A Model of Expenditure Shocks*

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

[1]International Monetary Fund and University of Queensland
[2]University of Virginia Darden School of Business
[3]ETH Zürich
[4]University of Virginia and Federal Reserve Bank of Cleveland

July 2023

Abstract

We introduce a new quantitative model of household expenditure shocks to rationalize the common anecdote of a low-income and low-liquidity household that uses additional income to save (repay debt) rather than consume. Our model also rationalizes key features of the joint dynamics of household-level consumption and income, including our finding that consumption is volatile yet disconnected from income. The key feature of our model is stochastic consumption thresholds that yield large utility costs if violated. The stochastic thresholds increase the welfare cost of income fluctuations by an order of magnitude.

JEL Codes: D14, E21

Keywords: Consumption, Expenditure Shocks, Household Debt

*This paper was previously circulated under the title "Saving-Constrained Households." We would like to thank Douglas Diamond, Peter Rupert, and seminar and workshop participants at 2019 LAEF/WAMS, 2019 SED, CERGE-EI, ETH Zürich, FRB Cleveland, FRB Richmond, Indiana University, Notre Dame, Reserve Bank of Australia, Texas A&M, UC Berkeley, UC Davis, UC Santa Cruz, University of New South Wales, University of Queensland, and University of Zürich for helpful comments. The opinions expressed are those of the authors and do not necessarily represent views of the Federal Reserve Bank of Cleveland, the Federal Reserve System, or the IMF.

†Email: jmirandap@bcentral.cl

‡Email: murphyd@darden.virginia.edu

§Email: kwalsh@ethz.ch

¶Email: ey2d@virginia.edu

1 Introduction

A common anecdote referenced by policymakers and politicians is of a 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. According to recent analysis of account-level data from JP Morgan Chase, these anectotal households are pervasive: consumption is extremely volatile, and the median household does not have sufficient liquid assets to buffer typical expenditure fluctuations (Farrell and Greig, 2015, 2017; Farrell et al., 2019). Debt for these households is burdensome in the sense that additional income is allocated to debt service rather than additional consumption. The poorest of these households are often considered especially vulnerable because they may forego medical care, food, or other basic necessities in order to service this unwanted debt burden.

In this paper, we formalize this story in a quantitative model that replicates key features of the joint dynamics of household-level consumption and income. The seminal work of Blundell et al. (2008) and its many successors combine panel data on food expenditure from the pre-1999 Panel Study of Income Dynamics (PSID) with the Consumer Expenditure Survey (CEX) to impute household consumption and use a standard consumption/income model to identify the response of consumption to transitory and permanent income shocks. Since 1999, the PSID has expanded its coverage to include comprehensive measures of consumption, thus permitting an analysis of the joint dynamics of consumption and income based on a broad measure of total consumption.[1] We exploit this comprehensive consumption data in the PSID to establish that household-level consumption is volatile yet disconnected from income, and we propose a quantitative expenditure shock model that explains this disconnect along with a variety of related time series and cross-sectional features of the microdata. Furthermore, the expenditure shock model rationalizes the behavior of the anecdotal low-income household that uses additional resources to repay debt or save rather than to consume. It also rationalizes recent evidence on excess sensitivity to anticipated income declines and marginal propensities to consume that are increasing in income (over some range).

We begin with four facts from the PSID on the joint dynamics of income and consumption. First, for the average household in the PSID, consumption is as volatile as income. In contrast, simple Permanent Income Hypothesis (PIH) models imply consumption should be much smoother than income (Fact 1). Two potential remedies are pervasive liquidity constraints (e.g., Kaplan and Violante, 2014) or very persistent income shocks, both of

[1] See, for example, Blundell et al. (2016), Blundell et al. (2018), and Commault (2022).

which closely tie consumption to current income.

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

Fact 3 is that household consumption growth is negatively autocorrelated. In a standard setting with households that follow the PIH, consumption is a random walk and consumption growth is independent of its lag, in contrast with the strong negative autocorrelation we estimate. As explained by Blundell et al. (2008), negatively autocorrelated consumption growth is indicative of shifters of observed consumption unrelated to income, such as measurement error or preference shocks.[3] And as Dynan (2000) shows, Fact 3 is also inconsistent with theories of habit in consumption, which induces positive autocorrelation in consumption growth.

Fact 4 is that the cross-sectional correlation between consumption growth and income growth is far smaller among households experiencing high consumption (relative to the within-household average) than among the full sample of households. So while Fact 2 establishes that consumption is relatively disconnected from income (within households), Fact 4 is that consumption is even less connected to income (across households) among households for which consumption levels are high. As we explain below, in the standard model with or without measurement error, on average the relative level of consumption barely matters for the relationship between income and consumption growth.

We propose a joint explanation of these facts based on a theory in which households face stochastic time-varying consumption thresholds that, if violated, yield substantial utility costs. These consumption thresholds represent unanticipated shocks such as medical emergencies, auto repairs, or even expenses associated with attending a wedding or funeral.[4] When an adverse shock hits a household, it chooses to accumulate debt (reduce wealth)

[2]Given that consumption is smoother than income in aggregate data, it is perhaps surprising that consumption growth is volatile yet disconnected from income at the household level. Although this pattern has not been emphasized in the literature, it is apparent in Farrell and Greig (2015), and the estimates in Tables 3 and 4 in Blundell et al. (2008) imply a ratio between consumption volatility and income volatility greater than one.

[3]To see this, suppose latent log consumption c_t follows a random walk: $c_{t+1} = c_t + \varepsilon_{t+1}$, where ε_{t+1} is mean zero i.i.d. We only observe, however, measured consumption: $c^*_{t+1} = c_{t+1} + u_{t+1}$, where u_{t+1} is mean zero i.i.d. It is natural to think of u_{t+1} as measurement error and c_{t+1} as actual consumption, but you can also interpret c_{t+1} as a frictionless tendency/benchmark and u_{t+1} as a consumption shifter. Since $\Delta c^*_{t+1} = \varepsilon_{t+1} + u_{t+1} - u_t$, it follows that $cov(\Delta c^*_{t+1}, \Delta c^*_t) = -var(u_t)$.

[4]Farrell and Greig (2017) emphasize that "extraordinary payments" from medical bills, auto repairs, and taxes are pervasive in their sample and suggest these "types of expenses that have a higher likelihood of being unexpected in timing or magnitude and are thus potentially more difficult to weather."

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 or dissave when faced with a large, unanticipated expenditure. Maintaining a low net asset position is optimal for households against their threshold consumption level, but it is costly in the sense that, in the event of another adverse shock, households may be forced to consume below the threshold level (e.g., move out of a home or forgo medical care), which is associated with a large utility cost. Therefore, households for which consumption is against the threshold use additional income to pay off debt (increase net assets) as a precautionary measure.

The average *level* of the consumption threshold represents, for example, consumption commitments in Chetty and Szeidl (2007) – goods and services that are infrequently adjusted and are subject to transaction costs. One can think of these goods as subject to stochastic maintenance costs, which are represented in our model by *shocks* to the consumption threshold. For example, the household commits to buy an automobile or take the children to a private school (average threshold) but also commits to cover the implied expenses of unexpected car repairs or school trips (shocks to the threshold). Even medical emergencies can be considered a form of maintenance, maintenance of a person's physical body. When a broad range of consumption is committed (Chetty and Szeidl, 2007), a health expenditure shock cannot be fully offset by reducing expenditure on other components of consumption, and hence total household expenditure is effectively subject to a minimum threshold.[5]

We refer to households against their consumption threshold 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.[6] 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.

To evaluate the ability of consumption thresholds to explain our facts, we develop a heterogeneous-agent model featuring stochastic consumption thresholds that are persistent. 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

[5]The consumption thresholds can also be interpreted as major expenses that Campbell and Hercowitz (2019) have documented are important for understanding household saving behavior.

[6]To 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: $k' - k \leq y - \underline{c}$.

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 the data. The estimation uncovers a consumption threshold process that is persistent, highly volatile, and (nearly) mean zero.[7] 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 the thresholds would deliver the large fit improvement we observe.

An alternative potential explanation for our facts is error in the measurement of household expenditure. Idiosyncratic and independently distributed measurement error, for example, would imply observed expenditure that is volatile and relatively uncorrelated with income. We evaluate the plausibility of these facts being driven by measurement error by incorporating potentially autocorrelated measurement error in the standard Bewley model and calibrating the parameters of the measurement error process to fit the data. The measurement error process necessary to match the targets is extremely volatile, accounting for nearly all of the variation of observed expenditure. Using the "noise" statistic from Alan and Browning (2010), in our calibrated measurement error model 89% of observed consumption growth is due to noise. This is above the 86% PSID estimate of Alan and Browning (2010), which the authors note is "somewhat higher" than in previous papers, and well above the 75% used in Alan et al. (2009), which they say is an "upper end" number for the literature. The implied measurement error process is also autocorrelated. These results alone raise doubt for a pure measurement error explanation of our facts. Furthermore, although both the expenditure shocks and measurement error models provide a strong fit to Facts 1 through 3, only the expenditure shock model generates Fact 4 (lower cross-sectional correlation between consumption and income growth conditional on high consumption). The superior fit of the expenditure shock model arises from asymmetry in the effects of the consumption threshold process on actual consumption. When the consumption threshold is lower than the (unconstrained) optimal consumption level, household consumption behavior resembles that of a standard PIH model. But high levels of consumption are driven by binding consumption thresholds, which delink consumption from income. In contrast, in the standard model with measurement error, observed consumption is relatively independent of income,

[7] We 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.

regardless of whether consumption is high or not.

Aguiar et al. (2020) also study the recent PSID and explain a variety of its puzzling features with heterogeneous preferences, specifically various combinations of the discount factor and elasticity of intertemporal substitution (EIS). In Appendix B, we solve a model based on their calibration and show that such preference heterogeneity does not help explain our four facts. While low discount factor types have high consumption volatility relative to income volatility, they are only around 16% of the population, so the model does not match Fact 1 on average. Furthermore, the heterogeneous preference model is not close on the other facts for any type.

In addition to explaining Facts 1 through 4, our model helps to explain a number of otherwise puzzling features of consumer behavior documented by recent empirical work. First, the expenditure shock model implies that a fraction of medium-to-low-income households are against their consumption threshold and hence exhibit marginal propensities to consume (MPCs) out of additional income of zero. As a consequence, the model-implied MPC distribution (by wealth or income) is U-shaped rather than monotonically decreasing, as implied by standard models in which low-income households are credit-constrained but not saving-constrained.[8] Our model-implied nonmonotonic relationship is consistent with a number of studies indicating that many low-to-medium-wealth households use additional income to pay down debt rather than spend. For example, Shapiro and Slemrod (2003) find that low-income and low-wealth households are more likely to save out of tax cuts. More recently, Kueng (2018) and Lewis et al. (2022) document stronger spending propensities out of new income among high-income households than low-income households, and Koşar et al. (2023) document that marginal propensities to repay debt out of transfer payments are decreasing in liquid wealth-to-income (and, relatedly, that marginal propensities to consume are increasing in wealth-to-income).[9]

Similarly puzzling is the evidence in Ganong and Noel (2019) 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 an-

[8]The model of Campbell and Hercowitz (2019) also generates the U-shape with respect to wealth. See Jeon and Walsh (2023) and Pistaferri and Saporta-Eksten (2012) for further discussion of the empirical evidence on the possible U-shape.

[9]The prevalence of medium-to-low income households with low MPCs 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. Relatedly, 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. We provide a more detailed discussion of the evidence of low-income households with low MPCs in Section 3.3.

ticipated 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.[10]

We also investigate the extent to which consumption threshold shocks amplify the welfare cost of income fluctuations. While in our standard Bewley model, the gain from shutting down labor income fluctuations is equivalent to a 2.84% permanent increase in consumption, with expenditure shocks the gain rises to 37.03%. Consumption thresholds make income fluctuations much more costly because they constrain saving and limit the ability of households to smooth out negative income shocks.

