

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

Rocio Madera[†]

November, 2024

Abstract

This paper characterizes the joint dynamics of earnings and consumption changes focusing on the tails of the distribution, motivated by new evidence that emphasizes the role of extreme earnings changes. In particular, this paper makes the following contributions. (i) I first propose an empirical method to discipline the outliers in survey data sampling using publicly available moments from administrative sources. (ii) I then apply this method to study the distribution of household disposable income and consumption, two measures that are not available in the U.S. social security data. An important result is that the empirical relation between 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) I build a life-cycle, incomplete markets model with lumpy durable consumption and non-gaussian earnings shocks. I parametrize the model using higher-order moments of earnings and consumption at the household level. The behavior of lumpy durables is crucial to rationalize the empirical findings in the event of non-gaussian shocks. (iv) I use the model to calculate the structural consumption response to extreme shocks and to understand the implications for the degree of self-insurance against higher-order income risk, the welfare cost of incomplete markets, and aggregate consumption dynamics.

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 and Ellen McGrattan, for invaluable inspiration and guidance. I thank Anmol Bhandari, Jonathan Heathcote, Kyle Herkenhoff, Loukas Karabarbounis, Morris Kleiner, and Fabrizio Perri for their time and feedback. I also thank all other participants at the University of Minnesota Macro-Labor Workshops, in particular Luis Diez-Catalan, Joaquin Garcia-Cabo, Eugenia Gonzalez-Aguado, and Sergio Salgado for comments and suggestions. A previous version of this paper was circulated as "How Shocking Are Shocks?".

[†]Southern Methodist University; romadera@smu.edu.

1 Introduction

What characterizes the joint distribution of income and consumption changes? This question is critical for measuring the extent to which households can smooth labor income fluctuations and how large uninsurable income risk is. It is also at the core of the transmission mechanism from fiscal and monetary policy to households' consumption behavior. Nevertheless, the answers to these questions so far have heavily relied on two assumptions that are increasingly challenged as access to data improves: that earnings shocks are well approximated by a normal distribution and that the consumption response to income changes is relatively linear on the size of the shock.

Few studies have inspected changes in income from a nonparametric perspective. Those who have ([Geweke and Keane, 2000](#); [Bonhomme and Robin, 2009](#); [Guvenen *et al.*, 2016](#)), however, reach a common conclusion: earnings changes are not normally distributed. In particular, they exhibit a high level of excess kurtosis (i.e. the importance of the tails and center of the distribution in relation to the area in between). Leaving the technical details to the main text, this means that the incidence of moderate-size changes is substantially smaller than under a log-normal distribution. In fact, 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. Yet, with few notable exceptions, studies that look at joint changes in income and consumption are biased towards average or moderate-size income changes – one standard deviation up or down – and disregard that (1) the average shock is very small and (2) not all income shocks are equally shocking. As a result, little attention has been devoted to the potentially most disruptive part of the distribution: the tails.

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 growth exhibit deviations from log-normality, especially the latter. This is surprising, as one would expect most of the large changes to go away with the inclusion of both spouses' incomes and, especially, government transfers. Whether or not that implies a higher risk for households in the form of larger consumption fluctuations is unclear. Larger changes usually trigger insurance

mechanisms that moderate changes do not. Acquiring insurance for disasters, moving back home with family, or buying a better car in the case of a positive change are all examples of mechanisms that households are likely to use in the case of extreme events. Because it is practically impossible to account for all these forms of risk-sharing and self-insurance, I follow the common practice of inspecting the joint distribution of consumption and earnings (Deaton and Paxson, 1994; Altonji *et al.*, 1996; Blundell *et al.*, 2008). Given my focus on the extremes of the distribution, I define new measures of comovement between earnings and consumption changes, which are borrowed from the finance literature in which the analysis of tails is typical. These measures will be defined in the main text but fall within the common denomination of tail dependence. While the overall covariance between earnings and nondurables is considerably higher than for durables, empirical measures of 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 the incorporation of durable consumption adjustments in a life-cycle, incomplete markets model with higher-order, idiosyncratic income risk in order to calculate the consumption response to structural tail shocks. In this case, a model is necessary as non-linearities are pervasive. Beyond the non-linearities implied by the presence of borrowing constraints¹, the size of the tail shocks creates jumps that usual, empirical identification strategies can hardly capture². 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, De Nardi *et al.* (2016) and Guvenen *et al.* (2016). 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.

¹See 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.

²A notable exception is Arellano *et al.* (2014).

To parameterize the model, I proceed in two steps. First, I use a simulated moments' method to estimate the parameters of the earnings process, targeting the second through fourth moments of four-year growth in after-transfers income. Additionally, I target the age profile of the variance of log income levels, which is important to discipline the persistence of income shocks. Next, I proceed to calibrate the remaining parameters of the model to match both aggregate and microeconomic targets. 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 separately.

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 insurance, I find that the average transmission coefficients (commonly known as BPP coefficients) are not very different from the current estimates in the literature; roughly a bit over half of the income shocks are transmitted to nondurable consumption. This implies that close to half of the income shocks are insured via risk-sharing and self-insurance. However, I show that this is masked by a large amount of heterogeneity in the size of the shock. Quantile regressions of consumption change on the structural income shocks show that there is a substantial response to extreme shocks.

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, 2010; Guvenen and Smith, 2014a). 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.

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 couple of years with the increasing availability of administrative data. [Guvenen *et al.* \(2016\)](#) 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 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 t10 ([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. While my sample is much smaller and my data is exposed to measurement error, the

rich set of covariates provides a different set of insights.

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 capturing the behavior at the tails and higher-order moments.

2.1 The *new* PSID

The PSID is a longitudinal study of a representative sample of U.S. households, tracking the socioeconomic status of US families from 1967 to 2014. Due to its length and panel structure, it has been extensively used for the study of income dynamics. In this section, I give a brief overview on the dataset itself and focus on the recent and less-explored waves that contain detailed data on consumption. I refer the reader to [Heathcote *et al.* \(2010\)](#), for example, for the basics and structure of the PSID.

The original focus of the PSID was the study of income dynamics and poverty. Hence, the coverage of socioeconomic and income covariates has been thorough since its inception, but questions on consumption behavior were limited to food and rent³. It was not until the 1999 wave, with information concerning tax year 1998, that questions on a wide set of expenditure categories were added, making it the first panel of income and consumption at a disaggregated level in the US that covers the vast majority of the information available in cross-sectional surveys such as the CEX⁴. As a result, most studies requiring panel data for income and consumption before 1999 were forced to use imputation strategies to incorporate CEX consumption categories into the PSID⁵. While these imputation strategies have been proven to be very successful for non-durable categories, adjustments on durable goods have been left out of the picture, which I show are fundamental in understanding and interpreting the joint dynamics of consumption and earnings changes.

