

w o r k i n g
p a p e r

20 04

A Model of Expenditure Shocks

Jorge Miranda-Pinto, Daniel Murphy,
Kieran James Walsh, and Eric R. Young



FEDERAL RESERVE BANK OF CLEVELAND

ISSN: 2573-7953

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

Working papers are available on the Cleveland Fed's website at:

www.clevelandfed.org/research.

A Model of Expenditure Shocks

Jorge Miranda-Pinto, Daniel Murphy, Kieran James Walsh, and Eric R. Young

We document four features of consumption and income microdata: (1) household-level consumption is as volatile as household income on average, (2) household-level consumption has a positive but small correlation with income, (3) many low-wealth households have marginal propensities to consume near zero, and (4) lagged high expenditure is associated with low contemporaneous spending propensities. Our interpretation is that household expenditure depends on time-varying consumption thresholds where marginal utility discontinuously increases. Our model with consumption thresholds matches the four facts better than does a standard model. Poor households in our model also exhibit “excess sensitivity” to anticipated income declines.

JEL classification codes: D14, E21.

Suggested citation: Miranda-Pinto, Jorge, Daniel Murphy, Kieran James Walsh, and Eric R. Young. 2020. “A Model of Expenditure Shocks.” Federal Reserve Bank of Cleveland, Working Paper No. 20-04. <https://doi.org/10.26509/frbc-wp-202004>.

*This paper was previously circulated under the title “Saving-Constrained Households.” Jorge Miranda-Pinto is at the University of Queensland (j.mirandapinto@uq.edu.au). Daniel Murphy is at the University of Virginia Darden School of Business (murphyd@darden.virginia.edu). Kieran James Walsh is at the University of California, Santa Barbara (kieranwalsh@ucsb.edu). Eric R. Young is at the University of Virginia, the Federal Reserve Bank of Cleveland, and Zhejiang University (ey2d@virginia.edu). The authors would like to thank Douglas Diamond; Peter Rupert; and seminar and workshop participants at the 2019 LAEF/WAMS; the 2019 SED; the Federal Reserve Bank of Cleveland; the Federal Reserve Bank of Richmond; Indiana University; Notre Dame; the Reserve Bank of Australia; Texas A&M; the University of California, Berkeley; the University of California, Santa Cruz; the University of New South Wales; and the University of Queensland for helpful comments. The collection of data used in this study was partly supported by the National Institutes of Health under grant number R01 HD069609 and R01 AG040213, and the National Science Foundation under award numbers SES 1157698 and 1623684.

1 Introduction

In this paper we establish that household-level consumption is volatile yet disconnected from income. We begin with three facts from microdata on the joint dynamics of income and consumption. First, for the average household in the Panel Study of Income Dynamics ([ISR \(2019\)](#)), consumption is as volatile as income. In contrast, simple permanent income hypothesis models imply that consumption should be much smoother than income. Two potential remedies are pervasive liquidity constraints (e.g., [Kaplan and Violante \(2014\)](#)) or very persistent income shocks, both of which closely tie consumption to current income (vs. permanent income).

However, Facts 2 and 3 raise doubts about these fixes. Our second fact is that for the average household, consumption and income growth have a low correlation of around 0.2. This suggests that explanations for volatile consumption cannot rely on a strong dependence of consumption on contemporaneous income.¹ Our third fact is that many low-wealth households exhibit marginal propensities to consume (MPCs) near zero. This fact is primarily based on our review of the literature, although we also replicate a recent study of MPC heterogeneity, [Misra and Surico \(2014\)](#), to document the large fraction of low-wealth households with MPCs near zero. Fact 3 also brings into question liquidity constraints as an explanation for volatile consumption: if liquidity constraints are to deliver highly volatile consumption, the poor (those close to the liquidity constraint) must have high MPCs.²

We propose a joint explanation of these facts based on a theory in which households face time-varying consumption thresholds that, if violated, yield substantial utility costs. These consumption thresholds represent unanticipated shocks such as medical emergencies and auto repairs. When an adverse shock hits a household, the household chooses to accumulate debt (reduce

¹Facts 1 and 2 hold even when focusing exclusively on nondurable consumption, which suggests that the durability of goods is not a driving force.

²Relatedly, [Luo et al. \(2017\)](#) argue that borrowing constraints have difficulty replicating the feature of the PSID that the cross-sectional dispersion in consumption relative to income is nearly constant across the income or wealth distribution.

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. And the low net asset position frontloads consumption suboptimally, relative to a world without expenditure shocks. Therefore, households for which consumption is against the threshold use additional income to pay off debt (increase net assets) as a precautionary measure.

We refer to these households as “saving-constrained” to capture the notion that, in the absence of binding consumption thresholds, households would save rather than reduce their asset position. Consumption thresholds effectively constrain households’ saving relative to a frictionless benchmark, just as credit constraints constrain borrowing relative to a frictionless benchmark.³ And just as “credit-constrained” households can presumably access credit only at exorbitantly high costs, “saving-constrained” households can save more (consume less than the consumption threshold) only by paying a large utility cost.

We develop a heterogeneous-agent model featuring stochastic consumption thresholds that are persistent and idiosyncratic income shocks. Income is exogenous and has both a persistent component and a near-permanent component. Households smooth consumption using a one-period risk-free asset. If a household chooses to consume below the threshold in a given period, it pays a utility cost proportional to the difference between consumption and the threshold. We calibrate the income process to the PSID and then estimate the model with time-varying consumption thresholds (which nests the model without thresholds) to attempt to match Facts 1–3 and other moments. The estimation uncovers a consumption threshold process that is persistent,

³To see this mathematically, let c be consumption, y be income, k be current assets, and k' be assets carried to the next period, so the budget constraint is $c + k' = y + k$. A consumption threshold $c \geq \underline{c}$ is equivalent to an upper bound on saving: $k' - k \leq y - \underline{c}$.

highly volatile, and (nearly) mean zero.⁴ The threshold model substantially improves the fit to our moments, relative to the standard “Bewley” model with no thresholds. That including the thresholds improves the fit is, on one hand, not surprising, as the model with thresholds has more free parameters. On the other hand, there was ex ante no guarantee that the thresholds would deliver the large fit improvement we observe, so we interpret the estimation results as evidence for our proposed explanation of Facts 1–3.

To attempt to cross-validate our consumption threshold model, we turn to an untargeted moment, Fact 4: in the PSID we document that contemporaneous spending propensities are strongly dependent on prior expenditure. In particular, lagged high expenditure (on broad consumption or exclusively nondurables) is associated with low current MPCs out of income, conditional on wealth and other covariates. A household that experienced high consumption in the prior period exhibits a contemporaneous spending propensity that is approximately 20 percent lower than the average (unconditional) spending propensity. The notion that prior 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.

The estimated model generates this path-dependence of spending propensities that we observe in the PSID, while the model without expenditure shocks does not. In the model with consumption thresholds, high lagged expenditure is associated with lower contemporaneous propensities to spend because high expenditure is a proxy for households being saving-constrained by high consumption thresholds. In summary, our framework is capable of explaining key moments of the household-level relationship between consumption and in-

⁴We constrain the consumption threshold process to be an AR(1). Since the estimated process is nearly mean zero, the implication is that about half of the realized thresholds are never binding, as the households never choose negative or near-zero consumption.

come (including the prevalence of medium to low-income households with high savings propensities) with a straightforward, intuitive, and easy-to-implement modification to standard theory.⁵

Our framework also helps to reconcile theory with otherwise puzzling consumer behavior documented by recent empirical work. In particular, [Ganong and Noel \(2019\)](#) show that, among low-wealth households receiving unemployment insurance (UI) benefits, spending drops precipitously upon the predictable expiration of UI benefits. This is puzzling from the perspective of standard theory since even low-income/poor households should be able to smooth over anticipated negative shocks. Since the income decline is predictable, households should cut consumption immediately, which is always feasible. Borrowing constraints would not matter for this response, since consumption is falling. Households in our model, in contrast, can exhibit the [Ganong and Noel \(2019\)](#) behavior if their consumption is at the threshold. In these cases, a decline in income (even if anticipated) leads households to consume below the threshold and pay the utility cost. While they *could* have smoothed over the shock, they choose not to because doing so entails paying the utility cost in the periods prior to the anticipated income decline.⁶

The consumption thresholds in our model are related to the notion of consumption commitments in [Chetty and Szeidl \(2007\)](#) – goods that are infrequently adjusted and are subject to transaction costs. In our theory, house-

⁵Our framework shares with standard Bewley models the prediction that average MPCs in response to unanticipated transitory income shocks are low (< 0.2) compared to the empirical evidence. We conjecture that integrating our theory of saving constraints with existing frameworks that deliver higher average MPCs (e.g., [Kaplan and Violante \(2014\)](#), [Carroll et al. \(2017\)](#)) will be fruitful for developing a comprehensive understanding of the MPC distribution.

