The Nonlinear Consumption Response to Income Risk: The Role of Tail Changes and Durables*

Rocio Madera†

September, 2025

Abstract

This paper characterizes the joint dynamics of earnings and consumption changes focusing on the tails of the distribution, in light of new evidence that finds large earnings changes to be empirically relevant. We make four contributions. (i) We discipline a survey-data sample with administrative-data moments and document departures from normality in the marginal distributions of household disposable income and consumption changes. (ii) We document that the comovement between household earnings and consumption changes is highly non-linear, with extreme events correlating strongly with durable consumption adjustments and less so with non-durable expenditures. (iii) We build a life-cycle, incomplete markets model with lumpy durable consumption and non-Gaussian earnings shocks. We parametrize the model using higher-order moments of earnings and consumption at the household level. The lumpy nature of durables is essential to rationalize the empirical patterns under non-Gaussian risk. (iv) We use the model to calculate the consumption response to shocks of different size and to measure the welfare cost of incomplete markets. Pass-through of persistent shocks is lower in the tails of the income changes distribution, especially at younger ages. Under non-Gaussian risk, the welfare cost of idiosyncratic risk increases 9pp., but the durable margin mitigates this loss.

JEL Codes: E21, D31, D91

Keywords: Idiosyncratic income risk, higher-order moments, consumption insurance, durable consumption, non-Gaussian shocks.

---

*I am grateful to Fatih Guvenen, Ellen McGrattan, Anmol Bhandari, and all other participants at the University of Minnesota Macro-Labor Workshops.

†Southern Methodist University; romadera@smu.edu.

1 Introduction

Households rely on a variety of formal and informal insurance mechanisms to cope with unexpected income changes, many of which are difficult, if not impossible, to observe at the relevant frequency. For this reason, the joint behavior of income and consumption has long served as a measure for the welfare cost of income risk and for the value of self- and public insurance (Deaton and Paxson, 1994; Blundell *et al.*, 2008). Yet, efforts to characterize the comovement of income and consumption beyond averages or specific natural experiments are limited. This paper measures the joint dynamics of the full distribution of income and consumption across a wider range of shocks and quantifies the consequences of incorporating this non-normal, tail behavior for estimates of consumption pass-through, partial insurance, and the welfare costs of idiosyncratic income risk.

Recent advances in access to administrative earnings data reveal a higher incidence of large changes than previously suggested by second-order measures based on survey data (Guvenen *et al.*, 2021). In a given year, most individuals see little or no change in their labor income and a few but non-negligible share experience extreme events. The implications of such tail shocks are ex ante ambiguous. On the one hand, extreme income realizations can tighten liquidity constraints and raise the welfare cost of idiosyncratic risk (e.g., Guvenen *et al.*, 2024; Nardi *et al.*, 2020). On the other hand, large shocks may trigger adjustments that small shocks do not–rebalancing illiquid assets, moving in with family, or undertaking (or postponing) lumpy purchases–thus activating buffering margins that mitigate welfare losses. At the aggregate level, the nonlinearities generated by the interaction of large shocks, liquidity constraints, and adjustment costs can meaningfully shape consumption dynamics and the transmission of macroeconomic policy.

Despite advances in income measurement, progress on the consumption side has been slower. Administrative sources provide precise earnings histories, but consumption is typically measured in surveys and often in repeated cross sections, making panels of income and consumption rare. Beyond these data constraints, *consumption* itself is a complicated flow measure. In only a few cases is consumption equal to expenditure, as many categories are durable and yield benefits over multiple years. Some goods and services are adjusted only infrequently due to high adjustment costs, creating lumpy patterns in observed behavior. Durable purchases often entail ongoing commitments through maintenance or debt service, and many consumption categories are intrinsically tied to the household life cycle, such as housing needs expanding with family size. These features complicate both the measurement and interpretation of consumption dynamics, especially when studying their joint evolution with income over time.

As a result, much of the evidence comes either from natural experiments (e.g., tax rebates; Johnson *et al.*, 2006) or from imputing cross-sectional consumption into panel income data (e.g., Blundell *et al.*, 2008); these approaches deliver clean identification for specific shocks but generally rely on approximate linearity of the consumption rule. Quantitative approaches within life-cycle model frameworks (e.g., Kaplan and Violante, 2010; Storesletten *et al.*, 2004) allow for nonlinear responses, yet the stochastic component of earnings is frequently limited to *normal* changes. More flexible empirical (e.g., Arellano *et al.*, 2017) and quantitative (e.g., Nardi *et al.*, 2020) approaches have recently inspected changes in income from a nonparametric perspective.

This paper makes two primary contributions. First, using panel data from the PSID on income, public transfers, and nondurable and durable consumption, I characterize the marginal and joint distributions of after-transfers household earnings and consumption changes. I show that excess kurtosis is not an exclusive feature of individual earnings, which fades out when transfers are considered. On the contrary, household earnings, nondurable consumption, and durable consumption changes exhibit deviations from log-normality, especially the latter. This, however, does not imply that large income changes correlate with large consumption fluctuations as discussed above. To inspect the joint distribution of income and consumption changes empirically, I define measures of co-movement and tail dependence between earnings and consumption changes. While the overall covariance between earnings and nondurables is considerably higher than for durables, empirical measures of quantiles and tail dependence show that the opposite is true in the case of extreme events. That is, earnings and nondurable consumption are correlated but less so in the tails. The opposite happens for durable consumption. This suggests that both types of goods be viewed as complementary in order to understand the response of consumption to income shocks.

Motivated by the empirical findings, the second contribution of this paper is to build a model that rationalizes these findings and provides a laboratory to inspect the mechanisms and to calculate the consumption response. I add two elements to the standard life-cycle incomplete markets model: durable consumption adjustments and higher-order, idiosyncratic income risk. In this case, a model is necessary as nonlinearities are pervasive. Beyond the nonlinearities implied by the presence of borrowing constraints,[1] the size of the tail shocks creates jumps that usual, empirical identification strategies can hardly capture.[2] I model earnings as a mixture of two normals plus a deterministic age profile, which is flexible enough to capture the excess kurtosis observed in the data. Richer forms of statistical processes have been proposed in the literature, starting with (Geweke and Keane, 2000) and, more recently, Nardi *et al.* (2020) and Guvenen *et al.* (2021). The mixture of two normals is enough to capture the differences with respect to a model without tail risk, which is the main point of my analysis. Durable consumption expenditures are exposed to non-convex adjustments' costs. This implies that the decision rule for durable consumption follows a Ss-type behavior, which will be the centerpiece of my mechanism.

To parametrize the model, I proceed in two steps. First, I use a simulated method of moments to estimate the parameters of the earnings process, targeting the second through fourth moments of two- and four-year growth in after-transfers income. Next, I proceed to calibrate the remaining parameters of the model to match both aggregate and microeconomic targets, including the distribution of consumption changes. In addition to the targeted moments, I evaluate the fit of my model comparing the Quantile-Quantile plots of the earnings and consumption changes distribution, both durable and nondurable.

Finally, I use the calibrated model to test a series of implications of tail income shocks for the response of both durable and nondurable consumption, as well as for the degree of self-insurance of households. Not surprisingly, large income shocks do have a strong impact on the probability of durable adjustment, and the response is of the Ss-type, as expected. That is, there is practically no change in the middle part of the distribution. There are two mechanisms that generate leptokurtosis in durable consumption changes in the model: One is the endogenous lumpiness in the adjustment of durable consumption as a result of adjustment costs, but there is also a delayed upward adjustment from the option value of durable goods. These two mechanisms are consistent with empirical evidence in (Chetty and Szeidl, 2007) and Browning and Crossley (2009), respectively.

Looking at the response of nondurable consumption and the degree of partial insur-

[1]See Kaplan and Violante (2010) for an in-depth discussion of the implications of borrowing constraints for empirical measures of self-insurance whose identification relies on the linearity of policy rules.

[2]A notable exception is Arellano *et al.* (2017)

ance, I show that it features a large amount of heterogeneity in the size of the shock, with larger shocks of either sign triggering smaller non-durable consumption responses than average-sized shocks. Finally, using the calibrated model, I calculate the fraction of consumption at every age that households would be willing to give up to move from a world with durables and non-Gaussian shocks. I find that the welfare costs attributable to higher-order risk are around 9% of the yearly household consumption bundle with durable adjustments. Additionally, in the absence of durable adjustment, this extra risk would be much larger, suggesting an insurance role of durable adjustment when shocks are non-Gaussian.

Related Literature

This paper is related to several streams of the literature, but mainly falls at the corner between the measurement of uninsurable income risk and the implications of higher-order moments in income changes for household consumption. The literature on consumption or *partial* insurance has a long list of reference papers (Blundell *et al.*, 2008; Primiceri and van Rens, 2009; Kaplan and Violante, 2014; Guvenen and Smith, 2014). All of them look at the response of nondurable consumption to unexpected income changes. The latter two estimate structural versions to account for nonlinearities in the consumption rule. My contribution to that literature is twofold: (1) I model the distribution of earnings in a way that potentially very large shocks of nonnegligible density can happen; and (2) I show the importance of studying nondurable consumption decisions in connection to durable to understand the substitution between the two at different parts of the income shocks distribution.

A more recent related line of work studies nonlinear consumption responses to income shocks within life-cycle models that feature incomplete markets and heterogeneous income risk. DeNardi (2002) show that allowing for skewed and leptokurtic earnings processes generates substantial nonlinearities in marginal propensities to consume (MPCs) across the wealth distribution, with the poorest households displaying markedly higher and more asymmetric responses. Similarly, Guvenen *et al.* (2024) provide evidence that nonlinearities in the transmission of income shocks are especially pronounced for large negative shocks. My paper complements these contributions by adding the durable dimension and quantifying the structural mechanisms behind these nonlinearities: rather than stemming solely from borrowing constraints or precautionary motives, I show that lumpy durable adjustment plays a key role in shaping heterogeneous consumption responses.

The only other paper, to the extent of my knowledge, that considers durable consumption in a life-cycle incomplete markets model for the purpose of evaluating the ability of households to self-insure is Cerletti and Pijoan-Mas (2012). There are three main differences in our frameworks: In their model, adjustment of durable goods is smooth, not subject to adjustment costs. This responds to the fact that their main focus is how durables provide a rebalancing option that alleviates borrowing constraints in the event of an unexpected shock. The second difference is their income process, which follows a standard random walk plus white noise. Lastly, our definition of durables differs in that I include housing as a durable good.

The empirical observation that individual earnings changes are leptokurtic is not new. Over a decade ago, Geweke and Keane (2000) characterize the distribution of male earnings in the PSID and find that a normal does poorly at approximating the observed numbers, which resemble a leptokurtic distribution. Bonhomme and Robin (2009), also making use of advances in nonparametric econometric methods, show that the same is true for France. This literature has become especially prolific in the last decade with the increasing availability of administrative data. Guvenen *et al.* (2021) study the dynamics of earnings over the lifecycle using social security records of millions of workers. Their sample size allows for a fully nonparametric analysis. Compared to the previous papers, they document that there is a large amount of heterogeneity in the higher-order moments over the life cycle and initial level of earnings. While previous papers have reported numbers of slightly below 10 (Bagger *et al.*, 2014), it ranges from 4 to 40 for different ages and income status. To this literature, my main contribution is to measure whether the tail changes implied by the higher-order moments in income, that could be potentially very disruptive if taken at face value, have any impact on consumption and household's welfare. First, by looking at data for households after government transfers, and second by moving forward to the response of consumption.

Chetty and Szeidl (2007) and Browning and Crossley (2009) look at the empirical relation between durable goods and income shocks. The former is closer to this paper in the sense that it focuses on household lumpy consumption responses to a large wage shock. The latter focuses on smaller durables, such as clothing, furniture, and the like. While methodologically different, their results are consistent with my findings. They both provide a theoretical framework that hints at a stronger response of durables in the event of an unemployment shock, dampening the transmission to nondurables. They conjecture an increase in welfare coming from the lower fluctuations in nondurable consumption, but their frameworks, unlike mine, do not allow for a welfare analysis of the value of durable consumption as a margin of adjustment in the event of income shocks.

The quantitative response of durable consumption to income shocks has been studied extensively in a business-cycle environment. Considering that recessions and expansions are times in which large negative and positive shocks, respectively, are more frequent, this paper is also related to this literature that includes, for example, Grossman and Laroque (1990); Flavin and Nakagawa (2004); Berger and Vavra (2015). The closest to my framework, but in an infinitely-lived households version, is the latter. Our problems are conceptually different, though. Their focus is in how positive durable expenditures respond more or less sluggishly to economic shocks. As a result, I set up the problem so that households can upgrade or downgrade the size of their durables. Considering the comparable case of upwards movements in my model, my results are consistent with theirs.

The remaining of the paper is structured as follows. Section 2 empirically inspects the marginal and joint distribution of earnings and consumption changes, the baseline model and its calibration are described in Section 3, Section 4 explains the main results and implications. Section 5 concludes.

