Expenditure Risk and Household Wealth Dynamics*

Luigi-Maria Briglia†

March 19, 2025

Abstract

A substantial fraction of U.S. households experience wealth loss, often transitioning into hand-to-mouth (H2M) status despite no clear income shocks. Standard models attribute these transitions to transitory income fluctuations, yet expenditure risk (ER)—persistent, unpredictable spending shocks—explains a larger share of wealth loss episodes. Using PSID data (1999–2021) and a flexible expenditure model, ER is identified as the primary driver of H2M transitions, surpassing transitory income risk. A heterogeneous-agent model incorporating ER successfully replicates these patterns, highlighting the role of expenditure shocks in shaping household wealth dynamics beyond standard income risk mechanisms.

Keywords: Income risk, Expenditure risk, Wealth dynamics, Precautionary savings
JEL Classification: D15, D31, D52, E21

*

†University of Tübingen: luigi.briglia@uni-tuebingen.de

1 Introduction

What drives households into wealth loss? A substantial fraction of U.S. households live hand-to-mouth (H2M), consuming most of their income with little or no accumulated wealth to buffer against financial shocks. According to the Panel Study of Income Dynamics (PSID), approximately 20% of households hold zero or negative net worth, and 10% of households with initially positive net worth transition into non-positive wealth between consecutive survey waves. Standard economic theory attributes these transitions to transitory income fluctuations, suggesting that households losing wealth are those experiencing short-term income setbacks. However, the frequency and persistence of these episodes indicate that transitory income risk alone cannot fully account for observed wealth dynamics.

This paper argues that expenditure risk (ER)—i.e. persistent, unpredictable spending shocks—plays a critical role in wealth loss among H2M households. Using PSID data from 1999 to 2021, I estimate a flexible expenditure model and define ER as the persistent component of household spending that remains unexplained after accounting for income risk, cash on hand, and demographic characteristics. I find that ER significantly predicts transitions into H2M status, contributing 6.73% to the explanatory power of the model in predicting these episodes, as measured from a random forest classification algorithm. In comparison, transitory income risk contributes 2.27%, while permanent income risk plays no role. These findings suggest that wealth loss is not merely a reaction to income fluctuations but also a consequence of persistent, unpredictable expenditure shocks that standard models fail to capture. Specifically, ER is driven by spending on transportation, education, and maintenance (house and car repairs), categories that impose sudden financial burdens, reduce savings, and increase the likelihood of wealth loss.

While much of the literature focuses on the persistence of low wealth, my approach highlights what triggers transitions into H2M status. Existing models explain the concentration of zero-wealth households through time-preference heterogeneity, suggesting that some individuals are persistently H2M due to a structural preference for low savings. In contrast, I show that expenditure risk acts as a distinct driver of wealth loss episodes, complementing time-preference heterogeneity. To investigate this mechanism, I extend a standard heterogeneous-agent model with idiosyncratic income risk to incorporate ER, modeled as persistent shocks to the marginal utility of part of the consumption basket. Unlike the standard model, which prescribe no direct correlation of wealth response to permanent income shocks near the borrowing constraint, my framework generates a marginal wealth response to these shocks, consistent with the empirical findings. Moreover, it successfully replicates the observed predictive power of ER and transitory income risk for wealth loss, while permanent income risk remains insignificant.

Standard models relying on persistent household time-preference heterogeneity[1] explain cross-sectional differences in net worth as a result of heterogeneity in discount factors, which generate different target wealth levels. In these models, wealth loss episodes are primarily driven by transitory income fluctuations rather than differences in discount factors. These frameworks suggest that low-discount-factor households accumulate little wealth and remain structurally vulnerable, with transitory income shocks acting as the proximate trigger for asset depletion, while permanent income shocks exert no direct effect near the borrowing constraint. This mechanism, rooted in the permanent income hypothesis (PIH) of Friedman (1957) and formalized in Hall (1978), remains a cornerstone of modern buffer-stock savings models (Deaton, 1991; Carroll, 1997).

However, empirical evidence challenges these predictions. While permanent income risk explains more of the cross-sectional variation in consumption growth than transitory risk, the pattern reverses for wealth accumulation, with transitory risk playing a dominant role. Strikingly, I find that permanent income risk, which should theoretically induce no direct wealth loss near the borrowing constraint, is correlated with wealth loss events, albeit negligible in predicting H2M transitions once financial controls are included. This suggests that financial vulnerability arises not only from time-preference heterogeneity but also from the dynamic interplay of income shocks and expenditure risk.

To better understand wealth transitions, I estimate an expenditure function that separates systematic spending patterns from unpredicted expenses. ER is identified as the persistent, unexplained component of household spending after controlling for income risk, cash on hand, and demographics—key determinants in standard consumption-savings models. Using PSID data (1999–2021), I implement a machine-learning-based expenditure model that accommodates nonlinear relationships across explanatory variables. Specifically, a random forest regression takes as inputs income risk measures identified following Blundell et al. (2008), available resources, demographics, and household types inferred using grouped fixed effects (Bonhomme et al., 2022). The model's residual—the gap between actual and predicted spending—captures unexpected expenses, which exhibit significant persistence. I exploit this persistence to identify ER, distinguishing systematic expenditure shocks from transitory fluctuations.

To investigate the sources of ER, I analyze its relationship with household spending across non-durable consumption categories. Cross-correlations between ER and expenditure shares reveal that ER is primarily driven by education costs and essential maintenance expenses, such as car and home repairs—categories imposing large, often unavoidable financial burdens. While education expenses reflect sustained commitments that adjust infrequently, maintenance shocks require large, immediate payments that can disrupt house-

[1] Examples include Krusell and Smith (1998) and Kaplan and Violante (2022), who model household discount rates as highly persistent stochastic processes, and Aguiar et al. (2024), who propose fixed discount factor heterogeneity as a fundamental driver of cross-sectional differences in households' average propensity to consume.

hold finances over multiple periods. In contrast, some spending categories remain stable regardless of ER fluctuations due to consumption commitments (Chetty and Szeidl, 2007) or their tight linkage to income and demographics. This heterogeneity underscores ER's role as a distinct financial stressor that standard models struggle to capture.

Empirical results confirm that ER significantly affects household financial dynamics alongside income risk. A one-standard-deviation increase in ER increases consumption growth by 3 percentage points, comparable to the effects of income risk. The impact on wealth accumulation is similarly substantial—households exposed to higher ER accumulate less wealth, with an effect size equal to 3% of their average annual lifetime income. Furthermore, ER explains a cross-sectional share of variation in consumption growth and wealth accumulation similar to that of income risk components. To assess the relative importance of different risk sources in explaining wealth loss, I estimate a random forest classification model. The results confirm that ER is the primary determinant of H2M transitions, contributing for 6.73% of the model explanatory power, compared to 2.27% for transitory income risk and a statistically negligible -0.05% for permanent income risk. These findings, robust across specifications, highlight ER's economic significance in shaping household financial trajectories.

To rationalize these findings, I develop a heterogeneous-agent model where households face idiosyncratic income and expenditure risk under incomplete markets. The model features two consumption goods that enter additively in household flow utility, with one subject to persistent expenditure shocks modeled as an AR(1) process on marginal utility, while the other remains unaffected. This structure introduces a novel channel through which expenditure risk (ER) affects household behavior, distinct from standard models that attribute wealth dynamics primarily to transitory income fluctuations and time-preference heterogeneity. In my framework, persistent and unpredictable expenditure shocks directly deplete household resources over time, leading to wealth loss episodes that cannot be explained by income risk alone.

The model is calibrated within a life-cycle framework to capture key features of household wealth accumulation. In both the benchmark (BE) and expenditure risk (ER) economies, the distribution of discount factors is separately calibrated to match empirical moments on household savings behavior. Specifically, the calibration targets the share of H2M households, as well as the median and third quartile of net worth at both mid-working life and retirement, ensuring consistency with observed cross-sectional wealth heterogeneity. The ER economy introduces additional parameters: (i) the fraction of the consumption basket exposed to ER, calibrated to match the average expenditure share of categories that positively co-move with ER in the data; and (ii) the persistence and volatility of the ER shock, which are jointly calibrated to match both the autocovariance of consumption expenditure in levels and in first differences. These additional features allow the model to replicate the empirical patterns of household expenditure risk and wealth loss dynamics.

An alternative explanatory mechanism for the observed behavior is presented by Hubbard, Skinner, and Zeldes (1995). They argue that asset-based means-tested transfers create moral hazard, leading to a fraction of zero-wealth households in their model. This mechanism generates the incentive to deplete their asset buffer for a household facing a permanent negative shock to income. However, while asset-based means-tested transfers have been significantly reduced over the past four decades, the proportion of H2M households has increased from 15% (as reported by Huggett in 1996, based on the 1983 Survey of Consumer Finances) to 20% in the last 20 years based on the PSID. Unlike asset-based transfers, income-based means-tested transfers remain a relevant feature of the U.S. safety net, as documented by Guner, Rauh, and Ventura (2024). My model accounts for these transfers in a reduced-form way by allowing for a floor on income realizations in the worst states of the world—specifically, during the lowest realizations of both permanent and transitory income shocks. This ensures that the model captures some of the insurance effects of income-based transfers while allowing for the distinct role of expenditure risk in driving wealth loss episodes.

The quantitative model results confirm that introducing expenditure risk fundamentally alters wealth dynamics. The ER economy generates a significant average marginal effect of permanent income risk on wealth loss episodes, a feature absent in the benchmark model, while still matching the qualitative effects of income risk on consumption growth and wealth accumulation. In particular, the ER economy preserves the empirical pattern where transitory income risk plays a dominant role in short-term consumption fluctuations, while permanent income risk has a stronger effect on wealth accumulation. Crucially, the ER economy replicates the predictive structure of risk sources on wealth loss transitions: expenditure risk emerges as the leading driver, transitory income risk has a secondary role, and permanent income risk—absent in the benchmark model—now exhibits a statistically significant but smaller marginal effect.

The key mechanism behind these results is the wedge introduced by expenditure risk in the intertemporal Euler equation. Persistent shocks to the marginal utility of part of the consumption basket act as a multiplicative term within the expectation over future consumption, effectively creating systematic deviations from standard consumption-savings behavior. While these shocks do not directly shift household expectations, they alter the realized covariance between marginal utility and consumption growth rates. Specifically, a higher realization of tomorrow's marginal utility shock implies increased consumption needs in the next period, generating an "as if" effect where more weight is placed on future states with higher expenditures. This covariance-driven mechanism reduces the optimal consumption growth rate relative to the benchmark economy. Consequently, households facing persistent expenditure shocks adjust their savings behavior, leading to amplified financial vulnerability and an increased likelihood of wealth depletion. These results underscore the importance of incorporating ER into heterogeneous-agent models to explain wealth loss episodes beyond standard income fluctuation mechanisms

This paper contributes to two distinct strands of literature. Firstly, it aligns with the applied macroeconomic literature that employs household dynamics to investigate consumption inequality and its relationship with income risk. My empirical approach can be viewed as a refinement of Hall and Mishkin (1982)'s seminal work, which examined the stochastic nature of the consumption-income relationship and used a second-order moving average process to model the component of non-durable consumption unexplained by the permanent-income hypothesis. Furthermore, I incorporate the framework proposed by Blundell et al. (2008) to capture the response of consumption to income shocks and estimate a flexible non-parametric expenditure model. By combining these two strategies, I use the Blundell et al. (2008) framework to identify an expenditure policy function based on income risk realizations, while extending the modeling strategy for unexplained consumption beyond the confines of the permanent income hypothesis presented by Hall and Mishkin (1982).