In summary, our framework is capable of explaining key moments of the household-level relationship between consumption and income (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.[11]

For comparison, we note here that alternative models (including heterogeneous preferences, which we explicitly addressed) 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., 2021). While these shocks can help the model address Fact 1, it is unclear if they will reconcile Facts 2 through 4, and they will not help explain the presence of moderately low-income households with low MPCs – in the stationary distribution low-income households will have low discount factors and therefore high MPCs. Similarly, idiosyncratic variation in returns (Hubmer et al., 2021)

[10] Our 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 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.

[11] Our framework shares with standard Bewley models the prediction that average MPCs in response to unanticipated transitory income shocks are low (< 0.2) compared to 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.

will deliver the same prediction, as poor households will have low returns.[12]

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 presents welfare calculations. Section 7 concludes.

2 The Joint Dynamics of Household-Level Income and Consumption

Here we highlight pieces of evidence on household-level income and consumption that, when taken together, are difficult to reconcile with existing theories of heterogeneous households with uninsurable idiosyncratic income risk. The relevant facts are the following: (1) Consumption growth is as volatile as income growth for the average household. (2) Household-level consumption growth is relatively uncorrelated with income growth. (3) Consumption growth is negatively autocorrelated. (4) The cross-sectional correlation between consumption and income growth is lower among households experiencing high consumption compared to their within-household average.

Our analysis of these facts 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 dataset 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 2005 the PSID added information on expenditure on clothing, travel, other recreation, telephone, internet, household repairs, and household furnishing. With these additions, the PSID captures nearly all of the expenditure categories measured by the CEX, which collects data at a higher frequency but does not

[12]There are other models of household expenditure shocks, Chatterjee et al. (2007) and Livshits et al. (2007) for example, 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.

have the longitudinal dimension that characterizes the PSID.[13] To maintain a consistent measure of expenditure across the panel, we exclude the categories added in 2005 from our analysis, although the three facts are apparent under alternative measures that incorporate the added expenditure categories (see Appendix Table A1). 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. We restrict our sample to respondents that are in each wave of the PSID from 1999 through 2017. Our resulting sample consists of 6,159 households.

2.1 Income Calibration for Comparison to Bewley Model

Below we present moments on the joint dynamics of income and consumption from the PSID. To benchmark these moments, we simulate data from a calibrated Bewley model. To calibrate the model, we simulate a quarterly labor income process, choosing the parameters such that when the model is aggregated to a biennial frequency and truncated to the number of waves in our PSID sample, the estimated process from the simulated data matches key moments from the PSID. Specifically, as in Storesletten et al. (2004), we assume that log labor income for each household i consists of a highly persistent component and a transitory component:

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

[13] As discussed by Andreski et al. (2014), the consumption data in the PSID closely correspond to that 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. $\mathcal{N}(0,1)$.[14] In the PSID, however, we observe only log annual income ($a_{i,\tau}$, $\tau = 1,2,3,\ldots$) sampled biennially. In the model, this corresponds to

$$
\begin{aligned}
a_{i,1} &= \log \sum_{t=1}^{4} \exp(y_{i,t}) \\
a_{i,2} &= \log \sum_{t=9}^{12} \exp(y_{i,t}) \\
a_{i,3} &= \log \sum_{t=17}^{20} \exp(y_{i,t}) \\
&\vdots
\end{aligned}
$$

In the PSID, we form labor income by subtracting capital income from total income. Our income process estimates are based on data from 2005 through 2017, since the income subcategory capital income (which we use to derive labor income) begins in 2005.[15] An alternative approach would be to use reported measures of wage income. We choose to focus on total-net-capital income, both because wage income is missing for 15% of households in our sample and because our measure covers other forms of labor income including self-employment income. Regardless, our estimates of the labor income process are relatively similar using either measure of labor income.

With labor income in hand, 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 $var(FI_i) \approx 1.06^2$, $\rho_a \approx 0.05$, and $\sigma_a \approx 0.96$. These moments – cross-sectional dispersion in fixed effects, autocorrelation, and residual standard deviation – are the moments we match. Setting $\sigma_x = 0.15$, $\sigma_z = 0.74$, and $\rho_z = 0.78$, when we run the fixed-effects panel regression on model-simulated $a_{i,\tau}$ for $\tau \in [1,...,7]$ (after burning 1000 quarters) and 50,000 households, the resulting values for $var(FI_i)$, ρ_a, and σ_a match what we see in the PSID.[16]

We then use a discretized version of the calibrated income process to compute a standard Bewley model (one asset, heterogeneous agents, and uninsurable idiosyncratic income risk).

[14] Our assumption of near-permanent income ($\rho_x = .99$) follows Carroll et al. (2017), who assume $\rho_x = 1$, and Krueger et al. (2016), who estimate an annual $\rho_x = .97$.

[15] Guvenen and Smith (2014) construct PSID labor income in a similar fashion.

[16] The 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 and similar to estimates based on standard methods that address the bias (Anderson and Hsiao, 1981; Arellano and Bond, 1991). In calibrating the model we apply the same OLS estimator used in the data to the simulated data, which has the same number of time periods but many more households.

The specifics of the model are outlined in Section 3 below. Our labor income process is convenient for our later estimations and simulations because while it is stationary and ex ante identical across agents, the highly persistent component $x_{i,t}$ resembles fixed effects (a feature of the PSID) with only 10 time periods (corresponding to 1999, 2001, ..., 2017 in the PSID).

While our income process matches the post-1999 panel data, it is more volatile and less persistent than others used in the literature (e.g., Aguiar et al., 2020; Karahan and Ozkan, 2013) that are based on pre-1999 data. In Appendix B we use the off-the-shelf income process from Krueger et al. (2016) and find that the discrepancy between the PSID data (discussed below) and predictions of the standard model are not specific to our particular income process.

2.2 Facts 1 and 2: Consumption is Volatile and Relatively Independent of Income

Table 1 presents key moments from the joint dynamics of total 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 over the same number of periods as covered in our PSID sample.

Panel A reports moments based on within-household joint movements of consumption growth and income growth. For each household, we compute the standard deviations of the change in log consumption ($d\log C$) and the change in log income ($d\log I$). We compute the within-household ratio of these standard deviations and report the cross-sectional summary average of $\frac{sd(d\log C)}{sd(d\log I)}$ in the first row of Panel A. We also compute the within-household correlation between the change in log consumption and the change in log income. The cross-sectional average of $corr(d\log C, d\log I)$ is reported in the second row of Panel A.

It is readily apparent that consumption is more volatile (relative to income volatility) in the data than in the Bewley model (Fact 1). The average of $\frac{sd(d\log C)}{sd(d\log I)} = 1.05$ (standard error .09) far exceeds the corresponding moment from the Bewley model. While not emphasized in many previous papers, this fact is consistent with other studies of the joint dynamics of income and consumption. In particular, Blundell et al. (2008) document for the pre-1999 PSID, that consumption growth volatility is as large income growth volatility (Tables 3 and 4). More recently, Cho et al. (2022) use the post-1999 PSID and show that shocks to consumption growth are relatively more volatile than shocks to income growth, especially for the period post-2007. Farrell and Greig (2015) analyze proprietary JP Morgan account data over 2013 to 2014 and conclude that "individuals experienced high levels of income volatility and higher levels of consumption volatility across the income spectrum.” In a follow-up paper that covers 2013-2018 but uses a different measure of volatility, Farrell and Greig (2017) find that consumption is less than volatile than income, but not by much (the median coefficient of variation for spending is 0.33 vs. 0.38 for income).

One possible explanation for volatile consumption in the data is that many 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 0.64, in the data it is only 0.23 (standard error 0.02).

It is important to note that household 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. This evidence is consistent with the analyses in Farrell and Greig (2015, 2017) and Farrell et al. (2019), who document a pervasive disconnect between consumption and income across the income distribution, despite the high relative volatility of consumption.

Nor are the summary statistics from the PSID in Table 1 driven by durable goods. In particular, the relative volatility of nondurable expenditure is 1.00 (slightly lower than for total expenditure), and the average correlation is 0.24 (see Appendix Table A1 for summary statistics based on various definitions of expenditure). Therefore, durability alone cannot account for consumption that is volatile and relatively independent of income.

Panel B reports estimates of the coefficients of autoregressions of log consumption (with household fixed effects - first row, and without household fixed effects - second row). Not only is consumption far less correlated with income than predicted by a standard model, it is also less persistent. We will not refer to these autoregressions as independent facts, but we will use them to estimate the model developed in Section 3.

2.3 Fact 3: Consumption Growth is Negatively Autocorrelated

While the moments for our first set of facts provide intuitive comparisons to predictions from the PIH, the moment for our third fact (autocorrelation of consumption *growth*) is perhaps less intuitive. As discussed in Blundell et al. (2008), the autocorrelation of consumption growth is indicative of consumption drivers that do not appear in a traditional Euler equation. In particular, negative autocorrelation is inconsistent with the standard PIH and, as highlighted by Dynan (2000), with models of habit formation. Consider an

Table 1
Consumption and Income Moments

		PSID	Bewley
		(1)	(2)
Panel A: Average across households			
Fact 1	$sd(d\log C)/sd(d\log I)$	1.05	0.29
		(0.09)	
Fact 2	$corr(d\log C, d\log I)$	0.23	0.64
		(0.02)	
Panel B: Panel autoregression coefficients			
	AR coefficient ($\log C$), FE	0.21	0.67
		(0.02)	
	AR coefficient ($\log C$), pooled	0.67	0.94
		(0.01)	
	AR coefficient (C growth), FE	-0.38	-0.078
		(0.02)	
Fact 3	AR coefficient (C growth), pooled	-0.36	0.019
		(0.02)	
Panel C: Average across years			
	Cross-sectional $corr(d\log C, d\log I)$	0.21	0.65
		(0.02)	
	Cross-sectional conditional $corr(d\log C, d\log I)$	0.073	0.67
		(0.03)	
Fact 4	ratio	0.35	1.03
		(0.12)	

Note: $sd(d\log C)/sd(d\log I)$ and $corr(d\log C, d\log I)$ in Panel A are the cross-sectional means of the listed statistics, which are first calculated at the household level. The model-based moments are calculated from a quarterly simulated panel of 20,000 households aggregated to ten biennial periods. Autoregressive (AR) coefficients in Panel B are calculated via regression with (FE) or without (pooled) household fixed effects. Time fixed effects are included in all PSID-based autoregressions. Estimation procedures that account for possible Nickell bias yield slightly larger estimates. The PSID statistics are nearly identical when focusing exclusively on nondurable goods. Panel C reports the average (across years) of cross-sectional correlations. Cross sectional conditional $corr(d\log C, d\log I)$ limits the sample to households experiencing high consumption, defined as household consumption exceeding the within-household average by 1.5 standard deviations. PSID standard errors are in parentheses. Fact 3 standard errors are from OLS, clustered at the household level. The standard errors for Facts 1, 2, and 4 are based on the stationary bootstrap of Politis and Romano (1994). We generate each of 10,000 bootstrapped panels by randomly selecting households with replacement and sampling common time periods for income growth, consumption growth, and $HighC$ by creating random blocks with tuning parameter $p = 1/3$, meaning the average block length is 3. See Politis and Romano (1994) for details.

increase in consumption. Under the PIH, one would expect this increase to be driven by a shock to permanent income such that consumption growth jumps and then returns to trend. Habit formation, as Dynan (2000) explains, implies a gradual adjustment of consumption and hence positively autocorrelated consumption growth. But if observed consumption instead jumps temporarily (for unclear reasons relative to the PIH), then consumption growth will be negatively autocorrelated (see Footnote 3).