2 Higher-Order Comovement of Earnings and Consumption

This section presents an empirical characterization of the marginal and joint distributions of household earnings and consumption changes, with a focus on the higher-order moments and the deviations from normality.

2.1 A Panel of Income and Consumption Changes

This section documents the construction of the PSID-based panel and the definitions of all variables used in the analysis.

2.1.1 The Consumption Panel in the PSID

The PSID is a longitudinal study of a representative sample of U.S. households, tracking a wide set of socioeconomic variables of US families from 1967. Due to its length and panel structure, it has been extensively used for the study of income, wealth, and consumption dynamics. In this section, I briefly describe the dataset and focus on the recent and less-explored waves that contain detailed data on consumption. For a general overview of the PSID's design and content, see, for example, Heathcote *et al.* (2010).

The PSID's original focus was income dynamics and poverty. Accordingly, coverage of socioeconomic and income covariates has been thorough since inception, whereas questions on consumption were long limited to food and rent.[3] Beginning with the 1999 wave (reporting tax year 1998), the survey added a broad set of expenditure categories, making the PSID the first U.S. panel with disaggregated income and consumption measures comparable to cross-sectional surveys such as the CEX.[4] Prior to 1999, most studies requiring panel data on both income and consumption relied on imputing CEX consumption into the PSID.[5] While these imputations perform well for nondurables, durable adjustments are largely absent.

Given our focus on consumption, this paper's main data of reference are precisely these later waves of the PSID, beginning in 1999 and spanning until 2015.[6] This poses two main challenges: The first one concerns the interviewing frequency starting with the 1999 wave, which changes from annual to biennial. We will therefore interpret the results accordingly. The second limitation of discarding nearly the first 30 years of data is the loss in observations for a relatively small survey. It is worth noticing, however, that the level of attrition is higher in the initial years of the survey, so the loss in the time dimension is compensated with a more balanced and stable panel, less likely to be affected by non-random entry and exit in the panel that can contaminate the estimation of the earnings distribution (Daly *et al.*, 2016). Additionally, the 1993 PSID wave underwent

[3] Some information on the value of owned homes and vehicles is available, but it is inconsistent over time, and there is no corresponding series on expenditures on these durables.

[4] Andreski *et al.* (2014) provide a detailed comparison of the post-1999 PSID expenditure modules with the CEX and conclude that the overlapping categories are broadly comparable in levels and over the life cycle.

[5] See, for example, Blundell *et al.* (2008); Kaplan and Violante (2010); Guvenen and Smith (2014).

[6] The consumption panel now runs through 2023. Because extending the sample leaves the core facts materially unchanged, I keep the earlier waves for the baseline estimates. See Guvenen *et al.* (2024) for an updated characterization of consumption moments.

a major revision in main variables concerning labor income; starting the analysis after that revision avoids spurious variation and the need for retrospective harmonization.

To sum up, this paper uses the 1999–2015 waves of the PSID, spanning 16 years (1998–2014) of information on income, consumption, and wealth, as well as a wide set of socioeconomic covariates, at a biennial frequency. These biennial waves of the PSID constitute the main dataset for the empirical analysis in this paper,[7] with measurement and sample design choices tailored to the goals of characterizing the joint distribution of income and consumption changes.

2.1.2 Sample

The baseline sample includes households whose heads are between 25 and 60 years old, have not retired, and that have not suffered major changes in their family structure in the past two years. I require at least three consecutive observations between 1998 and 2014. I retain both the representative SRC sample and the SEO oversample, applying PSID family weights throughout. At the end of Section 2.2 I describe the treatment of outliers, a key step in capturing tail changes.

The final sample with information on income and consumption is comprised of around 45,000-60,000 observations, corresponding to approximately 13,000 households over 18 years at biennial frequency. More details, including the number of observations left at each step of the sample selection, are given in Appendix A.

2.1.3 Main Variables: Definitions

Earnings. The reference measure of household earnings will be household earnings after taxes and transfers, which I will also refer to as post-government income. Post-government household earnings are defined as pre-government household labor income *plus* public transfers *minus* federal income taxes. Pre-government household labor income is composed of the head of household's labor income *plus* the spouse's labor income. Each member's labor income excludes self-employment. Transfers include unemployment insurance, welfare, and social security. Federal income taxes are calculated using TAXSIM. All amounts shown in dollars are in real 2010 dollars, deflated using the general PCE index for income.

[7]Earlier waves are used selectively when information on prior behavior is required.

Consumption. For the case of consumption, nondurable consumption includes food, utilities, nondurable transportation, and recreation. Durable consumption includes houses, cars, furnishings and repairs, and clothing. A detailed description of all consumption subcategories and the exact construction of each variable can be found in Appendix A. All amounts shown in dollars are in real 2010 dollars, deflated using the general PCE index for nondurable consumption categories, except for housing and vehicles. Housing and vehicle-related expenditures and adjustments are deflated using the corresponding PCE for housing and motor vehicles, respectively.

The remaining variables are defined in Appendix A.

2.1.4 Measures of Changes and Adjustments

Measuring consumption changes in income and nondurable goods and services is a relatively straightforward task. The case of durable consumption, however, requires some discussion.

It will be helpful to start by defining two measures of change that will be central in my analysis: Let $\log \Delta^s(x)$ and $\text{arc}\Delta^s(x)$ denote the log- and the arc-change in x from the the current period to s periods ahead, respectively. Formally:

$$\begin{aligned}
\log \Delta^s(x_t) &\equiv \log x_{t+s} - \log x_t \\
\text{arc}\Delta^s(x_t) &\equiv \frac{x_{t+s} - x_t}{(x_{t+s} + x_t)/2}.
\end{aligned}$$

The default measure will be log-changes. The log-change has the advantages that it is unbounded for both positive and negative changes and a very good approximation of the growth rate for most of the changes observed in income and consumption. The disadvantages include that it is not defined for changes that involve a zero and that it is a poorer approximation of growth for a few but very important large differences. Because large changes in income and infrequent changes in consumption are the focus of this paper, I will complementary use the two measures defined above, as well as dollar changes relative to current income, in the remaining of the paper. I next explain how I define adjustments in durable consumption.

For the case of the smaller durables, direct expenditure values are reported. For the case of vehicles and houses, I follow the definitions in Chetty and Szeidl (2007), who define an adjustment as the change in vehicles and houses beyond depreciation.

To minimize measurement error, I combine data on exchanges of vehicles and sales of houses with self-reported moves and value of the stock. If no move, purchase, or sale is reported, and the value of the good, as well as property taxes and home insurance, is within 20% of their value from last year, no adjustment is recorded. This approach to measuring durable adjustments is inspired by Berger and Vavra (2015). For the rest of the cases, I define different situations that are explained in Appendix A but, in general, an adjustment is considered. The value of the adjustment is an average between the self-reported value of the house or car and the value of the exchange, which very often coincide. Changes in durable consumption are calculated applying the measures described above directly on the value of the stock. If there is an adjustment, the value is converted to real values using the corresponding PCE for each category. I refer to changes in durables as adjustments, as a reminder that they are changes in the stock.

2.2 Tail Changes in Survey Data: Sample Selection

Measurement error is not uncommon in survey data. Therefore, a crucial step when selecting a sample for analysis is the treatment of outliers. The usual practice is to trim the tails of the earnings distribution somewhere between the 1 and 10 top and bottom percentiles (e.g., Blundell *et al.*, 2008; Heathcote *et al.*, 2010). This poses a trade-off when the purpose is to characterize the tails of the distribution or the full distribution for macroeconomic purposes: On the one hand, we want our sample to be as clean as possible from coding and response errors. On the other hand, the recent evidence from administrative sources mentioned above suggests that some–more than we thought–of those outliers can indeed be free of error and correspond to unusually large changes. This practice might explain why income changes were found to follow roughly a log-normal distribution. I propose a new method for *cleaning* the tails.

The method is very simple. It relies on the rich set of moments of the distribution of earnings levels and changes made publicly available by Guvenen *et al.* (2015) from administrative data in the US (SSA data). These moments are further available by age, sex, and past income levels.[8] I use the distribution of individual earnings changes in the SSA data to discipline the trimming of the tails in the PSID waves corresponding to information for the years 1992-1997. These waves are available at the annual frequency and explicitly separate labor earnings from incorporate business income, making it highly comparable to the SSA data.

[8]The Global Repository of Income Dynamics has expanded these statistics to a large set of countries.

We minimize the distance between an array of administrative-data and survey-data moments, trimming 0.01 percent from each time alternatively. In particular, let $m^{\text{PSID}}(c; p_L, p_R)$ be the same moments computed in PSID cell c after trimming the p_L left-tail and p_R right-tail percentiles of $\Delta \log y$ *within the cell.* Starting from $(p_L, p_R) = (0, 0)$, I iteratively adjust (p_L, p_R) in 0.01 percentage-point increments to minimize a quadratic distance

$$\mathcal{L}(p_L, p_R) \;=\; \sum_c \left(m^{\text{PSID}}(c; p_L, p_R) - m^{\text{SSA}}(c)\right)^\top W \left(m^{\text{PSID}}(c; p_L, p_R) - m^{\text{SSA}}(c)\right),$$

with W set to the identity.[9] The resulting cell-specific cutoffs are then averaged within year and applied to the main biennial sample (1998–2014). On average, this yields trimming of approximately 0.2% in the left tail and 0.5% in the right tail per year. Importantly, trimming is performed on earnings changes only; I do *not* trim on consumption or on joint outcomes to avoid selection on the dependent variable. Figure 1 shows the resulting distributions.

[9] Results are similar with diagonal W weighting tail probabilities more heavily.

Figure 1: SSA and the PSID: U.S. Males Annual Earnings 1-Year (top) and 5-Year (bottom) Log-Changes

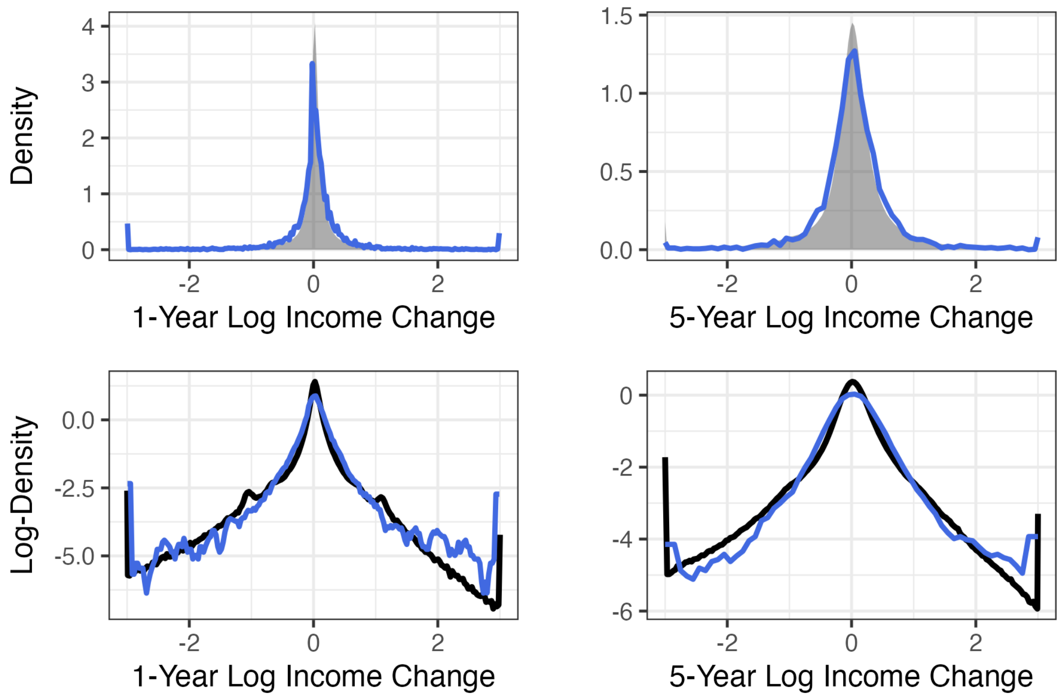

Notes: Blue curves plot PSID kernel densities of 1-year (top) and 5-year (bottom) log individual male earnings changes; gray shading shows the corresponding histograms (Epanechnikov kernel). Black curves reproduce the SSA densities reported by Guvenen *et al.* (2015). PSID moments are computed using the 1992–1997 annual waves to match the SSA frequency; tails are disciplined with administrative moments via within-cell trimming (approximately 0.2% left and 0.5% right per year; see Section 2.2 and Appendix A).

2.3 Marginal Distributions: The Tails of Household Earnings and Consumption Changes

This section analyzes the distribution of changes in labor income and consumption, with particular attention to deviations from normality and the tails. I use standard distributional graphs alongside summary measures—skewness, excess kurtosis, and tail probabilities.

