Essays in Idiosyncratic Income Risk

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Rocio Madera Holgado

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Fatih Guvenen, Adviser

June, 2017

Acknowledgements

I am grateful to my advisor Fatih Guvenen, for the inspiration and guidance. I am very thankful to Anmol Bhandari and Ellen McGrattan, for their constructive discussions and constant support. I am also appreciative of Morris Kleiner, for the thoughtful conversations and for participating in my dissertation committee. Additionally, I benefited from conversations with Luis Diez-Catalan, Joaquin Garcia-Cabo, Jonathan Heathcote, Kyle Herkenhoff, Loukas Karabarbounis, Fabrizio Perri, Sergio Salgado, and participants at the University of Minnesota Macro-Labor Workshops; as well as from the collaboration with Chris Busch, David Domeij, and Fatih Guvenen in the first chapter, and with Antonio Cabrales and Maia Guell in the last chapter. Finally, I acknowledge the financial support of the Bank of Spain and La Caixa Fellowships for Graduate Studies, in addition to the Hutcheson-Lilly and the Sacaroglu-Sargent Dissertation Fellowships.

Dedication

To my parents, Marisa and Miguel, who supported me unconditionally along this journey and taught me to take every step with enthusiasm and joy.

Abstract

The three chapters in this dissertation constitute an investigation of the determinants of idiosyncratic income risk faced by individuals and households over their lives and over time, as well as the consequences for individual behavior and welfare. Chapter 1 concentrates on the business-cycle variation in higher-order income risk and the extent to which such risk can be smoothed within households or with government social insurance policies. Chapter 2 studies the role of labor market institutions and employment protection in shaping individuals' exposure to earnings uncertainty over their lives. Chapter 3 measures the response of household consumption to large and unexpected earnings fluctuations that neither labor market institutions nor government policy can insure away.

Contents

List of Tables

List of Figures

Chapter 1

Business Cycle Earnings Risk and Government Insurance

1.1 Introduction

This paper studies how higher-order income risk varies over the business cycle as well as the extent to which such risks can be smoothed within households or with government social insurance policies. By higher-order income risk, we refer to risks that are captured by not only the variance of income shocks, but also their skewness and kurtosis. These higher order moments of the data can be a major source of risk for individuals as we show in this paper.

To provide a broad perspective on these questions, we study panel data on individuals and households from the United States, Germany, and Sweden, covering more than three decades of data for each country. It is useful to begin by putting our analysis in context. A broad range of empirical evidence indicates that idiosyncratic income risk rises in recessions. Earlier work in the literature was limited by the small sample size and time span on the available survey-based panel datasets, such as the Panel Study of Income Dynamics (PSID), forcing researchers to make parametric assumptions to obtain identification. One widely used assumptions that is common in the literature is that shocks to earnings are Gaussian and, therefore, its changes are bound to be symmetric. Restricting attention to the changes in the mean and variance of income shocks, earlier studies concluded that the variance of income shocks is countercyclical (e.g., Storesletten *et al.* (2004)). In recent work, Guvenen *et al.* (2014) used a very large panel dataset on earnings histories of a representative sample of 10% of all US working-age males from the U.S. Social Security Administration (SSA) records. Their large sample size, with millions of observations, allows for identification of changes in an unrestricted sense—without parametric assumptions.They found that the variance of income shocks is very stable over time and is robustly acyclical, whereas the skewness of shocks varies significantly over time in a procyclical fashion. This suggests that the changes in labor income risks in recessions and expansions are rather asymmetric.

Despite important advantages, the SSA data also have three shortcomings:

(i) earnings data are available only for individuals, and it is not possible to link household members to each other, (ii) no information is available on taxes and transfers (unemployment insurance, welfare payments, gifts, etc.), and (iii) no information is available on skills/education. Furthermore, Guvenen *et al.* (2014) focus on males with no corresponding information on women.