The last row of Panel B reports the autoregression coefficients for consumption growth. The autocorrelation is strongly negative in the PSID (Fact 3), whereas it is effectively zero in the Bewley model. This negative autocorrelation implies drivers of consumption that are independent of income, consistent with Facts 1 and 2.

In principle Facts 1 through 3 could arise from measurement error in consumption. Intuitively, measurement error would produce observed consumption that is volatile and disconnected from income. Below we formally evaluate this possibility and propose an alternative explanation for these facts.

2.4 Fact 4: Cross-Sectional Correlations between Consumption Growth and Income Growth

Our fourth fact is based on cross-sectional correlations of consumption growth and income growth among households experiencing episodes of high consumption. We define high consumption as a household having consumption in a period that exceeds the within-household average by 1.5 standard deviations.[17] Specifically, $HighC_{it} = 1$ if $C_{it} >$Mean(C_{it}) + 1.5$\times$ SD(C_{it}).

The reason for examining cross-sectional correlations by episodes of high consumption is based on intuition embedded in our proposed model of expenditure shocks. We defer a full explanation until we present the model in Section 3. Briefly, our model suggests that consumption thresholds are only binding when they are sufficiently positive. For low realizations of the consumption threshold, households are unconstrained and consumption is tied to income as in the standard model. But for high realizations of the threshold, the threshold binds and consumption is delinked from contemporaneous income. We choose to define high-consumption episodes using a threshold of 1.5 standard deviations based on the fact that in our calibrated model of expenditure shocks presented below, it roughly maximizes the share of households experiencing episodes for which consumption is delinked from income due to binding minimum consumption thresholds (see Table 5).

[17] We choose 1.5 standard deviations based on the fact that in our calibrated model of expenditure shocks presented below, it roughly maximizes the correlation between high-consumption episodes and binding minimum consumption thresholds (see Table 5).

To explore state-dependence in the relationship between consumption and income, we examine cross-sectional correlations between consumption growth and income growth, both for the full sample of households and for the subset that are experiencing high consumption episodes. In particular, we compute cross sectional correlations for each year in the sample. In Panel C of Table 1 we report the average correlation across years. The first row of Panel C reports the average correlation based on the full sample of households. The second row reports the same moment based on restricting the sample each year to households that experience a high consumption episode in that year. The last row reports the ratio between these correlations.

It is immediately apparent that in the PSID, consumption and income are relatively delinked among the households experiencing high expenditure episodes, whereas in the Bewley model the relationship between consumption and income is stable across sample restrictions.

Decomposing High Expenditure Episodes. Which categories of expenditure drive high expenditure (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 2 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.[18] Expenditure shares during episodes (column (2)) are generally similar to average expenditure shares (column (1)). Two categories are noticeably more prevalent during episodes: education and transportation. Food and housing are less prevalent during episodes.

Columns (3) and (4) offer an alternative approach to examining the relevance of different expenditure categories in driving episodes. Here, we identify episodes for each category and examine the extent to which category-specific episodes predict aggregate expenditure episodes. For example, the dummy for a transportation episode is set to unity whenever a household's transportation expenditure exceeds its within-household average by a standard deviation. We then regress the dummy for an episode on indicators for the category-specific episodes. The pattern that emerges from both OLS (column (3)) and probit (column (4)) models is that high sub-category expenditure is associated with a high-expenditure episode for each category of expenditure. Furthermore, the likelihood that any given category-

[18] To 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 category during episodes. We then do the same for total expenditure and take the ratio of the two.

Table 2
Decomposing High Expenditure Episodes

	Share of Total Expenditure (1)	Ratio of Category Expenditure Relative to Total Expenditure during High Expenditure Episodes (2)	Coefficient from Linear Probability Model (3)	Coefficient from Probit Model (4)
Food	0.24	0.13	0.15	0.96
Housing	0.38	0.19	0.22	1.27
Transportation	0.26	0.49	0.44	1.99
Education	0.03	0.10	0.13	0.94
Child Care	0.01	0.02	0.06	0.48
Health	0.08	0.08	0.10	0.76

Note: This table presents statistics for the broad categories that make up the measure of total expenditure. Expenditure on Clothing, Trips, Other Recreation, Household Repairs, Household Furnishings, and Telephone/Internet are not included in the measure of total expenditure since they were only recorded beginning in 2005. In column (2), expenditure during episodes is relative to within-household averages. In columns (3) and (4), the depicted statistics are the coefficients from a regression of an indicator for a high expenditure episode on indicator variables for high sub-category expenditure. All regression coefficients are significant at the 1% level.

specific episode is associated with a total expenditure episode is broadly proportional to that category's share of total expenditure and follows a similar pattern of relevance that is depicted in column (2). These results motivate us to model consumption thresholds as applying to aggregate consumption rather than to specific goods, although the slightly stronger role of education and transportation is consistent with our motivating examples in the introduction. An additional benefit of modeling threshold shocks that apply to aggregate consumption is to economize on state variables, which facilitates incorporating our proposed model extension (below) into large-scale models of household behavior.

Before turning to the model, we note that the discrepancy between the standard consumption/saving model and our PSID facts is not driven by our particular income process. Our process matches the estimated autocorrelation, volatility, and dispersion in fixed effects in the post-1999 PSID, but it is nonetheless more volatile and less persistent than standard processes from the literature (based on earlier data). In Appendix B, we instead use the off-the-shelf process from Krueger et al. (2016) and solve a Bewley model with Epstein-Zin preferences for different values of the EIS and discount factor. With a low quarterly discount factor of .9274, high EIS of 1.5, and risk aversion equal to 4, we can get the relative volatility of consumption to 0.95, which is within a standard error of the PSID estimate. But with a higher discount factor or an EIS of 0.5, the relative volatility drops considerably. Moreover, no configuration of preferences comes close on the other facts: as with our version of the Bewley model, there is no mechanism to delink consumption and income, regardless of the income process.

3 A Theory of Expenditure Shocks

Here we present a theory of expenditure shocks 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 (and other moments). Specially, we use the evidence in Facts 1 and 2 as calibration targets. We do not target Facts 3 or 4, which serve as external validation. Separately, we introduce autocorrelated measurement error in expenditure into the standard Bewley model and calibrate the measurement error process to replicate Facts 1 and 2. We then examine the ability of the expenditure shock model and the Bewley model with measurement error to replicate Facts 3 and 4.

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

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

subject to the budget constraint

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

where V is the value function, C is consumption (the numeraire), k is capital wealth (which exogenously depreciates at rate $\delta \geq 0$), z is persistent idiosyncratic 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 $\exp(z+x)\,\overline{h}$ 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:

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

where ϵ_j, $j \in \{z, x, c\}$, is an idiosyncratic mean-zero shock with standard deviation σ_j.

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 distribution Ω^* over $(k, z, x, \underline{c})$, and household value and policy functions V^*, c^*, and k'^* such that (1) the value and policy functions solve the household problem given prices, (2) K^* and L^* solve the firm problem:

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

(3) the labor market clears: $L^* = \overline{H} \equiv \overline{h} \int (z+x)\, d\Omega^*$, and (4) Ω^* is generated by k'^* and the exogenous processes for $(z, x, \underline{c})$.

Define

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

to be a household's income, that is, labor income plus net capital income, and let $\overline{K} = \int k'^* d\Omega^*$ be steady-state aggregate household capital.

3.2 Calibration and Estimation

We assume a period is one quarter and calibrate the model with a two-step procedure. First, we choose the productivity process parameters ($\rho_z = 0.74$, $\sigma_z = 0.78$, $\rho_x = 0.99$, and $\sigma_x = 0.15$) 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: $\phi = 0.03$ ($\approx 12\%$ annual premium on borrowing vs. saving), $\alpha = 0.36$ (standard in literature), $\delta = 0.0125$ (standard in literature), $r = 0.0225$ (so the net return on saving is 1% per quarter), and $b = -1$ (about 1/3 of quarterly net labor and capital income, with the normalization $\overline{h} = 0.33$). In the second step, we choose the remaining parameters (discount rate, utility cost, and $\underline{c}$ process) to target the following moments in the ergodic distribution of the stationary equilibrium:[19] (1) $corr\,(d\log I, d\log C) = 0.23$, $\frac{std(d\log C)}{std(d\log I)} = 1.05$, and $corr\,(\log(C_t), \log(C_{t-1})) = 0.67$ computed at the biennial frequency (PSID moments reported in Table 1), (2) $\frac{\overline{K}}{\overline{K}^{\alpha}\overline{H}^{1-\alpha}} = 12$ (quarterly frequency), which is a standard capital/income target, and (3) $Fraction\,(k<0) = 0.1$ (quarterly frequency), which is in the ballpark of commonly used values in the consumer credit literature (e.g., Athreya et al., 2009) for the proportion of borrowers.

Via global optimization, the best–fit parameter values are $\beta = 1/1.039$, $\lambda = 24.394$, $\rho_c = 0.587$, $\sigma_c = 3.077$, and $\mu_c = 0.053$. Further computational details are in Appendix D. The baseline Bewley model is computed analogously but without the $\underline{c}$ process and with β re-estimated to match $\frac{\overline{K}}{\overline{K}^{\alpha}\overline{H}^{1-\alpha}} = 12$.

We also investigate the extent to which our empirical observations can be explained by measurement error in consumption. For this exercise, we add measurement error

$$m_t = \rho_m m_{t-1} + \sigma_m u_t,\ u_t \sim iidN(0,1)$$

to the log of biennial consumption in the calibrated Bewley model simulations. Fixing the other parameters (including the estimated β), we find the values of (ρ_m, σ_m) that most closely match the Table 1 PSID moments in simulations of the Bewley model with measurement error. We have also explored more flexible measurement error processes in which the variance depends on a household's current wealth (or its square), but doing so hardly improve the fit to the three considered moments, so here we present the simplest specification.

Calibration, estimation, and fit for the three models (Bewley, expenditure shocks, and Bewley with measurement error) are summarized in Table 3. Further computational details regarding the simulations are included in the table's caption.

Note that in estimating the models we calculate moments, the relatively volatility of consumption and income for example, for a single household. This is done for computational convenience, and the assumption we are making is that our best guess of the true moments corresponds to our estimates from the PSID, which has a low number of time periods ($T = 10$) relative to the number of agents ($N = 6,159$). An alternative estimation approach would be to calculate model moments with low T and a large cross-section of households. Since the agents are ex ante identical, for high enough T we would get the same moments, but for $T = 10$ you get something a little different because our targeted moments have a time series dimension. In any case, for our main results in Tables 1 and 4 we show model numbers

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

coming from simulations with $T = 10$ and large N, which is the best comparison with the PSID. This is why those numbers differ slightly from the calibration and estimation ones (Table 3).