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. 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

³Some data on the value of owned houses and vehicles was provided, but this was generally inconsistent over time, and no data on actual expenditures on these durables was reported.

⁴[Andreski *et al.* \(2014\)](#) provide a detailed comparison between these new waves of the PSID and the CEX. They conclude that, overall, the overlapping categories are highly comparable on average and over the life cycle.

⁵See, for example, [Blundell *et al.* \(2008\)](#); [Kaplan and Violante \(2010\)](#); [Guvenen and Smith \(2014b\)](#), among many others.

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 a major revision in main variables concerning labor income; starting after that wave eliminates spurious variation and the need to make assumptions to homogenize those variables over time.

To sum up, the *new* PSID spans 16 years (1998–2014) and contains 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 paper⁶.

2.2 Tail Changes in Survey Data: The Importance of Sample Selection

The sample of reference 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 also impose that they have at least three consecutive observations between 1998 and 2014. I keep both the original representative SRC sample and the SEO sample from the PSID, with the appropriate weighting. At the end of Section 2.2, I elaborate on the treatment of outliers, since it is a major step in our sample selection.

The final sample with information on income and consumption is comprised of around 20000 observations, corresponding to approximately 5000 households over 16 years. More details, including the number of observations left at each step of the sample selection, are given in Appendix A.

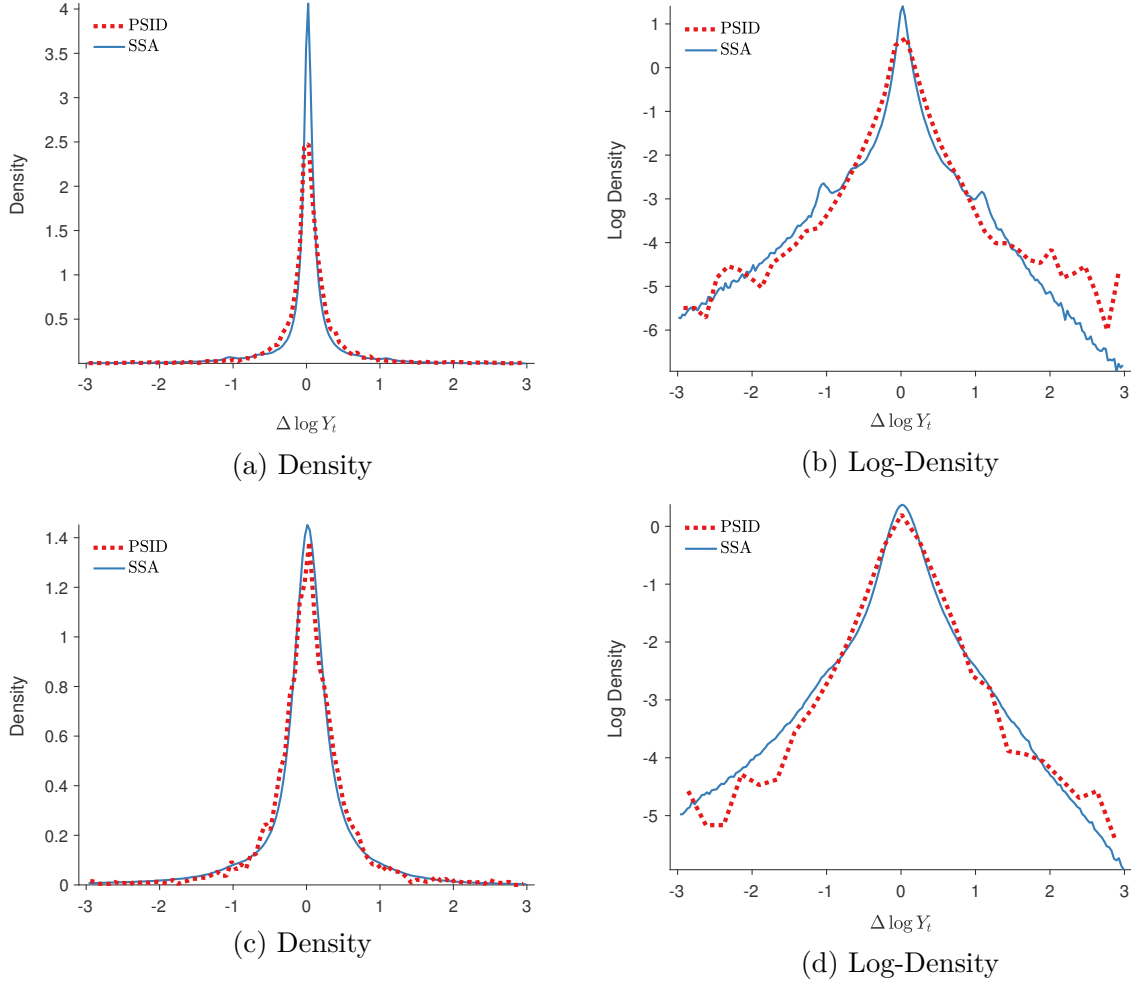
Trimming the Tails: A New Method

A crucial step when selecting a sample in survey data is the treatment of outliers. Due to the presence of measurement error, it is not uncommon to trim the tails of the distribution somewhere between the 1 and 10 top and bottom percentiles. This poses a problem when characterizing the tails of the distribution. In fact, in the past, one of the reasons why income changes were found to follow roughly a log-normal distribution has to do with the treatment of outliers when selecting a sample in the PSID. I propose a new method for cleaning the tails.

⁶For some exercises, I make use of previous waves whenever information of past behavior is necessary.

Guvenen *et al.* (2016) calculate a large number of moments of the distribution of earnings changes by age, sex, and past income levels from administrative data in the US (SSA data). I use the distribution of earnings changes in the SSA data to discipline the trimming of the tails in the PSID. Specifically, for this step, I use the PSID annual waves right before our sample (1992-1997) for comparison with the SSA frequency. I start from the left and right ends and trim 0.01 percentile at a time, until the distribution resembles that reported in the SSA data according to different moments of the distribution. After this method is executed, 0.2% of the left tail and 0.5% of the right tail has to be trimmed to mimic the distribution in the SSA data for every year of our sample. Figure 1 shows the resulting distributions. More details are provided in the appendix.

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



Note: SSA densities are those reported in [Guvenen *et al.* \(2014\)](#) and available in excel from the authors' websites. PSID density graphs are calculated with Epanechnikov kernel and 100 bins.

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

Measures of Income and Consumption

Several measures of earnings and consumption are used along the paper. For earnings, the reference measure 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.

For the case of consumption, nondurable consumption includes food, utilities, non-durable 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 income and 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.