Secondly, my theoretical approach relates to the literature on the spending behavior of low-wealth households. Hubbard et al. (1995) show how asset-based means-tested transfers create a 100% tax on savings below the given consumption floor, thus generating H2M households in their model. A different approach to match the observed wealth distribution imposes heterogeneity through discount factors or rates of returns. In a recent paper, Aguiar et al. (2024) propose a model where low-wealth households have specific target wealth levels determined by their discount factor and intertemporal elasticities of substitution, while Hubmer, Krusell, and Smith (2021) generate wealth inequality through idiosyncratic rates of returns. My theoretical approach expands on Aguiar et al. (2024)'s framework by adding a life-cycle dimension and delving into the nature of expenditure risk by calibrating non-income-related model ingredients[2] to the stochastic properties of unexplained consumption.

This paper closely relates to the approach of Miranda-Pinto, Murphy, Walsh, and Young (2020), motivated by the observation of medium-to-low-income households exhibiting low marginal propensity to consume (MPC) out of small unexpected income bumps in the form of tax rebates. They develop a theory of saving-constrained households generated by stochastic consumption thresholds that, when active, push household expenditure above its counterfactual level under the threat of a severe utility penalty. However, the mechanisms leave the household that is hit by a consumption threshold shock thirsty for unexploited savings opportunities, which is instrumental in delivering low MPCs but raises challenges for this modeling strategy to generate recurrent and persistent H2M status for a large fraction of households.

[2] I draw upon the literature on health shocks to inform the key components of my modeling strategies. Palumbo (1999) model health variations as shocks to the marginal utility of consumption, while French and Jones (2004) estimate a stochastic process for out-of-pocket healthcare expenditure; De Nardi, French, and Jones (2010) feature both.

The rest of the paper is structured as follows. Section 2 evaluates the standard framework's ability to capture wealth dynamics and motivates the role of expenditure risk. Section 3 examines empirical evidence on how income and expenditure risk shape household wealth transitions. Section 4 introduces the theoretical framework and quantifies the impact of expenditure risk. Finally, Section 5 concludes.

2 Income Risk and Wealth Dynamics

The impact of income risk on wealth dynamics has been extensively explored in economic studies, particularly within the framework of consumption-saving theory under conditions of uncertainty and incomplete markets. According to the theory, to a first-degree approximation, economic agents adjust their consumption levels in response to permanent shifts in income. Conversely, in the face of transitory income fluctuations, they should utilize their financial assets to maintain steady consumption levels. In this section, I study how standard theoretical prescriptions hold up when confronted with the data, particularly in the case of transitions toward low-wealth status.

To account for Aguiar et al. (2024)'s recent developments in the theory of low-wealth holdings, I augment the standard consumption-saving framework under income risk and incomplete markets with fixed household types along the time-preference dimension. The model is able to generate transitions toward low-wealth in response to negative transitory income shocks but cannot deliver a significant relationship between H2M transitions and permanent income risk realizations. The intuition is straightforward: particularly when approaching their borrowing constraint, households will cut down consumption expenditure when facing a permanent negative income shock.

On the other hand, in the data, I identify the differential effect of permanent and transitory income shocks on the binary outcome of transitioning into H2M status, exploiting the panel dimension of the PSID by following Blundell et al. (2008). I find that both types of income risk realizations, transitory and permanent, elicit H2M transitions. The relationship holds true and significant for permanent income shocks after controlling for household head age or educational attainment, and the effects are larger in magnitude when observing households that are closer to the H2M threshold.

2.1 Hand-to-Mouth Transitions in the Data

Table 1 shows the incidence and persistence of H2M status in PSID data. Notably, there is a 10% unconditional probability of falling into low-net-worth status. Once there, two out of three family units will fail, in the following two years, to accumulate enough assets to achieve positive net worth. The outcome of these transition probabilities is the observed share of more than 20% H2M households. Interestingly, the observed population share of H2M households is quite close to the unique stationary distribution[3] implied by the observed transition matrix, suggesting that average transition probabilities have not deviated much over the years from those measured in the data sample.

Table 1: Hand-to-Mouth Status Incidence & Persistence

	Transition Matrix		population share
	not-H2M$_{t+2}$	H2M$_{t+2}$	
not-H2M$_t$	.90	.10	80%
H2M$_t$	.34	.66	20%

Notes: H2M$_t$ is an indicator for households reporting non-positive net worth in a given survey wave t. Details on the definition of net worth can be found in section 2.1.1.

2.1.1 Data

The empirical analysis is conducted using the Panel Study of Income Dynamics (PSID) biennial surveys from 1999 to 2021. Since 1999, the PSID started measuring wealth and improved its expenditure data collection beyond food and housing. Specifically, my analysis centers on non-durable consumption expenditure, C, total household income, Y, and net worth, W. A set of controls X is used to partial out from the main interest variables the component explained by observable demographics. Flow variable values, although reported biennially, refer to their yearly realizations.

Consumption expenditure includes non-durable categories such as food, housing, childcare, healthcare, clothing, trips, recreation, house furnishing and repairs, utilities, vehicle payments and repairs, transportation, and schooling. As a measure of household income, I use the sum of all taxable income, transfer income, and social security income from the reference person, spouse (if present), and other family unit members (if present). For homeowners, housing costs are computed as 6 percent of the home value; on the other hand, I include the same value, subtracting associated mortgage interest and home insurance, as implicit rent.

In this paper, I identify households being in a low-wealth status with those reporting non-positive net worth, which I label as hand-to-mouth (H2M). Net worth is computed by subtracting liabilities from asset holdings. Assets include both liquid and non-liquid wealth, such as home equity, stocks, checking and savings balances, money market funds, certificates of deposit, treasury bills, and retirement accounts. Liabilities include mortgages, credit cards, medical and legal bills, and student loans. The net value of any business, farm, or vehicles is also added to net worth.

[3]The stationary distribution for H2M households, implied by Table 1 transition matrix, is: $\{.77, .23\}$

The choice of a measure of wealth status that is not relative to household income or earnings deviates from alternative hand-to-mouth-status approaches where H2M status is tipically assigned to household with net worth or liquidity below some arbitrary fraction of their annual earnings. Classifying households as hand-to-mouth, according to either net worth or liquidity, is instrumental in addressing questions related to their response to fiscal stimuli or monetary shocks, thus related to their marginal propensity to consume. This paper focuses on wealth dynamics at the bottom of the household wealth distribution; therefore, I believe the use of the proposed H2M definition using a zero-wealth threshold to classify low-wealth status is more fitting.

The set of controls used throughout the whole empirical analysis, X, features a third-order polynomial in household age, household dummies for education, race, and residential region, household size, and number of children, plus calendar year dummies. Household age, education, and race refer to the household representative person.[4] The three continuous variables of interest – log consumption, log income, and wealth – are residuals obtained after partialling out the component explained by X. When estimating the consumption expenditure function, the household representative person's health status is added to X. Therefore, the measure of expenditure risk identified in this paper, as confirmed by the cross-correlation analysis in 3.2 with respect to non-durable expenditure sub-categories, does not include health risk.

To estimate the role of idiosyncratic risk realization in determining transitions toward H2M, I run a linear probability model with the dependent variable $\text{FALL}_{i,t}$, representing a transition from positive to non-positive net worth between $t-2$ and t. The binary regression model uses lagged wealth, changes in marital status, and home ownership together with ΔX, a first-difference transformation of X, as controls. I refer to this set of controls as ΔX^{FALL}.

My sample draws only on PSID's nationally representative sample. The household representative person's age is restricted to between 25 and 64. Following Aguiar et al. (2024), households with less than $2000 in income or reporting food or housing expenditure above 90% or below 5% of total expenditure are dropped. All nominal values are converted to CPI-deflated 2001-$ using the Federal Reserve Economic Data (FRED) Consumer Price Index for the United States series.

2.1.2 Income risk identification

Income risk modeling, assumptions, and identification in this section closely follow Kaplan and Violante (2010), which builds on a simplified version of the seminal work by Blundell et al. (2008). Let y represent household income residuals, obtained after partialling out

[4] In older PSID waves, the representative person was referred to as the household head.

observable household characteristics[5] from observed log-income realizations $\tilde{y}$. Income residuals are then decomposed into a permanent component z^y and a transitory component ε^y. Given the sample timing characteristics, one period in the stochastic model corresponds to two years in the data.

$$\begin{aligned}\tilde{y}_{i,t} &= \Xi(X_{i,t}) + y_{i,t}, \quad y_{i,t} \perp\!\!\!\perp X_{i,t} \\ y_{i,t} &= z^y_{i,t} + \varepsilon^y_{i,t} \\ z^y_{i,t} &= z^y_{i,t-1} + \eta^y_{i,t} \\ \therefore \Delta y_{i,t} &= \eta^y_{i,t} + \Delta\varepsilon^y_{i,t} \\ \left(\varepsilon^y_{i,t}, \eta^y_{i,t}\right)' &\overset{\text{iid}}{\sim} \mathcal{N}\left(0_2, \quad \text{diag}\left(\sigma^2_{\varepsilon^y}, \sigma^2_{\eta^y}\right)\right)\end{aligned} \tag{1}$$

Given (1), income residual growth captures a noisy measure of income risk components and, under a pair of orthogonality conditions[6] on consumption growth $\Delta c_{i,t}$, there exists a set of measurable functions of income histories[7] $f^{\eta,\,\varepsilon}_t(y_i)$ with the following properties:

$$\begin{aligned}\text{cov}\left(\Delta c_{i,t},\ \eta^y_{i,t}\right) &= \text{cov}\left(\Delta c_{i,t},\ f^\eta_t(y_i)\right) \\ \text{var}\left(\eta^y_{i,t}\right) &= \text{cov}\left(\Delta y_{i,t},\ f^\eta_t(y_i)\right) \\ \text{cov}\left(\Delta c_{i,t},\ \varepsilon^y_{i,t}\right) &= \text{cov}\left(\Delta c_{i,t},\ f^\varepsilon_t(y_i)\right) \\ \text{var}\left(\varepsilon^y_{i,t}\right) &= \text{cov}\left(\Delta y_{i,t},\ f^\varepsilon_t(y_i)\right)\end{aligned} \tag{2}$$

The objective of the empirical strategy in (2) is to identify the relationship between permanent $\eta^y_{i,t}$ and transitory $\varepsilon^y_{i,t}$ income innovations and consumption growth $\Delta c_{i,t}$; to assess the pass-through of stochastic income innovations on consumption dynamics. Therefore, using the appropriate functions of income residuals y_i, we can construct pass-through coefficients $\beta_{\Delta c|\bigstar^y}$, for $\bigstar \in \{\eta,\ \varepsilon\}$, as:

$$\begin{aligned}\beta_{\Delta c|\eta^y} &= \frac{\text{cov}\left(\Delta c_{i,t}, \sum_{j=-1}^{1} \Delta y_{i,t+j}\right)}{\text{cov}\left(\Delta y_{i,t},\ \sum_{j=-1}^{1} \Delta y_{i,t+j}\right)} \\ \beta_{\Delta c|\varepsilon^y} &= \frac{\text{cov}\left(\Delta c_{i,t}, \Delta y_{i,t+1}\right)}{\text{cov}\left(\Delta y_{i,t},\ \Delta y_{i,t+1}\right)}\end{aligned} \tag{3}$$

Coefficients $\beta_{\Delta c|\bigstar^y}$ will identify how permanent or transitory income risk affects con-

[5]The Ξ function is estimated using a linear regression model, following (Blundell et al., 2008); the vector of observables $X_{i,t}$ contains a third-order polynomial in household age, household dummies for education, race, and residential region, household size, and number of children, plus calendar year dummies.