Robust Higher-Order Moments

To provide a definition of tail changes, I first look at the empirical distribution of labor income changes. The first panel of Figure 2 shows the distribution of household income after taxes and transfers. Table II reports the share and usual amount of income change for different sizes. This is an alternative and more intuitive representation of the same idea behind Figure 2.

The four first central moments of the distribution are useful descriptors of the underlying shape. Nonetheless, they are highly influenced by outliers and are sometimes hard to interpret. Therefore, I will complement the information contained in the central moments with their percentile-based counterparts. In addition to being robust to outliers, these measures have a clear interpretation in terms of easily identifiable parts of the distribution. Formally,

$$\mathsf{Pk} \equiv \mathsf{k}\text{th percentile}$$

$$\mathsf{Pk\ell} \equiv \mathsf{Pk} - \mathsf{P\ell}$$

$$\text{Kelley Skewness} \equiv \mathsf{S}^{\mathsf{K}} \equiv \frac{\mathsf{P9050} - \mathsf{P5010}}{\mathsf{P9010}} \tag{1}$$

$$\text{Crow-Siddiqui Kurtosis} \equiv \mathsf{K}^{\mathsf{CS}} \equiv \frac{\mathsf{P97.5} - \mathsf{P2.5}}{\mathsf{P7525}}. \tag{2}$$

Table I reports the values of the second through fourth moments of the distribution for different measures of earnings and consumption. I choose to include the robust measures and relegate the remaining moments to the appendix, in Table A.1. There are several important empirical results contained in this table. Because they are the centerpiece of my empirical analysis, I will discuss them in detail.

First, looking at the bold numbers referring to the whole sample, we can see that all variables exhibit deviations from normality, mostly in the form of excess kurtosis. This is a feature that is observed in administrative data for individual earnings, and it is thus important to observe it in my sample. More interesting is the fact that excess kurtosis remains high after including the spouse's earnings and government transfers, which we would expect to dampen the fluctuations in household income. Furthermore, changes in both measures of consumption are far from log-normal. This result has been pointed out by Toda and Walsh (2015) using the CEX data, but the fact that durable consumption changes are strongly leptokurtic is unexplored. Appendix A includes the histograms corresponding to these variables in log scale, to emphasize the size of the tails.

Table I: Higher-Order Moments of Earnings and Consumption: All Households

		Mean	Std. Dev.	Kelley Skewness	CS Kurtosis
Panel A: 2-Year Changes					
Individual Earnings	All	0.065	0.723	−0.027	9.409
HH Earnings Pre	All	0.059	0.704	−0.049	7.901
HH Earnings Post	All	0.064	0.635	−0.035	6.314
ND Consumption	All	0.047	0.489	−0.034	4.197
D Consumption	All	0.113	1.033	0.239	16.230
Panel B: 4-Year Changes					
Individual Earnings	All	0.121	0.822	−0.070	8.066
HH Earnings Pre	All	0.116	0.801	−0.065	6.715
HH Earnings Post	All	0.123	0.707	−0.059	5.341
ND Consumption	All	0.102	0.526	−0.031	4.033
D Consumption	All	0.228	1.228	0.261	13.554

Notes: The table reports distributional moments of changes in earnings and consumption over two horizons. For a variable x_t, the h-year change is defined as $\Delta_h x_t \equiv \log x_{t+h} - \log x_t$; "Mean" and "Std. Dev." are the sample mean and standard deviation of $\Delta_h x_t$. "Kelley Skew" is the quantile-based skewness statistic, which equals 0 under Normality, positive (negative) values indicate more extreme gains (losses). "Crow–Siddiqui" is the quantile-based kurtosis statistic, which equals 2.91 under a Normal distribution; larger values indicate heavier tails. "HH Earnings Pre" ("HH Earnings Post") denotes household earnings before (after) taxes and transfers. "ND Consumption" and "D Consumption" refer to nondurable and durable consumption, respectively. Panel A summarizes 2-year changes ($h = 2$) and Panel B summarizes 4-year changes ($h = 4$). The sample includes the full sample of households.

To provide a more intuitive characterization of how disturbing tail events in income can potentially be, Table II shows the share of households experiencing changes of different sizes in a given year, as well as the size of the change, both in log points and in dollars. For the moment, I pool positive and negative changes in the absolute value of the change. In order to define a relative measure of the size of the shock, I define thresholds depending on the number of standard deviations from the mean. For the purpose of understanding the significance of these numbers, it's important to remind a couple of features of the normal distribution so that we can understand its shortcomings. A normal distribution assumes that (1) all values in the sample will be distributed equally above and below the mean, and (2) only 0.3% of changes exceed three standard deviations in absolute value. This number is over 3% in my sample.

Table II: Incidence of Log Earnings Changes by Size

Size bin	Percent	Average Size (log Δ)	Average Size ($)
$0 \leqslant \lvert\Delta\rvert < 1\,\mathrm{SD}$	82.790	0.150	6,671.800
$1\,\mathrm{SD} \leqslant \lvert\Delta\rvert < 2\,\mathrm{SD}$	10.770	0.690	24,536.160
$2\,\mathrm{SD} \leqslant \lvert\Delta\rvert < 3\,\mathrm{SD}$	2.900	1.210	36,703.400
$3\,\mathrm{SD} \leqslant \lvert\Delta\rvert$	3.530	2.800	43,867.680
N		18,524	

Notes: Bins are defined by the absolute h-year log change in household earnings after taxes and transfers, $\lvert\Delta\rvert = \left|\log E_{t+h} - \log E_t\right|$ with $h = 2$. "Percent" is the share of observations in each bin. "Average Size (log Δ)" reports the mean of $\lvert\Delta\rvert$ (log points) within the bin. "Average Size ($)" converts each observation's log change to dollars using its baseline earnings level and averages within the bin; dollars are in 2010 USD.

Graphical Analysis and Deviations from Normality

Despite the strong evidence against normality shown in Tables I and II, it is still useful to provide a graphical description of how these numbers show up in the data. The upper panel in Figure 2 contains the histograms of all three main variables of interest: changes in income, nondurable consumption, and durable consumption, from left to right. The bars reflect the data, and the dashed line corresponds to a normal distribution with the same variance, which is approximately the distribution that would result from an estimated parametric specification that constrains shocks to income to be log-normal. It becomes evident that the majority of the changes within two standard deviations (approximately between -1 and 1) are very close to zero. However, it is very hard to extract conclusions on the tails based on the histograms. This happens mainly because the density function is bounded below by zero. Therefore, I complement the histograms with two other graphical constructs: (1) Log-densities, shown in Appendix A, and (2) Quantile-Quantile plots (QQ plots hereafter).

The lower panel in Figure 2 includes a set of QQ plots. QQ plots compare two distributions by plotting their quantiles against each other. They represent a particularly useful tool to assess the extent to which a variable is well approximated by a normal, or any given distribution. Both axes correspond to the x-axis in the histogram plot immediately above. We can thus think of the lower panel to be the two distributions in the upper panel against each other: the data is in the y-axis, and the normal is in the x-axis. As a result, the units are log changes of the corresponding variable. For illustration purposes, the axes in the case of earnings and durables are truncated at 3, but the conclusions do not change since the tails just keep diverging.

I will start describing the dashed line, which corresponds to the 45-degree line and coincides with the QQ plot if the variable in the y-axis was distributed exactly as the reference distribution. Next, the solid line in the left panel contains the sorted data. Notice that, particularly in the leftmost and rightmost graphs, it follows an S-shape. This is a sign of fat tails[10]. In the negative quadrant, points above (below) the 45-degree line are closer to (further from) the mean than their normal counterpart. The opposite happens in the positive quadrant. Moreover, the differences can reach 1 log point, despite being unnoticeable in the histograms.

With these concepts in mind, we can look at the three graphs and immediately infer both the middle part of the distribution and the tails, independently of the scale of the y-axis and the size of the bins, as opposed to the case of the histograms. In summary, both earnings and consumption exhibit deviations from normality. Nondurable consumption does so to a lesser extent, but a normal distribution would still miss the tails. The case of the durables is remarkable, mostly due to the fact that many households do not change their stock at all in a given year, but when they do, the change is large. Smaller adjustments correspond to furnishings and other smaller durables. The next question of interest is whether there is any relation between these tails of consumption and earnings changes.

[10] Appendix A.6 includes a stylized example of QQ plots for usual distributions.

Figure 2: Empirical Distributions: Histograms and Deviations from Normality

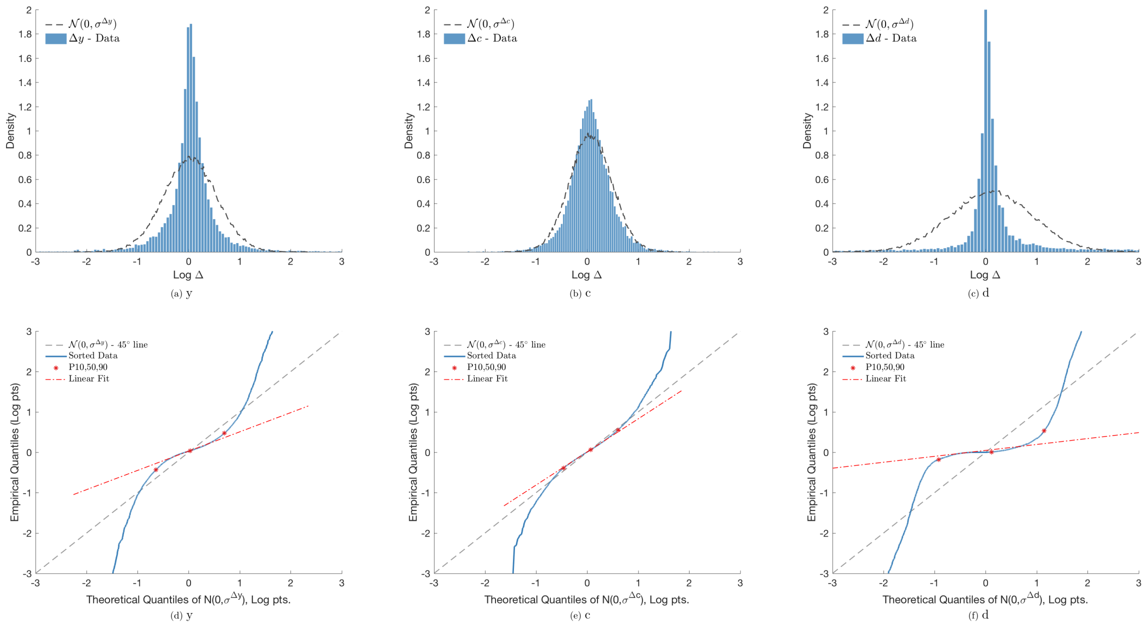

2.4 The Nonlinear Comovement of Earnings and Consumption Changes

The previous subsection showed that tail risk is pervasive even when private and public transfers are considered, and also in consumption, especially durable. Next, I assess the probability of these tail events that occur jointly in earnings and consumption. For that purpose, we proceed in three steps. First, we compare the *marginals* distributions of income changes with those of nondurable and durable consumption changes. Second, we estimate quantile regressions to trace out nonlinear conditional responses. And, finally, we measure the joint incidence of extreme events using tail dependence.

To compare the marginal distributions described in the previous subsection, we juxtapose the distributions of income changes with those of nondurable and durable consumption changes. With this graph, we want to emphasize the difference in densities of moderate changes. In particular, for the case of durable consumption, assuming a normal distribution in income would miss the important changes in durables. In the quantile–quantile style plots of consumption versus income (Figure 3), nondurables track income more closely near the center, while durables display a marked S–shape: the slope steepens in the tails, indicating relatively more mass in extreme durable adjustments than in income realizations.

To quantify heterogeneity and nonlinearities in the joint distribution, I estimate, for each $\tau \in \{0.10, \ldots, 0.90\}$, the weighted quantile regression

$$Q_{\Delta c}(\tau \mid \Delta y) \; = \; \alpha(\tau) + \beta(\tau)\, \Delta y,$$

separately for nondurable and durable consumption. The estimated $\beta(\tau)$ profiles (Figure 4) show that nondurable responses are largest around the middle of the distribution and attenuate toward the tails, whereas durable responses are U–shaped, with muted reactions at the center and much stronger sensitivity in the lower and upper quantiles.