3.3 Results

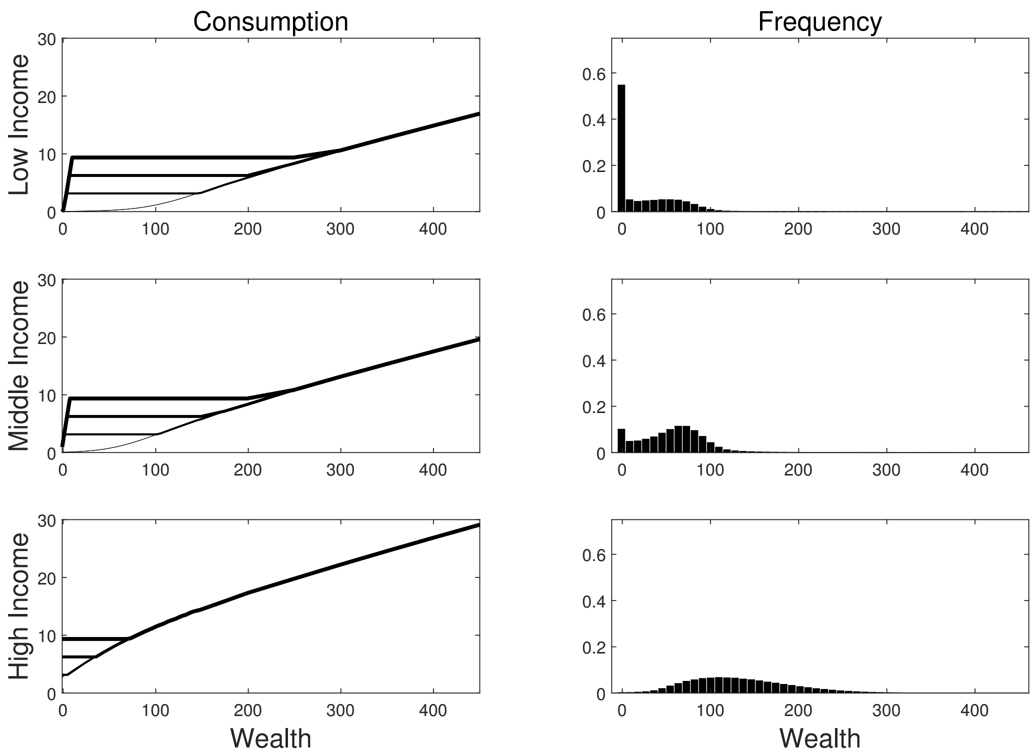

Figure 1

Note: The left column shows the expenditure shock 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 1 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.[20] However, not all low-wealth households with

[20]The 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, who face prohibitively large costs of accessing credit.

Table 3
Calibration and Estimation

		Bewley	Expenditure shock model	Bewley with meas. error	Notes
Calibrated Parameters	ρ_z	0.74	0.74	0.74	Inc. AR, pers.
	σ_z	0.78	0.78	0.78	Inc. vol., pers.
	ρ_x	0.99	0.99	0.99	Inc. AR, perm.
	σ_x	0.15	0.15	0.15	Inc. vol., perm
	ϕ	0.03	0.03	0.03	Borrowing cost
	α	0.36	0.36	0.36	Capital share
	δ	0.0125	0.0125	0.0125	Depreciation
	r	0.0225	0.0225	0.0225	Rental rate
	b	-1	-1	-1	Borrowing constr.
	$\bar{h}$	0.33	0.33	0.33	Labor normalization
	β			0.9889	Discount factor
Estimated Parameters	β	0.9889	0.9622		Discount factor
	λ		24.394		$\underline{c}$ utility cost
	ρ_c		0.5867		$\underline{c}$ AR
	σ_c		3.0767		$\underline{c}$ volatility
	μ_c		0.0529		$\underline{c}$ average
	ρ_m			0.2915	ME AR
	σ_m			0.8709	ME volatility
Moments	$\frac{\overline{K}}{\overline{K}^{\alpha}\overline{H}^{1-\alpha}}$	11.9377	14.0878	11.9377	Wealth/income
	$\frac{sd(d\log C)}{sd(d\log I)}$	0.3431	1.1089	1.0342	Ave. relative vol.
	$\rho(C,I)$	0.6518	0.2974	0.2175	$corr(d\log C, d\log I)$
	$k<0$	0.0365	0.0987	0.0365	Fraction borrowers
	$AR(C_t)$	0.9412	0.6852	0.6800	C AR
Targets	$\frac{\overline{K}}{\overline{K}^{\alpha}\overline{H}^{1-\alpha}}$	12	12		Wealth/income
	$\frac{sd(d\log C)}{sd(d\log I)}$		1.05	1.05	Ave. relative vol.
	$\rho(C,I)$		0.23	0.23	$corr(d\log C, d\log I)$
	$k<0$		.10		Fraction borrowers
	$AR(C_t)$		0.67	0.67	C AR

Note: The calibrated parameters are the same across models, except the measurement error exercise uses the estimated discount factor from the Bewley model. The income/productivity process parameters are chosen to approximate the PSID, as described in the text. The other calibrated parameters are reasonable values from the literature (see main text for details). For the Bewley model, with and without measurement error, the discount factor β is chosen to match an aggregate wealth/income ratio of 12. For the expenditure shock model, the 5 free parameters are estimated match the 5 targets described. For the measurement error model, after β is chosen, the remaining 2 free parameters are estimated to match the remaining 3 targets described. Reported moments are from a simulation of 2 million quarters (after a 20,000 quarter burn), yielding 250,000 biennial observations. See Appendix D for further details.

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. For comparison, Figure 2 shows the consumption functions and wealth distributions for the baseline Bewley model without expenditure shocks.

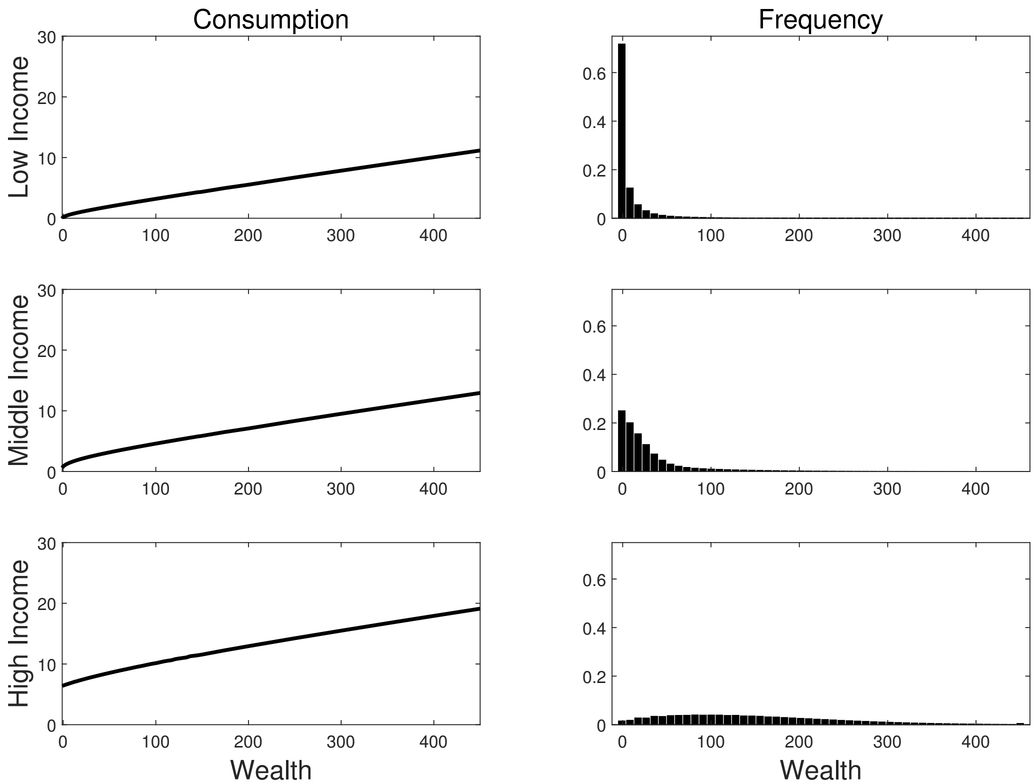

Figure 2

Note: The left column shows the Bewley model consumption policy functions (without expenditure shocks) at different income levels, and the right column shows the steady-state wealth distribution conditional on these income levels. Low (Middle, High) income means both the permanent and persistent components of productivity are at their low (middle, high) discretized values.

What does this imply for the MPC distribution across households? Comparing the wealth distributions in Figure 1 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 3 is the ergodic MPC histogram, drawn separately for households below and above median wealth (poor and rich).[21] Figure 3 shows that for low-wealth households,

[21]Figures 3 and 4 are constructed from model simulations of $2{,}000{,}000$ quarters with a burn-in of $20{,}000$ quarters. There are 20 bins with a bin size of ≈ 0.05. MPCs greater than 1 are set to 1. We calculate the

the largest mass of MPCs is around zero. While the model is lacking a substantial amount of intermediate MPC low-wealth households, there is a mass of very high MPC households. In contrast, about 40% of high wealth households are close to zero, the rest having moderate MPCs in the neighborhood of 8%. So while very high MPCs come from the poor, a large fraction of moderate MPCs come from the rich, and low MPCs are divided roughly equally amongst the rich and poor (defined relative to median wealth). For comparison, Figure 4 shows the corresponding graph for the Bewley model without expenditure shocks.

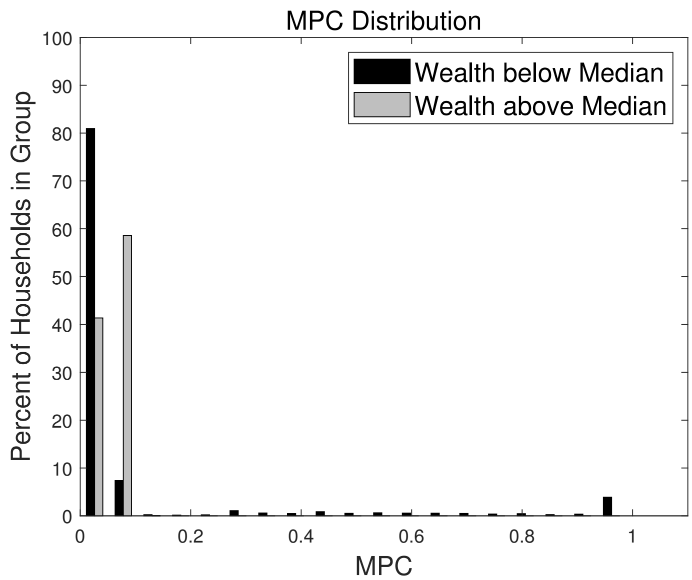

Figure 3

Note: The figure shows the expenditure shock model's ergodic MPC distribution for households, divided into low and high wealth groups. The bin size is 0.05.

This discussion suggests a possible U-shaped MPC distribution, which is indeed evident in Figure 5, which plots average MPC by wealth or income quintile. The lowest-MPC households are, on average, not rich. They are poor enough for the minimum consumption threshold to matter but not so poor that they violate it. For comparison, Figure 6 shows the corresponding graph in the Bewley model without expenditure shocks, which has the classic downward-sloping relationship.

Related evidence of non-decreasing MPCs in wealth/income: The U-shaped pattern in Figure 5 is reminiscent of discussions in Pistaferri and Saporta-Eksten (2012), Carroll et al.

MPC at a point in time as follows. First, we construct an MPC policy function from the optimal consumption policy function. The consumption function tells us optimal consumption on a wealth grid (for any values of the exogenous variables). We define the MPC at a point on the wealth grid to be $\Delta C/\Delta k$, where Δk is the value on the grid one higher minus current wealth (and ΔC is the difference between consumption at those points). Since wealth is not kept on the grid in the simulation, off-grid MPCs are calculated via cubic Hermite interpolation.