⁶Our theory also provides a simple explanation for the findings of [Shea \(1995\)](#). Using the PSID and data on union contracts, he shows that consumption responds to anticipated income declines but not anticipated income increases, which is the opposite prediction of a model with liquidity constraints. Our model can generate this asymmetry if saving constraints are more likely to bind than are borrowing constraints. In that case, households can smooth out income increases through borrowing/dissaving, while they are limited in their ability to smooth income declines by the possibility of contemporaneously hitting consumption thresholds. See [Jappelli and Pistaferri \(2017\)](#) and [Jappelli and Pistaferri \(2010\)](#) for surveys of the literature on “excess sensitivity” to anticipated income declines.

holds subject to a shock (e.g., a medical emergency) may, to avoid violating the threshold, not adjust downward other components of consumption and maintain relatively stable expenditure. The notion that many aspects of consumption are difficult to adjust is one way to interpret our assumption that threshold shocks affect consumption broadly. In the face of a medical emergency, a necessary education expense, or an auto repair, for example, households do not move out of their homes or downgrade household appliances. One can also interpret our consumption thresholds as a reduced-form way to overlay stochastic maintenance costs on top of committed consumption. For example, the need for households to spend to repair their vehicles arises from the fact that the vehicle represents committed consumption. Even medical emergencies can be considered a form of maintenance: maintenance of a person’s physical body. The consumption thresholds can also be interpreted as major expenses that [Campbell and Hercowitz \(2018\)](#) have documented are important for understanding household life-cycle saving behavior and the MPC distribution.

For comparison, we note here that alternative shocks will not replicate our facts. For example, many papers in the literature add shocks to discount factors in order to replicate the concentration of wealth observed in US data ([Krusell and Smith \(1998\)](#), [Hubmer et al. \(2019\)](#)). While these shocks can help the model address Fact 1, it is unclear if they will reconcile Fact 2, and they will not help with Fact 3 – in the stationary distribution poor households will have low discount factors and therefore high MPCs. Similarly, idiosyncratic variation in returns ([Hubmer et al. \(2019\)](#)) will deliver the same prediction, as poor households will have low returns.⁷

⁷There are other models of household expenditure shocks, for example, [Chatterjee et al. \(2007\)](#) and [Livshits et al. \(2007\)](#), although those frameworks do not have positive consumption thresholds. In these models expenditure shocks are equivalent to negative wealth shocks, and since they mechanically decrease net assets, these shocks raise the marginal propensity to consume, all else equal. In contrast, our expenditure shocks are an increase in the consumption threshold and endogenously decrease net assets. While the previous literature’s version of expenditure shocks is useful for understanding consumer bankruptcy, it does not generate the strong post-shock deleveraging motive present in our framework. In the previous literature, households want to save *less* after an expenditure shock.

The remainder of the paper proceeds as follows. Section 2 lists the set of facts on the joint dynamics of income and consumption. Section 3 presents the model of expenditure shocks and implications for the MPC distribution. Section 4 evaluates the model’s fit to the data by comparing moments on the joint dynamics of income and consumption from simulated data to the moments from the PSID. Section 5 simulates the dynamic effects of one-time transfers to households in our model and demonstrates how the model can be reconciled with recent evidence on the effect of anticipated income declines. Section 6 concludes.

2 The Joint Dynamics of Household-Level Income and Consumption

Here we highlight evidence on household-level income and consumption that, when taken together, is difficult to reconcile with existing theories of heterogeneous households with uninsurable idiosyncratic income risk. The relevant facts are the following: (1) Consumption is as volatile as income for the average household. (2) Household-level consumption is relatively uncorrelated with income. (3) A large share of low-wealth households have spending propensities of effectively zero. (4) Lagged high expenditure is associated with low current spending propensities.

Our analysis with respect to Facts 1, 2, and 4 relies on data from the 1999-2017 Panel Study of Income Dynamics (PSID), which is a biennial panel study of households that are representative of the US population. The PSID is the most comprehensive data set that tracks household-level expenditure and income over an extended period of time. Starting in 1999 the PSID began collecting 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. With these additions, the PSID captures nearly all of the expenditure categories measured by the Consumer Expenditure Survey (CEX), which collects data at a higher

frequency but does not have the longitudinal dimension that characterizes the PSID.⁸ For our analysis, we normalize each variable (consumption, income, wealth) by the personal consumption expenditure price index (PCE) for the year in which each measure is reported. To maintain a balanced panel, we restrict our sample to respondents who are in each wave of the PSID from 1999 through 2017. Our resulting sample consists of 6,159 households.

2.1 Facts 1 and 2: Consumption Is Volatile and Relatively Independent of Income

Table 1 presents key moments from the joint dynamics of income and consumption in the PSID. For each household, we compute the standard deviations of the change in log consumption and the change in log income. We compute the within-household ratio of these standard deviations and report cross-sectional summary statistics: the mean, median, and standard deviation of $\frac{sd(d\log C)}{sd(d\log I)}$. We also compute the within-household correlation between the change in log consumption and the change in log income and report summary statistics. Finally, we estimate the autoregressive coefficient on log consumption.

To benchmark these consumption moments, we simulate data from a calibrated Bewley model. To calibrate the model, we simulate a quarterly income process and choose the parameters of the data-generating process such that when the model is aggregated to a biennial frequency and truncated to the number of waves in our PSID sample, the estimated autoregressive process from the simulated data matches that in the data. Specifically, we assume that log income consists of a highly persistent component and a transitory component:

$$y_{i,t} = x_{i,t} + z_{i,t}, \quad x_{i,t} = \rho_x x_{i,t-1} + \sigma_x \varepsilon_{x,t}, \quad z_{i,t} = \rho_z z_{i,t-1} + \sigma_z \varepsilon_{z,t},$$

⁸As discussed by [Andreski et al. \(2014\)](#), the consumption data in the PSID closely correspond to those from the CEX. A number of features of the PSID help to improve upon the accuracy of the responses relative to the CEX. For example, the PSID offers respondents unfolding brackets when they cannot recall the exact amount spent on the subcategories of expenditure. This approach both improves response rates and improves data accuracy. The PSID also collects information at a more aggregated subcategory level than does the CEX.

where $\rho_x = .99$, and ε_x and ε_z are i.i.d. $\mathcal{N}(0,1)$. In the PSID, however, we observe only log annual income ($a_{i,\tau}$, $\tau = 1, 2, 3, \dots$) sampled biennially. In the model, this corresponds to

$$a_{i,1} = \log \sum_{t=1}^4 \exp(y_{i,t}), \quad a_{i,2} = \log \sum_{t=9}^{12} \exp(y_{i,t}), \quad a_{i,3} = \log \sum_{t=17}^{20} \exp(y_{i,t}), \quad \dots$$

In the 1999-2015 PSID, we estimate the panel regression model (with household and time fixed effects)

$$a_{i,\tau} = FI_i + \rho_a a_{i,\tau-1} + \gamma_\tau + \sigma_a \varepsilon_{a,\tau},$$

which yields estimates of $\text{var}(FI_i) \approx .87^2$, $\rho_a \approx .09$, and $\sigma_a \approx .86$. To calibrate the income process, we choose σ_x , σ_z , and ρ_z such that when we run the fixed-effects panel regression on model-simulated $a_{i,\tau}$ for $\tau \in [1, \dots, 9]$, the resulting values for $\text{var}(FI_i)$, ρ_a , and σ_a match what we see in the PSID.⁹ We then use the calibrated income process to compute a standard Bewley model (one asset, heterogeneous agents, and uninsurable idiosyncratic income risk). The specifics of the model are outlined in Section 3 below.

Table 1 presents key moments from the joint dynamics of income I and consumption C in the PSID alongside the same moments from the calibrated Bewley model. The model statistics are computed based on quarterly simulations that are transformed into biennial data. It is readily apparent that consumption is more volatile (relative to income volatility) in the data than in the Bewley model (Fact 1). Both the average of $\frac{sd(d \log C)}{sd(d \log I)}$ (1.05) and the median (0.85) far exceed the corresponding moments from the Bewley model. Furthermore, the ratio in the data exhibits a skewness that is not present in the model.

One possible explanation for volatile consumption in the data is that many

⁹The relatively small time dimension in our panel allows for the possibility of [Nickell \(1981\)](#) bias in the estimate of the autoregressive coefficient. Our estimated coefficient is relatively low (0.15) and similar to estimates based on standard methods that address the bias ([Anderson and Hsiao \(1981\)](#) and [Arellano and Bond \(1991\)](#)). In calibrating the model we apply the same OLS estimator used in the data to the simulated data truncated to nine periods.

households’ consumption tracks income due to very high MPCs. However, the second set of moments in Table 1 suggests that consumption is relatively independent of income (Fact 2). Whereas the average correlation between $d\log C$ and $d\log I$ in the model is 0.69, in the data it is only 0.24. Not only is consumption far less correlated with income than predicted by a standard model, it is also less persistent (bottom row of Table 1). Note that since the median consumption/income correlation in the data is also low, the low average correlation is likely not simply driven by a small number of rich households perfectly smoothing.¹⁰

The summary statistics from the PSID in Table 1 are nearly identical when focusing exclusively on nondurable goods (aggregate consumption net of purchases of new cars, furniture, appliances, and other household items such as floor coverings). Therefore, durability alone cannot account for consumption that is volatile and relatively independent of income.

2.2 Fact 3: Many Low-Wealth Households Have MPCs Near Zero

Much of the recent theoretical work on MPCs has focused on rationalizing the existence of high-MPC households. For example, [Kaplan and Violante \(2014\)](#) and [Carroll et al. \(2017\)](#) show that calibrated models with liquidity constraints and preferences exhibiting prudence can yield households with reasonably high MPCs through (i) precautionary savings, which generate a steep consumption function for lower wealth individuals, and (ii) binding liquidity constraints, which render some households “hand-to-mouth.”