Figure 3: The Joint Distribution of Income and Consumption: nondurables (left) vs. durables (right)

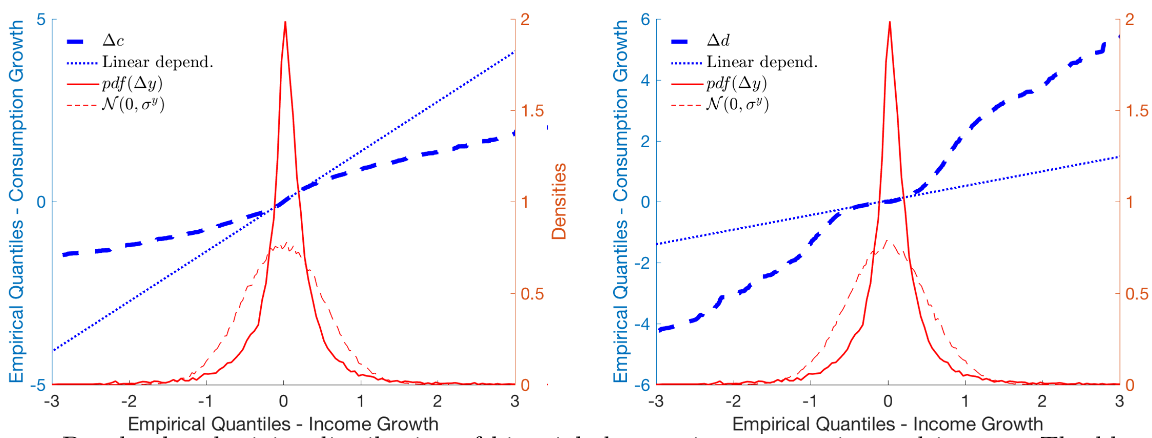

Notes: Panels plot the joint distribution of biennial changes in consumption and income. The blue dashed curve shows the empirical quantile map $Q_{\Delta c}(\tau)$ (left) and $Q_{\Delta d}(\tau)$ (right) against the income quantile $Q_{\Delta y}(\tau)$; the blue dotted line is the best linear fit ("Linear depend."). The red solid curve is the kernel density of Δy; the red dashed curve is a Normal $N(0, \sigma^2_{\Delta y})$ with the same variance; densities use the right axis.

Figure 4: Quantile Regression Coefficients of Income Change on Consumption Changes

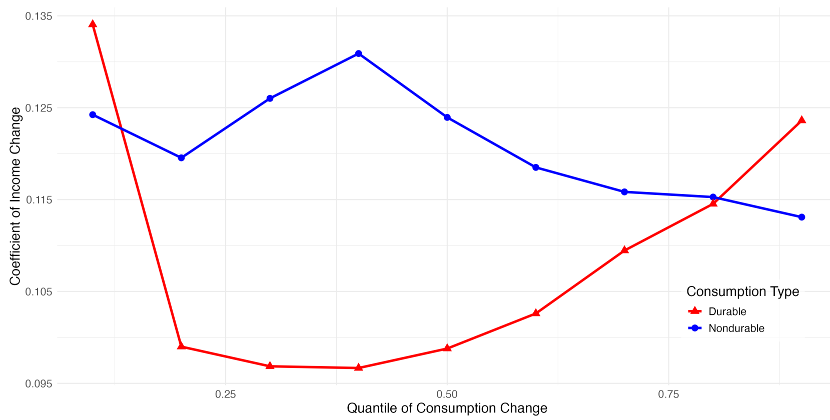

Notes: Points show τ-quantile coefficients from $rq(\Delta c_{h,t} \sim \Delta y_{h,t})$ by consumption type; lines connect τ.

Finally, we quantify the joint incidence of extreme events using tail dependence. We define tail dependence is as the *limiting probability that one random variable exceeds a certain threshold given that another random variable has already exceeded that specific threshold*. Formally, the so-called τ-measure for the dependence between the left tails of two random variables x and y is defined as

$$\tau_{y|x} = \lim_{p \to 0} \frac{\Pr(y < Q_y(p) \text{ and } x < Q_x(p))}{p} = \lim_{p \to 0} \Pr(y < Q_y(p) | x < Q_x(p)),$$

where $Q_y(p)$ denotes the quantile of the distribution of y at probability level p. It is very similar to the measure of correlation and does not imply causality. If $\tau = 1$, the tails of x and y are completely dependent, $\tau = 0$ denotes independence. There are several ways to estimate τ, I use the indicator proposed by van Oordt and Zhou (2012) for its non-parametric nature.

The estimator of $\tau_{y|x}$ is defined as the ratio between the number of observations in which both x and y are extreme and those in which only x is extreme. What *being extreme* means depends on the application. Formally:

$$\hat{\tau}_{y|x} = \frac{\sum_{i=1}^{n} I_{yi} I_{xi}}{\sum_{i=1}^{n} I_{xi} I_{xi}}$$
$$I_{xi} = 1(x_i < Q_x(k)),$$

where I choose k so that $Q_{y^{post}}(k) = 1.5$, 3 standard deviations for household labor income.

Table III: Correlation and Tail Dependence Between Income And Consumption

	Correlation	Tail Dependence
Nondurable c	0.136	0.048
Durable c	0.084	0.213

Table III shows the empirical measure of tail dependence, as well as the usual Pearson's correlation coefficient. The Pearson's correlation estimator averages deviations from the mean and does thus not distinguish between extreme or moderate outcomes or the sign of the returns. It is interesting to see that τ is very close to the Spearman's rank correlation. Despite weaker than tail dependence estimates, the Spearman's rank correlation has often been used as an alternative measure of joint tail behavior, and equals 0.068 and 0.187 for nondurables and durables, respectively.

Taken together, these facts point to pervasive nonlinearities in the joint distribution of income and consumption. Moreover, the role of durable consumption in measuring the response to income shocks seems crucial. Identifying insurance measures using linear empirical designs is therefore challenging. In the next section, I develop a life-cycle incomplete-markets model with lumpy durables to compute structural responses and explore their implications.

3 Quantitative Analysis

In this section, I build a life-cycle consumption-savings model with income uncertainty and incomplete markets, with two additions to the standard Bewley model in Kaplan and Violante (2010): (1) households are allowed to adjust durable consumption, subject to non-convex adjustment costs; and (2) shocks follow a distribution that is flexible enough to exhibit the excess-kurtosis observed in the data.

I first describe the specifics of the household problem, with an emphasis on the choice of durable consumption, which is the core of the model. Next, I begin discussing the parametrization of the model with the details of the idiosyncratic shocks specification and the adjustment costs of durable consumption. Finally, I proceed to the calibration of some parameters and estimation strategy of others.

3.1 A Model with Non-Gaussian Income Shocks and Lumpy Consumption

The economy is populated by a continuum of finitely-lived households. Each household works for T_R periods, lives as a retiree for $T - T_R$ periods, and dies with certainty at age T. During the retirement years, households have a probability of surviving from age t to the next age $t+1$ equal to $\xi_t < 1$. Perfect annuity markets are available. A period corresponds to two years.

Timing. The timing of events within a period is as follows. At the beginning of the period households observe their realizations of the idiosyncratic shocks ε and η. Next, households collect the capital income from the returns on their savings and make their consumption-savings decision, including durable consumption. Durables are chosen one period in advance, similarly to financial assets. This means that, even when agents sell durables in the current period, the service flow is received on the beginning-of-period stock.

3.1.1 Household Problem

Preferences Households have standard CRRA preferences over a consumption bundle of non-durable and durable consumption, denoted by c and d, respectively. Both types of goods conform the consumption aggregate following a Cobb-Douglas specification[11]. Future utility is discounted at the rate $\beta \in (0,1)$ and, after retirement, households have a probability of surviving $\xi_t \in (0,1)$. Formally,

$$E_0 \sum_{t=1}^{T} \beta^{t-1} \xi_t \frac{\mathcal{C}\left(c_t, s_t\right)^{1-\gamma}}{1-\gamma} \tag{3}$$

$$\mathcal{C}(c_t, s_t) \equiv c_t^{\alpha} (s_t)^{1-\alpha}, \tag{4}$$

where $\mathbb{E}_0$ is the expectation operator with respect to the stochastic processes intro-

[11] Piazzesi and Schneider (2007) provide evidence in favor of the Cobb-Douglas aggregation of both consumption goods.

duced in subsection 3.1.3, conditional on information available at time 0.

Borrowing and Saving Households can borrow and save issuing a risk-free bond. At every age, agents choose how much to save for the next period, a_{t+1}, and earn capital gains ra_t on currently held bonds, where $r > 0$ is the risk-free rate of return. Borrowing is constrained to a fraction λ^y of minimum labor income, $\underline{y}_t$, and a fraction λ^d of the chosen stock of durables, which can be understood as collateralized borrowing or a down payment requirement in the case of adjustment:

$$a_{t+1} \geqslant -\lambda^y \underline{y}_t - \lambda^d d_{t+1} \tag{5}$$

In the baseline case, I assume $\lambda^y = 0$, meaning that borrowing other than collaterized or for down payments is ruled out.

Old Age Income at retirement mimics the US system. Pensions are a function of lifetime average gross earnings.[12] Let $\bar{Y}_i^R$ denote the average labor income over the working life of a household and $\bar{Y}$ the average labor income in the economy. Then, pension income is defined as:

$$P\left(\bar{Y}_i^R\right) = \begin{cases} 0.9\bar{Y}_i^R & \text{if} \quad \bar{Y}_i^R \leqslant 0.3\bar{Y} \\ 0.27 + 0.32\left(\bar{Y}_i^R - 0.3\right) & \text{if} \quad 0.3\bar{Y} < \bar{Y}_i^R \leqslant 2\bar{Y} \\ 0.81 + 0.15\left(\bar{Y}_i^R - 2\right) & \text{if} \quad 2\bar{Y} < \bar{Y}_i^R \leqslant 4.1\bar{Y} \\ 1.13\bar{Y} & \text{if} \quad 4.1\bar{Y} < \bar{Y}_i^R \end{cases}, \tag{6}$$

where

$$\bar{Y}_i^R = \frac{1}{T_w}\sum_{t=1}^{T_w} Y_{it}$$

[12] For computational purposes, I follow Guvenen and Smith (2014) and estimate average labor earnings $\bar{Y}_i^R$ as the fitted value of

$$\bar{Y}_i = a_o + a_1 Y_{i,T_R},$$

where $\bar{Y}_i$ is the simulated individual average earnings and Y_{i,T_R} is income at retirement age. This avoids having to keep track of average earnings at each age.

Recursive problem of a working household For ages $t = 1, \cdots, T_r - 1$

$$\begin{aligned} V_t(a_t, d_t, z_t) &= \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} \\ \text{s.t.} \quad & c_t + a_{t+1} + d_{t+1} + A(d_t, d_{t+1}) = Y_t + (1+r)a_t + (1-\delta)d_t \\ & Y_t \text{ given by equations (9)} - (13) \\ & a_{t+1} \geqslant -\lambda^y \underline{Y}_t - \lambda^d d_{t+1}, \quad c_t \geqslant 0 \end{aligned}$$

Recursive problem of a retiree household For ages $t = T_R, \cdots, T$

$$\begin{aligned} V_t(a_t, d_t, z_t) &= \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \xi_t \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} \\ \text{s.t.} \quad & c_t + a_{t+1} + d_{t+1} + A(d_t, d_{t+1}) = P(\bar{Y}) + (1+r)a_t + (1-\delta)d_t \\ & P(\bar{Y}) \text{ given by equation (6)} \\ & a_{t+1} \geqslant -\lambda^y \underline{Y}_t - \lambda^d d_{t+1}, \quad c_t \geqslant 0 \\ & V_{T+1} = 0 \end{aligned}$$

3.1.2 Durable Consumption Choice

Durable Choice Set The choice to adjust durables is discrete. At each age t, households choose whether to keep the undepreciated portion of their durable stock or to adjust to one of the n^d sizes in set $\mathcal{D} = \{d_0, \cdots, d_{n_d}\}$.

Prices The relative price of durables regarding non-durables is normalized to one.

Adjustment Costs The adjustment of durable consumption is subject to the non-convex adjustment cost A. A is a function of the current and the next period's stock of durables:

$$A(d_{t+1}, d_t) = \chi_1 d_t + \chi_2 \Big| \underbrace{d_{t+1} - (1-\delta) d_t}_{c_t^d} \Big|, \tag{7}$$

where $c_t^d \equiv d_{t+1} - (1-\delta)d_t$ denotes expenditures on durable adjustments at age t. The first component in (7) is fairly standard in the literature (Berger and Vavra, 2015; Luengo-Prado, 2006) and it is responsible for an inaction region in durable adjustment as it is a fixed cost with respect to the adjustment. An empirical interpretation of this linear component includes sales agent fees, taxes, or repairs in preparation for selling. The second component, linear in the size of the adjustment, is added for the purpose of capturing the differences in size, in either direction, but it is quantitatively smaller than the first part of the adjustment cost function. A way to think about this second element is as new furniture needed to fill a new house or old furniture that needs to be disposed of, which is increasing in the difference between the size of the old and the new house. This formulation resembles that in Fella (2014), with the difference that, in his framework, downgrades are less costly than upgrades. In other words, the second term does not appear in absolute value. Appendix

Service Flow As opposed to the case of non-durables, expenditures on durable goods and the consumption services derived from them do not coincide. To obtain the latter, I assume that the service flow from durables, s^d, is proportional to its stock at the beginning of every period:

$$s_t^d = \kappa d_t, \kappa > 0 \tag{8}$$

3.1.3 Earnings Process and Idiosyncratic Shocks

During the working years, households receive an exogenous stream of labor income exposed to idiosyncratic fluctuations. To avoid confounding *private* consumption insurance with *public* government insurance, my income measure of reference is post-government households' earnings–that is, after transfers and taxes. I will then make use of a tax function to recover pre-government earnings, following Kaplan and Violante (2010).