This paper makes two contributions. First, applying non-parametric techniques and using robust statistics, we document that the overall dispersion of individual labor earnings growth is flat and acyclical in all three countries, whereas the left-skewness of shocks is strongly countercyclical. These findings are robust across gender, skill groups, private/public sector employment and occupation. Furthermore, we show (using data for Germany) that these results are primarily driven by changes in wages and not in hours. Moreover, we show that applying the same method to both survey and administrative data (the PSID and SSA for the U.S. and SOEP and SIAB for Germany) yields the same substantive conclusions.

Second we find that insurance provided within households or by the government plays an important role in reducing downside risk, but that how and to what extent differs between the countries. Within-household provided insurance reduces the countercyclicality in the skewness of earnings in Sweden, but evidence of within-household insurance is much weaker in United States and in Germany. Government provided insurance, in the form of unemployment insurance, welfare benefits, aid to low income households, and the like, plays a more important role in all three countries; the effectiveness is weakest in the United States, and strongest in Germany.

The paper is organized as follows. The next section discusses the data sources, and Section 1.3 describes the empirical approach. Section 1.4 presents the results for gross (before-government) individual earnings and examines how the patterns of cyclicality vary by gender, education, and type of employment. Section 1.5 expands the analysis to households and includes various types of government social insurance policies to examine their impact on the cyclicality of higher-order risk. Section **??** uses a structural life-cycle consumption-savings model with partial insurance to quantify the welfare benefits of governments' social insurance policies in the three countries under study. Section 1.6 concludes.

1.1.1 Related Literature

This paper is primarily related to two streams of the literature: the investigation of the cyclical properties of income risk and the design and analysis of government policy over the business cycle.

The question of how idiosyncratic income risk varies over the business cycle is essentially empirical, yet its answers have been dominated by parametric choices as a way to overcome the data limitations. Storesletten *et al.* (2004) addressed those difficulties and proposed an identification strategy that allowed exploiting a longer time span compared to that available in the micro data. Similar to our results, they find that income risk increases in recessions. However, as discussed in the introduction, they find this risk to be driven by countercyclical variance. Our paper is closer to Guvenen *et al.* (2014), also briefly discussed above, both in our non-parametric methods and in the findings. We also find the increase in income risk to be driven by higher-order moments. Moreover, compared to the latter paper, we extend the empirical analysis to different samples, demographic groups, income measures, and countries. We, together with contemporaneous research discussed below, confirm that the increase in downside risk in recessions is a pervasive phenomenon.

Despite the importance of this empirical question for policy analysis in general, there are few applications to countries other than the US. A notable exception is **?**, who extend the framework in Storesletten *et al.* (2004) and apply it to household data in Germany, the UK, and the US. Limiting their analysis to symmetric business-cycle risk of wages, they find mixed evidence for the cyclicality of the variance of shocks; namely that variance is procyclical in Germany, acyclical in the UK, and countercyclical in the US. They attribute this result to the differences in institutions between the three economies. In our paper, we compare three economies with very different institutions, and after allowing for the possibility of asymmetric changes find that all three countries exhibit similar cyclical patterns in higher-order labor income risk. Closer to our work, **?** analyze the cyclicality of labor income risk in Germany, explicitly allowing for time variation of the skewness. Extending the identification approach of Storesletten *et al.* (2004) to the third moment, they come to the same substantial conclusion as we do; namely, that variation of income risk over the business cycle is asymmetric. Finally, **?**

Our focus on higher-order moments, in addition to the recent empirical evidence, is motivated by a number of theoretical and quantitative papers that highlight the importance of these for household consumption behavior to be empirically plausible. In particular, Constantinides and Ghosh (2014) allow labor income shocks to exhibit procyclical skewness in an asset-pricing model. Their model is able to match the cross-sectional distribution of market returns, resolving several puzzles in the finance literature, including the equity premium and excess volatility puzzles. These results are in line with earlier theoretical results shown in Mankiw (1986).