To date there has been limited theoretical work that addresses an otherwise elusive set of findings from the empirical literature: a substantial frac-

¹⁰Household consumption’s high volatility and low correlation with income are not driven by particular parts of the income or wealth distribution. Dividing households into quartiles based on either average household wealth or income, the within-quartile average relative volatility ranges from 0.92 to 1.23, and the within-quartile average consumption/income correlation ranges from 0.17 to 0.29.

Table 1
Joint Dynamics of Consumption and Income: PSID and Bewley Model

		PSID	Bewley model
$sd(d\log C)/sd(d\log I)$	Mean	1.05	0.50
	Median	0.85	0.50
	St Dev	0.84	0.06
$corr(d\log C, d\log I)$	Mean	0.24	0.69
	Median	0.27	0.69
	St Dev	0.40	0.05
AR coefficient (C)		0.22	0.47

Note: The AR coefficient estimates from the PSID are based on an OLS regression with household and time fixed effects. Estimation procedures that account for possible Nickell bias yield slightly larger estimates. The PSID statistics are nearly identical when focusing exclusively on nondurable goods.

tion of households have an MPC of zero, and many of these households are not wealthy or high-income. Here we discuss evidence that many low-wealth households have MPCs of zero. We first survey and synthesize recent work that documents zero-MPC households. We then extend a prominent study of MPC heterogeneity, [Misra and Surico \(2014\)](#), to offer additional details on the prevalence of low MPCs among low-wealth households.

Most recently, [Bunn et al. \(2018\)](#) report an MPC of zero for 77 percent of surveyed British households with respect to positive shocks. Furthermore, the probability of reporting an MPC of zero is significantly higher for households with a mortgage loan-to-value ratio of 75–90 percent (vs. less than 75 percent). Prior studies have likewise documented a substantial (but smaller) share of households with MPCs near zero. Using the 2010 Italian Survey of Household Income and Wealth, [Jappelli and Pistaferri \(2014\)](#) find that around 22 percent of households would have an MPC of zero in response to a hypothetical income shock equal to a typical month of earnings. While the fraction of households with an MPC of zero increases in the cash-on-hand percentile, still around 10

percent of the *most* cash poor households have a zero MPC.

[Shapiro and Slemrod \(2003\)](#) and [Sahm et al. \(2015\)](#) survey US households around tax changes and present evidence related to zero-MPC behavior. Instead of soliciting precise MPCs, they ask households whether tax cuts (hikes) lead them “mostly” to increase (decrease) spending, mostly to increase (decrease) saving, or mostly to pay off (pay off less) debt. With respect to the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) [Shapiro and Slemrod \(2003\)](#) report that 78.2 percent of households say they would mostly increase saving or pay off debt. The authors find higher saving/delevering response rates for lower income households (vs. higher income households), nonstockholders (vs. stockholders), and households with small amounts of stock (vs. households with more stock). They also show households that “use credit” to “pay for unexpected expense” have higher saving/delevering rates than ones that “use savings” or “cut back spending.” In the [Sahm et al. \(2015\)](#) 2013 retrospective survey concerning the 2011-2012 payroll tax holiday, 65 percent of households say they mostly adjusted saving and debt/borrowing (the corresponding number was 86 percent in the prospective survey from March/April 2011). In the retrospective survey, more than 50 percent of these saving/debt adjusters have household income less than \$75,000 and around 20 percent make less than \$35,000.

Much of the above evidence is based on surveys of what households did or would do in response to an income shock. Given that self reports may not correspond to what households actually do in response to an income transfer, it is informative to examine direct estimates of MPCs. In a recent study of MPC heterogeneity in response to observed transfers, [Misra and Surico \(2014\)](#) exploit randomness in the timing of tax rebate transfers along with data from the Consumer Expenditure Survey (CEX). Their quantile regressions imply that for 40–50 percent of households, the spending response to tax rebates is not statistically different from zero. Inspection of Figures 1 and 4 in [Misra and Surico \(2014\)](#) suggests many of these households have moderate to low income (less than around \$50,000). Since the transfers in their study were potentially anticipated, the estimates are not necessarily equivalent to the MPC out of

a purely unanticipated transfer shock. Nonetheless, heterogeneous responses to anticipated shocks are informative about the constraints faced by different households.

Our survey of the existing evidence suggests that many low-wealth, high-debt households (the type of households typically associated with high MPCs) exhibit MPCs near zero. This echoes the survey of the literature in [Carroll et al. \(2017\)](#), who write, “much of the empirical work . . . does not find that the consumption response of low-wealth or liquidity constrained households is statistically higher.” Relatedly, [Kueng \(2018\)](#) finds, with respect to the Alaska Permanent Fund, that average MPCs are higher among households with higher income. [Lewis et al. \(2019\)](#), applying a “fuzzy C-means-based estimator” to the 2008 Economic Stimulus Act to capture household heterogeneity, also show a positive relationship between income (or liquid wealth) and MPCs.¹¹

To obtain a more complete understanding of the prevalence of low MPCs among low-wealth households, we redo the [Misra and Surico \(2014\)](#) analysis of MPC heterogeneity. Their quantile regression approach estimates MPCs for subgroups of households, where the subgroups are based on household quantiles of consumption changes during the period of a tax rebate. With their quantile estimates in hand, we then examine the relationship between the incidence of near-zero MPC and household wealth.

2.2.1 Short-Term Spending Propensities from the 2001 Tax Rebate: [Misra and Surico \(2014\)](#)

The mailing of the 2001 tax rebate was randomized based on the penultimate number of the tax filer’s Social Security number. Hence, the rebate receipt was exogenous to individual characteristics. Using household-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

¹¹[Lewis et al. \(2019\)](#) present mixed evidence on near-zero MPCs. In their full sample, they estimate that most households have a moderate but positive MPC. Restricting the sample to nondurables, however, 71 – 92 percent of households have MPCs “not statistically distinguishable from zero.”

main specification in [Johnson et al. \(2006\)](#) is

$$\Delta C_{it+1} = \sum_s \alpha_{0s} \times M_s + \alpha_1' X_{it} + \alpha_2 R_{it+1} + u_{it+1}, \quad (1)$$

where ΔC_{it+1} is household i 's 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-month period $t+1$.

[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

$$\Delta C_{it+1} = q(R_{it+1}, X_{it}, M_s, \lambda_{it+1}) \quad \text{with} \quad \lambda_{it+1} | R_{it+1}, X_{it}, M_s \sim U(0, 1), \quad (2)$$

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

$$\Delta C_{it+1} = q(R_{it+1}, X_{it}, M_s, \tau) = \sum_s \alpha_{0s}(\tau) \times M_s + \alpha_1(\tau)' X_{it} + \alpha_2(\tau) R_{it+1}. \quad (3)$$

The estimated consumption responses are common within a quantile τ but are heterogeneous across quantiles, representing unobserved heterogeneity.

2.2.2 Extending the [Misra and Surico \(2014\)](#) Analysis

We start by replicating the estimation of α_2 in equation (5) in [Misra and Surico \(2014\)](#). We focus on the estimated tax rebate coefficient of nondurable consumption for the 2001 tax rebate. Then, we gather additional information from the CEX on households' wealth. We define wealth as the sum of the value of holdings of checking accounts, saving accounts, US bonds, stocks, and property minus outstanding mortgage and nonmortgage debt. Non-mortgage debt is composed of credit card debt, bank loans, credit union debt, and dentist and hospital debt. We also define liquid wealth as the value of holdings of checking accounts, saving accounts, US bonds, and stocks minus nonmortgage debt, when nonmortgage debt is available.¹²

Figure 1 shows the distribution of MPCs for the group of low-wealth households (below median wealth). We classify MPC estimates below zero as zero. Some low-wealth households exhibit very large (≈ 1) MPCs, but around 40 percent of low-wealth households have an MPC estimate of zero or below. The distributions are nearly identical whether examining total wealth (left panel) or liquid wealth (right panel).

In Figure 2 we plot, for each quantile, the MPC estimate (with standard error bands) along with median wealth (total and liquid) for households in that quantile. With the total wealth sample, of the twenty quantiles eight have MPC estimates below zero (and are classified in Figure 1 as zero-MPC). Notably, the median wealth for most of these quantiles is low, but the lowest-wealth quantiles have higher MPC estimates (between 0.1 and 0.5). Further, three of the zero-MPC quantiles have moderate levels of median wealth (between \$45,000 and \$65,000). Therefore, zero-MPC behavior appears to be prominent among the low-but-not-lowest wealth levels and appears even at moderate wealth levels. The pattern is qualitatively similar in the liquid wealth sample (right panel of Figure 2). The fact that many households with low liq-

¹²The original [Misra and Surico \(2014\)](#) database has 13,606 household-time observations. We obtain total wealth data for 11,981 observations and liquid wealth for 5,608 observations. As noted by [Johnson et al. \(2006\)](#), a large fraction of households is missing information on liquid assets and non-mortgage debt.

uid wealth have near-zero MPCs strongly suggests that Fact 1 is not driven by liquidity-constrained households with high spending propensities.¹³

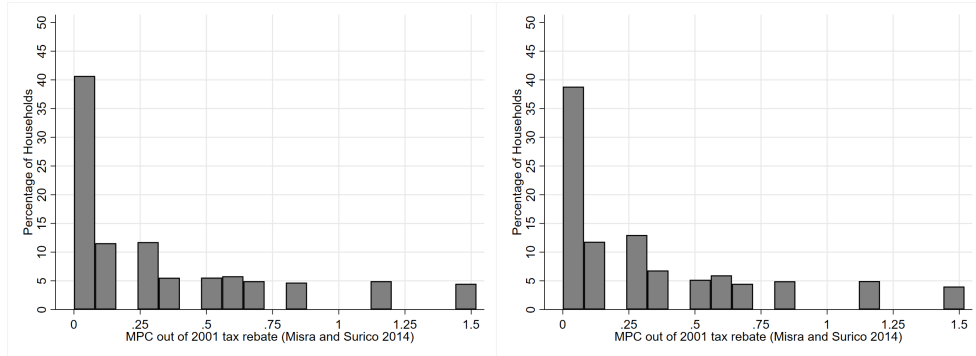


Figure 1

Note: This figure plots the distribution of estimated MPCs for US households with wealth below the median. We classify below-zero MPC estimates as having an MPC of zero. The left panel considers total wealth and the right panel considers liquid wealth.