Specifically, log labor income is the sum of a common deterministic age profile g_t^a and a household-specific stochastic component y_{it}. The latter has two elements: a transitory and a persistent element, with autoregressive coefficient ρ. Transitory shocks are normally distributed[13] with mean 0 and standard deviation σ_ε. Equations (9)-(12) formally summarize these relations.

[13] As pointed out in the introduction, the focus of this paper is the impact of potentially large shocks of persistent nature. I, therefore, model transitory fluctuations in the standard fashion.

Finally, equation (13) specifies the distribution of shocks to the persistent component. This is a crucial element of my analysis. In particular, η_{it} follows a mixture of two normals: with probability p, η_{it} will be drawn from a normal distribution with mean μ_1 and standard deviation σ_1; and with probability $1-p$ from a normal distribution with mean μ_2 and standard deviation σ_2. This type of distribution is simple but flexible enough to match the higher-order moments observed in the data.

$$\log Y_{it} = g_t^a + y_{it} \tag{9}$$

$$y_{it} = z_{it} + \varepsilon_{it} \tag{10}$$

$$z_{it} = \rho z_{it-1} + \eta_{it} \tag{11}$$

$$\varepsilon_{it} \sim \mathcal{N}(0, \sigma_\varepsilon) \tag{12}$$

$$\eta_{it} \sim \begin{cases} \mathcal{N}(\mu_1, \sigma_1) & \text{with prob.} \quad p \\ \mathcal{N}(\mu_2, \sigma_2) & \text{with prob.} \quad 1-p \end{cases} \tag{13}$$

3.1.4 Solution

I solve the model numerically, proceeding by backward induction and using the Endogenous Grid Method (Carroll, 2006; Barillas and Fernández-Villaverde, 2007). I apply the variant of the method developed in Fella (2014) to solve for the value and policy functions of both the continuous consumption-savings choice and the discrete decision of upgrading, downgrading, or not adjusting the stock of durables. Because my solution algorithm is an application of Fella (2014) in a life-cycle environment, I relegate the details to Appendix B.

3.2 Parametrization

The first period corresponds to age 25, retirement happens at age 60, and everybody dies at age 95, that implies $T_R = 35$ and $T = 70$. For the parametrization, I proceed in two steps: First, I estimate the income process characterized in equations (9)-(13) using Simulated Method of Moments. The targets are primarily second and higher order moments of the distributions of four-year income changes, as well as life-cycle restrictions on the level of income. The complete list is provided in Table VI. Second, to parametrize the rest of the model, I externally measure a subset of the parameters that have straightforward data counterparts or reliable evidence and then calibrate the remaining to target moments of the cross-sectional distribution o

3.2.1 Externally Calibrated Parameters

Preferences The coefficient of relative risk aversion is fixed at $\gamma = 2$.

Utility. The interest rate is fixed at $r = 4\%$, based on empirical evidence on the risk-free rate of U.S. Treasury Bonds in McGrattan and Prescott (2000)[14]. Given the choice for r, I then calibrate β to target the empirical value for the median wealth to median income ratio of households, which is equal to 1.35.

Share of nondurables in total consumption Given the Cobb-Douglas specification chosen for the consumption bundle, I measure α as the share of nondurable goods in total consumption in my household sample. This parameter is often found to be around 0.8 (Luengo-Prado, 2006) or even larger (Berger and Vavra, 2015). I find it to be closer to 0.7, given the consumption categories included in my benchmark sample. Table IV includes the different values for typically used definitions of nondurable consumption.

Table IV: Share of Nondurable Consumption in Total Consumption

Definition	α
C_1 = Food + Utilities + Nondurable Transportation + Recreation	**0.704**
$C_2 = C_1$+ Rent	0.725
$C_3 = C_2$+ Health + Education + Child Care	0.803

Notes: α is the Cobb–Douglas share in eq. (4). Durables are D = Cars + Houses + Furnishings + Repairs + Clothing. Boldface denotes the benchmark α used in calibration. All series are annual household expenditures.

Depreciation of durable goods To calculate the depreciation rate of durables, I use data from the BEA's NIPA and Fixed Assets and Consumer Durable Goods. In particular, I compute a weighted average of the depreciation for stock of durables and housing, where the weights are given by the relative size of each group. This gives an annual depreciation rate of $\delta = 0.072$.

[14] 4% is also around the average of the values used in related literature. I test robustness to changing this value to $r = 3\%$, as in Kaplan and Violante (2010), and $r = 5\%$, as in Berger and Vavra (2015).

Service flow of durable goods The flow of services derived from the stock of durable consumption, κ, is similarly calculated using aggregate data from the Flow of Funds and the BEA. It is measure to be $\kappa = 0.035$. This is, a car worth \$10000 provides yearly services for the value of \$350.

Survival Probabilities Conditional survival probabilities from the U.S. Life Tables.

Deterministic age profile This series is obtained as the predicted value of a regression of income after transfers on a quadratic on age and a set of education and year dummies.

Initial distribution of assets and durables Distribution of assets and durables, relative to income in the sample, respectively.

3.2.2 Internally Calibrated Parameters

Given the parameters described in the previous section and the exogenous process for idiosyncratic labor income, the critical parameters that determine the how households adjust durable consumption are the share of collateralized borrowing λ^d, the discount rate β, and the two parameters of the adjustment cost function χ_0 and χ_1. I choose these parameters in a second SMM to match seven moments of the distribution of household consumption and wealth: the median of household wealth in my sample, the share of households with negative or zero assets, the C-S measure of Kurtosis of 2- and 4-year changed of both nondurable and durable consumption.

Table V: Internally Calibrated Parameters

		Value
β	(Discount factor)	0.976
λ^d	(Collateralized borrowing)	0.720
χ_1	(Adjustment costs parameter)	0.108
χ_2	(Adjustment costs parameter)	0.008

.

Notes: Parameters are chosen by Simulated Method of Moments to match seven distributional targets for consumption and wealth (median wealth/income, share with nonpositive assets, and Crow–Siddiqui kurtosis and standard deviations of 2- and 4-year changes for nondurables and durables). β is the annual discount factor; λ^d is the collateral (down-payment) share that tightens the borrowing constraint $a_{t+1} \geqslant -\lambda^y \underline{y}_t - \lambda^d d_{t+1}$ with $\lambda^y = 0$ in the baseline; χ_0 and χ_1 parameterize the non-convex adjustment cost $A(d_{t+1}, d_t) = \chi_1 d_t + \chi_2 |d_{t+1} - (1-\delta) d_t|$ (eq. (7)), where the first term generates an inaction region and the second scales with the size of adjustment. Units: β and λ^d are unit-free; χ_0 loads on the beginning-of-period stock d_t; χ_1 loads on the absolute adjustment $|d_{t+1} - (1-\delta) d_t|$.

Income Process I use Simulated Method of Moments to estimate the parameters controlling the dynamics of the stochastic income component Θ^y, which include:

$$\Theta^y \equiv \{p, \rho, \mu_1, \sigma_1, \sigma_2, \sigma_\varepsilon, \sigma_{z_0}\}$$

I make the assumption that $\mu_2 = \frac{-(1-p)}{p}\mu_1$, which simply makes sure the mean of Δy is zero. This assumption allows the method of moments to focus on targeting higher-order moments without much loss, since matching the average of changes is relatively easy.

The targeted moments include the variance, Kelley Skewness, and Crow-Siddiqui Kurtosis of the moments of two- and four-year income and consumption changes.

Table VI: Income Process Estimates

$\boldsymbol{\sigma_\varepsilon}$	(Std. Dev. of transitory shock)	0.053
$\boldsymbol{p}$	(Probability of drawing from normal 1)	0.930
$\boldsymbol{\rho}$	(Persistence)	0.913
$\boldsymbol{\mu_1}$	(Mean of 1 persistent shock)	0.008
$\boldsymbol{\mu_2}$	(Mean of 2 persistent shock)	-0.106
$\boldsymbol{\sigma_1}$	(Std. Dev. of 1 persistent shock)	0.075
$\boldsymbol{\sigma_2}$	(Std. Dev. of 2 persistent shock)	1.189
$\boldsymbol{\sigma_{z_0}}$	(Std. Dev. of initial distribution)	0.753

.

Notes: Estimates of the stochastic component of earnings in eqs. (9)–(11). $y_{it} = z_{it} + \varepsilon_{it}$ with $\varepsilon_{it} \sim \mathcal{N}(0, \sigma_\varepsilon)$ (transitory) and $z_{it} = \rho z_{it-1} + \eta_{it}$ (persistent). The shock η_{it} follows a two-component Normal mixture (eq. (13)): with probability p, $\mathcal{N}(\mu_1, \sigma_1)$, and with probability $1 - p$, $\mathcal{N}(\mu_2, \sigma_2)$. The mixture captures higher-order moments (skewness and tail weight) in income changes. σ_{z_0} is the variance of the initial persistent component. Parameters are estimated by SMM targeting the variance and higher-order moments of 4-year income changes and the age profile of the variance of income levels. To focus on tails rather than mean shifts, we impose $\mu_2 = -\frac{1-p}{p}\mu_1$ so the expected change is approximately zero.

4 Results

In this section, I first evaluate the performance of the model in replicating the higher-order moments described in the empirical section, as well as the mechanisms at work. Next, I measure to what extent income shocks pass-through to consumption, sepaaretly for shocks of different size. I conclude with a welfare calculation.

4.1 Model Fit

Table VII compares the data moments with those implied by the model. The data corresponds to the moments of the change in the residual income and consumption after controlling for a quartic in age, and education and year dummies. Hence, the moments are not identical to those in Table I. The model is over-identified, hence it is not surprising that it does a good job matching targets. What is more interesting, though, is the fact that durable upgrade and downgrade probabilities, which are non-targeted, are matched quite closely. This is because the fourth moment in durable adjustment is a measure of lumpiness, and hence captures the probability of adjustment.

Table VII: Model Fit

	Data	Model
Cross-sectional moments (Income)		
SD($\log\Delta^2 y_t^{post}$)	0.585	0.589
KS($\log\Delta^2 y_t^{post}$)	-0.055	-0.005
CS($\log\Delta^2 y_t^{post}$)	6.890	6.711
SD($\log\Delta^4 y_t^{post}$)	0.660	0.532
KS($\log\Delta^4 y_t^{post}$)	-.086	-.113
CS($\log\Delta^4 y_t^{post}$)	5.921	6.010
Cross-sectional moments (Consumption)		
SD($\log\Delta^2 c_t$)	0.481	0.211
SD($\log\Delta^2 d_t$)	0.813	0.903
CS($\log\Delta^2 c_t$)	3.523	2.822
CS($\log\Delta^2 d_t$)	23.687	30.687
SD($\log\Delta^4 c_t$)	0.530	0.519
SD($\log\Delta^4 d_t$)	1.004	1.192
CS($\log\Delta^4 c_t$)	3.375	3.002
CS($\log\Delta^4 d_t$)	16.836	18.281
Non-targeted model implications		
% Households adjusting/year	15.212%	16.942%
% Households upgrading/year	8.028%	10.102%
% Households downgrading/year	7.084%	6.840%

Notes: The table compares data targets to the model under the calibrated parameterization. All moments are computed from residualized series: income and consumption are regressed on a quartic in age and education and year dummies, and moments are taken on the residuals. SD($\log \Delta^h x_t$) is the standard deviation of h-year log changes; KS is Kelley skewness $\big(Q_{0.9} + Q_{0.1} - 2Q_{0.5}\big)/\big(Q_{0.9} - Q_{0.1}\big)$ (0 under Normality); CS is the Crow–Siddiqui tail ratio $\big(Q_{0.975} - Q_{0.025}\big)/\big(Q_{0.75} - Q_{0.25}\big)$ (2.91 under Normality). "post" denotes earnings after taxes and transfers. Bold-labeled rows are targeted in the SMM; others are non-targeted implications. Durable adjustment rates (adjust/upgrade/downgrade per year) are non-targeted and closely matched, reflecting that kurtosis in durable changes is informative about the frequency and size of lumpy adjustments.