This paper is also related to the literature on economic stabilizers and cyclical government policy. Similar to our work, **?** focus on the stabilization power of taxes and transfers. Their model allows for different channels through which fiscal policy can interact with the business cycle. Our exercise is related to their *social insurance* channel; that is, how these policies alter the risks households are exposed to and their subsequent consumption response. In line with our results, they find that transfers and taxes help reduce the welfare costs of recessions. **?** study the design of optimal policy—transfers, taxes, and government debt---in response to aggregate shocks in a model with incomplete markets and redistribution concerns. They calibrate the model to US data, capturing the asymmetric variation in the tails of the distribution of earnings shocks. They find that it is optimal for the government to increase all three instruments as a hedging device against aggregate shocks.

1.2 The Data

This section provides an overview of the data sets we use in our empirical analysis, the sample selection criteria, as well as the variables used in the subsequent empirical analyses. Given the diversity of our data sources, we relegate the details to Appendix A. Briefly, we employ four longitudinal data sets corresponding to three different countries: the Panel Study of Income Dynamics (PSID) for the United States, covering 1976 to 2010;[1] the Sample of Integrated Labour Market Biographies (SIAB[2]) and the German Socio-Economic Panel (SOEP) for Germany, covering 1976 to 2010 and 1984 to 2011, respectively; and the Longitudinal Individual Data Base (LINDA) for Sweden, covering 1979 to 2010. The PSID and the SOEP are survey-based data sets. The PSID has a yearly sample of approximately 2000 households in the core sample, which is representative of the U.S. population; the SOEP started with about 10,000 individuals (or 5,000 households) in 1984 and, after several refreshments, covers about 18,000 individuals (10,500 households) in 2011.[3]

The SIAB is based on administrative social security records and our initial sample covers on average 370,000 individuals per year. It excludes civil servants, students and self-employed, which make about 20% of the workforce. From the

[1] The PSID contains information since 1967. We choose our benchmark sample to start in 1976 due to the poor coverage of income transfers before the 1977 wave. We complement our results using a longer period whenever possible and pertinent.

[2] We use the factually anonymous scientific use file SIAB-R7510, which is a 2% draw from the Integrated Employment Biographies data of the Institute for Employment Research (IAB).

[3] These numbers refer to observations after cleaning but before sample selection. Only the representative SRC sample is considered in the PSID. The immigrant sample and high income sample of the SOEP are not used, because they cover only sub-periods.

perspective of our analysis, the SIAB has two caveats: (i) income is top-coded at the limit of income subject to social security contributions, and (ii) individuals cannot be linked to each other, which prohibits identification of households. We deal with (i) by fitting a Pareto distribution to the upper tail of the wage distribution[4] and with (ii) by using data from SOEP for all household-level analyses. Throughout the analysis we focus on West Germany, which for simplicity we refer to as Germany. LINDA is compiled from administrative sources (the Income Register) and tracks a representative sample with approximately 300,000 individuals per year.

For each country, we consider three samples: two at the individual level–one for males and one for females–and one at the household level. The samples are constructed as revolving panels: for a given statistic computed based on the time difference between years t and $t+k$, the panel contains individuals who are aged 25 to 59 in periods t and $t+k$ ($k = 1$ or 5) and have yearly labor earnings above a minimum threshold in both years. This threshold is defined as the earnings level that corresponds to 520 hours of employment at half the legal minimum wage, which is about $1885 US dollars for the United States in 2010.[5] To avoid possible outliers, we exclude the top 1% of earnings observations in the PSID and SOEP, but not in LINDA (which is from administrative sources). For each individual, we record age, gender, education, and gross labor earnings. By gross earnings we mean a worker's compensation from his/her employer before any kind of government intervention in the form of taxes or transfers.

The household sample is constructed by imposing the same criteria on the

[4]The imputation is done separately for each year by subgroups defined by age and gender. For workers with imputed wages, across years, we preserve the relative ranking within the age specific cross-sectional wage distribution. The procedure follows Daly *et al.* (2014): see Appendix A.1.3 for details.

