augment their analysis by gathering data on households' total wealth and non-mortgage debt. With the additional data at hand, we study whether the micro evidence supports the existence of low-wealth and indebted households with MPCs near zero.

The advantage of using the quantile regression approach in Misra and Surico (2014) is that we estimate marginal propensities to consume (MPC) that are heterogeneous across households, even for households with similar observed characteristics. Therefore, taking the MPC estimates we are able to relate the MPC heterogeneity to proxies of saving-constrained households implied by our model.

4.1 Short-Term Spending Propensities from The 2001 Tax Rebate

The mailing of the 2001 tax rebate was randomized based on the penultimate number of the tax filer social security number. Hence, the rebate receipt was exogenous to individual characteristics. Using households level data on consumption from the CEX and individual tax records, Johnson et al. (2006) estimate the short-term consumption responses to the tax rebate receipt. The main specification in Johnson et al. (2006) is

$$\Delta C_{it+1} = \sum_{s} \alpha_{0s} \times M_s + \alpha_1' X_{it} + \alpha_2 R_{it+1} + u_{it+1}, \tag{2}$$

where ΔC_{it+1} is household i change in nondurable consumption in the three month period when the tax rebate was received. M_s is a set of time controls that capture seasonal effects and aggregate shocks. The matrix X_{it} contains household controls, in particular average

age and the change in the number of family members. The main variable R_{it+1} is the total dollar amount of tax rebate received by households i in the three months period $t+1$.

The authors find that between 20 to 40 percent of the tax rebate in 2001 was spent on non-durable goods. To capture the heterogeneity in the consumption responses, Johnson et al. (2006) interact the tax rebate variable R_{it+1} with indicators variables describing households with low (high) income and low (high) level of liquid assets. The results suggest that households with low income and low assets' liquidity have larger responses.

Misra and Surico (2014) amend the approach in Johnson et al. (2006) to account for the possibility that consumption responses may be heterogeneous, even within subgroups based on income. The authors estimate a version of equation (1) using quantile regression and find that high-income households are likely to have very low and very high consumption responses to the 2001 (and 2008) tax rebate.

As we are interested in understanding the determinants of low MPCs, we exploit the Misra and Surico (2014) approach. The main specification is a linear quantile model of the form

$$\Delta C_{it+1} = q(R_{it+1}, X_{it}, M_s, \lambda_{it+1}) \quad with \quad \lambda_{it+1} | R_{it+1}, X_{it}, M_s \sim U(0,1), \tag{3}$$

where λ_{it+1} captures the unobserved heterogeneity in households with similar observed characteristics (R_{it+1}, X_{it}, M_s). Let $q(R_{it+1}, X_{it}, M_s, \tau)$ the conditional τ-th quintile of ΔC_{it+1}, given observables, for each $\tau \in (0,1)$, the linear quantile model is[5]

$$\Delta C_{it+1} = q(R_{it+1}, X_{it}, M_s, \tau) = \sum_s \alpha_{0s}(\tau) \times M_s + \alpha_1(\tau)' X_{it} + \alpha_2(\tau) R_{it+1}. \tag{4}$$

As mentioned above, this model does not impose common MPCs among households with similar observable characteristics. The estimated consumption responses are only common within a quantile τ but are heterogeneous across quantiles, representing unobserved heterogeneity.

4.2 Extending the Misra and Surico (2014) Analysis

We start by replicating the estimation of α_2 in equation (3) in Misra and Surico (2014). We focus on the estimated tax rebate coefficient of non-durable consumption for the 2001 tax rebate. Then, we gather additional from the CEX on households' total wealth and non-mortgage debt. We define wealth as the sum of the balance in checking accounts, saving accounts, U.S. bond, value of stocks, and the value of properties, minus the outstanding mortgage debt and outstanding non-mortgage debt. Non-mortgage debt is composed of

[5] Misra and Surico (2014) extend the set of households controls X by adding squared age and changes in the squared number of family members.

credit card debt, bank loans, credit union debt, dentist and hospital debt, finance companies debt, and excludes business debt and car loans.

Figure 4 shows the distribution of MPC for the group of low-wealth households (below median wealth). Consistent with our theory, there is large fraction of low-wealth households with very low MPCs.