[6]Consumption growth, $\Delta c_{i,t}$, must be orthogonal to given leads and lags of income stochastic component innovations. Given the income process detailed in (1), the following No-Foresight and Short-Memory conditions need to hold:

$$\begin{aligned}\text{(NF)}: &\quad \text{cov}(\Delta c_{i,t},\ \eta^y_{i,t+1}) = \text{cov}(\Delta c_{i,t},\ \varepsilon^y_{i,t+1}) = 0 \\ \text{(SM)}: &\quad \text{cov}(\Delta c_{i,t},\ \eta^y_{i,t-1}) = \text{cov}(\Delta c_{i,t},\ \varepsilon^y_{i,t-2}) = 0.\end{aligned}$$

[7]The set of measurable functions depends on the stochastic process specified for income residuals, given stated assumptions: $f^\eta_t(y_i) = \Delta y_{i,t+1} + \Delta y_{i,t} + \Delta y_{i,t-1}$ and $f^\varepsilon_t(y_i) = \Delta y_{i,t+1}$ deliver the desired result.

sumption dynamics $\Delta c_{i,t}$. The intuition for this empirical specificaiton follows a two-step instrumental variable framework where $f_t^{\bigstar^y}(y_i)$ are employed as instruments to extract, from observed (residual) income dynamics, the $\bigstar$ income-process-component effect on the dependent variable.

To measure the marginal effect of idiosyncratic risks on ZW transition probabilities, I run a linear probability model that uses the same two-step instrumental-variable framework from the pass-through coefficient identification analysis. Adding lagged wealth, changes in marital status, and changes in home ownership together with ΔX, a first-difference transformation of X, as controls. Therefore, for income risk, the regressors are the projections of income growth on the appropriate instrument for each component. The coefficients associated with these regressors are labeled as $\beta_{\text{FALL}|\bigstar^y}$, for $\bigstar \in \{\eta,\ \varepsilon\}$.

2.1.3 Income risk and transitioning into hand-to-mouth status

As predicted by the theory, income risk affects consumption dynamics. Table 2 reports estimated pass-through coefficients from (3). The difference in the number of observations that enter permanent versus transitory risk estimates depends on the different identification data requirements between the two risk sources. Results are in line with estimations from the literature.[8] As expected, permanent shocks pass through consumption expenditure more than transitory ones, and the explanatory power of permanent shocks, as measured by the reported R^2s, is ten-fold that of transitory ones.

Table 2: Income Risk Effects: PSID vs. Benchmark Economy

	PSID			**Benchmark Economy**		
	Δc	$\Delta W/Y$	FALL	Δc	$\Delta W/Y$	FALL
y^P	0.04***	0.07***	-0.010***	0.05***	0.01***	0.001
	(0.004)	(0.011)	(0.0026)	(0.0003)	(0.0004)	(0.0007)
y^T	0.02***	0.09***	-0.007***	0.03***	0.05***	-0.016***
	(0.004)	(0.011)	(0.0025)	(0.003)	(0.004)	(0.007)
R^2	0.10	0.02	0.04	0.36	0.55	0.09
households		3,012			10,000	
hh waves		10,656			100,000	

Notes: Estimates for permanent (y^P) and transitory (y^T) income risk effects on consumption growth (Δc), wealth accumulation ($\Delta W/Y$), and H2M transitions (FALL). FALL is estimated using a linear probability model (LPM). Clustered standard errors in parentheses. Benchmark Economy results come from the calibrated model. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 2 reports the pass-through coefficient estimates for both income risk components for the binary $\text{FALL}_{i,t}$ event. The causal relationship between income risk and fall events is

[8] See Crawley and Theloudis (2024) for a detailed summary of structural-methods estimates.

present in the whole sample, as shown in the first column, and is robust to several cuts of the data aimed to isolate households who are more likely to experience a fall event.

2.2 Hand-to-Mouth Transitions in the Canonical Model

Here, I review some of the predictions of the canonical consumption-savings model under idiosyncratic income risk and incomplete markets augmented with ex-ante time-preferences heterogeneity, following Aguiar et al. (2024). Particularly, I focus on consumption dynamics and transitions towards H2M status, extending the analysis in Kaplan and Violante (2010) to an economy with household types tied to preference heterogeneity.

The benchmark economy (BE) used to represent the predictions of the canonical model is detailed in Section 4.1. To summarize, BE features both ex-ante and ex-post heterogeneity among households. Ex-ante heterogeneity is represented by an independent bivariate distribution of income level fixed effects and discount factors. Ex-post heterogeneity is generated by idiosyncratic permanent and transitory income shock realizations and the consequently diverse income histories.

As already shown by Kaplan and Violante (2010), the heterogeneous agent consumption-saving model with idiosyncratic income risk under incomplete markets can generate partial insurance levels, as measured by the pass-through coefficient introduced by Blundell et al. (2008), similar to the data when we allow for household borrowing. Results in Table 2 confirm their findings in an economy with ex-ante heterogeneity in time preferences. Moreover, as the permanent income hypothesis postulates, permanent shocks demand a larger pass-through to consumption than transitory ones.

What's the story the BE tells about falling into a low net worth position? The distribution of time-preference across households is calibrated to match the share of H2M households in the economy. The mechanism acts by creating heterogeneous optimal wealth-to-income ratios; given the assigned time-preferences, some households are content with holding little-to-no wealth, as these target dynamics adjust along the life-cycle. Therefore, as shown by Table 2, ZW transitions are dominated by transitory income risk realizations. For those households with target wealth-to-income ratios close to the threshold, fluctuations in transitory income greatly influence which side of the threshold they stand, as pointed out by the last column of Table 2.

Summarizing, in the canonical framework, permanent income shocks do not significantly affect the financial vulnerability of households at the bottom of the wealth distribution. The optimal dynamic response dictated by the Euler equation demands an appropriate cut in household expenditure. This relationship between income risk pass-through and wealth can be seen clearly in Figure 1. As their wealth decreases, households' consumption responses to income shocks become larger. The effect is present in both components of income risk but is much stronger for transitory shocks, whose pass-through to consumption approaches the level of permanent shocks as household wealth decreases.

Figure 1: Income Pass-Through Coefficients by Wealth Decile in BE

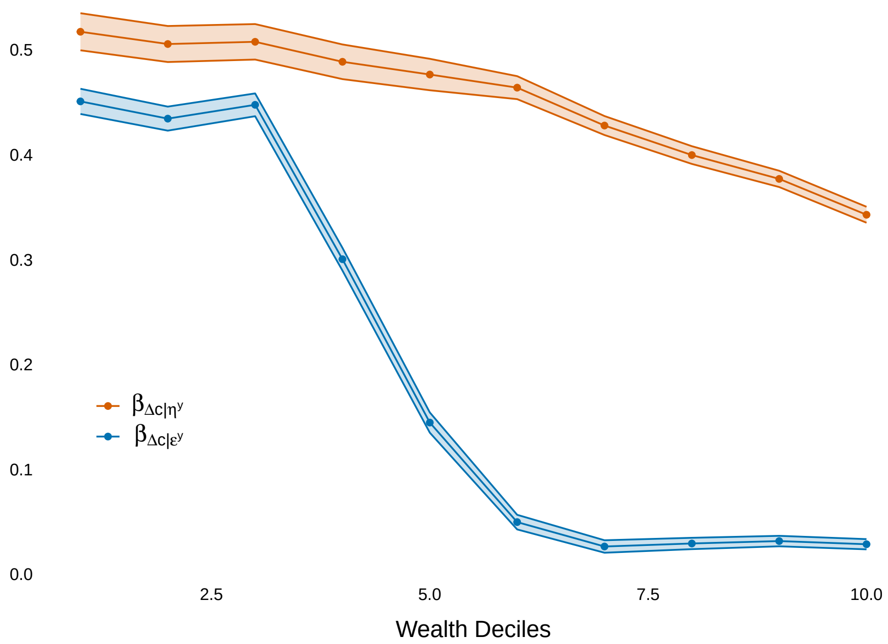

Notes: Sample cuts are generated using deciles of the beginning-of-period net worth variable. The pass-through coefficients for permanent and transitory income shocks are then computed within each data cut.

3 Expenditure Risk

How do households deal with unexpected expenses? Is there any evidence of persistence across periods? What are the consequences for wealth dynamics? I use income risk measures for permanent and transitory shocks, identified using observed histories of income dynamics as detailed in Section 2.1.2, together with a measure of fixed income level and other observable demographics to estimate a flexible expenditure model. The arguments of the consumption function are modeled after the same canonical consumption-savings framework adopted in Section 2. I use histories of changes in the deviation of observed expenditure from its predicted counterpart to identify expenditure shocks.

To identify expenditure risk alongside income risk, I build a stochastic-process dependent identification strategy that hinges on two core assumptions: (i) mean independence between permanent income and expenditure shock innovations, and (ii) household preferences[9] consistent with a consumption policy function featuring multiplicative independent

[9]For instance, CRRA preferences with homogeneous curvature across consumption goods over time and states of the world would imply (5).

arguments with respect to income and expenditure shock realizations. This framework allows for a simple two-step procedure, where expenditure risk is measured using consumption residuals obtained by subtracting from observed expenditure its predictable component derived from income types, asset holdings, income shocks, and household demographics.

3.1 Empirical Consumption Model

Consider a consumption policy function $\tilde{G}(\cdot)$ featuring both income $(\eta^y,\ \varepsilon^y)$ and expenditure η^c risks alongside household income type τ, log available resources[10] z, and observable demographics X as arguments; expressed for log consumption as:

$$c_{i,t} = \tilde{G}\left(\tau_i, \eta^y_{i,t}, \varepsilon^y_{i,t}, \eta^c_{i,t}, z_{i,t}, X_{i,t}\right) + \epsilon_{i,t}, \quad \epsilon_{i,t} \perp\!\!\!\perp \eta^y_{i,t}, \eta^c_{i,t}, z_{i,t}, X_{i,t} \tag{4}$$

Assume household preferences allow for the decomposition of observed expenditure into two components: what can be predicted using the income risk framework $\hat{c}$ plus a residual part ξ.

$$c_{i,t} = G\left(\tau_i, \eta^y_{i,t}, \varepsilon^y_{i,t}, z_{i,t}, X_{i,t}\right) + h\left(\eta^c_{i,t}\right) + \epsilon_{i,t}, \quad \eta^y_{i,t} \perp\!\!\!\perp \eta^c_{i,t} \tag{5}$$

To estimate $G(\cdot)$, I run a supervised machine learning[11] algorithm allowing for unstructured non-linearity along and across all argument dimensions. Once $G(\cdot)$ is estimated using two-fold cross-validation, predicted expenditure $\hat{c}$ is computed using observed realizations of persistent income shocks, available cash on hand, and household demographics. Figure A.1 in Appendix A shows the marginal effect of cash on hand and permanent income risk on log consumption predictions, obtained by computing a partial dependence function.

The estimated log consumption function is concave over log available resources and increasing in permanent income innovations. Figure A.2 displays the contribution of each argument used in the random forest regression framework to predict log consumption. Household income types, cash on hand, family size, and income risk innovations are leading determinants of prediction accuracy, as measured by the decrease in node impurity. Node impurity quantifies the homogeneity of the target variable within the subsets created by a split during the creation of each decision tree. Lower impurity means that the subset is more homogeneous (i.e., the target variable values are more similar).

[10] Cash on hand $Z_{i,t}$ is measured as the sum of household predetermined net worth $W_{i,t-2}$ and current family income $Y_{i,t}$.