Lumpiness and deviations from normality

Figure 5 shows the model counterpart of Figure 2. Besides replicating the data well, it is interesting to notice the behavior of durables adjustment. Coming from the model with durable adjustments, it is easy to see how the lumpiness translates into the QQ plot. With the intermediate quantiles all equal to zero. In other words, households only downgrade their durables when they receive a tail shock. It is interesting to see that there is some asymmetry between positive and negative changes, with the negative side being more lumpy. This is because of depreciation and semi-durable purchases. While the only reason why a household would downgrade their stock of durables, on top of age effects which are removed from this picture, is because of an income shock, households may choose to repair their current home or upgrade to a new one due to depreciation.

Figure 5: Deviations from Normality: Durables vs. Nondurables

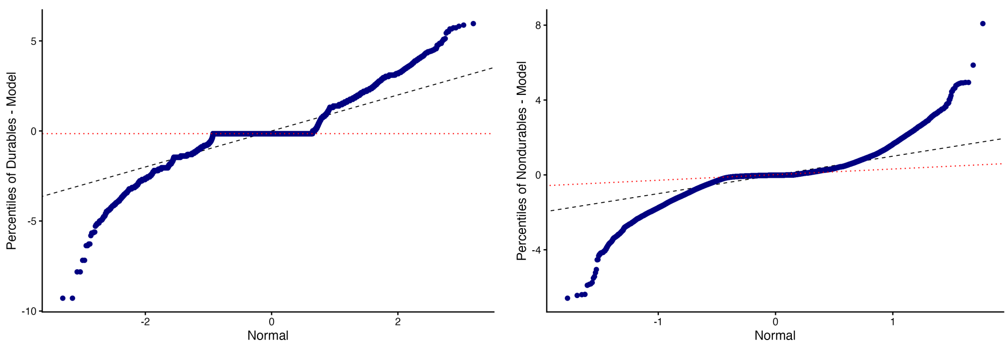

Notes: Model counterpart to Figure 2. Durables feature infrequent, lumpy adjustments, which show up as a flat segment around zero in the empirical quantile plot (left): most households make no durable change in typical years, and downgrades cluster after large negative income shocks. Nondurables adjust smoothly with no mass at zero. The QQ lines (right) depart from the 45°—leptokurtic tails for both series and stronger left–tail nonlinearity for durables—reflecting depreciation/semi-durable replacements and asymmetric responses to adverse shocks.

4.2 The Micro Consumption Response to (Tail) Income Shocks

Traditional measures of pass-through have relied on covariances between changes in income and consumption. Specifically,

$$\phi^{c,\eta} = \frac{\mathrm{Cov}(\Delta c_t, \eta_t)}{\mathrm{Var}(\eta_t)}, \qquad \phi^{c,\varepsilon} = \frac{\mathrm{Cov}(\Delta c_t, \varepsilon_t)}{\mathrm{Var}(\varepsilon_t)},$$

where η_t and ε_t denote the persistent and transitory income innovations, respectively, in my structural model. In practice, an empirical analogue instruments η_t (or ε_t) with functions of observed income changes. Under the linear specification $\Delta c_t = \beta^\eta \eta_t + \beta^\varepsilon \varepsilon_t + u_t$, orthogonal shocks, and u_t mean-independent of the shocks, these covariance ratios recover the slope coefficients: $\phi^{c,\eta} = \beta^\eta$ and $\phi^{c,\varepsilon} = \beta^\varepsilon$.

Following Blundell *et al.* (2008), I map these slopes into *partial insurance* via one minus pass-through:

$$\pi^\eta \equiv 1 - \beta^\eta, \qquad \pi^\varepsilon \equiv 1 - \beta^\varepsilon.$$

These measures have proven to be informative about the amount of insurance on top of self-insurance under certain linearity assumptions. They provide, however, a limited measure of partial insurance when income changes deviate from the standard case. In particular, two are the limitations that are crucial for my analysis: First, $\phi^{c,\eta}$ is a relative measure of the impact of the shock since it is divided by the variance. That is, even if we were to limit risk to second order variation, the size of the shock is irrelevant. Second, when higher-order moments of income shocks are nontrivial, the empirical analysis has shown that covariances can be misleading to represent the joint dynamics at the tails (Guvenen *et al.*, 2024).

In the presence of heterogeneity and nonlinearity, the same mapping applies pointwise—for example, by age a or along the distribution of outcomes:

$$\pi^\eta(a,\tau) = 1 - \beta^\eta(a,\tau), \qquad \pi^\varepsilon(a,\tau) = 1 - \beta^\varepsilon(a,\tau),$$

where $\beta^\cdot(a,\tau)$ comes from quantile regressions $Q_\tau(\Delta c \mid \Delta y)$.

The top row of Figure 6 shows nondurable responses to the persistent component. Two patterns stand out. First, coefficients decline steeply with age—high in the late 20s and falling toward zero by retirement—consistent with buffer-stock accumulation and rising insurance against permanent shocks. This is common in life-cycle models (e.g., Kaplan and Violante, 2010). Second, conditioning on extreme realizations (P1 for very negative changes or P99 for very positive changes) widens the wedge around the average only under the non-Gaussian process: tail shocks elicit substantially smaller nondurable adjustments than moderate shocks, whereas the Gaussian benchmark yields little heterogeneity across tails. Because the calibration is held fixed across columns, these gaps isolate the role of higher-order risk in amplifying state dependence. Note too that this difference decreases with age.

The bottom row turns to the transitory component. Levels are much smaller throughout and flatten quickly after age 35, indicating that short-lived shocks are smoothed primarily through liquid assets. Tails matters far less here, and again the Gaussian specification compresses heterogeneity even further. Taken together, the figure implies that (i) higher-order risk is quantitatively first-order for responses to persistent shocks, (ii) age gradients in MPCs emerge endogenously from wealth accumulation, and (iii) assuming Gaussian income innovations would materially understate the dispersion and tail sensitivity of nondurable adjustment over the life cycle.

Figure 6: Non-Durable Consumption coefficients under non-Gaussian (left column) and Gaussian (right column) shocks.

Non-Gaussian **Gaussian**

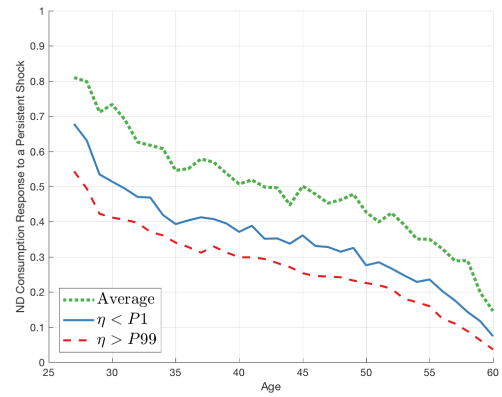

(a) η shock

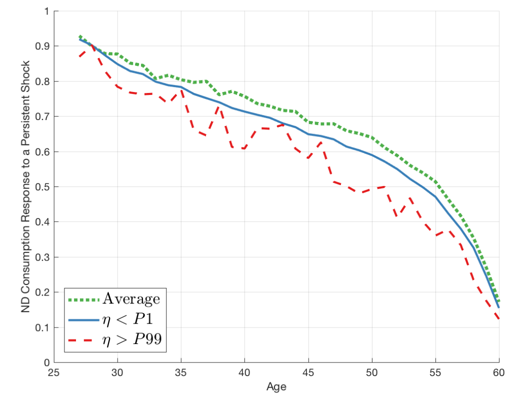

(b) η shock

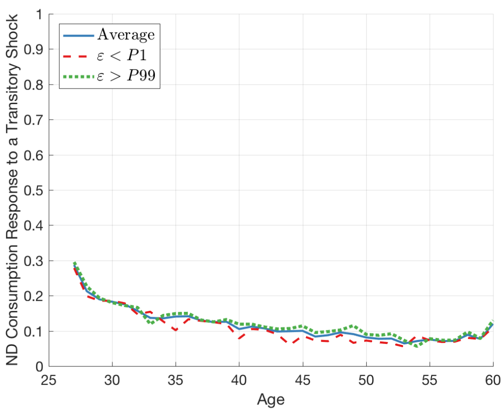

(c) ε shock

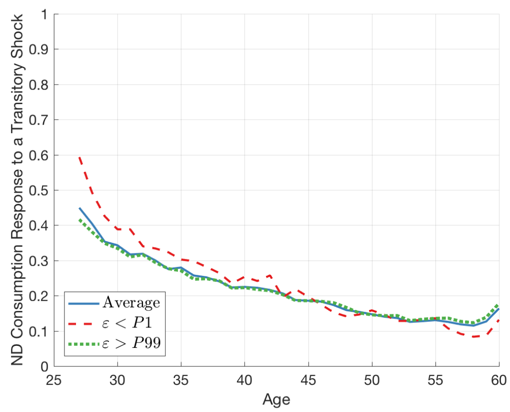

(d) ε shock

Notes: Panels report linear consumption coefficients from the model under non-Gaussian (left column) and Gaussian (right column) innovations. The top row shows responses to the persistent shock η; the bottom row to the transitory shock ε. Lines plot coefficients across horizons (biennial units); when present, shaded bands are model-simulated uncertainty bands. Parameters and calibration are held fixed across columns so differences reflect only the shock distribution. Values are in log changes and comparable across panels.

Through the lens of the BPP insurance coefficients, two implications emerge from Figure 6. First, $\beta^{\eta}(a)$ declines with age, so insurance against persistent shocks, $\pi^{\eta}(a)$, rises over the life cycle, consistent with buffer–stock accumulation. Second, under non–Gaussian risk the tail coefficients $\beta^{\eta}(a, \tau \in \{0.01, 0.99\})$ are markedly smaller than at the center, implying *more* partial insurance for extreme persistent shocks. By contrast, for transitory shocks $\beta^{\varepsilon}(a)$ is small and fairly flat, so $\pi^{\varepsilon}(a)$ is high and only weakly state dependent. Hence, assuming Gaussian innovations would understate insurance against persistent tail shocks and its heterogeneity.

To complement the analysis and overcome the limitations present in covariance measures, I calculate the response functions at different horizons to different sizes of income shocks for both consumption measures.

Figure 7: Response to Income Shocks of Different Size

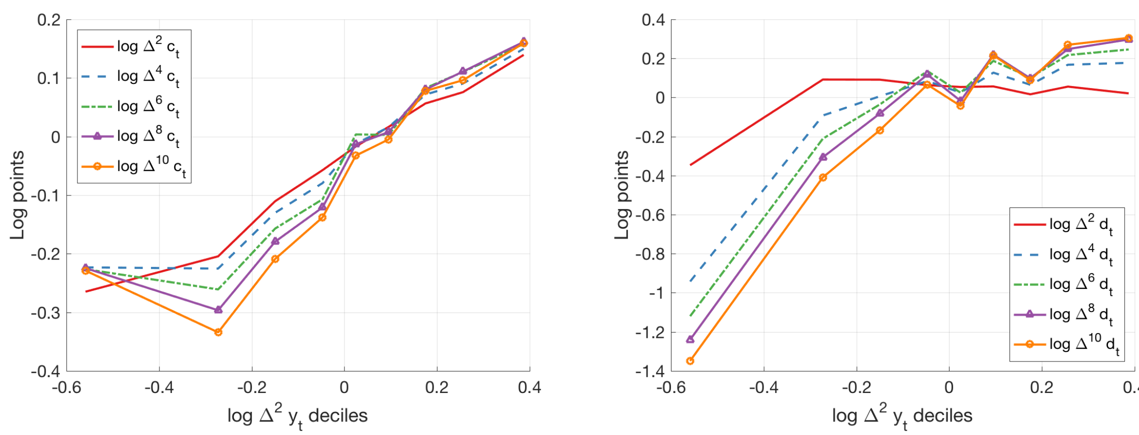

Notes: Each panel plots average consumption responses by the size of an income change. The x-axis groups observations into 10 equal-sized bins (deciles of $\log \Delta^2 y_t$); the y-axis reports the mean of $\log \Delta^h c_t$ within each bin, for horizons $h \in \{2, 4, 6, 8, 10\}$ years (lines). Left: nondurable consumption; right: durable consumption.

The x-axis shows the deciles of the income changes distribution. The y-axis contains the average response to this size of shock. The different lines correspond to different horizons, from 1 to 10 years ahead.

Notice that these figures are consistent with the idea that, in the event of large shocks, households adjust their durable stock in an unproportionate size. Moreover, the negative effect on consumption from negative income shocks follows a U-shape in the negative quadrant, getting closer to zero as the size of the durable adjustment increases. This effect is stronger for negative changes and the longer the horizon.

4.3 The Welfare Costs of Idiosyncratic Income Risk with Tail Shocks

In this section, I answer two questions: (1) How much worse is income uncertainty over the life-cycle when we consider the impact of tail shocks?, and (2) What is the insurance value–or cost–of durable consumption?

To answer these two questions, I perform the following welfare exercise: Using the calibrated structural model, I calculate the fraction of consumption at every age that households would be willing to give up to move from a world with durables and non-Gaussian shocks ($\mathsf{NG-D}$) to alternative worlds. Before getting into more details, it is useful to define my measure of welfare and consumption equivalent variation.