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 $\log \Delta^s(\mathbf{x})$ and $\text{arc}\Delta^s(\mathbf{x})$ denote the log- and the arc-change in \mathbf{x} from the the current period to s periods ahead, respectively. Formally:

$$\log \Delta^s(\mathbf{x}_t) \equiv \log \mathbf{x}_{t+s} - \log \mathbf{x}_t$$

$$\text{arc}\Delta^s(\mathbf{x}_t) \equiv \frac{\mathbf{x}_{t+s} - \mathbf{x}_t}{(\mathbf{x}_{t+s} + \mathbf{x}_t)/2}.$$

The default measure will be log-changes. 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. 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.

Next, I turn to analyze the distribution of changes in labor income and consumption. To inspect the extent to which these changes deviate from normality and, in particular, exhibit fat tails, I make use of two descriptive tools: higher-order moments and Quantile-Quantile plots.

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,

$$P_k \equiv \text{kth percentile}$$

$$P_{k\ell} \equiv P_k - P_\ell$$

$$\text{Kelley Skewness} \equiv S^K \equiv \frac{P_{9050} - P_{5010}}{P_{9010}} \quad (1)$$

$$\text{Crow-Siddiqui Kurtosis} \equiv K^{CS} \equiv \frac{P_{97.5} - P_{2.5}}{P_{7525}}. \quad (2)$$

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.

The second point to notice in Table [I](#) is the life-cycle effect. While all measures of income become increasingly leptokurtic over time, the opposite happens to consumption. The age effect on nondurable changes is not very strong, but it is striking because a standard consumption-savings model with a high degree of heterogeneity predicts an increasing level of kurtosis over time ([Guvenen *et al.*, 2016](#)). That is, there is empirical evidence of higher consumption insurance beyond self-insurance through savings against higher-order income risk if judged by fluctuations in nondurable consumption.

Table I: Higher-Order Moments of Earnings and Consumption

	Standard Dev.	Kelley Skewness (0 under a Normal)	C-S Kurtosis (2.91 under a Normal)
2-Year Changes			
$\log \Delta^2 y_t^{\text{ind}}$	0.654	-0.045	9.899
Y	0.660	-0.027	8.701
O	0.649	-0.074	11.059
$\log \Delta^2 y_t^{\text{hh}}$	0.571	-0.043	7.018
Y	0.584	-0.048	6.542
O	0.559	-0.040	7.444
$\log \Delta^2 y_t^{\text{post}}$	0.585	-0.055	6.890
Y	0.580	-0.035	6.515
O	0.589	-0.071	7.355
$\log \Delta^2 c_t$	0.471	0.004	4.000
Y	0.488	-0.002	4.141
O	0.455	0.006	3.937
$\log \Delta^2 d_t$	0.813	0.447	23.687
Y	0.936	0.510	26.671
O	0.692	0.339	19.182
4-Year Changes			
$\log \Delta^4 y_t^{\text{ind}}$	0.725	-0.126	8.151
Y	0.746	-0.121	7.797
O	0.703	-0.164	8.755
$\log \Delta^4 y_t^{\text{hh}}$	0.632	-0.095	5.838
Y	0.654	-0.090	5.854
O	0.611	-0.089	5.782
$\log \Delta^4 y_t^{\text{post}}$	0.660	-0.086	5.921
Y	0.683	-0.079	5.974
O	0.636	-0.092	5.827
$\log \Delta^4 c_t$	0.514	-0.041	3.796
Y	0.537	-0.052	3.925
O	0.490	-0.032	3.775
$\log \Delta^4 d_t$	1.004	0.451	16.836
Y	1.141	0.498	15.663
O	0.849	0.409	14.900

Note: **Columns** refer to the standard deviation and robust measures of skewness and kurtosis. See equations (1) and (2) for definitions. **Rows** include y^{ind} : individual earnings (heads), y^{hh} : household pre-gov. earnings, y^{post} : households post-gov. earnings, c : nondurable consumption, d : durable consumption. See Section 2.3 for detailed definitions. *Y*: Age group 25-44, *O*: Age group 45-60. See Table A.1 for extra moments.

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	Percent	Average Size (log Δ)	Average Size (\$)
$0 \leq \text{abs}\Delta < 1\text{SD}$	82.79	0.15	6671.80
$1\text{SD} \leq \text{abs}\Delta < 2\text{SD}$	10.77	0.69	24536.16
$2\text{SD} \leq \text{abs}\Delta < 3\text{SD}$	2.90	1.21	36703.40
$3\text{SD} \leq \text{abs}\Delta$	3.53	2.80	43867.68
N	18,524		

Note: $\text{abs}\Delta$ denotes the absolute value of $\log \Delta$. Earnings correspond to household earnings after transfers, therefore one standard deviation is equal to 0.58 log points. Dollars are in \$2010.

Graphical Analysis and Deviations from Normality

Despite the strong evidence against normality shown in 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).

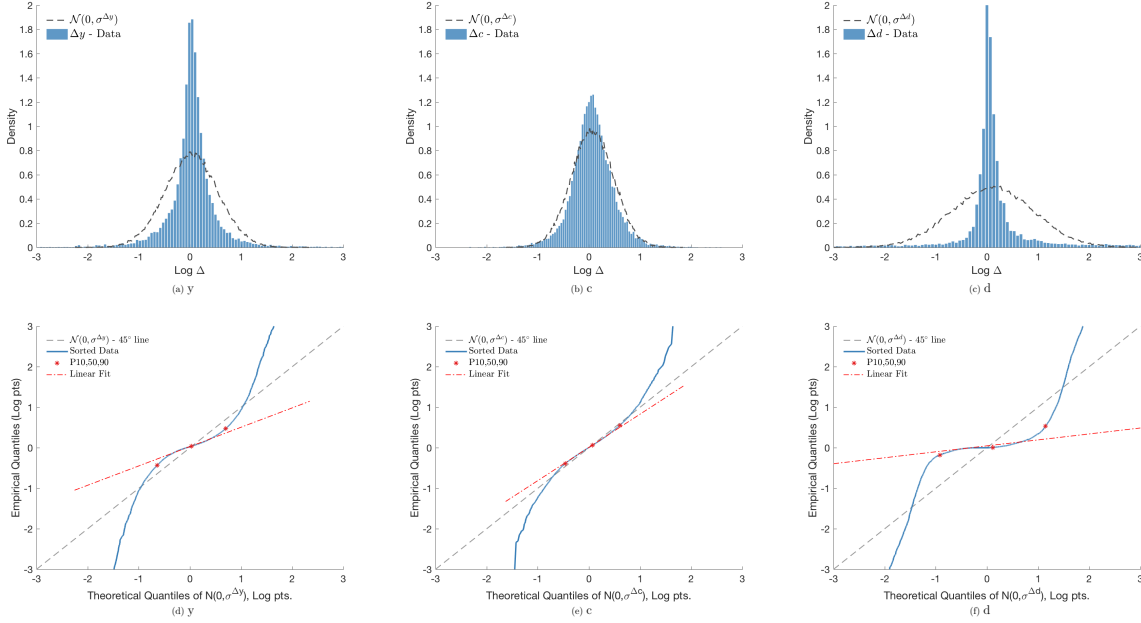
The lower panel in Figure 2 includes a set of QQ plots. QQ plots compare two distributions by plotting their quantiles against each other. They represent a particularly useful tool to assess the extent to which a variable is well approximated by a normal, or any given distribution. Both axes correspond 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 tails⁷. 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.