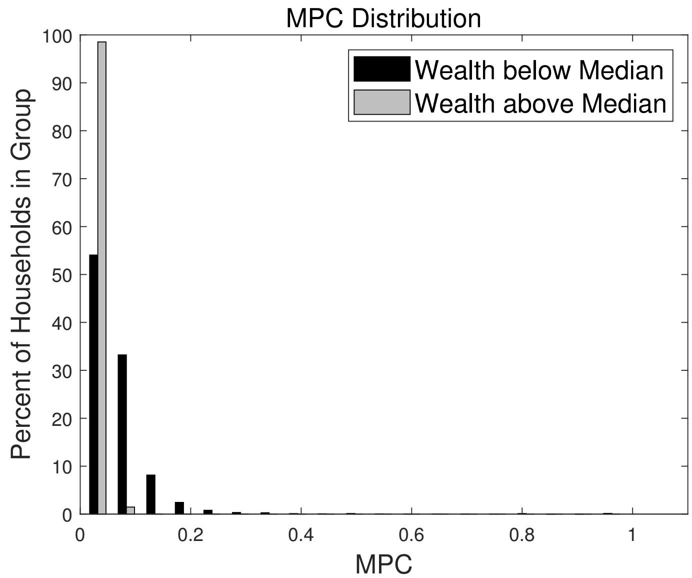

Figure 4

Note: The figure shows the baseline Bewley model's ergodic MPC distribution for households, divided into low and high wealth groups. The bin size is 0.05.

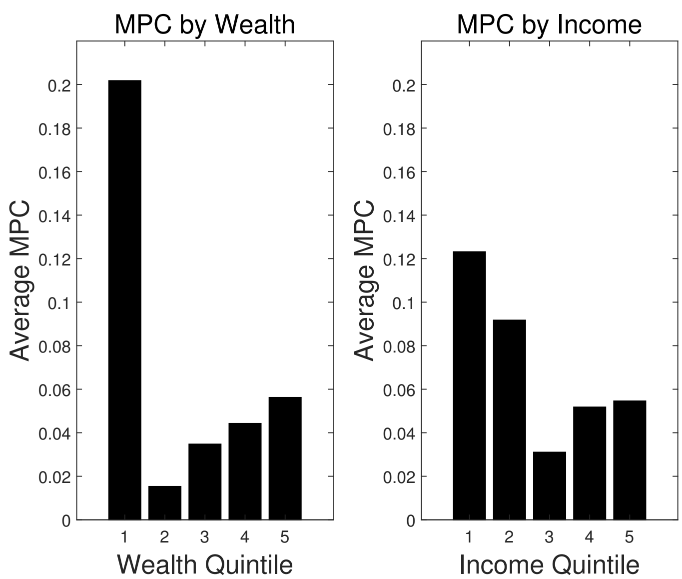

Figure 5

Note: The figure shows average MPC by wealth or income quintile in the ergodic distribution of the expenditure shock model.

(2017), and Campbell and Hercowitz (2019). These papers mention the possibility for a U-shape, but none takes a firm stance, noting the high standard errors in the literature. For the 2021 U.S. fiscal stimulus, Jeon and Walsh (2023) find a statistically significant U-shape by income in the likelihood of households reporting they "mostly spent" the stimulus.

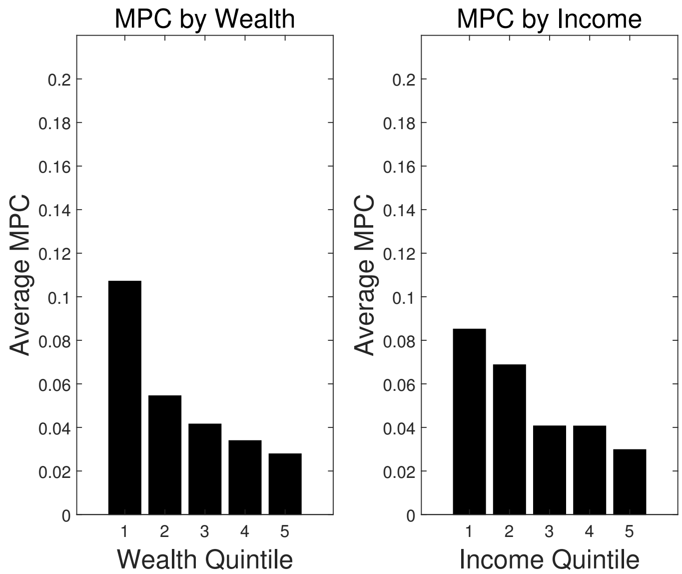

Figure 6

Note: The figure shows average MPC by wealth or income quintile in the ergodic distribution of the Bewley model without expenditure shocks.

Bunn et al. (2018) report an MPC of zero for 77% 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 of ratio of 75–90% (vs. less than 75%). 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% 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 cash-on-hand percentile, still around 10% 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% 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), non-stockholders (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% of households say they mostly adjusted saving and debt/borrowing (the corresponding number was 86% in the prospective survey from March/April 2011). In the retrospective survey, more than 50% of these saving/debt adjusters have household income less then $75,000 and around 20% make less than $35,000. More recently, Koşar et al. (2023) exploit the New York Fed's Survey of Consumer Expectations to estimate marginal propensities to repay debt (MPRD) and MPCs out of the 2020 stimulus payments. They find that MPRDs are substantially decreasing in liquid wealth-to-income ratios (and MPCs are increasing in wealth-to-income).[22]

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 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 CEX. Their quantile regressions imply that for 40–50% of households the spending response to tax rebates is not statistically different from zero. Inspection of Figures 1 and 4 in of 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, low-income, or high-debt households (the types 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. (2022), applying Gaussian mixture linear regression to the 2008 Economic Stimulus Act to capture household heterogeneity, also show a positive relationship between income and MPCs. In short, a multitude of studies call into question the notion that MPCs are falling in wealth/income and instead support the prediction of our model that MPCs may increase with wealth over part of the wealth distribution.

[22] Koşar et al. (2023) offer a complementary explanation —interest rates that rise with debt —for MPCs that are increasing in wealth.

4 Model Fit: Comparison to PSID

Here we examine the model's ability to improve the fit to the PSID with respect to Facts 1 through 4. To make the results comparable with the PSID, we simulate a large number of households over 80 quarters and then convert the quarterly data to a panel of biennial data that spans 10 biennieal periods (as in the PSID). We also examine the fit of the Bewley model and the Bewley model with measurement error. In Appendix B we consider a heterogeneous preference calibration based on Aguiar et al. (2020).

Table 4 shows that including the minimum consumption shock substantially improves the fit to the PSID. Whereas expenditure in the standard Bewley model is insufficiently volatile, too correlated with income, and too persistent, in the expenditure shock model expenditure is slightly more volatile than income (on average), relatively uncorrelated with income, and less persistent. The expenditure shock model also delivers a close fit to the negative autocorrelation of consumption growth observed in the PSID, even though this fact was not targeted in the estimation. Finally, the expenditure shock model replicates the conditional cross-sectional correlation between consumption growth and income growth: among households experiencing a high consumption episode the correlation is a third the size of the correlation across all households.

Introducing measurement error into the Bewley model also yields a strong fit to Facts 1 through 3, with a relative volatility of expenditure that is nearly identical to what we observe in the data. However, the measurement error process required to achieve this strong fit exhibits an extreme degree of volatility. Following Alan and Browning (2010), we evaluate the extent of noise by calculating $\frac{var(d\log C^{meas.\ error})-var(d\log C^{actual})}{var(d\log C^{meas.\ error})} = .89$. So in our Bewley model with measurement chosen to match Facts 1 and 2 and the autocorrelation of consumption, we see 89% noise in consumption growth. As noted in the introduction, this is beyond previous upper bounds from the literature. Furthermore, unlike the expenditure shock model, measurement error cannot explain Fact 4, and our best-fit measurement error process exhibits autocorrelation ($\rho_m \approx .3$).[23]

Column 5 reports the facts based on simulated data from a model similar to that of Aguiar et al. (2020), which assumes a less volatile income process and households with heterogeneous preferences. The details of the model are in Appendix B, and Table B1 reports the facts separately for each type of household. Their model produces a moderate improvement in relative consumption volatility (compared to our standard Bewley simulation), but

[23] Measurement error and expenditure shocks are not mutually exclusive theories. Including both and jointly estimating the $\underline{c}$ and measurement error processes could allow the model to fit a broader set of moments, although we have not attempted to do so, and the degree to which the models are complements vs. substitutes in matching the data is ex ante uncertain.

it increases the correlation between income growth and consumption growth. In short, the Aguiar et al. (2020) model generates consumption that is too connected to income (relative to the data), it generates close-to-zero negative autocorrelation in consumption growth, and high consumption in the model is irrelevant for the cross-sectional relationship between consumption and income growth.

Table 4
Consumption and Income Moments

		PSID	Bewley	Exp. Shock	Meas. Error	Het. Pref
		(1)	(2)	(3)	(4)	(5)
Panel A: Average across households						
Fact 1	$sd(d\log C)/sd(d\log I)$	1.05	0.29	1.10	1.13	0.55
Fact 2	$corr(d\log C, d\log I)$	0.23	0.64	0.34	0.18	0.82
Panel B: Panel autoregression coefficients						
	AR coefficient (log C), FE	0.21	0.67	0.12	0.26	0.76
	AR coefficient (log C), pooled	0.67	0.94	0.68	0.68	0.98
	AR coefficient (C growth), FE	-0.38	-0.078	-0.46	-0.34	-0.036
Fact 3	AR coefficient (C growth), pooled	-0.36	0.019	-0.44	-0.31	0.11
Panel C: Average across years						
	$corr(d\log C, d\log I)$	0.21	0.65	0.29	0.21	0.92
	Conditional $corr(d\log C, d\log I)$	0.073	0.67	0.098	0.22	0.92
Fact 4	ratio	0.35	1.03	0.34	1.05	1.00

Note: See caption for Table 1. "Exp. shock" refers to our estimated expenditure shocks model, and "Meas. error" refers to the Bewley model with measurement error. Het. Pref refers to the model similar to Aguiar et al. (2020) described in Appendix B.

4.1 High Consumption Episodes Analysis Based on Quarterly Data

Our interpretation of Fact 4 is that households that exhibit high consumption (relative to their within-household average) are constrained by high consumption thresholds. Here we examine this hypothesis using simulated quarterly data from the expenditure shock model. The advantage of examining quarterly data is that we can directly observe whether households are saving-constrained at any point in time. We can also aggregate the data to the biennial frequency while keeping track of how many times within a year households are saving-constrained.

Table 5 reports relevant summary statistics from the expenditure shock model, both at the quarterly frequency (column 1) and aggregated to the biennial frequency (column 2). In any given quarter, 24% of agents are on their saving constraint (row 1), and 5% are paying the utility cost of letting consumption fall below the threshold (row 7). Therefore, nearly 30% of agents are directly affected by the consumption threshold in any given quarter. At the biennial frequency, nearly 50% of agents experience a saving constraint for at least part of the year, and 10% of agents pay the utility cost for at least part of the year.

How well does a high consumption indicator capture saving constraints? According to column 4, over 60% of agents that have high consumption are indeed saving-constrained. That share increases to over 70% if we redefine high consumption as consumption exceeding the within-household average by *two* standard deviations (row 6). However, higher thresholds for the high consumption classification also imply that fewer households are included in the criteria. Under our baseline threshold of 1.5 standard deviations, 9% of agents experience high consumption in a quarter or year (row 3). That share falls to 5% when using a threshold of 2 standard deviations (row 5). 1.5 standard deviations roughly maximizes the correlation between episodes and being saving-constrained.

The strong correspondence between high consumption episodes and binding saving constraints from Table 5 explains why the expenditure shock model provides a strong fit to Fact 4. Periods of high consumption tend to indicate households that are saving-constrained and hence have consumption that is disconnected from income. Saving constraints are far less prevalent among those not experiencing high consumption, and for these households consumption and income are more closely connected. This asymmetry is not present in the Bewley model, with or without measurement error, which is why it cannot generate Fact 4.