[11] The chosen random forest algorithm outperforms both a simple OLS linear framework and a penalized LASSO classifier with five-degree spline interactions among all function arguments.

3.1.1 Unexpected expenses are not i.i.d.

Based on an income-risk driven theory of consumption and savings over the life cycle, any deviations from the household-policy-function's determined consumption expenditure are expected to be attributable solely to noise. However, I find that the residual component of log-consumption expenditure exhibits distinct patterns in both levels and its first difference.

$$\begin{aligned} c_{i,t} &= \hat{c}_{i,t} + \xi_{i,t} \\ \hat{c}_{i,t} &\equiv G\left(\tau_i, \eta^y_{i,t}, z_{i,t}, X^c_{i,t}\right) \\ \xi_{i,t} &\equiv h\left(\eta^c_{i,t}\right) + \epsilon_{i,t} \end{aligned} \tag{6}$$

Following (6), observed log-consumption expenditure, c, is decomposed into two components: one, $\hat{c}$, represents the best prediction we have of consumption expenditure based on household demographics and income dynamics, while the residual component ξ captures those expenses that go unpredicted using the model of idiosyncratic income risk from section 2.1.2.

Table 3: Stochastic Properties of log-Expenditure Components

	(a) Levels			(b) Growth Rates	
	autocorr.	variance		autocorr.	variance
$c_{i,t}$	.80	.37	$\Delta c_{i,t}$	-.34	.25
$\hat{c}_{i,t}$	.86	.18	$\Delta \hat{c}_{i,t}$	-.24	.05
$\xi_{i,t}$	.55	.14	$\Delta \xi_{i,t}$	-.40	.13

Notes: Log expenditure is decomposed into two components, one predictable and one unpredicted, according to equations (6). Correlation with its first lag and variance are then computed for all three objects both in levels and in first differences.

Although predicted expenditure captures the observed autocorrelation in consumption, its residuals ξ display autocorrelation as well, indicating that a household identified as an over-spender in a particular wave is more likely to maintain the same status in the subsequent one. As already discussed in Fernández-Villaverde and Krueger (2007), and documented in the first row of panel (a) in Table 3, non-durable consumption expenditure is serially correlated across periods. Its predicted component $\hat{c}$ is serially correlated as well, consistent with serially correlated income risk series used as an estimation argument, but cannot capture all of it. In fact, after accounting for the degree of serial correlation we can expect based on household characteristics and permanent income risk realizations, the residual component ξ still exhibits a non-negligible level of autocorrelation. This hints at the presence of some persistent component in expenditure[12] that is missed by the standard

[12] Hall and Mishkin (1982) estimate a moving average component to non-durable expenditure in order to fit the data.

income-risk model.

Furthermore, if we look at growth rates of consumption expenditure and its components, we can see in Table 3, panel (b), how consumption growth is negatively autocorrelated. Part of it can be explained by the standard income-risk driven framework, but the rest remains unexplained. Moreover, the growth rates of the unexplained component present higher volatility, pointing, as already postulated in (Blundell et al., 2008), again at the presence of some latent factor relevant for the household when formulating their optimal consumption-savings dynamic allocation.

Figure 2: Income and Consumption Residuals Impulse Response Functions

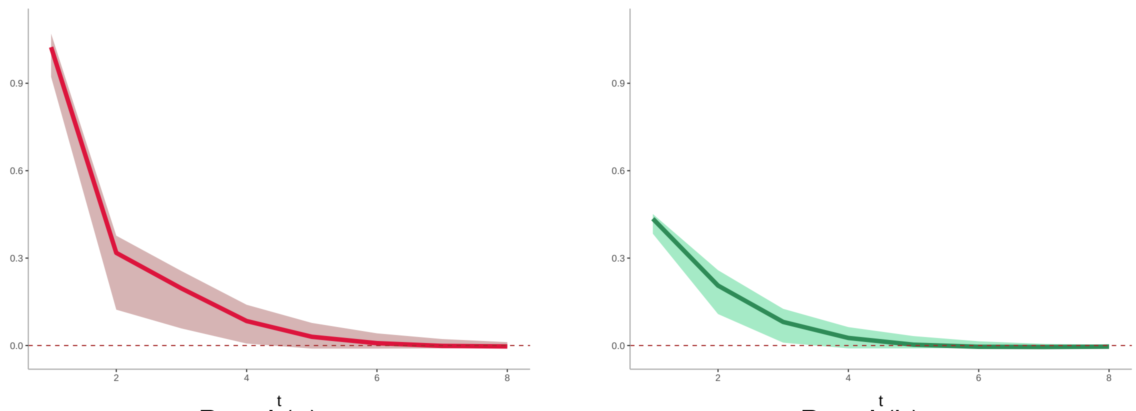

Panel (a)
income, corr$_{t,t+2}$ = 1

Panel (b)
consumption, corr$_{t,t+2}$ = 0.54

Notes: Side by side comparison of impulse response functions obtained by estimating a simultaneous panel-VAR regression of income and consumption residuals on its lags. One-period lags on the horizontal axis correspond to a two-year period in the survey data.

Finally, to complete the exploration of consumption residuals' stochastic properties, Figure 2 displays side-by-side the impulse response functions for income and consumption residuals computed by estimating a simultaneous Panel-VAR regression for both quantities. Autocorrelation estimates seem to exclude a random-walk modeling assumption for consumption residuals. Therefore, in the next section, I will estimate an autoregressive parameter for consumption residuals to be able to compute a persistence-adjusted first difference measure, as shown in equation set (7), necessary to identify persistent innovations to non-durable expenditure.

3.1.2 Expenditure risk identification

Based on their stochastic properties and generalizing on the income risk identification framework, consumption residuals ξ can be decomposed into a persistent component z^c and a transitory component ε^c. Following evidence about the stochastic properties of consumption residuals, an AR(1) assumption fits the data better than the martingale assumption used for income residuals. Therefore, to identify the relationship between consumption risk innovations η^c and any variable of interest, I will use a persistence-adjusted first-difference operator $\Delta_\rho x_t \equiv x_t - \rho^c x_{t-1}$.

$$\begin{aligned} \xi_{i,t} &= z_{i,t}^c + \varepsilon_{i,t}^c \\ z_{i,t}^c &= \rho^c \cdot z_{i,t-1}^c + \eta_{i,t}^c \\ \therefore \Delta_\rho \xi_{i,t} &= \eta_{i,t}^c + \Delta_\rho \varepsilon_{i,t}^c \\ \begin{pmatrix} \varepsilon_{i,t}^c \\ \eta_{i,t}^c \end{pmatrix} &\overset{\text{iid}}{\sim} \mathcal{N}\left(0_2, \quad \text{diag}\left(\sigma_{\varepsilon^c}^2, \sigma_{\eta^c}^2\right)\right) \end{aligned} \tag{7}$$

According to the stochastic modeling outlined in (7), the variance-covariance of persistent expenditure risk with respect to consumption growth is identified as:

$$\begin{aligned} \text{cov}\left(\Delta c_{i,t},\ \eta_{i,t}^c\right) &= \text{cov}\left(\Delta c_{i,t},\ g_t^\eta(\rho^c, \xi_i)\right) \\ \text{var}\left(\eta_{i,t}^c\right) &= \text{cov}\left(\Delta_\rho \xi_{i,t},\ g_t^\eta(\rho^c, \xi_i)\right) \end{aligned} \tag{8}$$

with $g_t^\eta(\rho^c, \xi_i) = \Delta_\rho \xi_{i,t+1} + \Delta_\rho \xi_{i,t} + \Delta_\rho \xi_{i,t-1}$. Therefore, we can construct the expenditure risk pass-through coefficient $\beta_{x|\eta^c}$ as:

$$\beta_{\Delta c|\eta^c} = \frac{\text{cov}\left(\Delta c_{i,t},\ \sum_{j=-1}^{1} \Delta_\rho \xi_{i,t+j}\right)}{\text{cov}\left(\Delta_\rho \xi_{i,t},\ \sum_{j=-1}^{1} \Delta_\rho \xi_{i,t+j}\right)} \tag{9}$$

The $\beta_{\Delta c_{i,t}|\eta^c}$ parameter measures the effect of the persistent component of unexplained expenditure on consumption growth $\Delta c_{i,t}$, controlling for all arguments featured in the estimated consumption policy function, especially permanent and transitory income innovations, household income type, and net worth. Beyond the proposed measure of expenditure-risk pass-through, a test for randomness of consumption residuals can be performed by imposing a random-walk modeling assumption, which is rejected. Details of the random-walk-modeling pass-through coefficient and related estimates can be found in Appendix B.

Finally, following the linear probability regression framework used to assess the marginal effect of income risk on ZW transition probabilities, as described in section 2.1.2, I instrument persistence-adjusted consumption-residuals growth rates $\Delta_\rho \xi_{i,t}$ with the proposed $g_t^\eta(\rho^c,\ \xi_i)$ instrument. The parameter associated with the described IV regression for expenditure risk is labeled as $\beta_{\text{FALL}|\eta^c}$.

3.2 Drivers of expenditure risk among non-durable categories

To investigate the drivers of expenditure risk, I look at the correlation between non-durable consumption sub-categories and measured persistent expenditure risk. The idea is to examine changes in expenditure shares in each consumption category and their comovement with expenditure risk; a positive correlation would indicate an expenditure-category which increases its expenditure share when expenditure risk hits. On the other hand, consumption categories that are not driving observed expenditure risk will mechanically show a weakly negative correlation between their expenditure shares and expenditure risk realizations.

Table 4 shows the results of this analysis based on the non-durable consumption categories available in the PSID. Only a subset of categories, together accounting on average for 23% of household expenditure, drive expenditure risk; it is expenditure in transportation, education, and maintenance (both house and car repairs) that drives measured expenditure risk.

Table 4: Consumption Categories Correlation with Expenditure Risk

			non-durable categories				
cross correlation	share	st.dev	negative corr.	positive corr.	cross correlation	share	st.dev
−0.07	0.13	0.06	food	transportation	0.03	0.15	0.08
−0.05	0.32	0.10	housing	school	0.06	0.03	0.07
−0.02	0.01	0.03	childcare	house repairs	0.03	0.02	0.04
−0.01	0.05	0.05	healthcare	car down pay	0.02	0.01	0.04
−0.05	0.02	0.02	clothing	car repairs	0.03	0.02	0.04
−0.02	0.03	0.04	trips				
−0.04	0.02	0.02	recreation				
−0.02	0.02	0.03	house furnish				
−0.06	0.08	0.04	all utilities				
−0.00	0.09	0.05	oth. car pay				
−0.06	0.77	0.77	c_1	c_2	0.06	0.23	0.27

Notes: Non-durable consumption categories have been split into c_1 and c_2 columns according to the sign of their cross-sectional correlation with expenditure shocks.

The addition of household representative's health-status information controls, in a way, for healthcare expenditure. Thus, it is not surprising that the identified persistent expenditure risk measure does not correlate with healthcare spending. Moreover, among the spending categories with the strongest negative correlation, we can list food, housing, clothing, and utilities, which can be considered among those consumption commitments as theorized by Chetty and Szeidl (2007).

The analysis of the drivers of expenditure risk not only allows for a more precise understanding of the phenomenon, it also provides instructions for any framework aimed to model of expenditure risk. Households subjected to persistent shocks affecting a minor but relevant share of their consumption basket have two degrees of freedom to react to these shocks: adjusting their expenditure in the rest of their consumption basket or their savings.