Figure 2

Note: The x-axis displays median wealth for each quantile. The left panel considers total wealth and the right panel considers liquid wealth. The y-axis displays the MPC estimate (along with bootstrap confidence intervals) for each quantile from [Misra and Surico \(2014\)](#).

¹³When interpreting Figure 2, it is relevant that the rebates may have been anticipated by households. If the rebates were anticipated, then positive MPCs reflect deviations from the permanent income hypothesis (PIH) and the results suggest that many households with low wealth (or low liquid assets) did not deviate from the PIH. Low-wealth households appear to save additional income, either because their MPC out of unanticipated income is approximately zero, or because they are not liquidity constrained and hence do not respond to anticipated income transfers.

2.3 Fact 4: Lagged High Expenditure Is Associated with Low Contemporaneous Spending Propensities

In the standard heterogeneous-agent model 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 in the data lagged expenditure in fact 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.¹⁴ 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 one standard deviation. Our baseline specification is

$$\log C_{it} = \beta_0 \log I_{it} + \beta_1 HighC_{it-1} + \beta_2 \log I_{it} \times HighC_{it-1} + \gamma X_{it} + \epsilon_{it}, \quad (4)$$

where $HighC_{it-1}$ is a dummy variable that equals one when the expenditure of household i exceeds the within-household average by one standard deviation at period $t - 1$. I_{it} is the income of household i in period t and X_{it} includes $\log(\text{wealth})$, the interaction of the wealth term with $\log I$, 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 0.16 (column (1)). Column (2) shows that this spending propensity varies with lagged high expenditure, and in particular that having high lagged expenditure reduces the spending propensity by 0.030, or 18.4 percent of the average effect (0.163) of income, indicating a large state dependence of spend-

¹⁴One limitation of the PSID is that it is impossible to distinguish between anticipated and unanticipated changes in income. 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.

ing 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 the interaction of (log) wealth and current (log) income. The coefficient on the interaction of income with lagged high expenditure is slightly attenuated but remains economically and statistically significant.

Our main specification is presented in column (4). Here, we isolate incidents of high lagged expenditure that are not associated with high income. The specific reason for doing so will be apparent from the theory in Section 3. Briefly, in standard 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. We classify these episodes as “high-expenditure episodes.” In particular, a high-expenditure episode is a dummy variable that equals unity when household consumption exceeds its within-household mean by one standard deviation and household income does not exceed its within-household mean by one 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).

Table 2
Lagged Expenditure and Spending Propensities in the PSID

Dependent variable: $\log(C_t)$					
	(1)	(2)	(3)	(4)	(5)
$\log(I_t)$	0.163*** (0.008)	0.170*** (0.009)	0.181*** (0.014)	0.183*** (0.014)	0.168*** (0.014)
High C_{t-1}		0.409*** (0.121)	0.308** (0.113)		
$\log(I_t) \times$ High C_{t-1}		-0.030*** (0.011)	-0.021** (0.010)		
High C_{t-1} & not high I_t				0.486*** (0.125)	
$\log(I_t) \times$ (High C_{t-1} & not high I_t)				-0.039*** (0.011)	
High C_{t-1} & high I_{t-1}					0.122*** (0.011)
$\log(I_t) \times$ (High C_{t-1} & high I_{t-1})					0.046*** (0.001)
Control for $\log(\text{wealth}) \times \log(I_t)$	NO	NO	YES	YES	YES
N	61590	55431	55431	55431	55431
R^2	0.65	0.66	0.67	0.67	0.69

Note: A household has high income, High I_t , (expenditure, High C_t) in periods in which income (expenditure) is over one 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 time and household fixed effects, $\log(\text{wealth})$, and a quadratic in age of the head of household. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Column (4) 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 spend-

ing elasticity that is lower by 0.039, or 23.9 percent of the average elasticity of spending with respect to income.¹⁵ 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 flips signs (column (5)). 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.

Table 3
Indicator Variable Definitions and Incidence

Indicator Variable	Description	Incidence
High Expenditure	$C_t > \text{Mean}(C_t) + \text{SD}(C_t)$	0.16
High Income	$I_t > \text{Mean}(I_t) + \text{SD}(I_t)$	0.26
High Expenditure Episode	High Expenditure and Not High Income	0.10
High Expenditure & High Income	High Expenditure and High Income	0.07

Decomposing High=Expenditure Episodes. Which categories of expenditure drive high-expenditure episodes (henceforth referred to as “episodes”)? Are episodes primarily driven by subsets 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 shares during episodes.¹⁶ Expenditure shares during episodes (column (2)) are generally similar to average expenditure shares

¹⁵Estimates of the coefficient on the interaction term based on the Anderson-Hsiao procedure and the Arellano-Bond procedure are -0.025 and -0.035 , respectively.

¹⁶To compute the statistics in column (2), we first demean category-specific expenditure for each household to obtain a measure of excess expenditure at any point in time. We then average over households experiencing an episode to obtain average excess expenditure for a

(column (1)). Two categories are noticeably more prevalent during episodes: education and transportation. Food and housing are less prevalent during episodes.

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	0.14	0.87
Housing	0.44	0.36	0.31	1.60
Transportation	0.27	0.37	0.39	1.84
Education	0.04	0.10	0.14	0.92
Child Care	0.01	0.01	0.05	0.36
Health	0.08	0.07	0.11	0.75

Note: This table presents statistics for the broad categories that make up the measure of total expenditure. Expenditure on Clothing, Trips, and Other Recreation is 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 subcategory expenditure. All regression coefficients are significant at the 1 percent level.

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 household's transportation expenditure exceeds its within-household average by one 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 subcategory 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 category during episodes. We then do the same for total expenditure and take the ratio of the two.

to that category's share of total expenditure and follows a similar pattern of relevance that is depicted in column (2).

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 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.¹⁷ We also examine an alternative measure of expenditure that excludes purchases of durables. Specifically, we identify high-expenditure episodes based on extreme realizations of expenditure net of purchases of automobiles, furniture, appliances, and other household items such as floor coverings.

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 is not necessary for understanding the effect of lagged expenditure on spending propensities.

¹⁷Note that anticipation effects should, if anything, reduce our coefficient estimates toward zero. The fact that spending propensities are positive suggests that households behave as if changes in income are to some extent unanticipated.

Table 5
Specification with Alternative Expenditure and Income Measures

Dependent Variable: $\log(C)$				
Income measure:	$\log(I)$		High Income	
	All	Nondurables	All	Nondurables
Expenditure measure used to identify high-expenditure episodes:	(1)	(2)	(1)	(2)
	Coefficient on interaction term	-0.039*** (0.011)	-0.041*** (0.011)	-0.112*** (0.015)
Fraction of average effect of income measure	-0.24	-0.25	-0.43	-0.44

Note: Nondurable expenditure is defined as total expenditure net of purchases of automobiles, furniture, appliances, and other household items. 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 one standard deviation above average expenditure (income) for the household. All regressions include the income measure, the episode indicator, a quadratic in age, $\log(\text{wealth})$ 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.4 Summary and Interpretation

Facts 1 through 4 are puzzling from the perspective of standard heterogeneous-agent models with idiosyncratic income risk. It might seem that volatile consumption (Fact 1) could be attributed to households having higher MPCs than in the benchmark Bewley model. But a model featuring a stronger effect of income on consumption would exhibit a higher correlation between consumption and income, contrary to Fact 2. It would also not help resolve Facts 3 and 4.

It is possible that measurement error in consumption contributes to consumption that is volatile and independent of income (Facts 1 and 2). But significant measurement error in consumption would also prevent us from de-

testing Fact 4 by attenuating toward zero the estimate of the coefficient on the interaction between lagged high consumption and contemporaneous income. Nor can measurement error in the PSID account for the existence of low-wealth, low-MPC households (Fact 3). Therefore, we conclude that a joint explanation of the facts is not readily apparent, and we turn to a theory of expenditure shocks to help rationalize the evidence.

3 A Theory of Expenditure Shocks

Here we present a theory to rationalize the evidence presented above. We introduce random consumption thresholds into a standard Bewley model with capital and calibrate it to match the consumption and income dynamics from the PSID. We use the evidence in Facts 1 through 3 as calibration targets. We do not target the regression evidence from Fact 4, which serves as external validation.