Define ex-ante welfare as the expected lifetime utility of a household i at time 0:

$$\mathcal{W}_i(\mathcal{C}) \equiv U_i\left(\{\mathcal{C}_{it}\}_{t=1}^{T}\right) \equiv \mathbb{E}_0 \sum_{t=1}^{T} \beta^{t-1}\xi_t u(\mathcal{C}_{it}), \tag{14}$$

where u is CRRA and $\mathcal{C}$ is given by equation 4 and denotes the Cobb-Douglas consumption bundle of non-durable consumption and durable consumption services.

Let $\mathcal{W}_i^{ng,d}$ be the ex-ante welfare from consuming the stream of $\mathcal{C}$ chosen in the benchmark case with non-Gaussian shocks and durable adjustment. Equivalently, I define $\mathcal{W}_i^{g,d}$ and $\mathcal{W}_i^{ng,nd}$ as the alternative worlds when shocks are Gaussian (g) or households cannot adjust durables (nd), respectively. Finally, let $\mathcal{W}^c$ be the complete-markets reference.

To answer the first question, I solve for the $\theta^j \in [0,1]$, for $j = g, ng$ in

$$\mathcal{W}_i^{ng,d}((1-\theta^{ng})\,\mathcal{C}) = \mathcal{W}^c \tag{15}$$

$$\mathcal{W}_i^{g,d}((1-\theta^{g})\,\mathcal{C}) = \mathcal{W}^c \tag{16}$$

Table VIII reports consumption–equivalent variations (CEV) for eliminating idiosyncratic risk. In the baseline non-Gaussian economy with durables, the certainty–equivalent fraction is about $\theta^{ng} \approx 0.15$: households would give up roughly 15 percentage points of consumption at every age and state to live in a riskless world. This is larger than under a

	No Durables	Durables
Gaussian	5%	6%
Non-Gaussian	19%	15%

Table VIII: Welfare Calculations

Gaussian calibration, reflecting the heavier tails documented in the data. The difference between the two, $\theta^{ng} - \theta^{g} \approx 0.09$, is the welfare cost attributable to higher–order risk.

$$\mathcal{W}_i^{ng,nd}((1-\theta^{ng})\,c) = \mathcal{W}^{c,nd} \tag{17}$$

$$\mathcal{W}_i^{g,nd}((1-\theta^{g})\,c) = \mathcal{W}^{c,nd} \tag{18}$$

In the nondurable–only environment the gap $\theta^{ng,nd} - \theta^{g,nd}$ is substantially larger than $\theta^{ng} - \theta^{g}$ (14pp. compared to 9pp.), showing that durable adjustment mitigates the welfare cost of tail risk: when the durable margin is shut down, nondurable volatility and pass–through both rise, and the incremental cost of moving from Gaussian to non-Gaussian risk increases.

Quantitatively, higher–order risk accounts for about 9% of annual consumption in the baseline with durables. Without the durable margin, this welfare loss would be considerably larger. This pattern is consistent with evidence that individual income changes feature excess kurtosis and asymmetric tails and that large adverse shocks are closely associated with larger consumption adjustments, in line with Guvenen *et al.* (2024).

5 Conclusion

This paper measures the joint dynamics of earnings and consumption changes with a focus on the tails and shows that departures from normality are a central feature of both series. Using household panel data, I document that excess kurtosis persists even after taxes and transfers and that durable consumption changes are especially leptokurtic. The joint evidence indicates that nondurable spending tracks income more closely in the center of the distribution, while extreme income changes are more strongly associated with durable adjustments. These facts motivate a life-cycle, incomplete-markets model with lumpy durables and higher-order income risk.

In the model, a flexible earnings process and a lumpy durable margin reproduce the main patterns in the data: little movement for typical changes and large, infrequent adjustments in the tails. Quantile-based measures make clear that the transmission of income shocks to consumption is not well summarized by average pass-through alone. In particular, nondurables absorb moderate fluctuations, whereas extreme shocks are accompanied by discrete movements in the durable stock. Standard insurance statistics remain informative on average but conceal the role of tail changes for both measurement and interpretation.

The model delivers two implications. First, allowing for tail risk and lumpy durables helps reconcile the nonlinear response to shocks of different size observed in the data. Second, durable adjustment plays an insurance role by absorbing part of the impact of very large shocks; removing that margin raises the cost of risk. A welfare calculation indicates that higher-order risk increases the consumption-equivalent cost of idiosyncratic income risk, and that the presence of a durable margin mitigates this loss.

Taken together, the results suggest that incorporating tail behavior and durable adjustment is essential for measuring insurance, evaluating the welfare cost of income risk, and understanding the link between household risk and aggregate consumption dynamics.

References

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 ce. *American Economic Review*, **104** (5), 132–35.

ARELLANO, M., BLUNDELL, R. and BONHOMME, S. (2017). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica*, **85** (3), 693–734.

BAGGER, J., FONTAINE, F., POSTEL-VINAY, F. and ROBIN, J. (2014). Tenure, experience, human capital and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review*, **104** (6), 1551–1596.

BARILLAS, F. and FERNÁNDEZ-VILLAVERDE, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, **31**, 2698–2712.

BERGER, D. and VAVRA, J. (2015). Consumption Dynamics During Recessions. *Econometrica*, **83**, 101154.

BLUNDELL, R., PISTAFERRI, L. and PRESTON, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, **98** (5), 1887–1921.

BONHOMME, S. and ROBIN, J. (2009). Assessing the equalizing force of mobility using short panels: France, 1990 a 2000. *Review of Economic Studies*, **76** (1), 63–92.

BROWNING, M. and CROSSLEY, T. F. (2009). Shocks, stocks, and socks: Smoothing consumption over a temporary income loss. *Journal of the European Economic Association*, **7** (6), 1169–1192.

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

CERLETTI, E. and PIJOAN-MAS, J. (2012). *Durable Goods, Borrowing Constraints and Consumption Insurance*. CEPR Discussion Papers 9035, C.E.P.R. Discussion Papers.

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

Daly, M., Hryshko, D. and Manovskii, I. (2016). *Improving the Measurement of Earnings Dynamics*. NBER Working Papers 22938, National Bureau of Economic Research, Inc.

Deaton, A. and Paxson, C. (1994). Intertemporal choice and inequality. *Journal of Political Economy*, **102** (3), 43767.

DeNardi, M. (2002). *Wealth inequality and intergenerational links*. Staff Report 314, Federal Reserve Bank of Minneapolis.

Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, **17** (2), 329–344.

Flavin, M. and Nakagawa, S. (2004). *A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence*. NBER Working Papers 10458, National Bureau of Economic Research, Inc.

Geweke, J. and Keane, M. (2000). An empirical analysis of earnings dynamics among men in the PSID: 19681989. *Journal of Econometrics*, **96**, 293–356.

Grossman, S. J. and Laroque, G. (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica*, **58** (1), 2551.

Guvenen, F., Karahan, F., Ozkan, S. and Song, J. (2015). *What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?* Tech. rep., National Bureau of Economic Research, Cambridge, MA.

—, —, — and — (2021). What do data on millions of U.S. workers reveal about lifecycle earnings dynamics? *Econometrica*, **89** (5), 2303–2339.

—, Ozkan, S. and Madera, R. (2024). Consumption dynamics and welfare under non-gaussian earnings risk. *Journal of Economic Dynamics and Control*, **169**, 104945, special Issue: Heterogeneity and Macroeconomics of Labor Markets.

— and Smith, A. A. (2014). Inferring Labor Income Risk and Partial Insurance from Economic Choices. *Econometrica*, **82** (6), 2085–2129.

Heathcote, J., Perri, F. and Violante, G. L. (2010). Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 19672006. *Review of Economic Dynamics*, **13** (1), 1551.

Johnson, D. S., Parker, J. A. and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, **96** (5), 1589–1610.

Kaplan, G. and Violante, G. L. (2010). How much consumption insurance beyond selfinsurance? *American Economic Journal: Macroeconomics*, **2** (4), 5387.

— and — (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, **82** (4), 1199–1239.

Luengo-Prado, M. J. (2006). Durables, nondurables, down payments and consumption excesses. *Journal of Monetary Economics*, **53** (7), 1509–1539.

McGrattan, E. R. and Prescott, E. C. (2000). Is the stock market overvalued? *Quarterly Review*, (Fall), 2040.

Nardi, M. D., Fella, G. and Paz-Pardo, G. (2020). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association*, **18** (2), 890–926.

Piazzesi, M. and Schneider, M. (2007). Asset Prices and Asset Quantities. *Journal of the European Economic Association*, **5** (23), 380–389.

Primiceri, G. E. and van Rens, T. (2009). Heterogeneous lifecycle profiles, income risk and consumption inequality. *Journal of Monetary Economics*, **56** (1), 20 39, carnegieRochester Conference Series on Public Policy: The Causes and Consequences of Rising Economic Inequality April 2526, 2008.

Shin, D. and Solon, G. (2011). Trends in men's earnings volatility: What does the panel study of income dynamics show?,. *Journal of Public Economics*, **95** (78), 973–982.

Storesletten, K., Telmer, C. I. and Yaron, A. (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, **51** (3), 609–633.

Toda, A. A. and Walsh, K. J. (2015). The double power law in consumption and implications for testing euler equations. *Journal of Political Economy*, **123** (5), 1177–1200.

van Oordt, M. R. and Zhou, C. (2012). The simple econometrics of tail dependence. *Economics Letters*, **116** (3), 371 373.

Appendices

A Data Appendix

This section describes the variables used in the analysis. The majority of the analysis is done with PSID, so the description is more detailed for this dataset.

A.1 The PSID

A.2 Structure and weights

Four different household samples compose the current version of the PSID from 1968 to 2013: (1) the Survey Research Center (SRC), (2) the Survey of Economic Opportunity (SEO), (3) the Latino sample, and (4) the Immigrant sample. The SRC, usually referred to as core sample, corresponds to a representative sample of the U.S. population in 1967 and their offsprings in later years. Most studies based on the PSID use this subsample only. The SEO also begins with the first available wave and included an additional set of low-income households. In 1990, 2000 Latino families were added and then dropped in 1995. Due to its short span, this sample is rarely used. Finally, a nationally representative sample of immigrant households that were not eligible in 1968 starts being surveyed in 1997.

All of these samples are probability samples with equal weights. Their combination, however, has unequal selection probabilities. I make use of the cross-sectional weights for the core, SEO, and immigrant samples. I do not use the Latino sample.

A.3 Variables

Head and Relationship to Head. I identify *current* heads and spouses as those individuals within the family unite with `Sequence Number` equal to 1 and 2, respectively. In the PSID, the man is labelled as the household head and the woman as his spouse. Only when the household is headed by a woman alone, she is considered the head. If the family is a split-off family from a sampled family, then a new head is selected.

Post-Government Household Labor Earnings. Pre-government household earnings *minus* taxes *plus* public transfers, as defined below. I construct an alternative version by subtracting `household capital income` from `family money` (i.e. disposable income) and the correlation is 0.98.

Taxes. Federal and state labor income taxes after credits. Estimated using TAXSIM.

Public Transfers. Transfers are considered at the family unit level, when possible. Broadly, the transfers included are unemployment benefits, welfare, and disability insurance. They are defined as in **?**, an extensive discussion and specific description is given in their Data Appendix.

Pre-Government Household Earnings. Head and spouse earnings, without self-employment.

Individual Head Labor Earnings. `Annual Total Labor Income` includes all income from wages and salaries, commissions, bonuses, overtime> I remove the labor part of self-employment (farm and business income)[15].

Individual Spouse Labor Earnings. Same definition as head's earnings for the spouse.

Variables not used in the main analysis for sample selection or controls include:

Education Level. Highest education level that an individual ever reports.

Annual Hours. Sum of annual hours worked on main job, extra jobs and overtime. It is computed using usual hours of work per week times the number of actual weeks worked in the last year.

[15] Self-employment income is split between asset and labor income in a somewhat arbitrary manner. See Shin and Solon (2011) for a detailed discussion.