⁷Appendix A.6 includes a stylized example of QQ plots for usual distributions.

Figure 2: Empirical Distributions: Histograms and Deviations from Normality



2.4 Joint Distribution

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, I introduce one tool that will prove to be useful in this context: tail dependence, or dependence of extreme events.

Tail Dependence and the Joint Distribution of Earnings and Consumption

Figure 3 confirm these findings. Notice, in particular, the S shape of the right panel. The steeper slope only indicates that the distribution is more disperse, as in the illustrative example of the middle panel in figure A.1. The S shape, however, is evidence of *fatter tails* of durable consumption changes as compared to earnings changes.

Finally, I plot both the empirical distribution of income changes (in solid red) and the corresponding normal distribution with the same variance (in dashed thin red). With

this, I 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.

At this point, it is evident that different measures point to a joint distribution of consumption and income that diverts at the tails. Moreover, the role of durable consumption in measuring the response to tail shocks seems necessary. An important problem with extreme events is the potential non-linear behavior in income and consumption. It is thus very hard to identify the measures of insurance in an empirical fashion, as many of the identification assumptions would be violated. As a result, in the next section, I develop a life-cycle incomplete-markets model, which will allow me to compute structural responses within a non-linear framework, as well as a wider set of implications.




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

3 Quantitative Analysis

In this section, I present a life-cycle consumption-savings model with income uncertainty and incomplete markets, with two additions to the standard⁸ case: (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 higher-order moments observed in the data.

⁸Aiyagari (1994); Kaplan and Violante (2010).

In the remaining of this section, 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 comment the calibration and estimation strategy and results.

3.1 Model

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.

3.1.1 The Household Problem

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.

Preferences. Households have standard CRRA preferences over a consumption bundle of non-durable and durable consumption, denoted by \mathbf{c} and \mathbf{d} , respectively. Both types of goods conform the consumption aggregate following a Cobb-Douglas specification⁹. 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,

$$E_0 \sum_{t=1}^T \beta^{t-1} \xi_t \frac{\mathcal{C}(\mathbf{c}_t, \mathbf{s}_t)^{1-\gamma}}{1-\gamma} \quad (3)$$

$$\mathcal{C}(\mathbf{c}_t, \mathbf{s}_t) \equiv \mathbf{c}_t^\alpha (\mathbf{s}_t)^{1-\alpha}, \quad (4)$$

⁹Piazzesi and Schneider (2007) provide evidence in favor of the Cobb-Douglas aggregation of both consumption goods.

where \mathbb{E}_0 is the expectation operator with respect to the stochastic processes introduced in subsection 3.1.2, 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, \mathbf{a}_{t+1} , and earn capital gains $\mathbf{r}\mathbf{a}_t$ on currently held bonds, where $\mathbf{r} > 0$ 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 downpayment requirement in the case of adjustment:

$$\mathbf{a}_{t+1} \geq -\lambda^y \underline{y}_t - \lambda^d \mathbf{d}_{t+1} \quad (5)$$

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

Pensions. Income at retirement mimics the US system. Pensions are a function of lifetime average gross earnings¹⁰. 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:

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

¹⁰For computational purposes, I follow Guvenen and Smith (2014a) and estimate average labor earnings \bar{Y}_i^R as the fitted value of

$$\bar{Y}_i = \alpha_0 + \alpha_1 Y_{i,T_R},$$

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.

where

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

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

$$\begin{aligned} V_t(a_t, d_t, z_t) &= \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} \\ \text{s.t.} \quad &c_t + a_{t+1} + d_{t+1} + A(d_t, d_{t+1}) = Y_t + (1+r)a_t + (1-\delta)d_t \\ &Y_t \text{ given by equations (7) -- (11)} \\ &a_{t+1} \geq -\lambda^y y_t - \lambda^d d_{t+1}, \quad c_t \geq 0 \end{aligned}$$

Recursive problem of a retiree household. For ages $t = T_r, \dots, T$

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

3.1.2 Idiosyncratic Shocks and Labor Income

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 distributed¹¹ with mean 0 and standard deviation σ_ε . Equations (7)-(10) formally summarize these relations.

Finally, equation (11) 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.

$$\log Y_{it} = g_t^a + y_{it} \quad (7)$$

$$y_{it} = z_{it} + \varepsilon_{it} \quad (8)$$

$$z_{it} = \rho z_{it-1} + \eta_{it} \quad (9)$$

$$\varepsilon_{it} \sim \mathcal{N}(0, \sigma_\varepsilon) \quad (10)$$

$$\eta_{it} \sim \begin{cases} \mathcal{N}(\mu_1, \sigma_1) & \text{with prob. } p \\ \mathcal{N}(\mu_2, \sigma_2) & \text{with prob. } 1 - p \end{cases} \quad (11)$$

3.1.3 Durables

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 $\mathcal{D} = \{d_0, \dots, d_{n^d}\}$.

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:

$$A(d_{t+1}, d_t) = \chi_1 d_t + \chi_2 \underbrace{|d_{t+1} - (1 - \delta)d_t|}_{c_t^d}, \quad (12)$$

¹¹As 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.

where $c_t^d \equiv d_{t+1} - (1 - \delta)d_t$ denotes expenditures on durable adjustments at age t . The first component in (12) 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 as of 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:

$$s_t^d = \kappa d_t, \kappa > 1 \quad (13)$$

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 Calibration

One period in the model is one year of life. 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 (7)-(11) 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 III. 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 of nondurable and durable consumption.

3.2.1 Estimation of the Income Process with SMM

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 and higher-order moments of four-year income changes, as well as the life cycle profile of the variance of income levels. Targeting the life-cycle profile of the variance of income levels is important to discipline the persistence parameter.

Table III: Income Process Estimates

σ_ε	(Variance of transitory shock)	0.053
p	(Probability of drawing from normal 1)	0.930
ρ	(Persistence)	0.913
μ_1	(Mean of 1 persistent shock)	0.008
μ_2	(Mean of 2 persistent shock)	-0.106
σ_1	(Variance of 1 persistent shock)	0.075
σ_2	(Variance of 2 persistent shock)	1.189
σ_{z_0}	(Variance of initial distribution)	0.753

3.2.2 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 \(2000\)](#)¹². 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 consumptions. Given the Cobb-Douglass 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.