[5]For the United States, we use the federal minimum wage. There is no official minimum wage in Sweden or Germany during this period. For Germany, we a take a minimum threshold of 3 Euros (in year 2000 Euros) for the hourly wage. For Sweden, the effective hourly minimum wage via labor market agreements was around SEK 75 in 2004 (Skedinger, 2007). For other years, we adjust the minimum wage by calculating the mean real earnings for each year, estimating a linear time trend for these means and removing that time trend from the SEK 75 minimum wage.

household head and adding specific requirements at the household level. More specifically, a household is included in our sample if it has at least two adult members, one of them being the household head,[6] that satisfy the age criterion and household income that satisfies the income criteria. At the household level, we will analyze pre- and post government earnings and disposable income. Pre-government earnings defined as the sum of gross labor earnings earned by the adults in the household. Post-government earnings is constructed by adding taxes and transfers, and disposable income by in addition adding capital income.

Classifying Expansions and Recessions

For the United States, the classification of expansionary and recessionary episodes is based on the NBER peak and trough dates, with small timing variations. Given the time span covered by our sample, we classify the following years as recessions: 1980-1983, 1991-92, 2001-2002, and 2008-2010. The main difference compared to the NBER list is that we treat the 1980-1983 period as a single "double-dip" recession because of the short duration of the intervening expansion and the lack of recovery in the unemployment rate. Based on this classification, there are four expansions and four recessions during our sample period.

[6] In PSID and SOEP the head of a household is defined within the data set. In LINDA, the head of a household is defined as the sampled male.

Figure 1.1: Cyclical Component of Quarterly GPD Growth: U.S., Germany, and Sweden

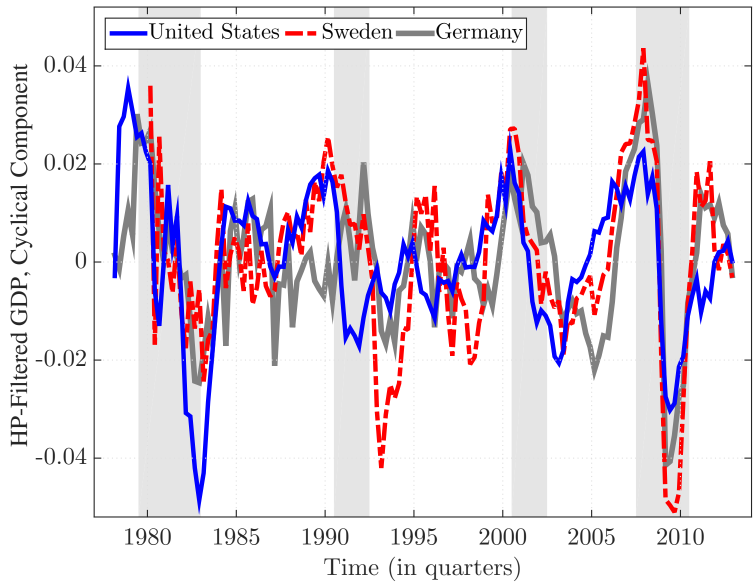

Note: The shaded areas in the figure indicate U.S. recessions according to our classification described in the text. The series for Germany corresponds to West Germany up to and including 1990Q4, and to (Unified) Germany from 1991Q1 on. The cyclical component is obtained by HP-filtering the series for GDP per capita from 1970Q1 to 2014Q1.

For both Germany and Sweden, we base the dating of expansions and recessions on data from the Economic Cycle Research Institute (ECRI), which applies the NBER methodology to OECD countries since 1948. The classification is consistent with various aggregate measures of the German and Swedish economies, respectively. In the time period covered by the panel data, recession periods for Germany (peak to trough) are from January 1980 to October 1982, January 1991 to April 1994, January 2001 to August 2003, and April 2008 to January 2009. Our sample period hence covers four recessions and four expansions. For Sweden, ECRI recession periods are from February 1980 to June 1983, June 1990 to July 1993, and April 2008 to March 2009. This leaves us with three recessions and three expansions during our sample period.

1.3 Empirical Approach

Measuring Income Volatility over the Business Cycle

For each year, we calculate robust statistics of log s-year changes in income. We consider different choices of s in order to distinguish between earnings growth over short and long horizons, and interpret these as corresponding to "transitory" and "persistent" earnings shocks.