3.3 Empirical Results: Household Dynamics vis-à-vis Expenditure Risk

How "big" is expenditure risk? Expenditure risk, estimated as the projection of persistence-adjusted consumption-residuals growth between $t-2$ and t on its $t-4$ to $t+2$ counterpart, accounts for 35% of cross-sectional variability in observed log consumption, measured by the explained sum of squares of the dependent variable in a linear regression. A relative measure of expenditure risk can be produced by comparing its standard deviation with those of the other two sources of idiosyncratic risk identified in this paper. Expenditure risk volatility of 0.16 sits between permanent income shocks volatility, 0.10, and transitory shocks volatility, 0.25. All three quantities are measured in log deviations.

Table 5: Income and Expenditure Risk Effects: PSID

	Δc	$\Delta W/Y$	Fall
y^P	0.04***	0.07***	$-$0.011***
	(0.004)	(0.011)	(0.0025)
y^T	0.02***	0.09***	$-$0.007***
	(0.003)	(0.004)	(0.0024)
ER	0.03***	$-$ 0.03***	0.008***
	(0.003)	(0.004)	(0.0024)
R^2	0.11	0.02	0.04
households		3,012	
hh waves		10,656	

Notes: Estimates for permanent (y^P), transitory (y^T), and expenditure risk (ER) effects on consumption growth (Δc), wealth accumulation ($\Delta W/Y$), and H2M transitions (FALL). FALL is estimated using a linear probability model (LPM). Standard errors in parentheses. Benchmark Economy results come from the calibrated model. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Unpredicted household expenditures not only exhibit persistence, as discussed in Section 3.1.1, but their persistent component (ER) significantly affects consumption growth, wealth accumulation, and the probability of falling into H2M episodes. Table 5 presents the average marginal effects of permanent income risk, transitory income risk, and expenditure risk on these three key household dynamics.

The estimated coefficients reveal economically meaningful effects across all three outcomes. Permanent and transitory income shocks have positive and significant effects on consumption growth, consistent with standard theory, although their impacts on wealth accumulation and ZW transitions differ. Transitory income risk negatively affects wealth accumulation, suggesting households smooth temporary income shocks through their wealth, whereas permanent income risk has a positive correlation with wealth changes. Expenditure risk (ER), capturing persistent unexpected expenses, shows a positive effect on the probability of falling into ZW and negatively affects household wealth accumulation.

The magnitudes of these effects underscore the importance of ER: a one-standard-deviation increase in expenditure risk increases the probability of transitioning into zero or negative wealth by 0.8 percentage points. Given an unconditional fall probability of approximately 10%, this represents a substantial marginal increase. These findings highlight a crucial distinction between expenditure and income risks, emphasizing the unique and substantial impact of ER on financial vulnerability among low-wealth households.

The predictive importance of the three sources of idiosyncratic risk differs markedly. Expenditure risk emerges as the most significant predictor of transitions into hand-to-mouth (H2M) status, accounting for 6.7% of these transitions. In contrast, transitory income risk explains only 2.3%, while permanent income risk contributes negligibly. Figure 3 summarizes these relative magnitudes visually, clearly underscoring the dominant predictive role of expenditure risk. A comprehensive variable-importance analysis from a random forest classification model, which further corroborates these findings, is reported in Appendix C.

Figure 3: PSID Fall-Event Prediction, Predictors Relative Importance

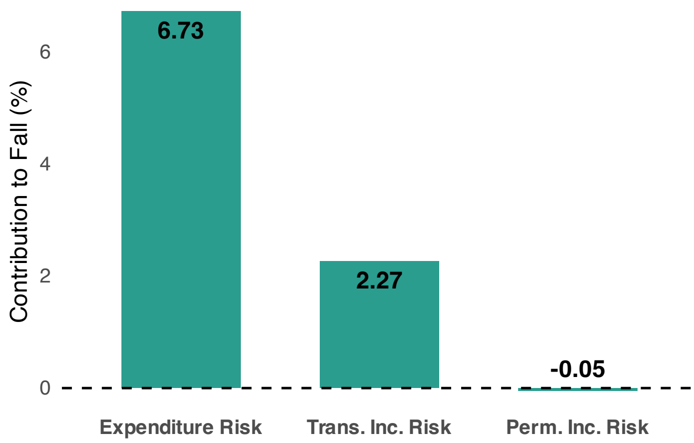

Notes: The bars measure the total decrease in node impurities from splitting on each variable, averaged over all trees, with respect to out-of-bag predictors classification. The node impurity is measured by the Gini impurity index, which measures the average misclassification probability.

Taking stock, I have jointly identified three distinct sources of risk shaping household consumption and wealth dynamics. The novel channel here, expenditure risk, explains 35% of the cross-sectional variability in observed log consumption. Additionally, the persistent component of unpredicted expenses significantly affects consumption dynamics and predicts H2M episodes, outperforming income risk measures. These results underscore the critical role of expenditure risk in financial vulnerability and highlight the need to incorporate it into models of household wealth transitions.

4 A Model of Income and Expenditure Risk with Time-Preference Heterogeneity Across Households

I will start with the description and calibration of the benchmark economy with idiosyncratic income risk and fixed preference heterogeneity to match the life cycle dynamics of asset accumulation. Aguiar et al. (2024) show how household heterogeneity in the discount factor can account for the observed fraction of households holding little wealth in the data, and I confirm it also delivers low-wealth status persistence across survey waves. Afterward, I will introduce the model ingredient for the expenditure-risk economy, which will be calibrated on the same set of targets as the benchmark economy plus those needed to discipline the added ingredient.

4.1 Benchmark Economy: Income Risk and Household Types

Since I focus on net worth, one asset is enough to characterize household wealth dynamics around the zero-wealth threshold. I will thus use a one-asset life-cycle adaptation of Aguiar et al. (2024) economy as the benchmark economy (henceforth BE) featuring preference heterogeneity by means of the discount factor. Households live $J^D = 60$ years, as workers until $J^R = 40$ and as retirees afterward. Agent i values their stream of non-durable consumption from age j onwards according to

$$
\begin{aligned}
V_{i,j}\left(x_{i,j}, \eta_{i,j}\right) &= \max_{a_{i,j+1} \in [-b,\ x_{i,j}]} u_j(x_{i,j} - a_{i,j+1}) + \beta_i \mathbb{E}_j V_{i,j+1}\left(x_{i,j+1}, \eta_{i,j+1}\right) \\
x_{i,j+1} &= R \cdot a_{i,j+1} + y_{i,j+1} \\
y_{i,j} &= \begin{cases} \exp\left(\mu_j + \tau_i + \eta_{i,j}\right), & j \leq J^R; \quad \text{at } j=1,\ \tau_i \overset{\text{iid}}{\underset{i}{\sim}} \mathcal{N}(0, \sigma_\tau^2) \\ rr_{y_{i,J^R}} \cdot y_{i,J^R}, & j > J^R \end{cases} \\
\eta_{i,j+1} &= \eta_{i,j} + \nu_{i,j}, \qquad \nu_{i,j} \overset{\text{iid}}{\underset{i,j}{\sim}} \mathcal{N}(0, \sigma_\nu^2) \\
u_j(x) &= n_j \frac{(x/n_j)^{1-\gamma} - 1}{1-\gamma}
\end{aligned}
\tag{10}
$$

where x is net worth, η is the persistent income level, $b \geq 0$ is the borrowing limit, a is savings, β_i is the household-specific discount factor drawn from a Gaussian-mixture distribution discussed in the next section, R is the economy gross rate of return on savings, μ_j is a deterministic age income profile, τ_i is the household-specific income fixed effect,

$rr_{y_{i,JR}}$ are income-dependent replacement ratios, and $\nu_{i,j}$ are iid Gaussian innovations to persistent income. Flow utility exhibits constant relative risk aversion γ over per-person consumption computed with family number weights n_j estimated from the data.

4.1.1 Benchmark economy calibration

The benchmark economy features income risk, incomplete markets, and discount factor heterogeneity. Households draw their initial asset holdings, income fixed effect, and discount factor types from mutually independent distributions at the beginning of their working life; permanent income shocks are drawn every period until retirement.

In a life cycle model, target wealth-to-income ratios dynamically determine household asset holdings together with retirement preparation needs. Discount factor heterogeneity will be disciplined by a truncated discretized mixture of two Gaussian distributions, thus its distribution will be calibrated using 5 parameters.

$$F_\beta = \pi \cdot \mathcal{N}(\mu_{\beta,1}, \sigma^2_{\beta,1}) + (1-\pi) \cdot \mathcal{N}(\mu_{\beta,2}, \sigma^2_{\beta,2}) \tag{11}$$

Table 6: Externally Calibrated Parameters

parameter(s)	value	source
μ_j	match life cycle profile	PSID data sample
σ^2_τ	0.3764	Income risk identification
σ^2_η	0.0118	Income risk identification
σ^2_ε	0.0265	Income risk identification
R	1.03	
b	0.2 $\cdot\mathbb{E}$(discounted future income flow)	
γ	2	
rr vec	[.778, .591, .484, .422, .333]	IRS Tables

Notes: Externally calibrated parameters with their assigned value and data source. Missing data source indicates the parameter value has been arbitrarily set.

Table 6 shows the parameters that are determined outside the simulated-method-of-moments (SMM) calibration loop. The deterministic income profile μ_j is estimated from the data, income risk components variance parameters are computed using the observed variance of identified permanent and transitory income shocks from section 2.1.2, variance of the household-specific income fixed effect is calibrated to minimize the distance between simulated and observed income volatility at the beginning of the household working life, annual interest rate is fixed at 3%, households can borrow against 20% of their annualized expected discounted future income flow, the coefficient of relative risk aversion is set to 2, and the lifetime-income dependent replacement ratios are taken from US Internal Revenue Service (IRS) statistics.

Table 7: Internal Calibration Benchmark Economy

$\mathcal{F}_\beta$ parameter	value	target	PSID	BE
$\mu_{\beta,1}$	0.96	share H2M	0.23	0.24
$\sigma^2_{\beta,1}$	0.09	middle working life Q50	1.81	1.54
$\mu_{\beta,2}$	0.94	retirement Q50	5.26	4.87
$\sigma^2_{\beta,2}$	0.11	middle working life Q75	5.58	5.09
π_β	0.61	retirement Q75	14.68	15.74

Notes: Table shows parameter values determining the ex-ante time-preference heterogeneity in the BE. Data target values used for calibration and their model-simulated counterpart are provided.

Internal calibration of the discount factor distribution is obtained by matching the average H2M share in the economy and the net worth holdings of the second and third quartiles of the wealth distribution[13] both in the middle of their working life and upon retirement. Table 7 shows the calibrated values and calibration data fit. The benchmark economy can deliver a wealth accumulation pattern compatible with the data on average and on both sides of the wealth distribution. Model predictions on consumption pass-through coefficients and H2M transitions are discussed in detail in Section 2.2.

The benchmark economy effectively matches the share of H2M households in the economy, both on average and across different ages, as illustrated in Figure 4. The heterogeneity in time preferences, along with the retirement-saving incentives built into the life-cycle model, results in a distribution of the timing when these incentives become increasingly significant with respect to other optimal consumption-saving forces. This creates a smooth, decreasing age profile that aligns closely with the observed data.

4.2 A Theory of Expenditure Risk: Persistent Shocks to Marginal Utility

The empirical analysis in Section 3 overall highlights the role of expenditure risk in explaining wealth dynamics at the bottom of the distribution, specifically around the zero-wealth threshold. Moreover, Section 3.2 indicates there exists a subset of non-durable consumption categories whose expenditure shares respond positively to the identified measure of expenditure shocks. I thus model expenditure risk as persistent innovations to the marginal utility of a fraction of the household's non-durable basket of goods, as expressed in (12).