Table 5
Correspondence between Saving Constrained and High Consumption in Expenditure Shock Model

	Share of agents who	Quarterly (1)		Biennial (2)
(1)	Are on saving constraint in a given quarter	0.24	Are on saving constraint for at least a quarter out of the year	0.48
(2)	Are on saving constraint for four consequtive quarters	0.05	Are on saving constraint for a full year	0.05
(3)	Have high consumption	0.09		0.09
(4)	Are saving constrained if experiencing high consumption	0.61	Are saving constrained for at least a quarter if experiencing high consumption	0.64
(5)	Have very high consumption	0.05		0.04
(6)	Are saving constrained if experiencing very high consumption	0.72	Are saving constrained for at least a quarter if experiencing very high consumption	0.73
(7)	Are paying the utility cost (consumption $<$ minimum threshold)	0.05	Are paying the utility cost for at least a quarter	0.10

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.

Figure 7 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. On impact, the average consumption of the stimulus is 30% higher in the expenditure shock model, but less is consumed after a few quarters. This is likely because the expenditure shock model has many more of both high MPC households (below the threshold) and close-to-zero MPC households (at the threshold). The high MPC households immediately consume a large fraction of the transfer, but the remaining households respond on average less than in the Bewley model. So the expenditure shock model generates a more frontloaded aggregate response to a homogeneous transfer. That said, Figure 7 suggests the aggregate difference between the two models is not particularly large with respect to a uniform stimulus. However, as we will see, the models are more different with respect to targeted transfers.

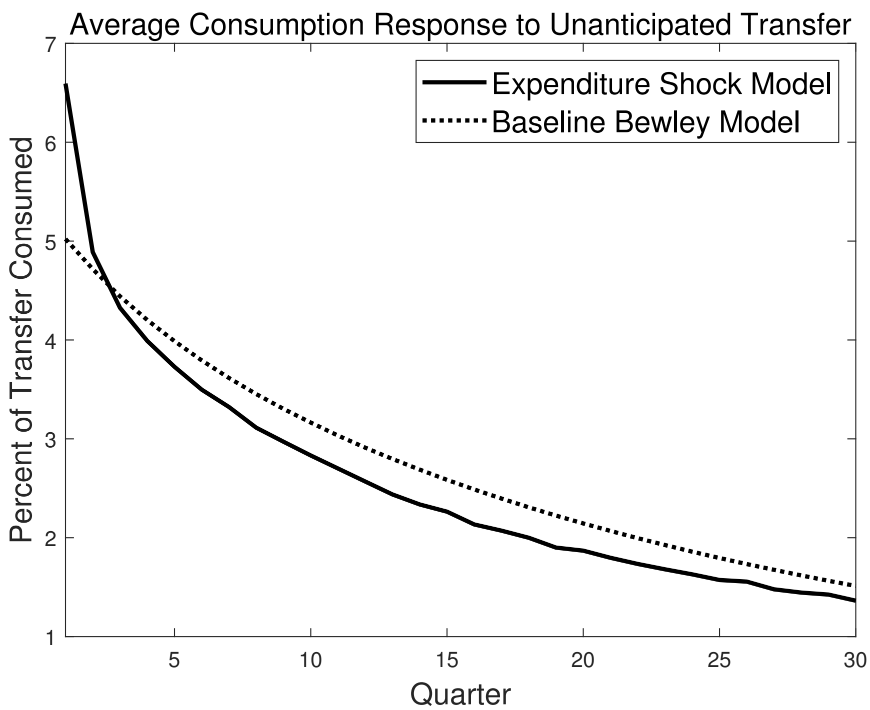

Figure 7

Note: The figure shows the effect on aggregate consumption of a one-time wealth transfer of 0.5 to all households (a little over 13% of average quartly income), starting at the ergodic distribution. Each line is the difference between the average consumption path with and without the transfer, divided by the transfer and then multiplied by 100. Therefore, the y-axis is the amount of the transfer consumed on average in the corresponding quarter.

In the expenditure shock model, what is the heterogeneous effect of the transfer? Figure 8 shows the average response by saving constraint status, starting at the ergodic distribution.[24] For each group (saving-constrained, unconstrained, and paying-the-utility cost), the consumption response to the transfer is expressed as the percent of the transfer consumed in e ach quarter. Paying-the-utility-cost households have a large initial response that rapidly decays 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/long-term prospects.[25] Saving-constrained households, on the other hand, initially save *all* of the transfer, allowing them to increase consumption over a longer horizon. That is, their consumption response is hump-shaped. Unconstrained households have a moderate initial increase in consumption, which dissipates slowly. By about 4 quarters, the consumption response of saving-constrained households is similar to 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 9 repeats the exercise of Figure 8 but defines groups by wealth tercile instead of saving constraint status. Echoing the U-shaped MPC by wealth from Figure 5, the consumption response to the transfer is non-monotonic in wealth. While the bottom 33% poorest households have the largest initial response, the middle group has the weakest initial response, as they are more likely to be saving-constrained. The richest 33% 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 by 15 quarters the poorest households have the lowest response. For comparison, Figure 10 is the corresponding graph for the baseline Bewley model without expenditure shocks. In the Bewley model, households in the second tercile of wealth have much stronger initial consumption responses than households in the third tercile (in contrast to the model with expenditure shocks).

[24] Figures 7, 8, 9, and 10 are constructed by drawing $200{,}000$ households from the stationary distribution and simulating the economy for 30 quarters (after a pre-shock burn-in of 100 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. The series are normalized by the transfer size and multiplied by 100 so they can be interpreted as the percent of the transfer consumed in each quarter.

[25] This implication from our model could rationalize the large spike in consumption at the time of the transfer documented in Aladangady et al. (2022) for tax refunds in the US and in Hamilton et al. (2023) for pension withdrawals in Australia.

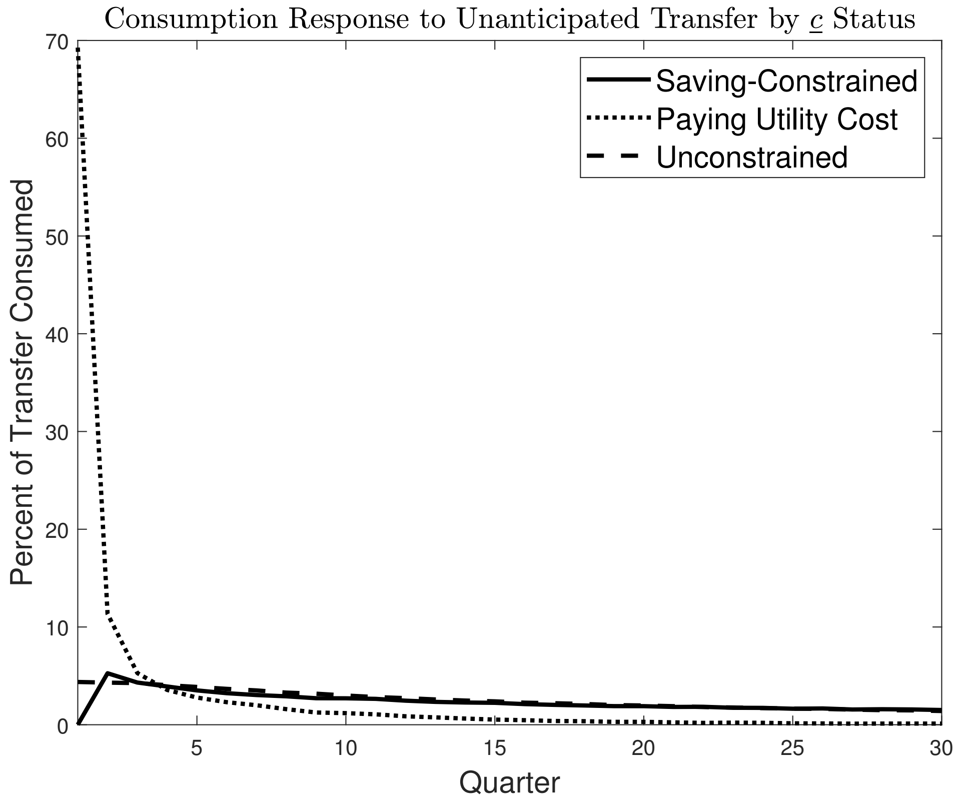

Figure 8

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, $c \leq \underline{c}-0.00001$), starting at the ergodic distribution. For each group, the line is the difference between the average consumption path with and without the transfer, divided by the transfer and then multiplied by 100. Therefore, the y-axis is the amount of the transfer consumed on average in the corresponding quarter.

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 saving-constrained households maintain consumption at

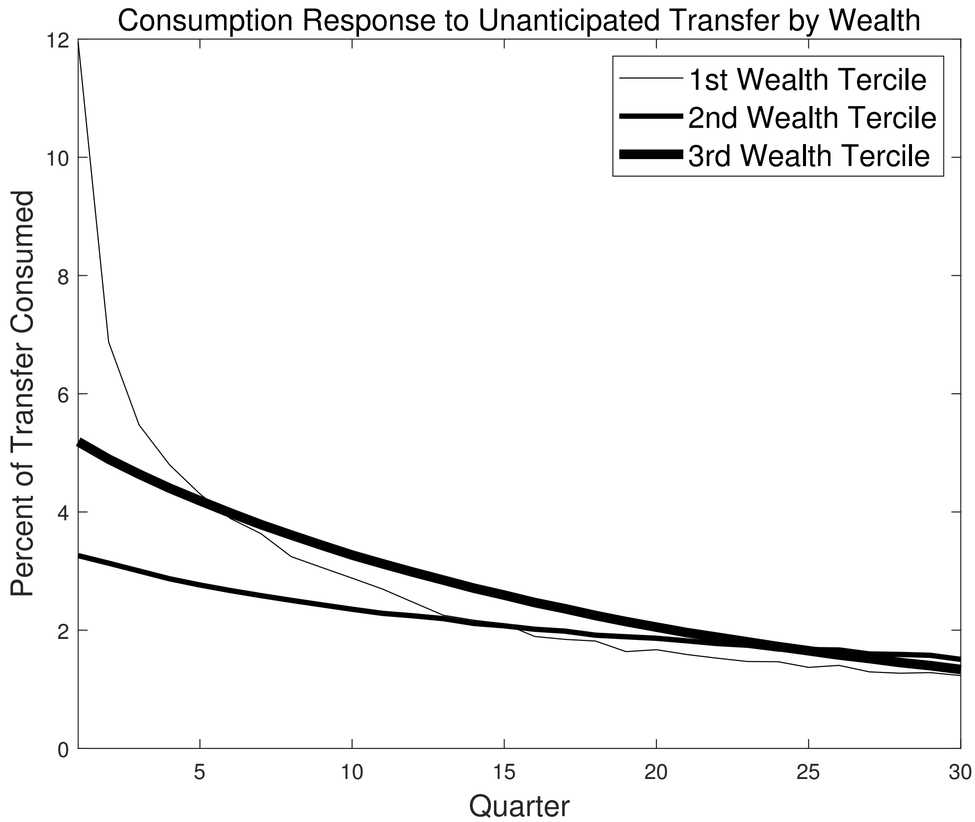

Figure 9

Note: The figure shows the average consumption response, in the expenditure shock model, to a one quarter unanticipated transfer of 0.5 by wealth, starting at the ergodic distribution. Line thickness corresponds to wealth tercile, so the thinnest line represents the bottom 33% by wealth in the ergodic distribution. For each group, the line is the difference between the average consumption path with and without the transfer, divided by the transfer and then multiplied by 100. Therefore, the y-axis is the amount of the transfer consumed on average in the corresponding quarter.