More specifically, we compute moments $m[\Delta_s y_t]$, where $y_t \equiv \ln Y_t$ (natural logarithm) and $\Delta_s y_t \equiv y_t - y_{t-s}$. The moments m we consider are: the log differential between the 90th and 10th percentiles (L9010), the (Kelley) skewness, and the top (L9050) and bottom (L5010) tails. For Germany and Sweden, s refers to 1- and 5-year changes. Due to the biennial structure of the PSID from the 1997 wave, our analyses of earnings for the United States refer to 2- and 4-year changes instead.[7]

We do not impose any parametric assumption on the dynamics of income but instead analyze the behavior of the tails of the distribution of earnings changes. We think this is important since interpretations when using the variance as a summary statistic of the distribution alone can be misleading. To see this point, consider a widening of both the upper and lower tails of a normally distributed variable. This is, P90 is shifted to the right and P10 is shifted to the left. This certainly implies an increase in the variance; the opposite, however, is not necessarily true. Think of the case in which only the lower tail shifts to the left. Notice how the overall dispersion of the distribution increases here as well, but if we were to interpret this increase in isolation we would wrongfully conclude that not only one tail, but both of them expand. Similarly, unchanged overall dispersion does not imply an unchanged distribution, but can be observed when both tails move together (i.e., one tail shrinks while the other expands). Both of these last two scenarios imply a change of the relative size of the tails–a feature

[7]We calculate overlapping s-year differences up to $\Delta_s y_{1996}$, and non-overlapping s-year differences from then and up to $\Delta_s y_{2010}$, for $s = 2, 4$.

summarized by the skewness of the distribution. In our empirical analysis, these are the two scenarios we observe when considering cyclicality: either overall dispersion does not change while skewness does, or dispersion is cyclical, caused by one tail expanding and the other shrinking.

We conclude that, when measuring income volatility, the tails should be explicitly analyzed. Furthermore, when relying on summary statistics of the distribution, limiting the analysis to the variance cannot possibly identify the nature of the change, yielding misleading results. Higher-order moments, like skewness, should be then considered. Note how any assumption on the distribution of income shocks would drive our results: a (log-) normal distribution cannot capture changes in skewness, for example. This is why, and in light of recent evidence on male earnings growth using administrative data for the United States (Guvenen *et al.*, 2014), we take a skeptical–non parametric–point of view.

Broadening the Definition of Business Cycles

Some of the important macroeconomic variables do not perfectly synchronize with expansions and recessions, but their fluctuations might have an impact on earnings. For example, the U.S. stock market experienced a significant drop in 1987, during an expansion, and we can see in the time series analysis how the third moment falls sharply in that year. Similarly, the U.S. economy displayed an overall weakness in 1993-1994, which is evident in a range of economic variables, but these years are technically classified as part of an expansion by the NBER dating committee. Other examples are easy to find for Germany and Sweden (e.g., 1996). Therefore, the main focus of our analysis will be on the co-movement of higher-order moments of earnings changes with a continuous measure of business cycles. We use the (natural) log growth rate of GDP—i.e., $\Delta_s GDP_t \equiv \ln(GDP_t) - \ln(GDP_{t-s})$—as our measure of aggregate fluctuations. More specifically, we regress each moment m of the log income change between

$t-s$ and t on a constant, a linear time trend, and the log growth rate of GDP between year $t-s$ and t :

$$m\left(\Delta_s y_t\right) = \alpha + \gamma t + \beta^m \times \Delta_s(GDP_t) + u_t. \tag{1.1}$$

For a quantitative interpretation of the results reported in the next sections, Table I reports the short- and long-run volatility of GDP growth for each country and year sample considered along the paper.

Table I: Short- and Long-Run GDP Growth Volatility: United States, Germany, and Sweden

		Std. Dev. of GDP Growth	
	Data period	short-run	long-run
United States	1976-2010	3.34%	4.44%
Germany	1976-2010	2.01%	3.95%
Sweden	1976-2010	2.36%	5.42%

Note: Short-run is 1-year difference for Germany and Sweden, and 2-year difference for the United States. Long-run is 5-year difference for Germany and Sweden and 4-year difference for the United States.