	α
D = Cars + Houses + Furnishings + Repairs + Clothing	
C ₁ = Food + Utilities + Nondurable Transportation + Recreation	0.7039
C ₂ = C ₁ + Rent	0.7249
C ₃ = C ₂ + Health + Education + Child Care	0.8034

Table IV: Share of Nondurable Consumption in Total Consumption

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 $\delta = 0.072$.

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.

¹²4% 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 \(2015\)](#).

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.

3.2.3 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_0	(Adjustment costs parameter)	0.108
χ_1	(Adjustment costs parameter)	0.008

4 Results

In this section, I first evaluate the performance of the model in replicating the tail behavior described in the empirical section, as well as the mechanisms at work. Next, I measure to what extent income shocks pass-through to consumption, comparing with previous estimates in the literature. A novel component of my analysis is that, instead of calculating the OLS response, I estimate quantile regressions to obtain heterogeneous effects by the size of the shock. I conclude with a welfare calculation.

4.1 Model Fit Discussion

Higher-order moments in the model

Table VI: Model Fit

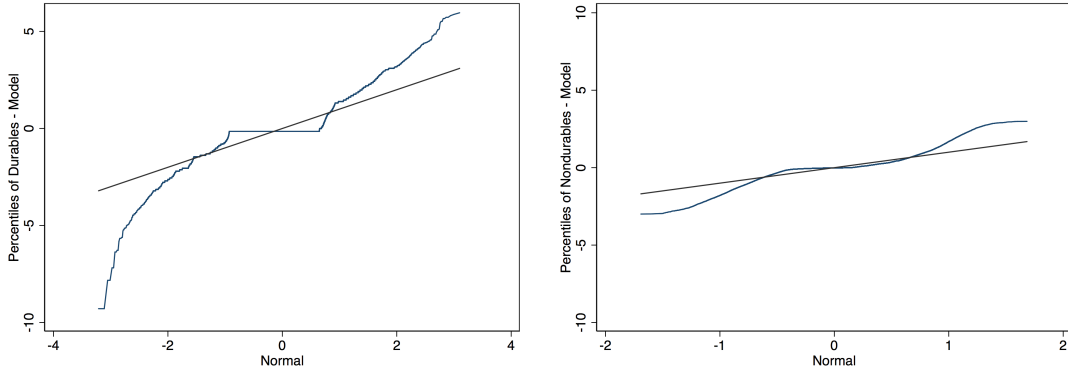
	Data	Model
Cross-sectional moments (Income)		
SD($\log\Delta^2 y_t^{\text{post}}$)	0.585	0.589
KS($\log\Delta^2 y_t^{\text{post}}$)	-0.055	-0.005
CS($\log\Delta^2 y_t^{\text{post}}$)	6.890	6.711
SD($\log\Delta^4 y_t^{\text{post}}$)	0.660	0.532
KS($\log\Delta^4 y_t^{\text{post}}$)	-.086	-.113
CS($\log\Delta^4 y_t^{\text{post}}$)	5.921	6.010
var(y_t^{post})	See fig. ??	
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%

Note: Names in bold are targeted. **Abbreviations:** SD - standard deviation, KS - Kelley Skewness, CS - Crow-Siddiqui Kurtosis.

Lumpiness and deviations from normality

Figure 4 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 4: Deviations from Normality: Durables vs. Nondurables



4.2 The Micro Consumption Response to Tail Shocks

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

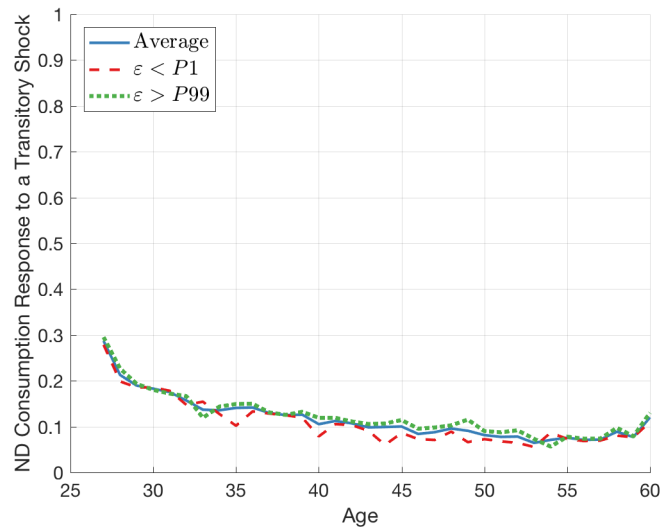
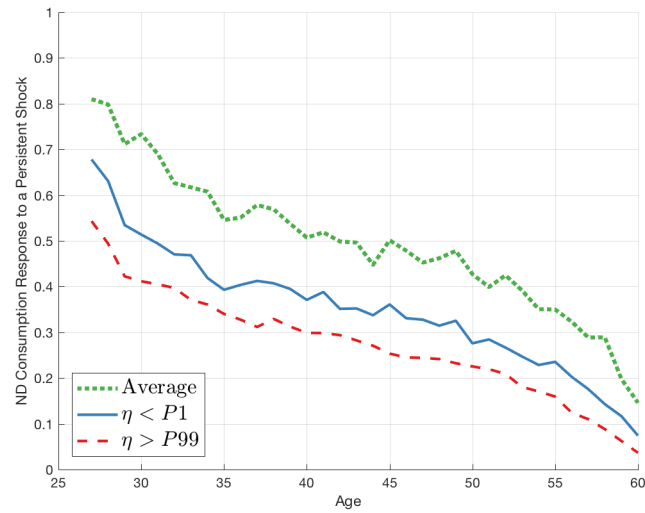
$$\phi^{c,\eta} = \frac{\text{Cov}(\Delta c_t, \eta_t)}{\text{Var}(\eta_t)},$$

where η here stands for the persistent shock in the case of my structural model. An empirical measure can be obtained by instrumenting η with a function of income changes in the data.

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.

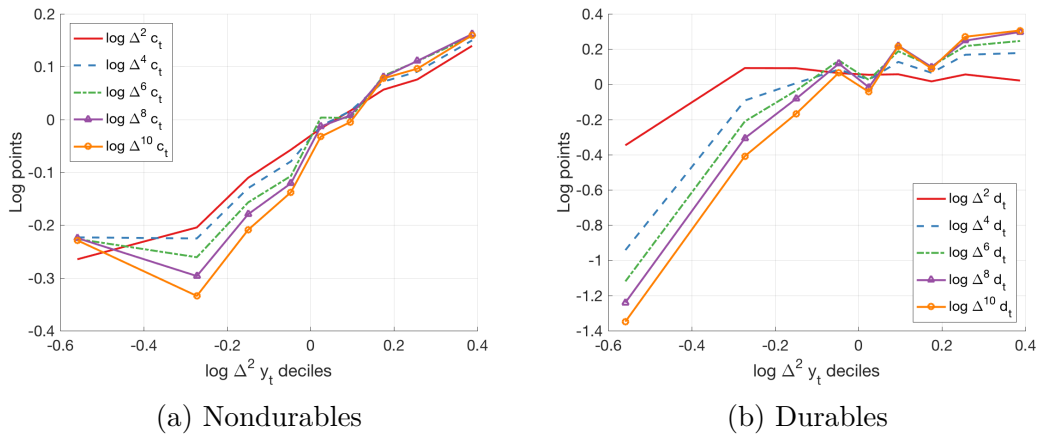
For comparison with the literature, and to gain insights on the effects of Non-Gaussian shocks and durable consumption, I first report the ϕ coefficients:



I will comment on the estimates for the persistent component, as the transitory component is highly insurable. It is interesting to notice that the non-Gaussian case exhibits a larger degree of self-insurance than the Gaussian case. Moreover, considering durable adjustments reduces the amount of partial insurance. However, as I have mentioned above, these estimates have to be interpreted with caution in the presence of very large shocks, as the average effect does not necessarily coincide with the behavior at the tails. Overall, this provides a second check on the fit of the model, since the estimates for models without nondurables are strongly consistent with previous literature, namely [Kaplan and Violante \(2010\)](#) and [Guvenen *et al.* \(2016\)](#) for the Gaussian and Non-gaussian case, respectively.

To complement the analysis and overcome the limitations present in covariane measures, I calculate empirical impulse response functions to different sizes of income shocks.

Figure 5: Response to Income Shocks of Different Size



The x-axis shows 20 percentiles 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 the longer the horizon. Again, notice how the upward changes do

not feature this result. This coincides with the lack of upper tail dependence between durable consumption and earnings.

4.3 The Macro Consumption Response to Tail Shocks

4.4 The Welfare Cost of Tail Shocks

In this section, I ask and 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 (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:

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

where \mathcal{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 $\mathcal{W}_i^{\text{ng,d}}$ 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 $\mathcal{W}_i^{\text{g,d}}$ and $\mathcal{W}_i^{\text{ng,nd}}$ 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 $j = \text{g}, \text{ng}$ in

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

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

I find $\theta^{\text{ng}} = 9\% > 5\% = \theta^{\text{g}}$, which means that households in the non-Gaussian

world are willing to give up an extra 4% of their consumption at every age and state in order to live in a world without uncertainty. This result is not surprising, since in the non-Gaussian world the fluctuations are larger. It is more interesting, however, when combined with the next exercise.

Next, I turn to the second question. I repeat the previous calculations in a world where households only get utility from non-durables.

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

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

It is interesting that, in this case, the gap between the Gaussian and non-Gaussian worlds is much larger. Specifically, $\theta^{ng} = 15\% > 5\% = \theta^g$.

I thus conclude that the welfare costs attributable to higher-order risk – tail events – are around 4% of the yearly household consumption bundle. Additionally, in the absence of durable adjustment, this extra risk would be much larger.

Finally, I use the calibrated model to ask a simple question: What fraction of consumption at each age and state would a household be willing to give up to live in a world without tail shocks? The answer to this question provides insights on whether previous welfare experiments that assumed normal shocks were missing anything from ignoring non-Gaussian shocks in previous quantitative studies. To calculate the consumption equivalent variation I proceed as follows:

Let household i 's lifetime utility be given by:

$$U_i \left(\{\mathbf{c}_{it}, \mathbf{d}_{it}\}_{t=1}^T \right) \equiv \sum_{t=1}^T (\beta \xi)^{t-1} u(\mathbf{c}_{it}, \mathbf{d}_{it}),$$

for series of nondurable and durable consumption $\{\mathbf{c}_{it}, \mathbf{d}_{it}\}$ as obtained in the model.

I am going to assume a utilitarian welfare function of the form:

$$W(\lambda) \equiv \sum_{i=1}^N \pi_i U_i \left(\{(1 - \lambda)\mathbf{c}_{it}, \mathbf{d}_{it}\}_{t=1}^T \right),$$

where π_i denotes the population size at age i .

I simulate the model economy for 10000 households in two scenarios: one in which shocks are as in the benchmark model, and one in which shocks are drawn from a Gaussian distribution, re-estimated to match all but the higher order moments in Table ???. The details of this model are in Appendix B. The simulation yields series $\{c_{it}^G, d_{it}^G\}$, for the Gaussian world, and $\{c_{it}^B, d_{it}^B\}$ for the benchmark world. And corresponding welfare functions

$$W^B(\lambda) \equiv \sum_{i=1}^N \pi_i U_i \left(\{(1-\lambda)c_{it}^B, d_{it}^B\}_{t=1}^T \right)$$

$$W^G(\lambda) \equiv \sum_{i=1}^N \pi_i U_i \left(\{(1-\lambda)c_{it}^G, d_{it}^G\}_{t=1}^T \right)$$

The welfare cost is λ such that

$$W^B(\lambda) = W^G(0).$$

and equal to

$$\lambda = 12\%.$$

In other words, a household would be willing to give up 12% of their consumption per year to live in a world without tail shocks. This is a smaller number compared to similar studies. [Guvenen *et al.* \(2016\)](#) find that it ranges from 12.3%, in their simplest specification, to 27.6%, when the income processes exhibits a high degree of heterogeneity in age and recent earnings. In their simplest specification they include a safety net parameters of a minimum annual income of \$2000. This is suggestive evidence that durable consumption adjustments are providing a similar level of insurance compared to that minimum level of transfers.

5 Conclusions

Using data from the Panel Study of Income Dynamics, I have documented that extreme changes in household income are more pervasive in the data than usually assumed in parametric assumptions. In particular, I show that excess kurtosis is not an exclusive

feature of individual earnings that fades out when transfers are considered. On the contrary, all three of households earnings, nondurable consumption, and durable consumption growth exhibit deviations from log-normality, especially the latter. This is surprising, as one would expect most of the large changes to go away with the inclusion of both spouses' income and, especially, of government transfers.

To gain insight on the joint distribution of consumption and income, I have defined new measures of comovement between earnings and consumption changes that are borrowed from the finance literature, where the analysis of tails is typical. While the overall covariance between earnings and nondurables is considerably higher than for durables, empirical measures of tail dependence show that the opposite is true for 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.

I estimated a quantitative model where non-convex adjustment costs in durable adjustments can endogenously generate fat tails in durable consumption. I test the implications of the model for usual estimates of partial insurance and find that they are indeed very similar, and slightly smaller than in the general case. This would be surprising if it weren't for the fact that there is a large amount of heterogeneity in the size of the shock. Inspecting the responses of consumption change on the structural income shocks show that there is a substantial response to extreme negative shocks, especially of durable consumption.