A.4 Detailed Summary Statistics and Extra Moments

Table A.1: Tails and Higher-Order Moments of Earnings and Consumption

	Std. Dev.	**L9010**	Skewness	**Kelley Sk.**	Kurtosis	**C-Siddiqui K.**	**L9050**	**L5010**
				2-Year Changes				
$\log \Delta^2 y_t^{ind}$	0.654	0.875	-0.475	-0.045	20.926	9.899	0.418	0.457
Y	0.660	0.922	-0.355	-0.027	22.471	8.701	0.449	0.473
O	0.649	0.816	-0.587	-0.074	19.527	11.059	0.378	0.438
$\log \Delta^2 y_t^{hh}$	0.571	0.846	-0.193	-0.043	24.045	7.018	0.405	0.441
Y	0.584	0.909	0.107	-0.048	25.110	6.542	0.433	0.476
O	0.559	0.793	-0.487	-0.040	22.852	7.444	0.380	0.412
$\log \Delta^2 y_t^{post}$	0.585	0.760	-0.980	-0.055	43.711	6.890	0.359	0.401
Y	0.580	0.772	-0.026	-0.035	39.629	6.515	0.372	0.400
O	0.589	0.753	-1.763	-0.071	46.930	7.355	0.350	0.403
$\log \Delta^2 c_t$	0.471	0.999	0.025	0.004	9.426	4.000	0.501	0.498
Y	0.488	1.007	-0.032	-0.002	10.729	4.141	0.502	0.505
O	0.455	0.987	0.077	0.006	7.874	3.937	0.496	0.491
$\log \Delta^2 d_t$	0.813	0.771	0.829	0.447	16.164	23.687	0.558	0.213
Y	0.936	1.125	0.857	0.510	12.477	26.671	0.849	0.276
O	0.692	0.537	0.575	0.339	21.330	19.182	0.360	0.178
				4-Year Changes				
$\log \Delta^4 y_t^{ind}$	0.725	1.082	-0.551	-0.126	19.349	8.151	0.473	0.609
Y	0.746	1.135	-0.822	-0.121	19.109	7.797	0.499	0.636
O	0.703	1.018	-0.237	-0.164	19.603	8.755	0.425	0.592
$\log \Delta^4 y_t^{hh}$	0.632	1.044	-0.306	-0.095	21.097	5.838	0.473	0.572
Y	0.654	1.069	-0.456	-0.090	19.377	5.854	0.487	0.583
O	0.611	1.006	-0.119	-0.089	23.104	5.782	0.459	0.548
$\log \Delta^4 y_t^{post}$	0.660	0.914	-1.297	-0.086	37.590	5.921	0.418	0.496
Y	0.683	0.910	-1.095	-0.079	32.716	5.974	0.419	0.491
O	0.636	0.918	-1.539	-0.092	43.565	5.827	0.417	0.501
$\log \Delta^4 c_t$	0.514	1.107	-0.464	-0.041	10.883	3.796	0.531	0.576
Y	0.537	1.127	-0.942	-0.052	11.435	3.925	0.534	0.592
O	0.490	1.089	0.144	-0.032	10.032	3.775	0.527	0.562
$\log \Delta^4 d_t$	1.004	1.235	0.824	0.451	10.885	16.836	0.896	0.339
Y	1.141	1.611	0.807	0.498	8.747	15.663	1.207	0.404
O	0.849	0.929	0.686	0.409	13.998	14.900	0.655	0.274

Note: **Moments.** Columns in dark gray denote the 2nd through 4th central moments of the distribution of each variable. Columns in black are the corresponding *robust* and other percentile-based measures. *P9010*: 90th/10th percentiles, *Kelley Sk.*: Kelley Skewness (**0 under a normal**), *C-Siddiqui K.*: Crow-Siddiqui Kurtosis (**2.91 under a normal**), *P9050*: 90th/50th percentiles, *P5010*: 50th/10th percentiles. **Variables.** y^{ind}: individual earnings (heads), y^{hh}: household pre-gov. earnings, y^{post}: households post-gov. earnings, c: nondurable consumption, d: durable consumption.

A.5 Detailed Sample Selection

I start with an initial sample of 105813 SRC and SEO households interviewed between 1992 and 2012. We then impose the next criteria every year. The number of individuals kept at each stage in the sample selection is listed in Table I. Previous to this selection process, I have cleaned the raw data and corrected duplicates. Outliers are considered bottom 0.2% and top 0.5% in order to obtain distributions for 1 year income changes that resemble those in ?. See the next appendix section for a comparison with their data.

	Observations Remaining
Start with	105813
No Major HH Composition Changes	89349
No Retirees	76627
Nonmissing Main Variables	70903
Drop Inconsistent Y and H	70806
Income Outliers	67413
Age in 25,60	58751
Not reliable $\widetilde{\mathsf{Y}}$	22373
Not enough consequtive obs	20954
Final # Observations	20954
Final # Households	4697

Table A.2: Detailed Sample Selection

A.6 Understanding QQ Plots

Figure A.1: Understanding QQ Plots

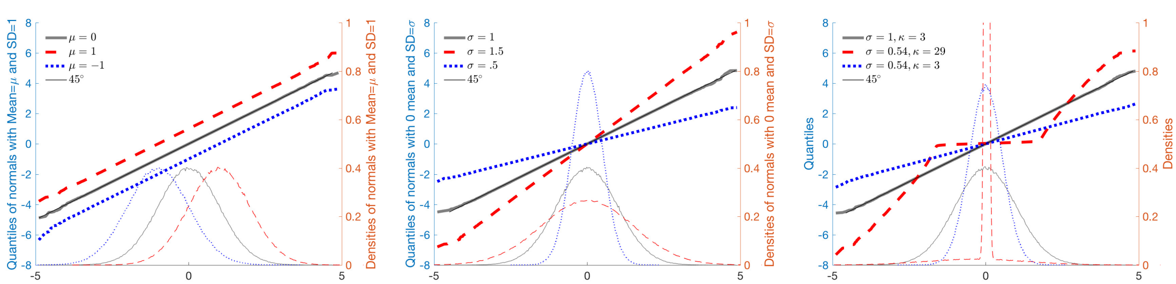

B Numerical Appendix

B.1 Solution Details

B.1.1 Optimality Conditions (Euler Equation)

Recall the maximization problem solved by working households at each age $t = 1, \cdots, T_r - 1$:

$$\begin{aligned}
V_t(a_t, d_t, z_t) &= \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} && (19) \\
\text{s.t.} \quad & c_t + a_{t+1} = Y_t + (1+r)a_t + (1-\delta)d_t - d_{t+1} - A(d_t, d_{t+1}) && (20) \\
& Y_t \text{ given by equations } (9) - (13) && (21) \\
& a_{t+1} \geqslant -\lambda^y \underline{y}_{t+1} - \lambda^d d_{t+1}, \quad c_t \geqslant 0, d_{t+1} \in D && (22)
\end{aligned}$$

To facilitate the solution of (19), I implement a set of transformations described in the next paragraphs.

First, let

$$b_{t+1} \equiv a_{t+1} + \lambda^y \underline{y}_{t+1} + \lambda^d d_{t+1}. \tag{23}$$

Then, the borrowing constraint can be rewritten as

$$\begin{aligned}
& c_t + (b_{t+1} - \lambda^y \underline{y}_{t+1} - \lambda^d d_{t+1}) = Y_t + (1+r)(b_t - \lambda^y \underline{y}_t - \lambda^d d_t) + (1-\delta)d_t - d_{t+1} - A(d_t, d_{t+1}) \\
& c_t + b_{t+1} = \underbrace{Y_t + (1+r)b_t + (1-\delta)d_t - (1-\lambda^d)d_{t+1} - A(d_t, d_{t+1}) - \lambda^d(1+r)d_t + \lambda^y(\underline{y}_{t+1} - (1+r)\underline{y}}_{m(d_{t+1})}
\end{aligned} \tag{24}$$

Notice that (24) defines a notion of cash in hand conditional on the choice of durables. $m(d_{t+1})$ denotes the total amount of resources available to be split between (non-durable) consumption and savings.

Next, I redefine V as follows

$$\tilde{V}_t = \max_{c_t, d_{t+1}, a_{t+1}} \left\{ \underbrace{\mathcal{C}(c_t, s_t)^{1-\gamma} + \beta \mathbb{E}_t \tilde{V}_{t+1}^{1-\gamma}(a_{t+1}, d_{t+1}; z_t)}_{\mathbb{V}_t} \right\}^{\frac{1}{1-\gamma}},$$

which yields an equivalent problem while reducing the curvature of the value function.

The dynamic programming problem to be solved is thus

$$\begin{aligned} \tilde{V}_t(b_t, d_t, z_t) &= \max_{d_{t+1}, b_{t+1}} \left\{ \mathcal{C}(m(d_{t+1}) - b_{t+1}, s_t)^{1-\gamma} + \beta \mathbb{E}_t \tilde{V}_{t+1}^{1-\gamma}(b_{t+1}, d_{t+1}; z_t) \right\}^{\frac{1}{1-\gamma}} \quad (25) \\ \text{s.t.} \quad & (b_{t+1}, d_{t+1}) \in \{ b_{t+1}, d_{t+1} : b_{t+1} \in [0, m(d_{t+1}; b_t, a_t, z_t)], d_{t+1} \in D \}. \end{aligned}$$

The FOC therefore are given by

$$\frac{1}{1-\gamma} \mathbb{V}_t^{\frac{1}{1-\gamma}-1} \left\{ -(1-\gamma)(\mathcal{C}_t)^{-\gamma} \underbrace{\left[\alpha c_t^{\alpha-1} s_t^{1-\alpha} \right]}_{\mathcal{C}_{c,t} \equiv \frac{\alpha}{c_t} \mathcal{C}_t} + \beta(1-\gamma) \widetilde{V}_{t+1}^{-\gamma} \widetilde{V}_{b,t+1} \right\} = 0$$

$$\mathcal{C}_t^{1-\gamma} \frac{\alpha}{c_t} = \beta \mathbb{E}_t \widetilde{V}_{t+1}^{-\gamma} \widetilde{V}_{b,t+1} \quad (\mathrm{FOC}_b)$$

And the envelope condition

$$(1-\gamma) \widetilde{V}_t^{-\gamma} \widetilde{V}_{b,t} = (1-\gamma)(1+r) \mathcal{C}_t^{-\gamma} \mathcal{C}_{c,t}$$

$$\widetilde{V}_t^{-\gamma} \widetilde{V}_{b,t} = (1+r) \mathcal{C}_t^{1-\gamma} \frac{\alpha}{c_t}. \quad (\mathrm{EC})$$

Combining (FOC_b) and (EC), I obtain the usual Euler Equation:

$$\mathcal{C}_t^{1-\gamma} \frac{\alpha}{c_t} = \beta(1+r) \mathbb{E}_t \mathcal{C}_{t+1}^{1-\gamma} \frac{\alpha}{c_{t+1}} \quad (\mathrm{EE})$$

B.1.2 Implementing the Endogenous grid Method

For a given choice of d_{t+1}, (EE) can be inverted to write optimal consumption c_t as a function of next period's assets and the income process specifics:

$$\left[c_t^{\alpha} s_t^{1-\alpha}\right]^{1-\gamma} \frac{\alpha}{c_t} = \underbrace{\beta(1+r)\mathbb{E}_t \mathcal{C}_{t+1}^{1-\gamma} \frac{\alpha}{c_{t+1}}}_{EMU_{c,t+1}}$$

$$c_t^{\alpha(1-\gamma)} s_t^{(1-\alpha)(1-\gamma)} \frac{\alpha}{c_t} = EMU_{c,t+1}$$

$$c_t^{\alpha(1-\gamma)-1} = \frac{EMU_{c,t+1}}{\alpha} s_t^{-(1-\alpha)(1-\gamma)}$$

$$c_t = \left[\frac{EMU_{c,t+1}}{\alpha} s_t^{-(1-\alpha)(1-\gamma)}\right]^{\frac{1}{\alpha(1-\gamma)-1}} \tag{iEE}$$

B.2 Details of the solution and calibration

I solve for the value and policy functions and each age using the method developed in Fella (2014) to apply the endogenous grid method with discrete choice variables. Starting from $V_{T+1} = 0$, I proceed backwards. The retirement problem is deterministic. For the worker's problem, I compute expected marginal and continuation utilities using a Gauss-Kronrod integration application. Interpolation is always linear, due to the discrete jumps associated with the durable decision. The size and bounds of the grids included in table A.3.

For the case of the income grid, the transitory shock is discretized using an equally spaced grid. The persistent component z grid is also equally spaced and the bounds are calculated via simulation of the income process. The bounds for η are also chosen by simulation. The estimates are found using a standard method of simulated moments, with weighting matrix that gives 0.8 weight to all the cross-sectional moments and 0.2 to the variance life-cycle profile.

Concerning the calibration algorithm. For a given β, I find λ^d and χ by simulating the economy until a criterium for the distance of the value function is met (relative tolerance of 10^{-6}). The algorithm used for both this and the previous step is the global method MSLS from the NLOPt library. The local search is done with Nelder Mead, also the version in the same library.

Table A.3: Numerical Parameters

		Value
Grids		
n_a	# Asset Grid Points	100
n_d	# Durable Grid Points	20
n_z	# Persistent Component Grid Points	41
n_ε	# Transitory Component Grid Points	21
$\underline{d}, \bar{d}$	Bounds Durable Consumption	0,1000000
$\underline{a}, \bar{a}$	Bounds Assets	0,5000000
Power a	Exponential Grid Power a	3
Power d	Exponential Grid Power d	2
N_{sim}	# Simulations	40000