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, announced in period 1, which remains until period 15, when it returns to normal. In Figure 11, we show the percent change in consumption between periods 4 and 5 for the lowest income households (lowest x_4 and z_4),

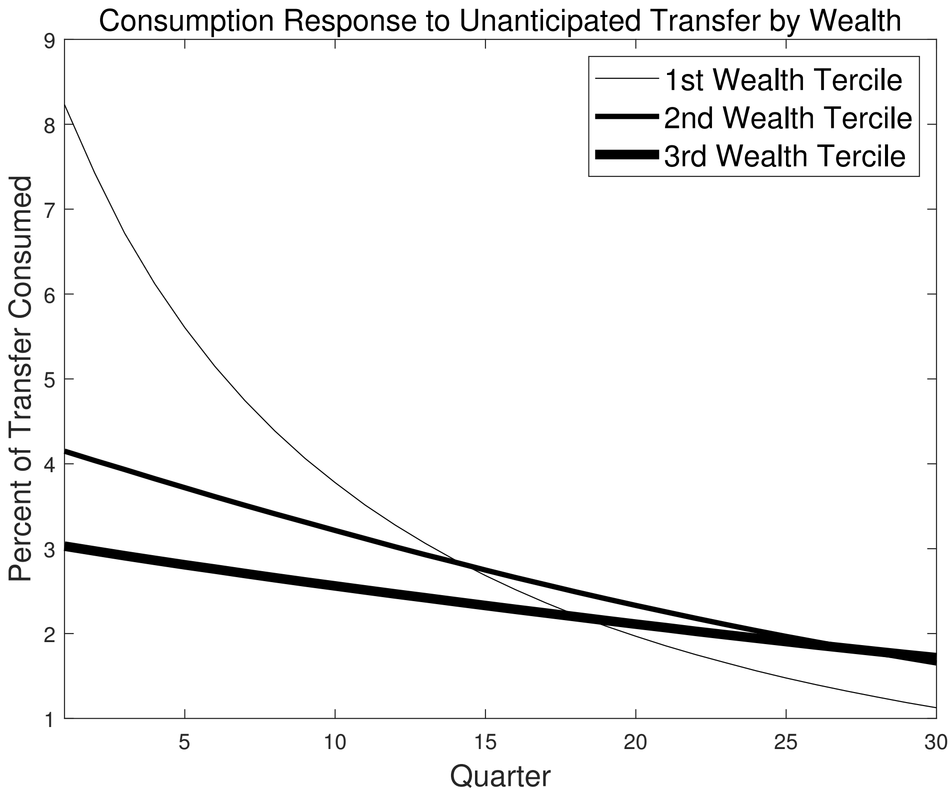

Figure 10

Note: The figure shows the average consumption response, in the baseline Bewley model without expenditure shocks, to a one quarter unanticipated transfer of 0.5 by wealth, starting at the ergodic distribution. Line thickness corresponds to wealth tercile, so the thinnest line represents the bottom 33% by wealth in the ergodic distribution. For each group, the line is the difference between the average consumption path with and without the transfer, divided by the transfer and then multiplied by 100. Therefore, the y-axis is the amount of the transfer consumed on average in the corresponding quarter.

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 their period 4 values. Considering the consumption functions in Figure 1, we see that many of the households with wealth less than around $k = 20$ are saving-constrained but near the point where they begin to violate the threshold. Consequently, 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. Without binding $\underline{c}$, there is no decline in consumption between periods 4 and 5: these house-

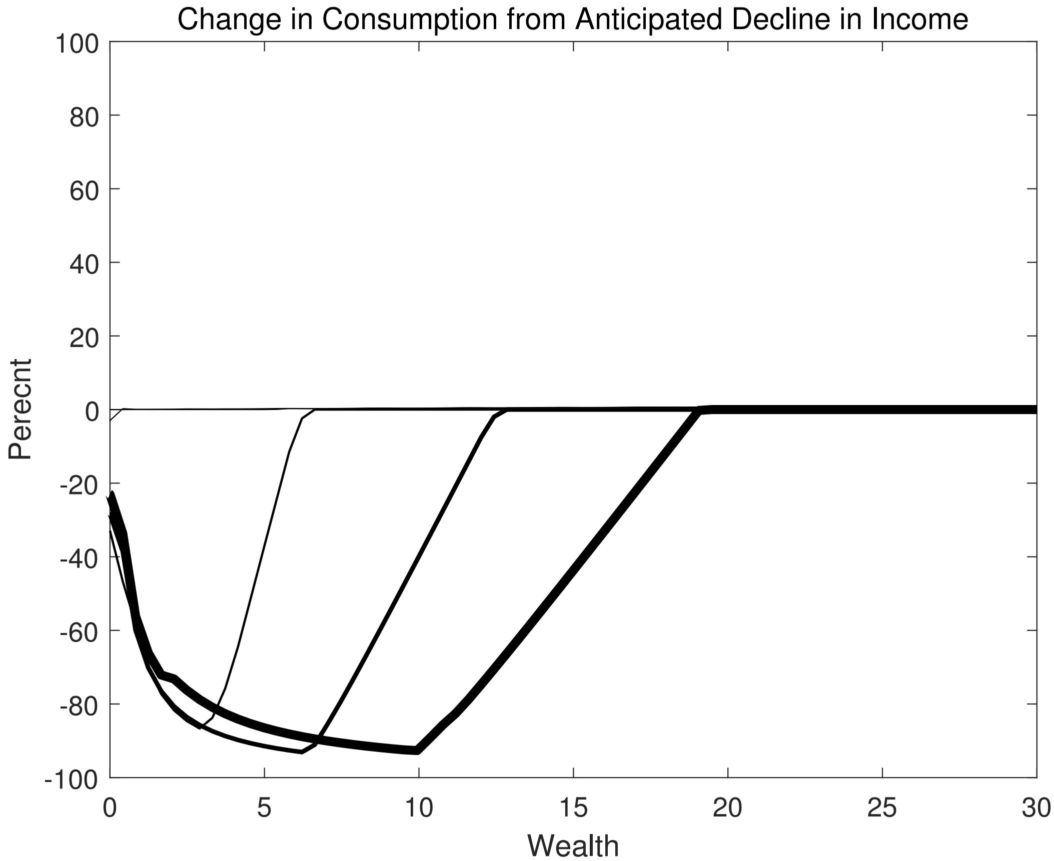

Figure 11
Note: The figure shows the percent change in consumption between periods 4 and 5 from a 25% decline in income in periods 5 through 15 that was announced in period 1, as a function of period 4 wealth, for households with the lowest realization of the income process in period 4. Line thickness corresponds to $\underline{c}_4$, with the thickest representing the highest realization of $\underline{c}$. While wealth changes endogenously between periods 4 and 5, we assume the exogenous shocks remain unchanged between 4 and 5.

holds 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 due to the fact that they can somewhat soften the blow of the income decline through winding down wealth before breaking the threshold.

6 Welfare

It is well-known that with standard preferences, income fluctuations in Bewley models are not particularly costly in welfare terms, despite the fact that the only insurance vehicle is a single asset subject to a borrowing constraint: agents are able to self-insure relatively well through precautionary savings (see, for example, Krusell and Smith, 1998). Consumption threshold shocks reduce the efficacy of such self-insurance by constraining saving and limiting the ability of households to smooth out the effects of random income declines.

To quantify the extent to which expenditure shocks amplify the welfare costs of income fluctuations, we perform the following exercise (described in detail in Appendix C): for each point $(k, z, x, \underline{c})$ in the state space, we freeze labor income (z, x) and resolve the model. We then calculate the welfare gain at each point in the state space by finding the % permanent increase in consumption from the original model that gives the same utility as freezing labor income at that point. The total welfare gain (from behind the veil of ignorance), integrates the gains with respect to the stationary distribution. The procedure is the same for the Bewley model, except the model becomes deterministic with labor income frozen.

In our Bewley model, the welfare gain is 2.84%, whereas in the expenditure shock model it rises by an order of magnitude to 37.03%. The welfare gain from turning off labor income fluctuations in the heterogeneous preference model (described in Appendix B) is 6.7%, larger than the gain from the Bewley model but far less than the gain in the expenditure shock model. In short, expenditure shocks greatly amplify the welfare cost of income fluctuations.

7 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 forego 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. First, household-level consumption is nearly as volatile as income. Second, household-level consumption is relatively uncorrelated with income. Third, household-level consumption growth is negatively autocorrelated. Fourth, the cross-sectional correlation between consumption growth and income growth is weaker among household with high levels of consumption (relative to the within-household average).

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 many lower-income 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. First, a number of studies document that many medium-to-low-income households exhibit low marginal propensities to consume out of transfer income. Our theory rationalizes such behavior because many medium-to-low-income households are against their minimum consumption threshold (and hence are saving-constrained). Second, 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 wealth is low and consumption is at 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.

Expenditure shocks drastically increase the value of insurance against income fluctuations (relative to a world without expenditure shocks), and they point to the potential benefits of alternative forms of social insurance including in-kind benefits that help cover expenditure shock expenses. Our theory also has important implications for the propagation of macroeconomic shocks. In Miranda-Pinto et al. (2023) we show that consumption thresholds are important for understanding the cross-country relationship between fiscal effects and credit markets. In particular, higher shares of proxies for savings-constrained households (during periods of normal-to-loose credit supply) are associated with stronger credit market relaxation in response to expansionary fiscal shocks.[26] With these applications in mind, we suggest that researchers view our expenditure shocks as an important add-on to the basic model with income shocks.

[26] See also Jeon and Walsh (2023), who use our theory to understand the differential effect of U.S. fiscal stimulus in 2020 vs. 2021.

References

Aguiar, M. A., Bils, M., and Boar, C. (2020). Who are the hand-to-mouth? Technical report, National Bureau of Economic Research.

Aladangady, A., Aron-Dine, S., David, C., et al. (2022). Spending responses to high-frequency shifts in payment timing: evidence from the earned income tax credit. *American Economic Journal: Economic Policy*.

Alan, S., Attanasio, O., and Browning, M. (2009). Estimating euler equations with noisy data: Two exact GMM estimators. *Journal of Applied Econometrics*, 24:309–324.

Alan, S. and Browning, M. (2010). Estimating intertemporal allocation parameters using synthetic residual estimation. *Review of Economic Studies*, 77:1231–1261.

Anderson, T. and Hsiao, C. (1981). Estimation of dynamic models with error components. *Journal of the American Statistical Association*, 76:589–606.

Andreski, P., Li, G., Samancioglu, M. Z., and Schoeni, R. (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):132–135.

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies*, 58:277–297.

Athreya, K., Tam, X. S., and Young, E. R. (2009). Unsecured credit markets are not insurance markets. *Journal of Monetary Economics*, 56:83–103.

Bernheim, B., Skinner, J., and Weinberg, S. (2001). What accounts for the variation in retirement wealth among U.S. households? *American Economic Review*, 91(4):832–857.

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

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

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy*, 126(S1):73–115.

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

Campbell, J. R. and Hercowitz, Z. (2019). Liquidity constraints of the middle class. *American Economic Journal: Economic Policy*, 11(3):130–155.

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

Chatterjee, S., Corbae, D., Nakajima, M., and Ríos-Rull, J.-V. (2007). A quantitative theory of unsecured consumer credit with risk of default. *Econometrica*, 75(6):1525–1589.

Chetty, R. and Szeidl, A. (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics*, 122(2):831–877.

Cho, Y., Morley, J., and Singh, A. (2022). Did marginal propensities to consume change with the housing boom and bust? *Working Paper*.

Commault, J. (2022). Does consumption respond to transitory shocks? Reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics*, 14(2):96–122.

Di Pillo, G., Liuzzi, G., S., L., Piccialli, V., and Rinaldi, F. (2016). A DIRECT-type approach for derivative-free constrained global optimization. *Computational Optimization and Applications*, 65(2):361–397.