1.4 Empirical Results: Gross Individual Earnings

In this section, we address four questions concerning higher-order risk for individual earnings. First, we ask if the countercyclical skewness and the acyclical dispersion is a US-only phenomenon or a robust feature of business cycles that can be observed in other countries whose labor markets differ greatly from that in the U.S.. For example, according to OECD (1993) 15 percent of U.S workers are unionized and 21 percent are covered by trade union agreements. In Germany the equivalent shares are somewhat higher; 30 and 44 percent respectively, but in Sweden the overwhelming majority are members (84 percent) or are covered (94 percent) by trade union agreements. Second, we ask if business cycle variation in higher-order income risk differs across observationally distinct groups, defined by gender, education, private/public sector employment and occupation. Third, we ask if cyclicality of earnings changes can be attributed mainly to changes in wages or to changes in hours worked. Fourth, we ask if the countercyclicality of skewness and the acyclicality of dispersion found in U.S. administrative earnings data is also borne out in U.S. survey data, e.g., the PSID. This question is important because earlier papers that used the PSID and adopted parametric methods found strongly countercyclical variance of shocks. This begs the question: is it the data set or is it the methodology that accounts for these different conclusions?

Table II: Cyclicality of Individual Earnings

	L9010	Kelley	L9050	L5010
		United States		
Males	–0.11	1.67***	0.57***	–0.68***
	(–0.51)	(5.00)	(3.71)	(–3.96)
Females	0.40***	0.62*	0.48**	–0.08
	(1.85)	(1.97)	(2.61)	(–0.52)
		Sweden		
Males	–0.11	3.74***	0.91***	–1.01***
	(–1.22)	(4.00)	(3.80)	(–3.74)
Females	0.43**	1.64***	0.67***	–0.24**
	(2.24)	(3.33)	(3.09)	(–2.67)
		Germany (SIAB)		
Males	0.15	5.48***	0.95***	–0.80***
	(0.36)	(5.80)	(3.14)	(–4.11)
Females	0.34	2.55**	0.80	–0.46*
	(0.48)	(2.05)	(1.25)	(–1.80)

Note: Each cell reports the coefficient on log GDP change of a regression of a moment of the distribution of changes in a income measure on log GDP change, a constant, and a linear time trend. Newey-West t-statistics are included in parentheses (maximum lag length considered: 3 for SIAB and LINDA, 2 for PSID). Asterisks ($*, **, ***$) denote significance at the (10%, 5%, 1%)-level.

Cyclicality of Dispersion

In Table II, we report the cyclicality of four key statistics computed from the distribution of earnings changes of individual workers. To provide a comparative discussion, we report the results for all three countries in the same table. For now, we focus on the first row of each panel, corresponding to the sample of male workers in each country. The first column reports the cyclicality of the L9010. In the United States, the L9010 for males is acyclical, as seen from the small (–0.11) and statistically insignificant (*t-stat* of –0.51) coefficient. Turning to Sweden and Germany, the L9010 for male earnings are also acyclical.[8]

[8] All regression results based on SIAB data are robust to various robustness checks that address issues of top-coding and a structural break in the wage variable. See appendix **??** for details.

Figure 1.2: Standard Deviation, Skewness, and Tails of Short-Run Earnings Growth: United States, Sweden, and Germany (SIAB): All Males

(a) United States, SD (left) and KS (right)