[13]Wealth is expressed in units of the sample median wealth measured in 2001$, approximately $33,000.

Figure 4: H2M Share Age Profile, Data vs BE

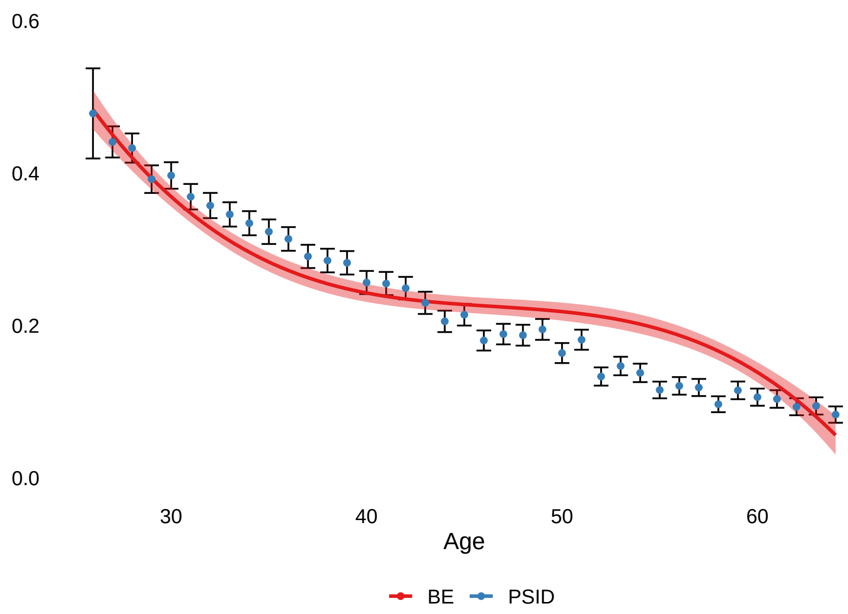

Notes: Hand-to-Mouth age-specific shares are computed both in the PSID and in the Benchmark economy (BE). The age-specific share, together with 95% error bars, are plotted for the PSID. A third order polynomial projection over age, with 95% confidence bands, is plotted for the benchmark economy.

$$\begin{gathered}u_j\left(c_1, c_2, \xi, \phi_0\right)=n_j\left(\phi_0 \cdot \frac{\left(c_1 / n_j\right)^{(1-\gamma)}-1}{1-\gamma}+\exp (\xi) \cdot \frac{\left(c_2 / n_j\right)^{(1-\gamma)}-1}{1-\gamma}\right) \\ \xi_{i, j+1}=\rho^c \xi_{i, j}+\phi_{i, j+1}, \qquad \phi_{i, j+1} \underset{i, j}{\overset{\text{iid}}{\sim}} \mathcal{N}\left(0, \sigma_\phi^2\right)\end{gathered} \tag{12}$$

where ϕ_0 pins down the average relative expenditure size of the riskless portion of the consumption basket. Two parameters characterize the stochastic properties of the marginal utility shocks law of motion, the autocorrelation parameter ρ^c and the innovations variance σ_ϕ^2.

On average, the household spends a given fraction on the risky portion of the basket. As expenditure shocks hit, households will adjust relative consumption expenditure shares and savings as characterized by the intertemporal Euler equation and the intratemporal expenditure shares tradeoff:

$$\left(\frac{c_{1,j}}{n_j}\right)^{-\gamma} \leq R\beta_i\, \mathbb{E}_j\left(\frac{c_{1,j+1}}{n_{j+1}}\right)^{-\gamma} \tag{13}$$

$$\frac{c_{2,j}}{c_{1,j}}=\left(\frac{\exp(\xi_j)}{\phi_0}\right)^{1/\gamma} \tag{14}$$

The introduction of a second good in utility, subject to risky marginal utility realizations, delivers one extra degree of freedom to the household's dynamic optimal allocation problem together with one extra source of idiosyncratic risk to insure against.

4.2.1 Expenditure risk economy calibration

On top of calibrating discount factor distribution parameters, the expenditure risk economy simulation needs to pin down the parameters regulating the magnitude and stochastic properties of expenditure shocks. The average share of expenditure in each consumption basket category is regulated directly by ϕ_0 and can be computed given other preference parameters.[14] Expenditure shocks autocorrelation ρ^c and variance σ^2_ϕ are calibrated to match the serial correlation properties of observed log expenditure, both in levels and in first differences. As discussed in Section 3.1.1, log consumption autocorrelation properties can be explained by the assumed autoregressive expenditure risk stochastic process, thus representing an indirect calibration approach.

Table 8: Internal Calibration Expenditure Risk Economy

	parameter	BE	ER	target	PSID	BE	ER
	$\mu_{\beta,1}$	0.96	0.98	share H2M	0.23	0.24	0.23
	$\sigma^2_{\beta,1}$	0.09	0.005	middle working life Q50	1.81	1.54	1.93
$\mathcal{F}_\beta$	$\mu_{\beta,2}$	0.94	0.28	retirement Q50	5.26	4.87	5.30
	$\sigma^2_{\beta,2}$	0.11	0.05	middle working life Q75	5.58	5.09	5.67
	π_β	0.61	0.80	retirement Q75	14.68	15.74	14.45
	ϕ_0		5.34	$\mathbb{E}c_2$ share	0.23		0.23
ξ	ρ^c		0.31	c autocorrelation	0.80		0.96
	σ^2_ϕ		0.14	Δc autocorrelation	-0.34		-0.27

Notes: Parameter estimates for the ER economy are attached to the information displayed in Table 7. Beyond calibrating the parameters of the discount factor distribution $\mathcal{F}_\beta$, the ER economy needs to pin down marginal-utility-risk coefficients ξ.

In the top panel of Table 8, we can see how matching similar life cycle accumulation dynamics for the two economies generates quite different patterns of discount factor heterogeneity. Figure 5 and Table 9 help us visualize the differences in the discount factor distributions resulting from the calibrated truncated Gaussian mixture in each economy. Unlike the benchmark economy, the expenditure risk calibration generates a bimodal ex-ante household heterogeneity with a smaller fraction of households having a relatively low discount factor.

[14]Taking expectation on the intra-temporal tradeoff from equation (14), we can solve for ϕ_0 as a function of the average expenditure share of good c_2, let's call it α, and of a geometric average of marginal-utility realization weighted by their unconditional probabilities: $\phi_0 = \left(\frac{1-\alpha}{\alpha}\right)^\gamma \cdot \left(\sum_{k=1}^{n_\xi} p_k \exp(\xi_k)^{1/\gamma}\right)^\gamma$ for ξ_k in the support of r.v. ξ pinned down by (12) and discretized using the Rouwenhorst method.

Figure 5: Discount Factor Probability Density Function

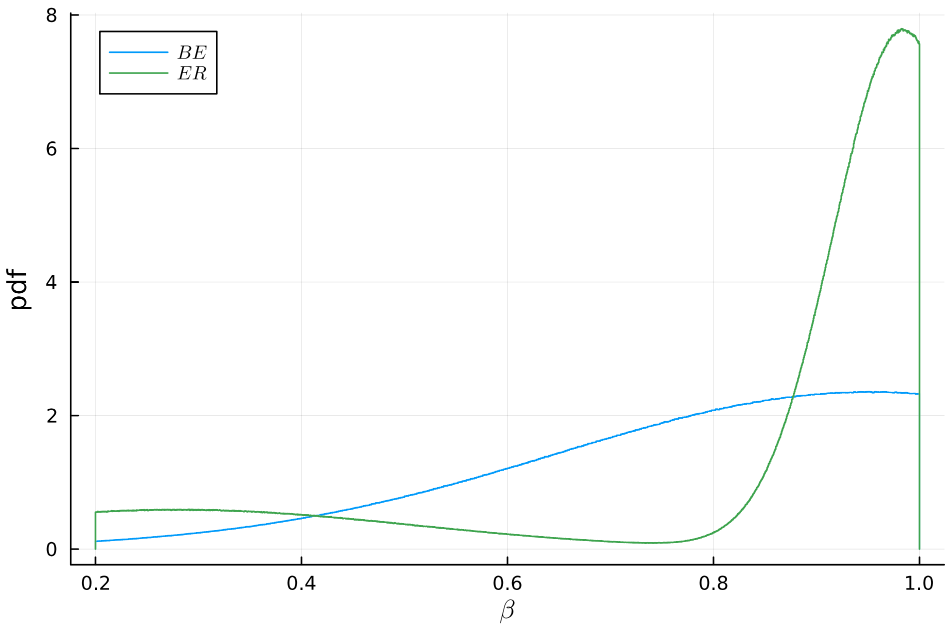

Notes: Using calibrated discount factor distributions $\mathcal{F}_\beta$ from Table 7, this figure compares the probability density function of the two continuous distributions.

Table 9: Discount Factor Quantiles

	quantiles				
	0.1	0.25	0.5	0.75	0.9
BE	0.48	0.63	0.78	0.89	0.96
ER	0.37	0.85	0.93	0.97	0.99

Notes: Using calibrated discount factor distributions $\mathcal{F}_\beta$ from Table 7, this table reports relevant quantiles of the populations of households used in the generation of the simulated economies.

From the bottom panel of Table 8, the calibrated marginal utility shocks exhibit mild autocorrelation and strong variance. The resulting autocorrelation is sufficient to generate a negative autocorrelation in first differences, while the large variance is consistent with Section 3.2 results on the observed volatility of the risky expenditure share. The calibration thus results in an ER economy with a larger share of households displaying stronger intertemporal-savings motives, represented by the big hump in the $\mathcal{F}_\beta$ pdf close to one, while affected by fairly large marginal utility fluctuations.

4.3 Quantitative Results: ER Matters for H2M transitions

First, let's examine how effectively the expenditure risk (ER) economy matches the empirical dynamics captured in the PSID data. Table 10 reports the average marginal effects of permanent income risk, transitory income risk, and expenditure risk on three key household outcomes: consumption growth (Δc), wealth accumulation ($\Delta W/Y$), and transitions into H2M status (FALL). Both the empirical estimates from the PSID and the results generated by the calibrated expenditure risk (ER) model are presented for comparison.

Table 10: Income and Expenditure Risk Effects: PSID vs. ER Economy

	PSID			**ER Economy**		
	Δc	$\Delta W/Y$	FALL	Δc	$\Delta W/Y$	FALL
y^P	0.04***	0.07***	-0.010***	0.05***	0.02***	-0.001**
	(0.004)	(0.011)	(0.0026)	(0.0003)	(0.0004)	(0.0004)
y^T	0.02***	0.09***	-0.007***	0.02***	0.06***	-0.009***
	(0.004)	(0.011)	(0.0025)	(0.003)	(0.004)	(0.004)
exp	0.03***	-0.03***	0.008***	0.02***	-0.02**	0.003***
	(0.004)	(0.009)	(0.0024)	(0.002)	(0.003)	(0.004)
R^2	0.11	0.02	0.04	0.34	0.54	0.07
households		3,012			10,000	
hh waves		10,656			100,000	

Notes: Estimates for permanent (y^P) and transitory (y^T) income risk effects on consumption growth (Δc), wealth accumulation ($\Delta W/Y$), and H2M transitions (FALL). FALL is estimated using a linear probability model (LPM). Clustered standard errors in parentheses. Benchmark Economy results come from the calibrated model. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The estimated magnitudes of the coefficients indicate that the ER economy is broadly successful in replicating the data patterns, though with some quantitative differences. For consumption growth, the empirical and simulated economies exhibit similar responses to both permanent and transitory income risks, but the empirical response to expenditure risk is slightly larger than in the theoretical model. The wealth-to-income responses also align qualitatively, though the empirical estimates show stronger sensitivity to both transitory income and expenditure risk compared to the model's predictions.