Dynan, K. E. (2000). Habit formation in consumer preferences: Evidence from panel data. *American Economic Review*, 90(3):391–406.

Farrell, D. and Greig, F. (2015). Weathering volatility: Big data on the financial ups and downs of U.S. individuals. *JPMorgan Chase Institute*.

Farrell, D. and Greig, F. (2017). Coping with costs: Big data on expense volatility and medical payments. *JPMorgan Chase Institute*.

Farrell, D., Greig, F., and Yu, C. (2019). Weathering volatility 2.0: A monthly stress test to guide savings. *JPMorgan Chase Institute*.

Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424.

Guvenen, F. and Smith, A. A. J. (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica*, 82(6):2085–2129.

Hamilton, S., Liu, G., and Sainsbury, T. (2023). Early pension withdrawal as stimulus. *Available at SSRN 4389699*.

Hubmer, J., Krusell, P., and Smith, A. J. (2021). Sources of U.S. wealth inequality: Past, present, and future. *NBER Macroeconomics Annual*, 35:391–455.

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

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

Jappelli, T. and Pistaferri, L. (2017). *The Economics of Consumption: Theory and Evidence*. Oxford University Press, New York.

Jeon, W. and Walsh, K. (2023). Heterogeneity in the spending response to stimulus: Evidence from the pulse survey. *Working Paper*.

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

Karahan, F. and Ozkan, S. (2013). On the persistence of income shocks over the life cycle: Evidence, theory, and implications. *Review of Economic Dynamics*, 16(3):452–476.

Koşar, G., Melcangi, D., Philossoph, L., and Wiczer, D. (2023). Stimulus through insurance: the marginal propensity to repay debt. *Working Paper*.

Krueger, D., Mitman, K., and Perri, F. (2016). Macroeconomics and household heterogeneity. In *Handbook of Macroeconomics*, volume 2, pages 843–921.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kueng, L. (2018). Excess sensitivity of high-income consumers. *Quarterly Journal of Economics*, 133(4):1693–1751.

Lewis, D., Melcangi, D., and Pilossoph, L. (2022). Latent heterogeneity in the marginal propensity to consume. *Working Paper*.

Livshits, I., MacGee, J., and Tertilt, M. (2007). Consumer bankruptcy: A fresh start. *American Economic Review*, 97(1):402–418.

Luo, Y., Nie, J., Wang, G., and Young, E. R. (2017). Rational inattention and the dynamics of consumption and wealth in general equilibrium. *Journal of Economic Theory*, 172:55–87.

Miranda-Pinto, J., Murphy, D., Walsh, K., and Young, E. (2023). Saving constraints, inequality, and the credit market response to fiscal stimulus. *European Economic Review*, 151(104355).

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

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426.

Pistaferri, L. and Saporta-Eksten, I. (2012). Changes in the income distribution and aggregate consumption. *Italian Ministry of Economy and Finance, Working Paper 11*.

Politis, D. and Romano, J. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.

Sahm, C. R., Shapiro, M. D., and Slemrod, J. (2015). Balance-sheet households and fiscal stimulus: Lessons from the payroll tax cut and its expiration. NBER Working Papers 21220, National Bureau of Economic Research, Inc.

Shapiro, M. D. and Slemrod, J. (2003). Consumer response to tax rebates. *American Economic Review*, 93(1):381–396.

Shea, J. (1995). Union contracts and the life-cycle/permanent-income hypothesis. *American Economic Review*, 85(1):186–200.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717.

Young, E. 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*, 34:36–41.

A Appendix Table

Table A1
Facts 1 through 4 under Alternative Expenditure Measures

Expenditure Measure:	Average reltive volatility (Fact 1)	Correlation with income growth (Fact 2)	Autocorrelation of consumption growth (Fact 3)	Ratio of cross-sectional correlations (Fact 4)
Baseline	1.05	0.23	-0.38	2.88
Broad	1.05	0.25	-0.38	1.48
Baseline Net Durables	0.99	0.23	-0.40	2.26
Broad Net Durables	1.00	0.24	-0.38	1.70

Note: This table presents, for various definitions of expenditure in the PSID, the average volatility relative to income (Fact 1), the average correlation with income growth (Fact 2), the autoregressive coefficient on expenditure growth (Fact 3), and the ratio of the cross-sectional correlation between consumption growth and income growth in the full sample relative to the subset of households experiencing high expenditure (Fact 4). The baseline measure of expenditure excludes all categories of expenditure that were added in 2005: clothing, travel other recreational expenses, telephone, internet, household repairs, and household furnishing. The Broad measure includes these categories. The Baseline Net Durable measure excludes purchases of automobiles, and the Broad Net Durable measure excludes purchases of household furnishings as well as purchases of automobiles.

B Simulations with Off-the-Shelf Income Process and Heterogeneous Preferences

Appendix Table B1 compares our four PSID facts with simulations from an alternate calibration of our Bewley model with heterogeneous preference and an off-the-shelf income process. It is similar to the one-asset model from Aguiar et al. (2020). Specifically, we solve the problem

$$V(a,y) = \max_{C\geq 0, a'\geq 0} \left[C^{1-1/\sigma} + \beta \left(E_{y'} \left[V\left(a', y'\right)^{1-\gamma} | a, y \right] \right)^{\frac{1-1/\sigma}{1-\gamma}} \right]^{\frac{1}{1-1/\sigma}}$$

subject to the budget constraint

$$C + a' \leq (1+r)\, a + \bar{h} y.$$

The annual process for income is the Krueger et al. (2016) one used in Section 2 of Aguiar et al. (2020):

$$\begin{aligned} \log y_t &= x_t + \epsilon_t, \\ x_t &= \rho_x x_{t-1} + \eta_t, \end{aligned} \tag{1}$$

in which $\rho_x = 0.97$, and the error terms ϵ_t and η_t are normal distributed with mean 0 and standard deviations $\sigma_\epsilon = 0.23$ and $\sigma_\eta = 0.20$. We convert to quarterly by taking $\rho_z^{1/4}$, $\sigma_\epsilon/2$, and $\sigma_\eta/2$, and approximate each process as a Markov chain using Rouwenhorst's method with 7 states each.

We set risk aversion $\gamma = 4$, solve the model for the four combinations of high and low $\beta \in \{0.9274, 0.9789\}$ (the discount factor) and $\sigma \in \{0.5, 1.5\}$ (the EIS), and use the population weights from Aguiar et al. (2020). We choose the high β to match the wealth/income ratio and then set the low β so that we have the same ratio of high to low β as in Aguiar et al. (2020). We normalize $\bar{h} = 2.5895$ and set $r = 0.0083$.

We report the moments weighted across all households in Column 2, and we report moments exclusively from each of the four types in the remaining columns. It is readily apparent that the (low-β) households have a higher ratio of consumption growth volatility to income growth volatility. But these households are not prevalent enough to yield an average (across-household) ratio that matches the PSID. Furthermore, these households' consumption growth is far too correlated with their income growth. And for each household type, consumption is too autocorrelated.

Table B1
Simulations with Off-the-Shelf Income and Heterogeneous Preferences.

		PSID	Heterogeneous Preferences				
			All	Low β Low σ	Low β High σ	High β Low σ	High β High σ
		(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Average across households							
Fact 1	$sd(\log C)/sd(\log I)$	1.05	0.55	0.82	0.95	0.46	0.54
Fact 2	$corr(d\log C, d\log I)$	0.23	0.82	0.92	0.97	0.79	0.81
Panel B: Panel autoregression coefficients							
	AR coefficient ($\log C$), FE	0.21	0.76	0.68	0.69	0.79	0.82
	AR coefficient ($\log C$), pooled	0.67	0.98	0.93	0.94	0.99	0.98
	AR coefficient (C growth), FE	-0.38	-0.04	0.063	-0.073	-0.018	0.027
Fact 3	AR coefficient (C growth), pooled	-0.36	0.11	0.12	0.04	0.16	0.21
Panel C: Average across years							
	Cross-sectional $corr(d\log C, d\log I)$	0.21	0.92	0.96	0.99	0.93	0.92
	Cross-sectional conditional $corr(d\log C, d\log I)$	0.073	0.92	0.98	0.99	0.94	0.94
Fact 4	ratio	0.34	1.00	1.02	1.00	1.01	1.03
	Fraction of households		1.00	0.001	0.16	0.73	0.11

Note: See the caption of main-text Table 1.

With regard to Fact 3, all household types display a positive autocorrelation of income growth in pooled regressions, failing to match the large negative autocorrelation in the data. With fixed effects, two groups have negative autocorrelation, but the magnitude is small relative to the PSID. Finally, the heterogeneous preference model fails to match our Fact 4 (across all types). The cross-sectional correlation of income growth and consumption growth is close to one, independent of episodes of high consumption.

C Welfare and Income Fluctuations

Denote the decision rules in the expenditure shock model by $(c, k') = g^*(k, z, x, \underline{c})$. The model without income changes has a value function that solves

$$\begin{aligned} V(k,z,x,\underline{c}) &= \max_{k' \geq b, c \geq 0} \left\{ \log(c) - \lambda \max\{\underline{c} - c, 0\} + \beta E_{\underline{c}'}\left[V\left(k', z, x, \underline{c}'\right)\right]\right\} \\ c + k' &\leq (1 + r - \delta + \phi \mathbf{1}(k \leq 0))\, k + w \exp(z + x)\, \overline{h}. \end{aligned}$$

We solve for the value of following the benchmark decision rules but receiving a $\psi\%$ increase in consumption permanently:

$$\begin{aligned} V^*(k,z,x,\underline{c};\psi) &= \log(1+\psi)+\log\left(g_c^*(k,z,x,\underline{c})\right)-\lambda\max\left\{\underline{c}-(1+\psi)\,g_c^*(k,z,x,\underline{c}),0\right\}+ \\ &\quad \beta E_{z',x',\underline{c}'}\left[V^*\left(g_k^*(k,z,x,\underline{c}),z',x',\underline{c}';\psi\right)\right]. \end{aligned}$$

For each state $(k,z,x,\underline{c})$, the welfare gain in consumption units solves the nonlinear equation

$$V^*(k,z,x,\underline{c};\psi(k,z,x,\underline{c}))=V(k,z,x,\underline{c}).$$

We solve this equation by approximating V^* in the ψ direction by a linear spline and using Brent's method. The total welfare gain is $\psi(k,z,x,\underline{c})$ integrated over the stationary distribution. The welfare gain for the heterogeneous preference model is calculated analogously.

In the Bewley model, the approach is similar, but the household problem is deterministic if income does not fluctuate, so we can compute the welfare change in closed form:

$$\psi(k,z,x)=\exp\left((1-\beta)\left(V(k,z,x)-v^*(k,z,x)\right)\right)-1,$$

where v^* is the original value function for the Bewley model. The expression is derived from the fact that in the Bewley model

$$V^*(k,z,x;\psi)=v^*(k,z,x)+\frac{\log(1+\psi)}{1-\beta}.$$

D Computational Appendix

The recursive problem for the household can be written as

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

subject to the budget constraint

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

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 $\left[b, \overline{k}\right]$ 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 which 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 studying anticipated income changes, we solve the sequential version of the household problem for $t \in \{1, ..., T\}$ 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

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

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

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

We assume $\tau_t = 0$ for $t \leq 4$ and $t \geq 15$ and set $\tau_t = .25$ for $t = 5, \ldots, 14$. 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. For the measurement error model, we fix both the Bewley simulation (of 250,000 biennial observations) and 250,000 standard normal draws. Then we use Matlab's fminsearch.m to find the persistence and volatility of measurement error that best match the targeted moments (with an intial condition of all measurement error parameters being zero). Specifically, we minimize the sum of squared deviations from the moments.