3.1 Model

The economy consists of a measure one of infinitely-lived households 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 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

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

subject to the budget constraint

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

where V is the value function, c is consumption (the numeraire), k is capital wealth (which exogenously depreciates at rate $\delta \geq 0$), z is persistent idiosyn-

cratic household productivity, x is (nearly) permanent idiosyncratic productivity, and \underline{c} is a persistent consumption threshold. If the household consumes less than \underline{c} , it must pay utility cost $\lambda(\underline{c} - c)$ for some $\lambda \geq 0$. 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 $\exp(z+x)\bar{h}$ at wage w . For any variable q , q' represents its value in the subsequent period.

We assume that household productivity and consumption thresholds evolve according to:

$$z' = \rho_z z + \epsilon'_z, \quad x' = \rho_x x + \epsilon'_x, \quad \underline{c}' = (1 - \rho_c)\mu_c + \rho_c \underline{c} + \epsilon'_c,$$

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 \bar{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 are denoted by stars and 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:

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

(3) the labor market clears: $L^* = \bar{H}$, and (4) Ω^* is generated by k'^* . Let \bar{K} denote aggregate household capital, define $y = (r - \delta + \phi \mathbf{1}(k \leq 0))k + w \exp(z+x)\bar{h}$ 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 ($\rho_z = 0.63$, $\sigma_z = 0.79$, $\rho_x = 0.99$, and $\sigma_x = 0.14$) to approximate household income from the PSID (as discussed in Section 2) and set the borrowing cost, capital share,

depreciation rate, global interest rate, and borrowing constraint to reasonable values: $\phi = 0.03$ (≈ 12 percent annual premium on borrowing vs. saving), $\alpha = 0.36$ (standard in the literature), $\delta = 0.0125$ (standard in the literature), $r = 0.0225$ (so the net return on saving is 1 percent per quarter), and $b = -1$ (about 1/3 of quarterly net labor and capital income, with the normalization $\bar{h} = 0.33$). In the second step, we choose the remaining parameters (discount rate, utility cost, and \underline{c} process) to target the following moments from the ergodic distribution of the stationary equilibrium:¹⁸

$$\begin{aligned} \frac{\bar{K}}{\bar{Y}} &= 12 \\ \text{mean}(\text{corr}(d\log y, d\log c)) &= 0.24 \\ \text{mean}\left(\frac{\text{std}(d\log c)}{\text{std}(d\log y)}\right) &= 1.05 \\ \text{Fraction}(k < 0) &= 0.1 \\ \text{corr}(\log(c_t), \log(c_{t-1})) &= 0.22 \\ \text{Fraction}(MPC < 0.01) &= 0.3 \\ \text{mean}(MPC) &= 0.2, \end{aligned}$$

where $\text{corr}(d\log y, d\log c)$, $\frac{\text{std}(d\log c)}{\text{std}(d\log y)}$, and $\text{corr}(\log(c_t), \log(c_{t-1}))$ are computed at a biennial frequency and are based on the values in Table 1 above. $\bar{K}/\bar{Y} = 12$ is standard, $\text{Fraction}(k < 0) = 0.1$ is in the ballpark of commonly used values in the consumer credit literature (e.g., Athreya et al. (2009)), and $\text{mean}(MPC) = 0.2$ is at the low end of the range suggested by Carroll et al. (2017). Finally, $\text{Fraction}(MPC < 0.01)$ governs the proportion of households that are saving-constrained. Thirty percent is in the range of estimates of the fraction of households with near-zero MPCs (from the papers cited in Section 2.2).

¹⁸Given 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 \bar{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.

Via global optimization, the best-fit parameter values are $\beta = 1/1.049$, $\lambda = 257.6$, $\rho_c = 0.65$, $\sigma_c = 1.48$, and $\mu_c = 0.0047$. The baseline Bewley model is computed analogously but without the \underline{c} process and with β re-calibrated to match $\bar{K}/Y = 12$.

3.3 Results

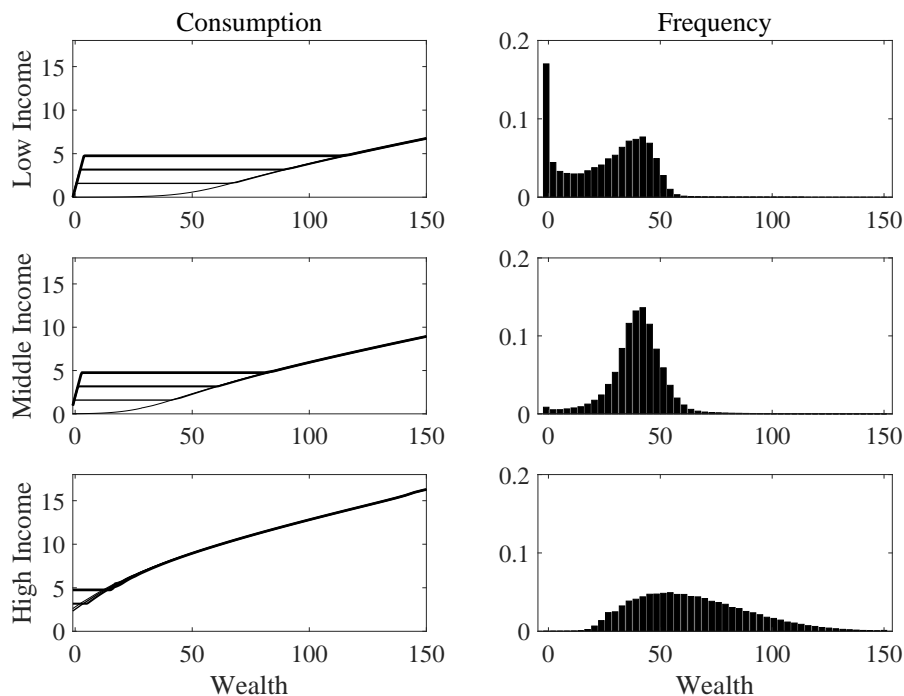


Figure 3

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.

Figure 3 shows the consumption functions for households with different realizations of \underline{c} . 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 are “saving-constrained” and use all additional wealth/income to save.¹⁹ However, not all low-wealth households with high realizations of \underline{c} have MPCs of zero. The poorest households cannot even achieve the consumption threshold. They consume all additional income, up until they are able to consume at the threshold.

What does this imply for the MPC distribution across households? Comparing the wealth distributions in Figure 3 with the consumption functions, 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: the former have moderate MPCs consistent with the permanent income hypothesis, while the latter are violating their threshold and anxious to not pay the utility cost. Figure 4 is the ergodic MPC histogram for below-median wealth households.²⁰ Echoing the corresponding Figure 1 from the data, Figure 4 shows that for low-wealth households, the largest mass of MPC is around zero. Relative to the data, the model is lacking intermediate MPC low-wealth households, although both the model and data histograms exhibit a mass of very high MPC households.

Figure 5 shows the model’s average MPCs by wealth decile, illustrating that MPCs are U-shaped in wealth. The lowest MPC households are, on average, not rich. They are poor enough for the minimum consumption level to matter but not so poor that they violate it. This pattern is reminiscent of [Misra and Surico \(2014\)](#), who uncover a non-monotonic U-shaped relationship between either income, homeownership (with a mortgage), or liquid assets and spending propensities. They write, “the largest propensity to consume out of the tax rebate tends to be found for households with both high levels of mortgage debt and high levels of income.”

¹⁹The saving-constrained (zero-MPC) households could save more (consume below the consumption threshold), but they chose not to do so because of the large utility cost associated with violating the threshold. This notion of saving-constrained households is analogous to the notion of credit-constrained households, which face prohibitively large costs of accessing credit.

²⁰Figures 4 and 5 are constructed from model simulations with 10,000 households and a burn-in period of 300 quarters. For Figure 4, there are 50 bins with a bin size of 0.0219. MPCs greater than 1 are set to 1.




Figure 4

Note: The figure shows the model's ergodic MPC distribution for households with below-median wealth. The bin size is 0.0219.




Figure 5

Note: The figure shows average MPC by wealth decile in the model's ergodic distribution.

4 Model Fit: Comparison to PSID

As discussed above, the model features a substantial fraction of zero-MPC households that have low-to-median wealth (Fact 3). Here we examine the

model’s ability to improve the fit to the PSID with respect to Facts 1, 2, and 4. To make the results comparable with the PSID, we transform the simulated quarterly data into biennial data.

With respect to Facts 1 and 2, Table 6 shows that including the minimum consumption shock substantially improves the fit to the PSID. Whereas consumption in the standard Bewley model is insufficiently volatile, too correlated with income, and too persistent, in the expenditure shock model consumption is nearly as volatile as income (on average), relatively uncorrelated with income, and less persistent. Consumption volatility in the expenditure shock model also exhibits skewness as in the PSID.

Table 6
Consumption and Income Moments

		PSID	Bewley model	Expenditure Shock Model
$sd(d\log C)/sd(d\log I)$	Mean	1.05	0.50	1.11
	Median	0.85	0.50	0.53
	St Dev	0.84	0.06	1.07
$corr(d\log C, d\log I)$	Mean	0.24	0.69	0.37
	Median	0.27	0.69	0.35
	St Dev	0.40	0.05	0.22
AR coefficient (C)		0.22	0.47	0.10

Note: The PSID AR coefficients are estimated after removing household and time fixed effects. All model moments are based on data simulated at a quarterly frequency and then converted to biennial data.