References

- AIYAGARI, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, **109** (3), 659–84.
- ALTONJI, J. G., HAYASHI, F. and KOTLIKOFF, L. (1996). *The Effects of Income and Wealth on Time and Money Transfers between Parents and Children*. NBER Working Papers 5522, National Bureau of Economic Research, Inc.
- 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. (2014). *Household Earnings and Consumption: A Nonlinear Framework*. mimeo, University College London.
- BAGGER, J., FONTAINE, F., POSTEL-VINAY, F. and ROBIN, J.-M. (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-VILLAYERDE, 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**, 101–154.
- 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.-M. (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.
- DE NARDI, M., FELLA, G. and PAZ PARDO, G. (2016). *The Implications of Richer Earnings Dynamics for Consumption and Wealth*. NBER Working Papers 21917, National Bureau of Economic Research, Inc.
- DEATON, A. and PAXSON, C. (1994). Intertemporal choice and inequality. *Journal of Political Economy*, **102** (3), 437–67.
- FELLA, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*.
- 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: 1968-1989. *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), 25–51.
- GUVENEN, F., KARAHAN, F., OZKAN, S. and SONG, J. (2014). *What Do Data on Millions of U.S. Workers Say About Labor Income Risk?* Working paper, University of Minnesota.
- , —, — and — (2016). *What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?* Working paper.
- and SMITH, A. A. (2014a). Inferring Labor Income Risk and Partial Insurance from Economic Choices. *Econometrica*, **82** (6), 2085–2129.

- and — (2014b). 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, 1967-2006. *Review of Economic Dynamics*, **13** (1), 15–51.
- KAPLAN, G. and VIOLANTE, G. L. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, **2** (4), 53–87.
- 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), 20–40.
- PIAZZESI, M. and SCHNEIDER, M. (2007). Asset Prices and Asset Quantities. *Journal of the European Economic Association*, **5** (2-3), 380–389.
- PRIMICERI, G. E. and VAN RENS, T. (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of Monetary Economics*, **56** (1), 20 – 39, carnegie-Rochester Conference Series on Public Policy: The Causes and Consequences of Rising Economic Inequality April 25-26, 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** (7-8), 973–982.
- 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.

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)¹³.

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.

¹³Self-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								
$\log \Delta^2 y_t^{\text{ind}}$	0.654	0.875	-0.475	-0.045	20.926	9.899	0.418	0.457
<i>Y</i>	0.660	0.922	-0.355	-0.027	22.471	8.701	0.449	0.473
<i>O</i>	0.649	0.816	-0.587	-0.074	19.527	11.059	0.378	0.438
$\log \Delta^2 y_t^{\text{hh}}$	0.571	0.846	-0.193	-0.043	24.045	7.018	0.405	0.441
<i>Y</i>	0.584	0.909	0.107	-0.048	25.110	6.542	0.433	0.476
<i>O</i>	0.559	0.793	-0.487	-0.040	22.852	7.444	0.380	0.412
$\log \Delta^2 y_t^{\text{post}}$	0.585	0.760	-0.980	-0.055	43.711	6.890	0.359	0.401
<i>Y</i>	0.580	0.772	-0.026	-0.035	39.629	6.515	0.372	0.400
<i>O</i>	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
<i>Y</i>	0.488	1.007	-0.032	-0.002	10.729	4.141	0.502	0.505
<i>O</i>	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
<i>Y</i>	0.936	1.125	0.857	0.510	12.477	26.671	0.849	0.276
<i>O</i>	0.692	0.537	0.575	0.339	21.330	19.182	0.360	0.178
4-Year Changes								
$\log \Delta^4 y_t^{\text{ind}}$	0.725	1.082	-0.551	-0.126	19.349	8.151	0.473	0.609
<i>Y</i>	0.746	1.135	-0.822	-0.121	19.109	7.797	0.499	0.636
<i>O</i>	0.703	1.018	-0.237	-0.164	19.603	8.755	0.425	0.592
$\log \Delta^4 y_t^{\text{hh}}$	0.632	1.044	-0.306	-0.095	21.097	5.838	0.473	0.572
<i>Y</i>	0.654	1.069	-0.456	-0.090	19.377	5.854	0.487	0.583
<i>O</i>	0.611	1.006	-0.119	-0.089	23.104	5.782	0.459	0.548
$\log \Delta^4 y_t^{\text{post}}$	0.660	0.914	-1.297	-0.086	37.590	5.921	0.418	0.496
<i>Y</i>	0.683	0.910	-1.095	-0.079	32.716	5.974	0.419	0.491
<i>O</i>	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
<i>Y</i>	0.537	1.127	-0.942	-0.052	11.435	3.925	0.534	0.592
<i>O</i>	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
<i>Y</i>	1.141	1.611	0.807	0.498	8.747	15.663	1.207	0.404
<i>O</i>	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. See Section 2.3 for detailed definitions.

Table A.2: Summary Statistics- Income

	Moments					
	mean	sd	skewness	kurtosis	min	max
Individual Labor Income - Males	68821.50	91776.56	12.37	279.95	0.00	3015425.75
Individual Labor Income - Females	35222.03	32104.80	2.55	17.55	0.00	370975.31
HH Pre-Government Labor Income	49437.46	56320.71	10.62	248.34	0.00	2038943.00
HH Post-Government Labor Income	41520.77	42060.71	1.41	355.23	-1.57e+06	1343660.38
Durable consumption	91908.48	124277.98	7.13	127.44	0.00	3786477.00
Non-durable consumption	14224.52	9455.32	3.41	36.67	0.00	187641.42
Semi-durable consumption	8246.56	15652.97	12.69	255.19	0.00	425631.84
	p10	p25	p50	p75	p90	p95
Individual Labor Income - Males	15638.53	31669.62	51727.05	81516.67	125709.75	168855.97
Individual Labor Income - Females	0.00	13232.61	30128.67	49186.45	71308.83	87944.59
HH Pre-Government Labor Income	7815.66	21668.51	39594.66	63260.14	93289.87	121559.35
HH Post-Government Labor Income	10822.88	21008.85	35789.11	53385.75	74759.59	94156.75
Durable consumption	1414.21	12322.44	65848.66	124733.64	200051.12	272782.53
Non-durable consumption	5314.89	8220.80	12316.20	17864.74	24844.35	30416.90
Semi-durable consumption	499.69	1952.96	4918.65	9793.83	17167.76	24674.76
Observations	23287					

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 [Guvenen *et al.* \(2016\)](#). 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 \tilde{Y}	22373
Not enough consecutive obs	20954
Final # Observations	20954
Final # Households	4697