The expenditure risk (ER) economy successfully replicates several key empirical patterns observed in household consumption and wealth dynamics. Table 10 shows average marginal effects for the three sources of risk, both in the data and the ER economy. The simulated economy closely matches the empirical consumption responses to permanent and transitory income risks, though the magnitude of the marginal effect of expenditure risk on H2M transitions is somewhat smaller in the model (0.003) compared to the data (0.008).

Figure 6 illustrates the relative predictive contributions of the three idiosyncratic risks for transitions into H2M status, comparing empirical results (PSID) with predictions from the ER economy. Expenditure risk accounts for 10.1% of the predictive power in the ER economy, compared to 6.73% in the data. Similarly, transitory income risk has a larger predictive contribution in the model (7.95%) relative to the empirical estimate (2.27%). Per-

Figure 6: Fall-Event Prediction, Predictors Relative Importance

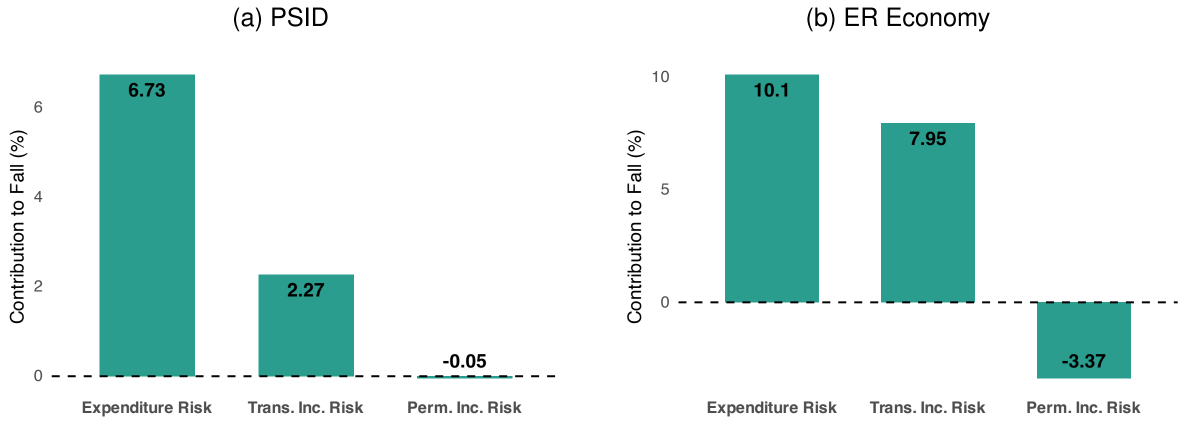

Notes: The bars measure the total decrease in node impurities from splitting on each variable, averaged over all trees, with respect to out-of-bag predictors classification. The node impurity is measured by the Gini impurity index, which measures the average misclassification probability.

manent income risk consistently shows negligible predictive power for H2M transitions in both the ER economy and empirical data.

To gain an intuition for the effects of expenditure risk on household response to permanent income shocks, it is useful to look at the Euler equation expressed in terms of the risky expenditure good c_2, obtained by combining equations (13) and (14):

$$\exp(\xi_j)\left(\frac{c_{2,j}}{n_j}\right)^{-\gamma} \leq R\beta_i\,\mathbb{E}_j\left[\exp(\xi_{j+1})\left(\frac{c_{2,j+1}}{n_{j+1}}\right)^{-\gamma}\right] \tag{15}$$

The multiplicative expectation term on the right-hand side of the equation implies a wedge in the household's abilities to adjust consumption to achieve income risk-dependent optimal consumption growth. It can be clearly seen by employing the marginal utility shocks law of motion from (12) and thus rewriting equation (15):

$$\mathbb{E}_j\left(\frac{c_{2,j+1}}{c_{2,j}}\right)^{-\gamma} - \frac{\exp\left\{(1-\rho^c)\,\xi_j\right\}}{R\beta_i} = -\text{Cov}_j\left(\exp(\phi_{j+1}),\left(\frac{c_{2,j+1}}{c_{2,j}}\right)^{-\gamma}\right) \geq 0 \tag{16}$$

In the absence of shocks to marginal utility, the right-hand side of equation (16) is equal to zero, and we are back to the benchmark economy consumption growth behavior. Expenditure risk, as defined, implies a negative covariance between next period's marginal utility realization and risky expenditure. A positive shock to marginal utility tomorrow generates a higher need to consume the risky good, therefore implying a lower marginal utility growth rate. This positive wedge dampens optimal consumption growth with respect to its benchmark economy counterpart.

5 Conclusions

This paper investigates the empirical and theoretical role of income and expenditure risk together in explaining household wealth dynamics. The analysis is motivated by an intriguing phenomenon: despite economic theory suggesting the importance of precautionary savings, a significant proportion of U.S. households, 22%, hold zero or negative wealth, exposing themselves to consumption fluctuations. Surprisingly, more than half of these households maintain this H2M position over multiple years, indicating a role for consumption dynamics.

Using the Panel Survey of Income Dynamics (PSID), I examine the joint dynamics of non-durable consumption expenditure, total household income, and wealth. To investigate the role of consumption fluctuations in the determination of wealth dynamics, I develop an econometric framework that allows for the identification of expenditure risk independently from income risk. The results demonstrate that both sources of risk matter for wealth dynamics and that, contrary to the standard theory of precautionary savings under idiosyncratic risk, permanent shocks to income can significantly lead to persistent depletion of family wealth and transition into non-positive net worth positions.

Expenditure risk, modeled as shocks to marginal utility hitting a fraction of the household consumption basket, is able to replicate household responses to permanent income variations by creating a wedge in the optimal intertemporal condition for consumption growth and savings dynamics. The covariance of marginal utility and future consumption choices hinders the household's ability to adjust expenditure when facing negative shocks to permanent income.

In summary, my study contributes to two strands of literature: the applied macroeconomic literature investigating consumption inequality and its relationship with income risk, and the theoretical literature on the spending behavior of low-wealth households. By refining empirical approaches and incorporating a flexible non-parametric consumption model, I provide insights into the dynamics of household wealth and the distinct role of consumption risk. The findings shed light on the nature of expenditure risk and its implications for household behavior, offering valuable insights for policymakers and researchers interested in understanding wealth dynamics and consumption inequality.

References

Mark Aguiar, Mark Bils, and Corina Boar. Who are the hand-to-mouth? *The Review of Economic Studies*, page rdae056, 23 May 2024. ISSN 0034-6527. doi: 10.1093/restud/rdae056. URL https://doi.org/10.1093/restud/rdae056. 3, 6, 7, 9, 12, 21

Eric Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. *American Economic Review*, 98(5):1877–1921, 2008. 3, 6, 7, 9, 10, 12, 16

Stéphane Bonhomme, Thibaut Lamadon, and Elena Manresa. Discretizing unobserved heterogeneity. *Econometrica*, 90(2):625–643, 2022. 3

Christopher D. Carroll. Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics*, 112(1):1–55, 1997. doi: 10.1162/003355397555109. URL `https://academic.oup.com/qje/article/112/1/1/1917095`. 3

Raj Chetty and Adam Szeidl. Consumption commitments and risk preferences. *The Quarterly Journal of Economics*, 122(2):831–877, 2007. 4, 18

Edmund Crawley and Alexandros Theloudis. Income shocks and their transmission into consumption. *arXiv preprint arXiv:2404.12214*, 2024. 11

Mariacristina De Nardi, Eric French, and John B Jones. Why do the elderly save? the role of medical expenses. *Journal of political economy*, 118(1):39–75, 2010. 6

Angus Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248, 1991. doi: 10.2307/2938366. URL `https://www.jstor.org/stable/2938366`. 3

Jesús Fernández-Villaverde and Dirk Krueger. Consumption over the life cycle: Facts from consumer expenditure survey data. *The Review of Economics and Statistics*, 89(3): 552–565, 2007. 15

Richard French and John Bailey Jones. On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19:705–721, 2004. 6

Milton Friedman. *A Theory of the Consumption Function*. National Bureau of Economic Research. Princeton University Press, Princeton, NJ, 1957. URL `https://press.princeton.edu/books/hardcover/9780691138862/a-theory-of-the-consumption-function`. 3

Nezih Guner, Christopher Rauh, and Gustavo Ventura. Means-tested transfers in the us: Facts and parametric estimates. Technical Report IZA DP No. 17551, IZA - Institute of Labor Economics, December 2024. URL `https://docs.iza.org/dp17551.pdf`. 5

Robert E. Hall. Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987, 1978. doi: 10.1086/260724. URL `https://www.journals.uchicago.edu/doi/10.1086/260724`. 3

Robert E. Hall and Frederic S. Mishkin. Consumption inequality and partial insurance. *Econometrica*, 50(2):461–481, 1982. URL `http://www.jstor.org/stable/1912638`. 6, 15

Glen R. Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995. 5, 6

Joachim Hubmer, Per Krusell, and Jr Anthony Smith. Sources of us wealth inequality: Past, present, and future. *NBER Macroeconomics Annual*, 35(1):391–455, 2021. 6

Greg Kaplan and Giovanni L Violante. How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87, 2010. 9, 12

Greg Kaplan and Giovanni L Violante. The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, 14(1):747–775, 2022. 3

Per Krusell and Anthony A Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, 106(5):867–896, 1998. 3

Jorge Miranda-Pinto, Daniel Murphy, Kieran James Walsh, and Eric R Young. A model of expenditure shocks. *Cleveland FED, Working Paper*, 2020. 6

Michael G Palumbo. Uncertain medical expenses and precautionary saving near the end of the life cycle. *The Review of Economic Studies*, 66(2):395–421, 1999. 6

A Consumption Model Estimates

In this Appendic section, I report two measures informing on the estimated consumption function from Section 3.1. Figure A.1 shows the partial dependence function for two of the main arguments used in the expenditure function estimation, cash on hand and income shocks.