With respect to Fact 4, we first demonstrate that the theory fits the empirical relationship between lagged expenditure and spending propensities at a biennial frequency. We then analyze simulated quarterly data (for which we can directly observe saving constraints) to assess how well the empirical specification captures the behavior of saving-constrained households. In assessing the model’s fit to Fact 4, it is important to note that the regression evidence was not used to calibrate the model.

Table 7 shows that spending propensities out of income are lower when households experienced high-expenditure episodes in the previous period (column (4)). In the Bewley model without consumption thresholds, average spending propensities are approximately independent of lagged high expenditure (column (6)).

The dependence of spending propensities on lagged expenditure in the expenditure shock model is based on the relationship between persistent consumption thresholds and MPCs. The theory implies that saving constraints (binding consumption thresholds) are associated with lower spending propensities. Saving constraints are persistent (due to the persistence of the consumption threshold), which implies that lagged saving constraints are also associated with lower spending propensities. High-expenditure episodes are a proxy for saving constraints: when households receive a high realization of \underline{c} , their expenditure increases even in the absence of corresponding income increases. Therefore, the theory implies that a high-expenditure episode, which is a proxy for a saving constraint, tends to be associated with lower future spending propensities. Below we examine these relationships in more detail using quarterly data generated by the model.

Table 7
Lagged Expenditure and Spending Propensities in the Model

Dependent variable:	Models					
	PSID		Expenditure Shock		Bewley	
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(C)$						
$\log(I)$	0.163	0.170	0.238	-0.126	0.228	0.202
L. High Exp. Episode		0.486		0.077		0.030
$\log(I) \times$ L. High Exp Episode		-0.039		-0.038		0.009
Control for $\log(\text{wealth})$	NO	YES	NO	YES	NO	YES
$\times \log(I)$						
R^2	0.65	0.67	0.68	0.69	0.99	0.99

Note: A high-expenditure episode is a dummy that equals unity when a household experiences high expenditure but not high income. All regressions control for $\log(\text{wealth})$. All regressions have household fixed effects, and the PSID ones also have time fixed effects. Standard errors are not shown since regression estimates are highly precise. Models are simulated at a quarterly frequency and then converted to biennial data before estimating regressions.

4.1 Analysis Based on Quarterly Data

The advantage of analyzing quarterly data is that we can observe whether households are saving-constrained at any point in time. If the data are aggregated to a 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.

Table 8 shows summary statistics and regression coefficients based on the quarterly data produced by the model. 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 17 percent of households are saving constrained in any given quarter (row (3)), which implies that a substantial share of households have MPCs near zero. One percent of households experience saving constraints for 8 consecutive quarters (row (4)). This persistence helps explain how the effects of savings constraints can be detected even in time-aggregated data.

Table 8
Model Summary Statistics and Regression Coefficients (Quarterly Data)

		Saving Constraint Model	Bewley Model
		(1)	(2)
(1)	Coefficient on interaction between $\log(\text{income})$ and saving constrained	-0.34	NA
(2)	Interaction coefficient as fraction of average spending elasticity	-2.43	NA
Share of agents who			
(3)	Are on saving constraint in a given quarter	0.17	NA
(4)	Are on saving constraint for eight consecutive quarters	0.01	NA
(5)	Have high consumption	0.19	0.18
(6)	Have high consumption but not high income (episode)	0.09	0.04
(7)	Pay the utility cost (consumption < minimum threshold)	0.04	NA
Correspondence between high consumption episode and saving constraint:			
(8)	Incidence of episode if saving constrained	0.42	NA
(9)	Incidence of saving constraint if episode	0.74	NA

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

Saving constraint episodes do not perfectly correspond to high-expenditure episodes. While the incidence of being saving constrained is 17 percent, the incidence of a high-expenditure episode is only 9 percent (row (6)) 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 three-fourths of households that have a high-expenditure episode are also saving-constrained (row (9)), suggesting that the proxy used in the regressions does indeed capture a large share of households that have MPCs of zero.

5 Consumption Responses to Temporary Income Shocks

Here we examine the implications of saving constraints for the dynamic effects of stimulus measures. We also demonstrate how our model can help explain otherwise puzzling evidence that for many low-income households, anticipated reductions in income are associated with large declines in consumption on impact.

The model with consumption thresholds and saving constraints features lower average spending propensities than would occur in an environment without thresholds. This implies that stimulus measures, such as one-time transfers, should be less effective in the short run. Figure 6 simulates the effect of a one-time unanticipated transfer to all households in the saving constraint (expenditure shock) model and in the standard Bewley model. Spending on impact is indeed much lower (by over half) in the saving constraint model. But despite stimulus being less effective in the short run, its effects are more persistent. Saving-constrained households eventually move away from the constraint and use additional income to increase their consumption.

In the expenditure shock model, what is the heterogeneous effect of the transfer? Figure 7 shows the average response by saving constraint status, starting at the ergodic distribution.²¹ For each group (saving-constrained, unconstrained, and paying-the-utility cost), the consumption response to the transfer is expressed as an elasticity with respect to average group income in the period of impact. Paying-the-utility-cost households have a large initial response, but within a year, consumption returns to the level of the control group (that does not receive transfer). In other words, these households immediately consume the transfer, without substantially changing their medium and long-term prospects. Saving-constrained households, on the other hand, initially save *all* of the transfer, allowing them to increase consumption over a

²¹Figures 7 and 8 are constructed by drawing 100,000 households from the stationary distribution and simulating the economy for 30 quarters. Given initial conditions and simulated shock paths, impulse responses are defined by differences between the consumption paths with and without an initial transfer of 0.5, all else equal.

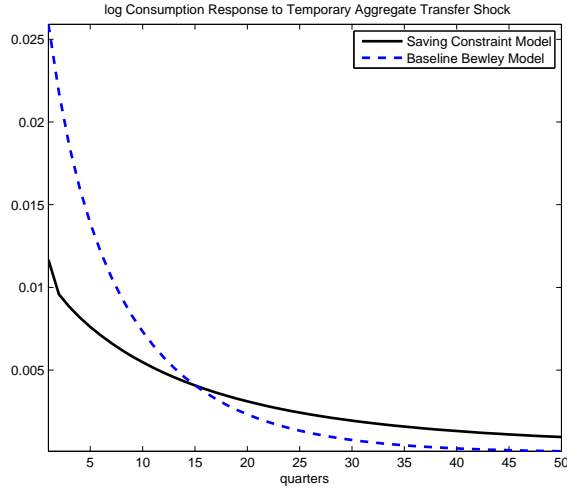


Figure 6

Note: The figure shows the effect on aggregate consumption of a one-time wealth transfer of 0.5 to all households. Each line is the log difference between the average consumption path with and without the transfer.

longer horizon. That is, their consumption response is hump-shaped. Unconstrained households have a moderate initial increase in consumption, which dissipates slowly. By about 20 quarters, the consumption response of saving-constrained households exceeds that of unconstrained ones. This illustrates a key consequence of binding consumption thresholds: they render households saving-constrained in the sense that they have pent-up demand for saving. Transfers help them save in the short run and consume over longer horizons.

Figure 8 repeats the exercise of Figure 7 but defines groups by wealth quintile instead of saving constraint status. Echoing the U-shaped MPC by wealth from Figure 5 and Misra and Surico (2014), the consumption response to the transfer is nonmonotonic in wealth. While the bottom 20 percent poorest households have the largest initial response, the 20-40 percent by wealth have the weakest initial response, as they are more likely to be saving-constrained. The richest 20 percent of households are unlikely to be saving-constrained and have the second highest initial response of consumption to the transfer. However, since many of the poorest households are effectively hand-to-mouth, the average response of that group quickly dissipates and within a few quarters

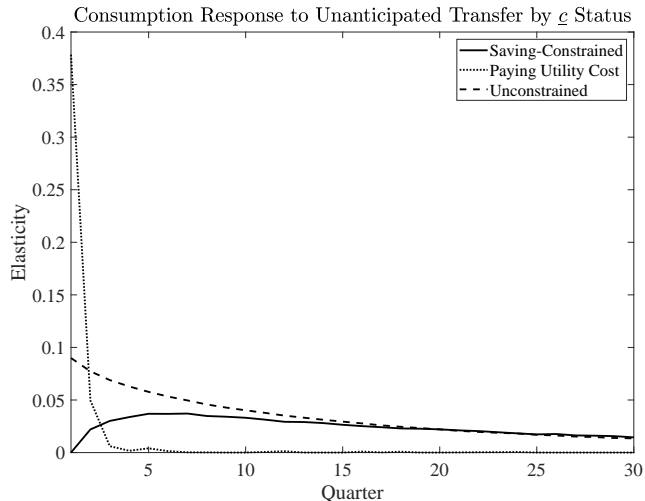


Figure 7

Note: The figure shows the average consumption response to a one quarter unanticipated transfer of 0.5 for saving-constrained households (solid, $|c - \underline{c}| < 0.00001$), unconstrained households (dashed, $c \geq \underline{c} + 0.00001$), and households paying the utility cost (dotted, $c \leq \underline{c} - 0.00001$), expressed as an elasticity with respect to initial income. For each group, the line is the log difference between the average response of households with and without the transfer, divided by .5 over average group income, starting at the ergodic distribution.

the poorest households have the lowest response.