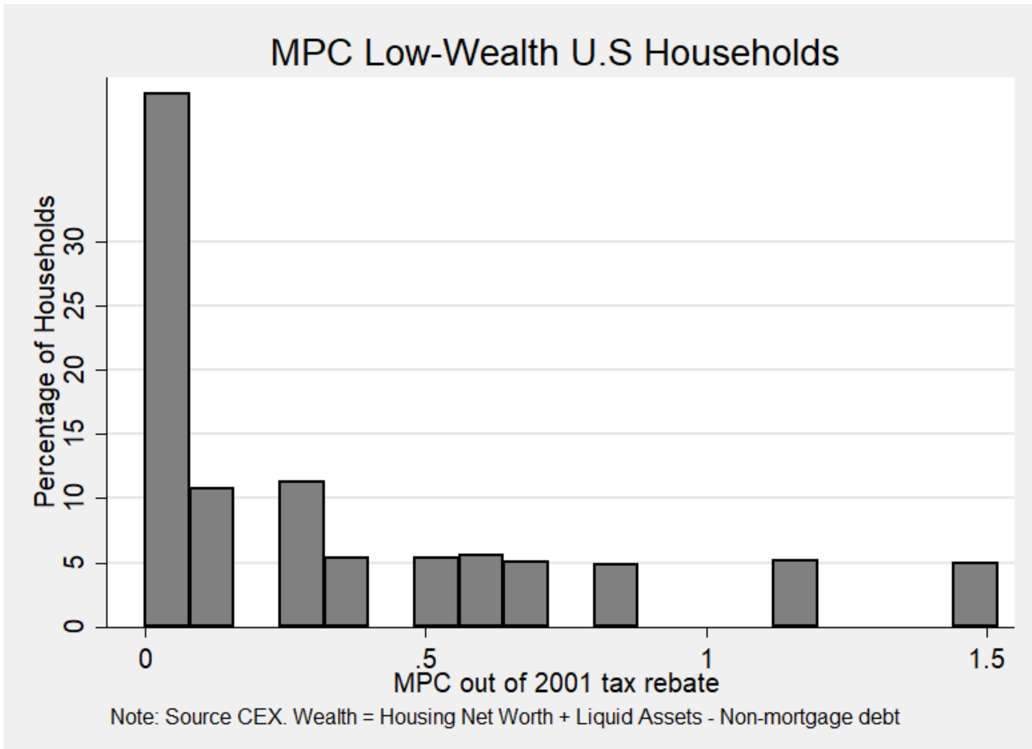

Figure 4
MPC distribution low-wealth Households

We now study who these zero-MPC households are. We define a dummy variable called zero-MPC that equals one when the MPC is below 0.05 and equal to zero otherwise. Figures 5 and 6 show the fraction of zero-MPC households across quantiles of non-mortgage debt and mortgage debt, respectively. We observe that, consistent with our theory in Figure 3, the fraction of zero-MPC households grows with debt.