Table A.3: 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 $t = 1, \dots, T_r - 1$:

$$V_t(\mathbf{a}_t, \mathbf{d}_t, \mathbf{z}_t) = \max_{\mathbf{c}_t, \mathbf{d}_{t+1}, \mathbf{a}_{t+1}} \{u(\mathbf{c}_t, \mathbf{s}_t) + \beta \mathbb{E}_t V_{t+1}(\mathbf{a}_{t+1}, \mathbf{d}_{t+1}; \mathbf{z}_t)\} \quad (19)$$

$$\text{s.t.} \quad \mathbf{c}_t + \mathbf{a}_{t+1} = Y_t + (1+r)\mathbf{a}_t + (1-\delta)\mathbf{d}_t - \mathbf{d}_{t+1} - A(\mathbf{d}_t, \mathbf{d}_{t+1}) \quad (20)$$

$$Y_t \text{ given by equations (7) -- (11)} \quad (21)$$

$$\mathbf{a}_{t+1} \geq -\lambda^y \mathbf{y}_{t+1} - \lambda^d \mathbf{d}_{t+1}, \quad \mathbf{c}_t \geq 0, \mathbf{d}_{t+1} \in D \quad (22)$$

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

First, let

$$\mathbf{b}_{t+1} \equiv \mathbf{a}_{t+1} + \lambda^y \mathbf{y}_{t+1} + \lambda^d \mathbf{d}_{t+1}. \quad (23)$$

Then, the borrowing constraint can be rewritten as

$$\begin{aligned} \mathbf{c}_t + (\mathbf{b}_{t+1} - \lambda^y \mathbf{y}_{t+1} - \lambda^d \mathbf{d}_{t+1}) &= Y_t + (1+r)(\mathbf{b}_t - \lambda^y \mathbf{y}_t - \lambda^d \mathbf{d}_t) + (1-\delta)\mathbf{d}_t - \mathbf{d}_{t+1} - A(\mathbf{d}_t, \mathbf{d}_{t+1}) \\ \mathbf{c}_t + \mathbf{b}_{t+1} &= \underbrace{Y_t + (1+r)\mathbf{b}_t + (1-\delta)\mathbf{d}_t - (1-\lambda^d)\mathbf{d}_{t+1} - A(\mathbf{d}_t, \mathbf{d}_{t+1}) - \lambda^d(1+r)\mathbf{d}_t + \lambda^y(\mathbf{y}_{t+1} - (1+r)\mathbf{y}_t)}_{\mathbf{m}(\mathbf{d}_{t+1})} \end{aligned} \quad (24)$$

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

Next, I redefine V as follows

$$\tilde{V}_t = \max_{\mathbf{c}_t, \mathbf{d}_{t+1}, \mathbf{a}_{t+1}} \left\{ \underbrace{\mathcal{C}(\mathbf{c}_t, \mathbf{s}_t)^{1-\gamma} + \beta \mathbb{E}_t \tilde{V}_{t+1}^{1-\gamma}(\mathbf{a}_{t+1}, \mathbf{d}_{t+1}; \mathbf{z}_t)}_{\mathbb{V}_t} \right\}^{\frac{1}{1-\gamma}},$$

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

The dynamic programming problem to be solved is thus

$$\begin{aligned} \tilde{V}_t(\mathbf{b}_t, \mathbf{d}_t, \mathbf{z}_t) &= \max_{\mathbf{d}_{t+1}, \mathbf{b}_{t+1}} \left\{ \mathcal{C}(\mathbf{m}(\mathbf{d}_{t+1}) - \mathbf{b}_{t+1}, \mathbf{s}_t)^{1-\gamma} + \beta \mathbb{E}_t \tilde{V}_{t+1}^{1-\gamma}(\mathbf{b}_{t+1}, \mathbf{d}_{t+1}; \mathbf{z}_t) \right\}^{\frac{1}{1-\gamma}} \quad (25) \\ \text{s.t.} \quad &(\mathbf{b}_{t+1}, \mathbf{d}_{t+1}) \in \{\mathbf{b}_{t+1}, \mathbf{d}_{t+1} : \mathbf{b}_{t+1} \in [0, \mathbf{m}(\mathbf{d}_{t+1}; \mathbf{b}_t, \mathbf{a}_t, \mathbf{z}_t)], \mathbf{d}_{t+1} \in \mathbf{D}\}. \end{aligned}$$

The FOC therefore are given by

$$\begin{aligned} \frac{1}{1-\gamma} \mathbb{V}_t^{\frac{1}{1-\gamma}-1} \left\{ -(1-\gamma)(\mathcal{C}_t)^{-\gamma} \underbrace{[\alpha \mathbf{c}_t^{\alpha-1} \mathbf{s}_t^{1-\alpha}]}_{\mathcal{C}_{c,t} \equiv \frac{\alpha}{\mathbf{c}_t} \mathcal{C}_t} + \beta(1-\gamma) \tilde{V}_{t+1}^{-\gamma} \tilde{V}_{b,t+1} \right\} &= 0 \\ \mathcal{C}_t^{1-\gamma} \frac{\alpha}{\mathbf{c}_t} &= \beta \mathbb{E}_t \tilde{V}_{t+1}^{-\gamma} \tilde{V}_{b,t+1} \quad (\text{FOC}_b) \end{aligned}$$

And the envelope condition

$$\begin{aligned} (1-\gamma) \tilde{V}_t^{-\gamma} \tilde{V}_{b,t} &= (1-\gamma)(1+r) \mathcal{C}_t^{-\gamma} \mathcal{C}_{c,t} \\ \tilde{V}_t^{-\gamma} \tilde{V}_{b,t} &= (1+r) \mathcal{C}_t^{1-\gamma} \frac{\alpha}{\mathbf{c}_t}. \quad (\text{EC}) \end{aligned}$$

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

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

B.1.2 Implementing the Endogenous grid Method

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

$$\begin{aligned}
[\mathbf{c}_t^\alpha \mathbf{s}_t^{1-\alpha}]^{1-\gamma} \frac{\alpha}{\mathbf{c}_t} &= \underbrace{\beta(1+r)\mathbb{E}_t \mathbf{c}_{t+1}^{1-\gamma} \frac{\alpha}{\mathbf{c}_{t+1}}}_{\text{EMU}_{\mathbf{c},t+1}} \\
\mathbf{c}_t^{\alpha(1-\gamma)} \mathbf{s}_t^{(1-\alpha)(1-\gamma)} \frac{\alpha}{\mathbf{c}_t} &= \text{EMU}_{\mathbf{c},t+1} \\
\mathbf{c}_t^{\alpha(1-\gamma)-1} &= \frac{\text{EMU}_{\mathbf{c},t+1}}{\alpha} \mathbf{s}_t^{-(1-\alpha)(1-\gamma)} \\
\mathbf{c}_t &= \left[\frac{\text{EMU}_{\mathbf{c},t+1}}{\alpha} \mathbf{s}_t^{-(1-\alpha)(1-\gamma)} \right]^{\frac{1}{\alpha(1-\gamma)-1}} \quad (\text{iEE})
\end{aligned}$$

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.4.

For the case of the income grid, the transitory shock is discretized using an equally spaced grid. The persistent component \mathbf{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.4: 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