5.1 Rationalizing Evidence that Consumption Responds to Anticipated Declines in Income

Motivated by empirical evidence from the prior literature, we also examine the effect of anticipated income declines in the model. [Ganong and Noel \(2019\)](#) document that anticipated declines in income cause consumption to plummet on impact, especially for low-income households, and [Shea \(1995\)](#) shows that consumption responds to anticipated income declines but not anticipated income increases. Similarly, [Bernheim et al. \(2001\)](#) document a large drop in consumption for new retirees, especially for the lowest income quartile.

These findings are puzzling from the perspective of standard theory, since even poor households should be able to smooth over expected income declines. Consumption thresholds help explain this behavior because some households

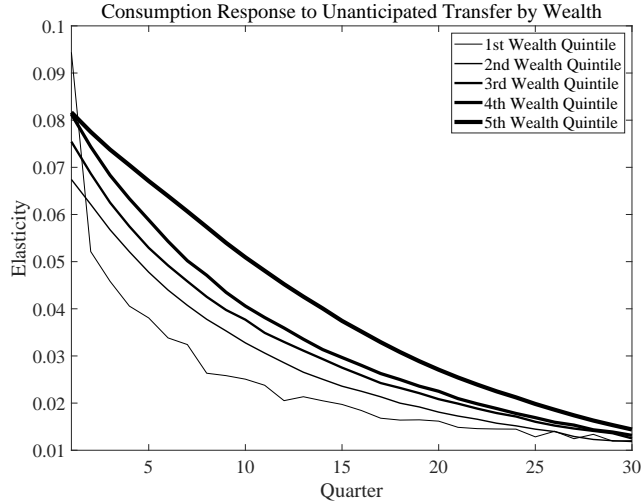


Figure 8

Note: The figure shows the average consumption response to a one-quarter unanticipated transfer of 0.5 by wealth, expressed as an elasticity with respect to initial income. Line thickness corresponds to wealth quintile, so the thinnest line represents the bottom 20 percent by wealth in the ergodic distribution. For each group, the line is the log difference between the average response of households with and without the transfer, divided by .5 over average group income, starting at the ergodic distribution.

maintain consumption at the threshold rather than pay a utility cost in advance of the income decline. Specifically, when these households receive news that income is expected to decline, they maintain consumption at the threshold (and therefore do not cut consumption) in hopes that future income may be higher than expected (that is, there is an option value of waiting until the realization of future income). Even if the future income decline is known with certainty, agents may maintain consumption at the threshold depending on how they trade off current utility costs with future utility costs.

To demonstrate this in the expenditure shock model, we assume that households experience a decline in their labor income of 25 percent in period 5, which remains until period 15, when it returns to normal. In Figure 9, we show the percent change in consumption between periods 4 and 5 for the lowest income households (lowest x_4 and z_4), as a function of wealth at period 4 (k_4), for different levels of \underline{c}_4 (represented by line thickness in the figure). While wealth can endogenously change between the periods, we assume x , z , and \underline{c} stay at

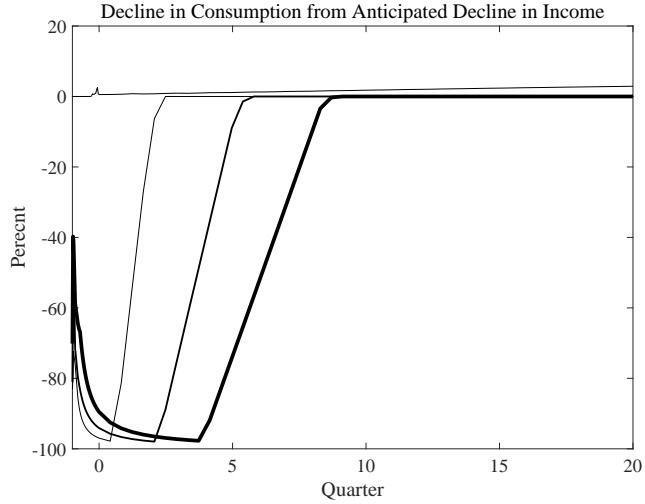


Figure 9

Note: The figure shows the percent change in consumption between periods 4 and 5 from a 25 percent decline in income in periods 5 through 15 that was announced in period 1, as a function of period 4 wealth, for households with the lowest realization of the income process in period 4. Line thickness corresponds to \underline{c}_4 , so that the thickest represents the highest realization of \underline{c} and the thinnest represents the lowest positive value of \underline{c} . While wealth changes endogenously between periods 4 and 5, we assume the exogenous shocks remain unchanged between 4 and 5.

their period 4 values. For sufficiently low wealth, consumption dramatically falls between periods 4 and 5 for households with high \underline{c} , even though the income decline was anticipated a year earlier. These households are willing to sacrifice smoothness in consumption to limit utility costs from breaking the threshold before the anticipated income decline, which they realize may eventually force them to pay the cost in the discounted future. However, as wealth increases, households have resources sufficient to continue to consume at the threshold, even with the expected income decline. With low or negative \underline{c} , there is no decline in consumption between periods 4 and 5: these households face no impediment to smoothing over the expected income decline. Note that the very lowest and intermediate wealth households have muted consumption declines. For the former, this is because they were already breaking the threshold. For the latter this is because they can somewhat soften the blow of the income decline through winding down wealth before breaking the threshold.

6 Conclusion

A common anecdote referenced by policymakers and politicians is of an American household with limited financial resources that is susceptible to an adverse shock, such as a health expense or a broken automobile, that causes the household to accumulate debt. This debt is burdensome in the sense that additional income is allocated to debt service (net asset accumulation) rather than additional consumption. The poorest of these households are often considered especially vulnerable because they may forgo medical care, food, or other basic necessities in order to service this unwanted debt burden.

In this paper, we explore this story. We first establish four empirical facts: three of which guide our theory and a fourth that serves as validation of our mechanism. First, household-level expenditure is nearly as volatile as income. Second, household-level expenditure is relatively uncorrelated with income. Third, a large share of low-wealth households have MPCs near zero. Fourth, lagged high expenditure is associated with low contemporaneous spending propensities.

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

The key implication of our theory is that a large share of high-debt households are saving-constrained rather than credit-constrained. That is, higher saving (lower consumption) is associated with large costs, just as “credit-constrained” households can access credit only at exorbitantly high costs. As a result, income transfers to low-income, high-debt households are less expansionary in the short term (a few months) than previous models of incomplete

markets with heterogeneous agents would predict, because savings-constrained households already spend too much.

Our theory also helps explain otherwise puzzling consumer behavior documented by recent empirical work. In particular, [Ganong and Noel \(2019\)](#) show that, among low-wealth households receiving unemployment insurance (UI) benefits, spending drops precipitously upon the predictable expiration of UI benefits. Households in our model exhibit similar behavior if their income is low and consumption is at or below the consumption threshold. In that case, a decline in income (even if anticipated) leads households to consume below the threshold and pay the utility cost.

Our theory has important implications for the propagation of macroeconomic shocks. In [Miranda-Pinto et al. \(2019\)](#) we show that consumption thresholds are important for understanding the cross-country relationship between fiscal effects on interest rates and consumer debt. In particular, higher debt (during periods of normal-to-loose credit supply) 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 debt. A general equilibrium extension of our model with consumption thresholds shows that the interest rate response to fiscal stimulus depends on consumer debt as in the data. It is likely that our theory also has implications for the effects of other shocks, particularly those that operate through interest rates (like monetary policy, bankruptcy, and macroprudential policy). With these applications in mind, we suggest that researchers should view our expenditure shocks as an important add-on to the basic model with income shocks.