Figure A.1: Predicted Expenditure, Partial Dependence Function

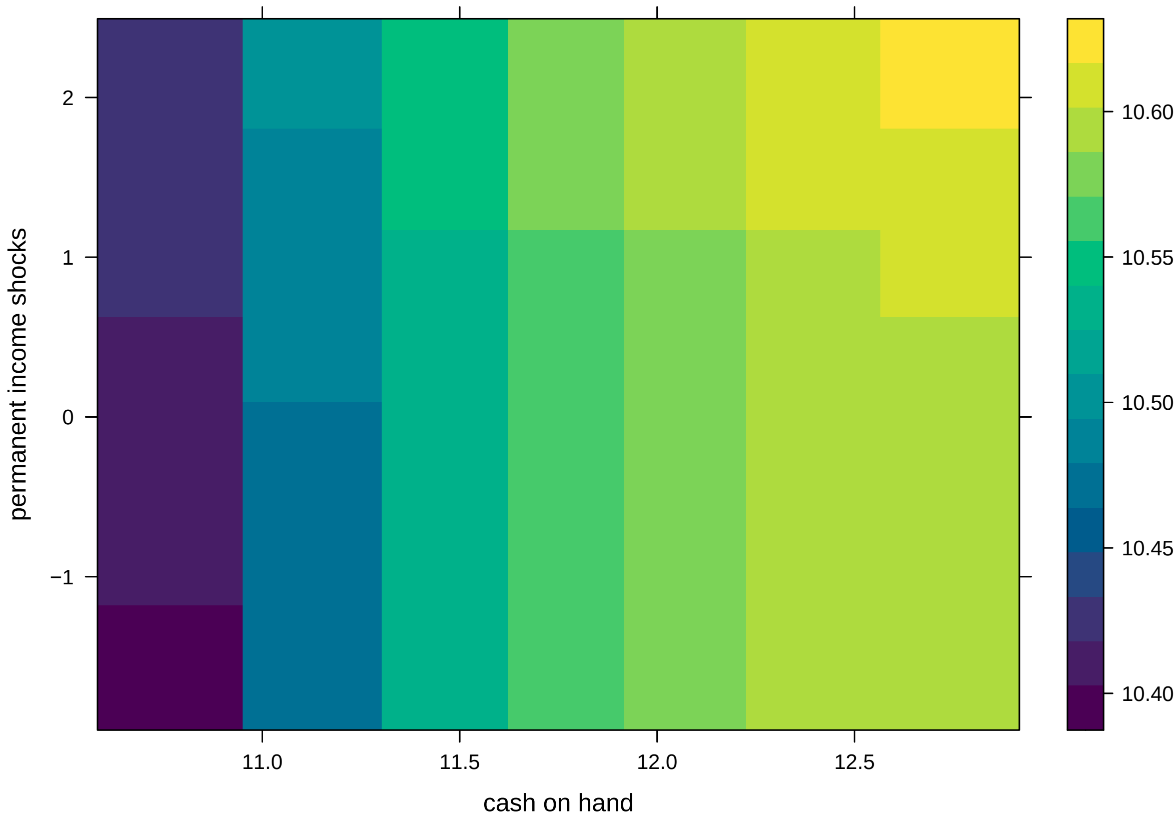

Notes: Function $G(\cdot)$ from eq. (5) is estimated using the full set of regressors as detailed in Section 3.1. A partial dependence function is then computed and evaluated on a quantile grid to extract the marginal effect of log cash on hand and permanent income shocks on predicting log consumption. Permanent income shocks are measured in units of its standard deviation. The colorkey legend on the right of the plot represents levels of log consumption.

Figure A.2 reports the mean decrease in node impurity obtained by each argument of the estimated consumption function through the random forest algorithm. Node impurity quantifies the homogeneity of the target variable within the subsets created by a split measured in units of residual sum of squares of the dependent variable. Lower impurity means that the subset is more homogeneous.

Figure A.2: Arguments Contribution to Consumption Function Estimation Accuracy

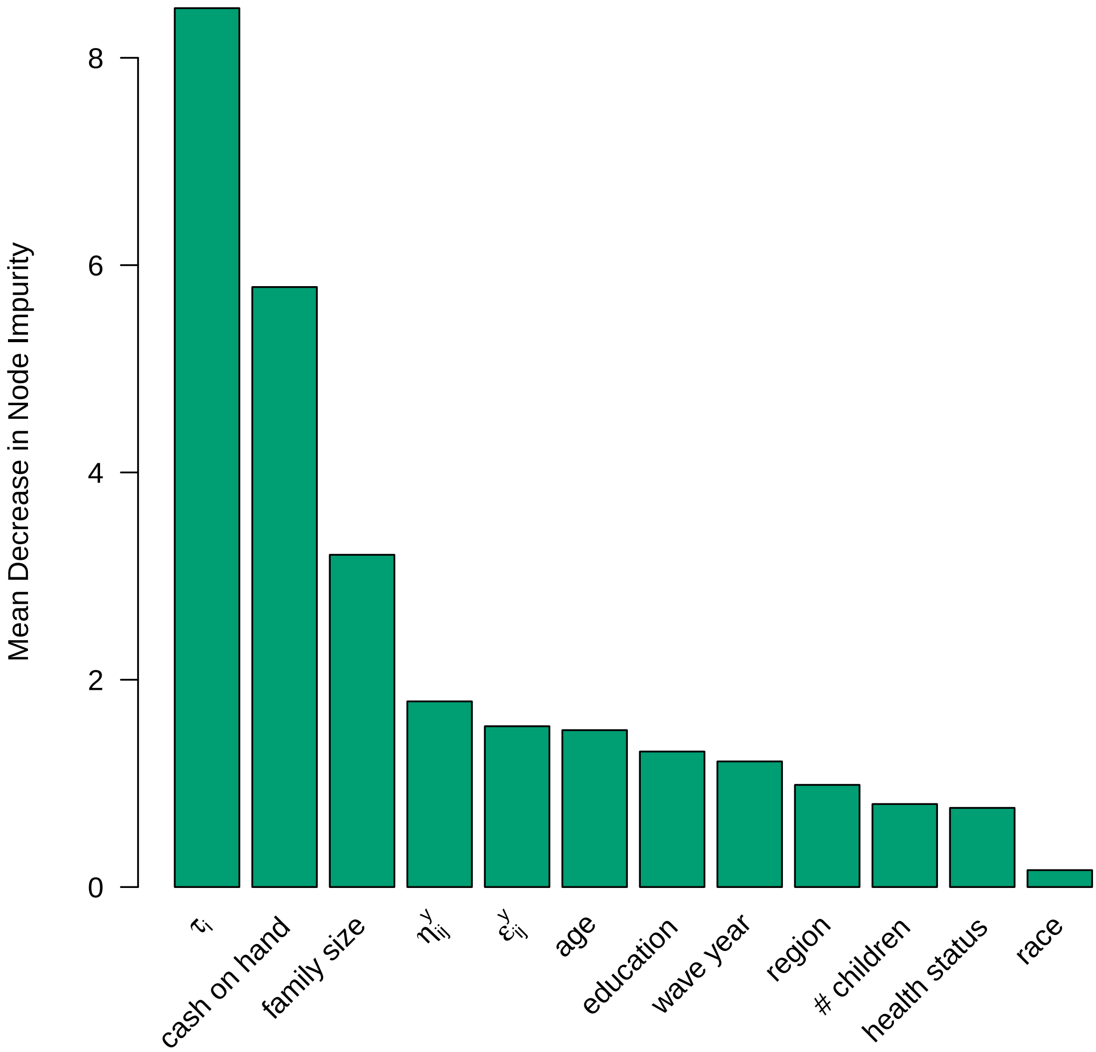

Notes:The measure is the total decrease in node impurities from splitting on the variable, averaged over all trees. Node impurity quantifies the homogeneity of the target variable within the subsets created by a split. Lower impurity means that the subset is more homogeneous (i.e., the target variable values are more similar). It is measured in units of residual sum of squares of the dependent variable, in this case log consumption.

B Alternative Modeling of Consumption Residuals

In this Appendix Section, I propose an alternative modeling assumption for consumption residuals. As shown in detail in Section 3.1.1, there's no evidence indicating consumption residuals might follow a random walk. However a random walk modeling assumption can be convenient to pursue. The following simple derivations show how, using consumption residuals in first differences, the modified expenditure-risk pass-through parameter can be used to test the null hypothesis of no correlation betwen consumption growth and persistent changes in consumption residuals.

Under the random walk assumption, i.e. $\rho^c = 1$, the pass-through coefficient defined by (9) boils down to:

$$\begin{aligned} \beta^{\text{alt}}_{\Delta c_{i,t}|\eta^c} &\overset{[\rho^c=1]}{=} \frac{\text{cov}\left(\Delta c_{i,t},\ \sum_{j=-1}^{1}\Delta\xi_{i,t+j}\right)}{\text{cov}\left(\Delta\xi_{i,t},\ \sum_{j=-1}^{1}\Delta\xi_{i,t+j}\right)} \\ &= \frac{\text{cov}\left(\Delta c_{i,t},\ \xi_{i,t+1}-\xi_{i,t-2}\right)}{\text{var}\left(\Delta\xi_{i,t}\right)} \end{aligned} \tag{B.1}$$

Therefore, considering the time structure of the PSID, $\beta^{\text{alt}}_{\Delta c_{i,t}|\eta^c}$ measures the covariance of 2-years (observed) consumption growth with 6-years residual consumption growth. As Table B.1 clearly shows, we can reject the null hypothesis of no permanent component left in consumption residuals once we've controlled for household income-level type, income risk realizations, cash on hand, and other demographics as detailed in Section 2.1.

Table B.1: Consumption Pass-Through Coefficients in PSID, Alternative Expenditure Risk Modeling

	Δc			
	(1)	(2)	(3)	(4)
$\boldsymbol{\beta_{\Delta c\|\eta^y}}$	0.43***			0.53***
	(0.037)			(0.094)
$\boldsymbol{\beta_{\Delta c\|\varepsilon^y}}$		0.06***		0.19***
		(0.013)		(0.040)
$\boldsymbol{\beta^{\text{alt}}_{\Delta c\|\eta^c}}$			0.86***	0.85***
			(0.107)	(0.106)
R^2	0.0130	0.0014	0.0479	0.0692
households	3,543	4,323	949	949
hh waves	12,733	17,056	2,099	2,099

Notes: Estimates here replicates those expressed in Table **??** with the random walk assumption for expenditure risk permanent component as detailed in Appendix B. Household-level clustered standard errors are displayed in parenthesis. ***, **, and * express significance at the 1%, 5%, and 10% level, respectively.

However, if we look at the ability to capture variability in log consumption, as expressed by explained sum of squares, this paper chosen AR(1)-specification expenditure risk measure explaines 35% of variability in log consumption, while the RW-specification can only account for about 1% of it.

C Fall into H2M Random Forest Classification

Figure C.1: Fall-Event Prediction, Predictors Relative Importance

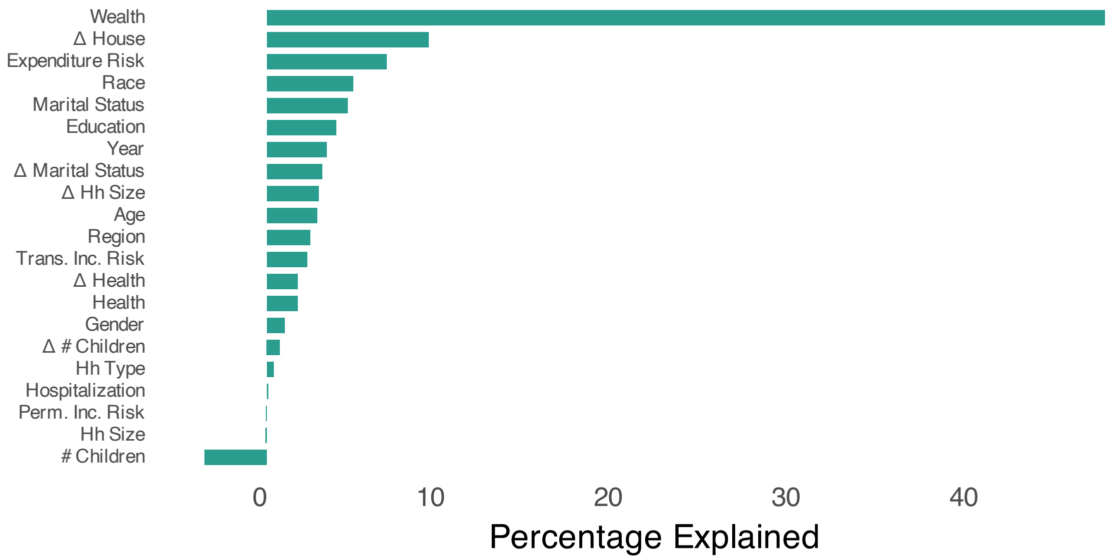

Notes: The bars measure the total decrease in node impurities from splitting on each variable, averaged over all trees, with respect to out-of-bag predictors classification. The node impurity is measured by the Gini impurity index, which measures the average misclassification probability.