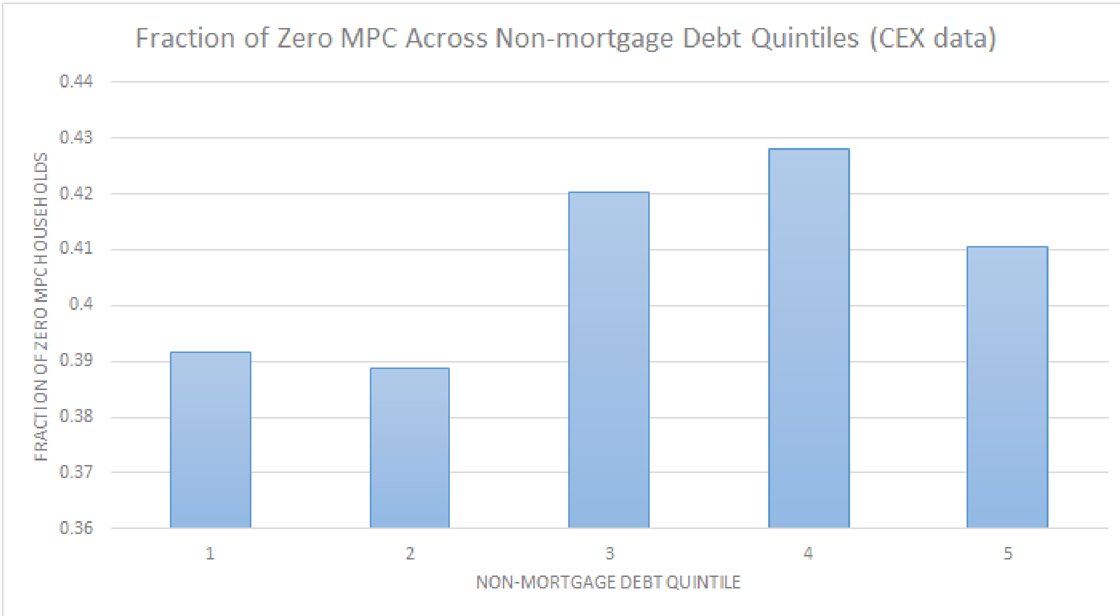

Figure 5
Fraction of low-MPC HH by non-mortgage debt quintile

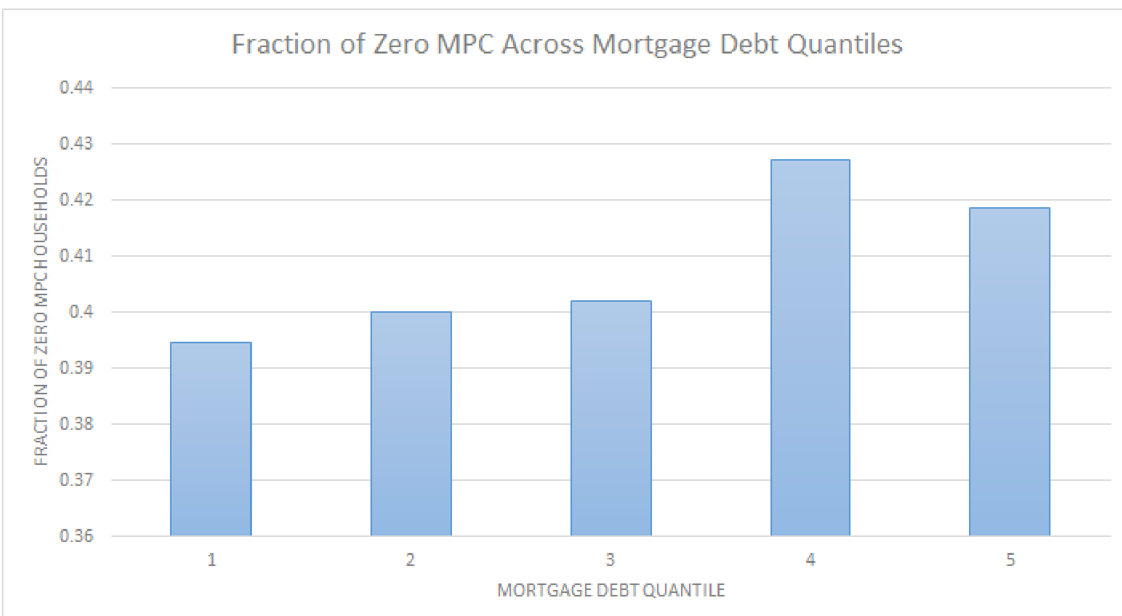

Figure 6
Fraction of low-MPC HH by mortgage debt quintile

5 Conclusions

We have presented a framework to reconcile a number of puzzling features of the microdata. First, household expenditure is nearly as volatile as income. Second, a large share of low-income, high-debt households have MPCs near zero. Third, we document that lagged high expenditure is associated with low contemporaneous spending propensities. Each category of expenditure contributes to this pattern.

To explain these facts, we develop a theory of saving-constrained households. The theory incorporates time-varying minimum consumption thresholds that, if violated, yield substantial utility costs. Households that experience a high minimum consumption threshold (relative to their wealth) increase consumption and debt. In order to avoid the potential utility cost of violating the minimum consumption threshold in the future (due to insufficient wealth/excess leverage) saving-constrained households buffer themselves by saving rather than spending out of additional income.

A striking implication of our theory is that a large share of high-debt households are saving-constrained rather than credit-constrained. As a result, income transfers to low-income, high-debt households are less expansionary in the short-tern (a few months) than previous models of heterogeneous agents would predict. Fiscal stimulus also tightens credit markets less than implied by prior theories. As a result, our theory has implications for the propagation of macroeconomic shocks. In Miranda-Pinto et al. (2019), we document cross-country differences in the macroeconomic response to fiscal shocks that are consistent with our theory's predictions.

References

Hall, R. E. (1978). Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy*, 86(6):971–987.

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

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

Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, 96(5):1589–1610.

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

Miranda-Pinto, J., Murphy, D., Walsh, K., and Young, E. (2019). Saving Constraints, Debt, and the Credit Market Response to Fiscal Stimulus: Theory and Cross-Country Evidence. *Mimeo University of Virginia*.

Misra, K. and Surico, P. (2014). Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs. *American Economic Journal: Macroeconomics*, 6(4):84–106.

Sahm, C. R., Shapiro, M. D., and Slemrod, J. (2015). Balance-Sheet Households and Fiscal Stimulus: Lessons from the Payroll Tax Cut and Its Expiration. NBER Working Papers 21220, National Bureau of Economic Research, Inc.