References

- Anderson, T.W. and Cheng Hsiao (1981). “Estimation of dynamic models with error components.” *Journal of the American Statistical Association*, 76, pp. 589–606. doi:[10.1080/01621459.1981.10477691](https://doi.org/10.1080/01621459.1981.10477691).
- Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni (2014). “Estimates of annual consumption expenditures and its major components in the PSID in comparison to the CEX.” *American Economic Review Papers and Proceedings*, 104(5), pp. 132–135. doi:[10.1257/aer.104.5.132](https://doi.org/10.1257/aer.104.5.132).
- Arellano, Manuel and Stephen Bond (1991). “Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations.” *Review of Economic Studies*, 58, pp. 277–297. doi:[10.2307/2297968](https://doi.org/10.2307/2297968).
- Athreya, Kartik, Xuan S. Tam, and Eric R. Young (2009). “Unsecured credit markets are not insurance markets.” *Journal of Monetary Economics*, 56, pp. 83–103. doi:[10.1016/j.jmoneco.2008.08.005](https://doi.org/10.1016/j.jmoneco.2008.08.005).
- Bernheim, B.Douglas, Jonathan Skinner, and Steven Weinberg (2001). “What accounts for the variation in retirement wealth among U.S. households?” *American Economic Review*, 91(4), pp. 832–857. doi:[10.1257/aer.91.4.832](https://doi.org/10.1257/aer.91.4.832).
- Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico (2018). “The consumption response to positive and negative income shocks.” *Journal of Monetary Economics*, 96, pp. 1–15. doi:[10.1016/j.jmoneco.2017.11.007](https://doi.org/10.1016/j.jmoneco.2017.11.007).
- Campbell, Jeffrey R. and Zvi Hercowitz (2018). “Liquidity constraints of the middle class.” *American Economic Journal: Economic Policy*, 11(3), pp. 130–155. doi:[10.1257/pol.20180070](https://doi.org/10.1257/pol.20180070).
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokunaka, and Matthew N. White (2017). “The distribution of wealth and the marginal propensity to consume.” *Quantitative Economics*, 8, pp. 977–1020. doi:[10.3982/QE694](https://doi.org/10.3982/QE694).
- Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull (2007). “A quantitative theory of unsecured consumer credit with risk of default.” *Econometrica*, 75(6), pp. 1525–1589. doi:[10.1111/j.1468-0262.2007.00806.x](https://doi.org/10.1111/j.1468-0262.2007.00806.x).
- Chetty, Raj and Adam Szeidl (2007). “Consumption commitments and risk preferences.” *Quarterly Journal of Economics*, 122(2), pp. 831–877. doi:[10.1162/qjec.122.2.831](https://doi.org/10.1162/qjec.122.2.831).
- Di Pillo, G., G. Liuzzi, Lucidi S., V. Piccialli, and F. Rinaldi (2016). “A DIRECT-type approach for derivative-free constrained global optimization.” *Computational Optimization and Applications*, 65(2), pp. 361–397. doi:[10.1007/s10589-016-9876-3](https://doi.org/10.1007/s10589-016-9876-3).
- Ganong, Peter and Pascal Noel (2019). “Consumer spending during unemployment: Positive and normative implications.” *American Economic Review*, 109(7), pp. 2383–2424. doi:[10.1257/aer.20170537](https://doi.org/10.1257/aer.20170537).
- Hubmer, Joachim, Per Krusell, and Jr. Anthony Smith (2019). “Sources of U.S. wealth inequality: Past, present, and future.” NBER Working Papers 23011, National Bureau of Economic Research.
- ISR (2019). *Panel Study of Income Dynamics*. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI, public use dataset edition.
- Jappelli, Tuillio and Luigi Pistaferri (2017). *The Economics of Consumption: Theory and Evidence*. Oxford University Press, New York. doi:[10.1093/acprof:oso/9780199383146.001.0001](https://doi.org/10.1093/acprof:oso/9780199383146.001.0001).
- Jappelli, Tullio and Luigi Pistaferri (2010). “The consumption response

- to income changes.” *Annual Review of Economics*, 2, pp. 479–506. doi:[10.1146/annurev.economics.050708.142933](https://doi.org/10.1146/annurev.economics.050708.142933).
- Jappelli, Tullio and Luigi Pistaferri (2014). “Fiscal policy and MPC heterogeneity.” *American Economic Journal: Macroeconomics*, 6(4), pp. 107–136. doi:[10.1257/mac.6.4.107](https://doi.org/10.1257/mac.6.4.107).
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles (2006). “Household expenditure and the income tax rebates of 2001.” *American Economic Review*, 96(5), pp. 1589–1610. doi:[10.1257/aer.96.5.1589](https://doi.org/10.1257/aer.96.5.1589).
- Kaplan, Greg and Giovanni L. Violante (2014). “A model of the consumption response to fiscal stimulus payments.” *Econometrica*, 82(4), pp. 1199–1239. doi:[10.3982/ECTA10528](https://doi.org/10.3982/ECTA10528).
- Krusell, Per and Anthony A. Smith (1998). “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106(5), pp. 867–896. doi:[10.1086/250034](https://doi.org/10.1086/250034).
- Kueng, Lorenz (2018). “Excess sensitivity of high-income consumers.” *Quarterly Journal of Economics*, 133(4), pp. 1693–1751. doi:[10.1093/qje/qjy014](https://doi.org/10.1093/qje/qjy014).
- Lewis, Daniel, Davide Melcangi, and Laura Pilossoph (2019). “Latent heterogeneity in the marginal propensity to consume.” *Working Paper*. URL <https://ideas.repec.org/p/fip/fednsr/902.html>.
- Livshits, Igor, James MacGee, and Michèle Tertilt (2007). “Consumer bankruptcy: A fresh start.” *American Economic Review*, 97(1), pp. 402–418. doi:[10.1257/aer.97.1.402](https://doi.org/10.1257/aer.97.1.402).
- Luo, Yulei, Jun Nie, Gaowang Wang, and Eric R. Young (2017). “Rational inattention and the dynamics of consumption and wealth in general equilibrium.” *Journal of Economic Theory*, 172, pp. 55–87. doi:[10.1016/j.jet.2017.08.005](https://doi.org/10.1016/j.jet.2017.08.005).
- Miranda-Pinto, Jorge, Daniel Murphy, Kieran Walsh, and Eric Young (2019). “Saving constraints, debt, and the credit market response to fiscal stimulus.” *Mimeo University of Virginia*. URL https://faculty.darden.virginia.edu/murphyd/documents/MMWY_IRRF.pdf.
- Misra, Kanishka and Paolo Surico (2014). “Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs.” *American Economic Journal: Macroeconomics*, 6(4), pp. 84–106. doi:[10.1257/mac.6.4.84](https://doi.org/10.1257/mac.6.4.84).
- Nickell, Stephen (1981). “Biases in dynamic models with fixed effects.” *Econometrica*, 49(6), pp. 1417–1426. doi:[10.2307/1911408](https://doi.org/10.2307/1911408).
- Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod (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. doi:[10.3386/w21220](https://doi.org/10.3386/w21220).
- Shapiro, Matthew D. and Joel Slemrod (2003). “Consumer response to tax rebates.” *American Economic Review*, 93(1), pp. 381–396. doi:[10.1257/000282803321455368](https://doi.org/10.1257/000282803321455368).
- Shea, John (1995). “Union contracts and the life-cycle/permanent-income hypothesis.” *American Economic Review*, 85(1), pp. 186–200. URL <https://www.jstor.org/stable/2118003>.
- Young, Eric R. (2010). “Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm and non-stochastic simulations.” *Journal of Economic Dynamics and Control*, 34, pp. 36–41. doi:[10.1016/j.jedc.2008.11.010](https://doi.org/10.1016/j.jedc.2008.11.010).

A Computational Appendix (For Online Publication)

The recursive problem for the household can be written as

$$v(k, z, x, \underline{c}) = \max_{k', c} \left\{ \log(c) - \lambda \max\{\underline{c} - c, 0\} + \beta E \left[v(k', z', x', \underline{c}') \right] \right\}$$

subject to the budget constraint

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

the borrowing constraint

$$k' \geq b,$$

and the processes for z , x , and \underline{c} . Assume each of these variables follows an AR(1):

$$q' = (1 - \rho_q) \mu_q + \rho_q q + \sigma_q \epsilon'_q$$

with ϵ a standard normal, for $q \in \{z, x, \underline{c}\}$. We then approximate each process as a Markov chain using Rouwenhorst's method with (3,3,7) states, respectively.

We iterate on the Bellman equation to solve the recursive problem, using Howard's improvement steps. We approximate v using piecewise-cubic Hermite polynomials in k over an irregularly-spaced grid on $[b, \bar{k}]$ that contains 0 and solve the maximization using feasible sequential quadratic programming (FSQP). Since the kink in the objective function slows down SQP methods substantially (they rely on local quadratic approximations that are inaccurate around the kink), we use a two-part procedure – we first ignore the $\lambda \max\{\underline{c} - c, 0\}$ term and compute the optimal decisions, then if the optimal c satisfies $c < \underline{c}$ we impose \underline{c} as an upper bound in FSQP and resolve.

To compute the distribution, we use the method from [Young \(2010\)](#): we linearly interpolate the decisions onto a dense evenly spaced grid and use histograms to approximate the distribution of k given (z, x, \underline{c}) . For model-based moments in [Tables 1, 6, 7, and 8](#), including the regression coefficients, we

use a stochastic simulation constructed with 4500 households tracked for 1600 quarters (200 biennial observations). The initial conditions for this sample are drawn from the stationary distribution, and we impose the condition that the value for x does not change over the sample. This is so that the simulations exhibit household fixed effects, a property of the PSID, even over long time horizons.

For studying anticipated income changes, we solve the sequential version of the household problem for $t \in \{1, \dots, T\}$ and anticipated labor income tax $\{\tau_t\}$:

$$v_t(k, z, x, \underline{c}) = \max_{k' \geq b, c} \left\{ \log(c) - \lambda \max\{\underline{c} - c, 0\} + \beta E \left[v_{t+1} \left(k', z', x', \underline{c}' \right) \right] \right\}$$

subject to

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

where

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

subject to

$$c + k' \leq (1 + r + \phi \mathbf{1}(k \leq 0))k + w \exp(z) \exp(x) \bar{h} (1 - \tau_T).$$

We assume $\tau_t = 0$ for $t \leq 4$ and $t \geq 15$ and set $\tau_t = .25$ for $t = 5, \dots, 14$. We set T very large so that the wealth dynamics have converged well before the horizon ends.

To estimate the model, we use DiRDFN, which is based on the DiRect (divided rectangles) optimizer and includes general constraints and local derivative-free searches (see [Di Pillo et al. \(2016\)](#)). The DiRect algorithm takes a feasible space of parameters (a hyperrectangle) and subdivides it iteratively, and then DiRDFN adds derivative-free local searches with active set methods to handle the constraints. The algorithm is globally convergent to the global minimum, but since the bounds matter (even if they end up not binding, they can affect

the search process if the global solution lies outside them) we check that the solution does not change if the bounds are increased. The parameters we estimate are $(\beta, \mu_c, \rho_c, \sigma_c, \lambda)$. For the standard Bewley model, we use a nonlinear root-finder (Brent's method) to find the β that matches the wealth-income ratio target.