(b) United States, Upper and Lower Tail

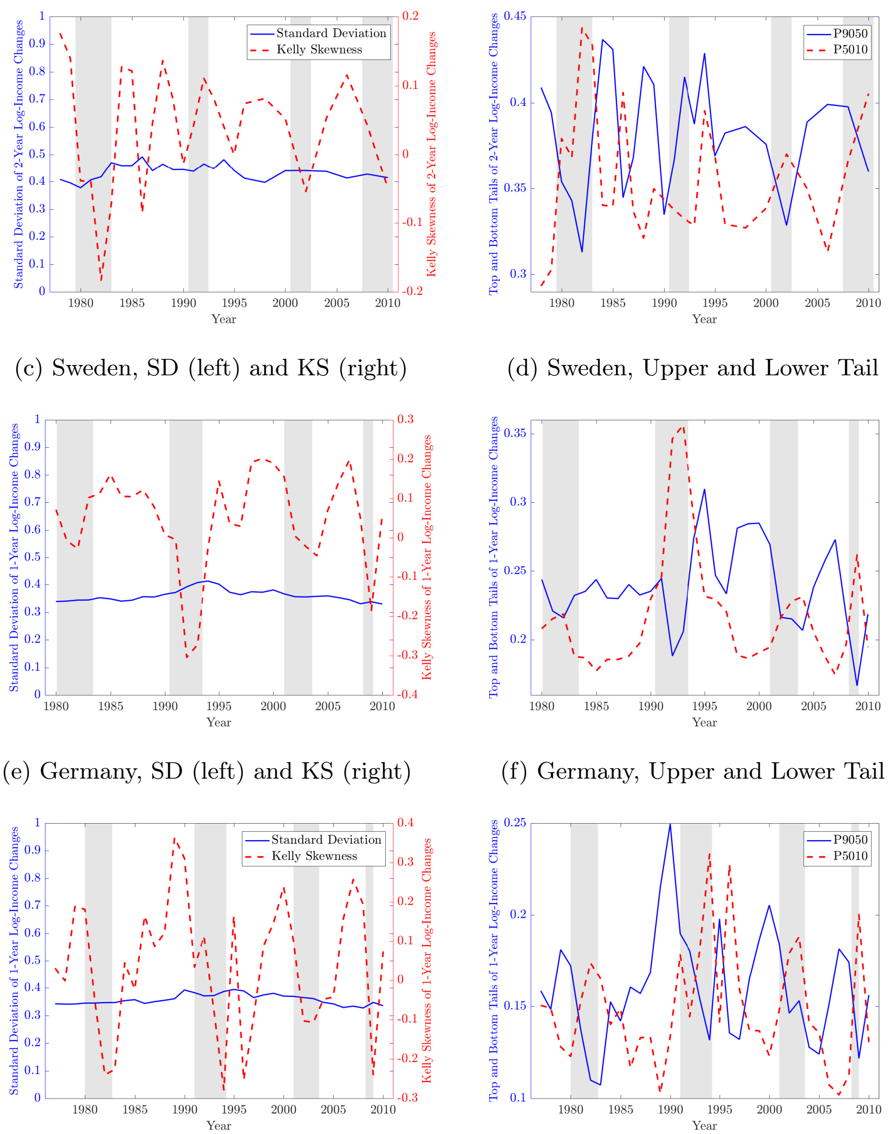

(c) Sweden, SD (left) and KS (right)

(d) Sweden, Upper and Lower Tail

(e) Germany, SD (left) and KS (right)

(f) Germany, Upper and Lower Tail

Note: Linear trend removed, centered at sample average.

Overall, we conclude that in all three countries the dispersion of earnings changes does not display any robust pattern of cyclicality, judging from these regressions. In addition to being acyclical, the dispersion of earnings changes is quite flat over time (left panels of Figure 1.2). These figures should be compared with typical calibrations in the literature that assume the volatility of earning shocks doubles or triples during recessions. Here the largest movements are on the order of 10% to 15%, and they show no signs of cyclicality.

Cyclicality of Skewness

We next turn to the cyclical behavior of skewness. Column 2 reports a measure of asymmetry, called Kelley's skewness, defined as:

$$\mathcal{S}_k = \frac{(P90 - P50) - (P50 - P10)}{(P90 - P10)}.$$

This measure has several attractive features compared with the third standardized moment. First, it is much less sensitive to extreme observations, since it does not depend on observations beyond the 90th and 10th percentiles of the distribution. This deals with the concern about potential outliers. It is therefore our preferred measure of skewness, especially when considering the U.S. and Germany (GSOEP) where measurement issues could be more important.[9] Second, the particular value of Kelley's skewness has a simple interpretation, in terms of the relative lengths of the top and bottom tails. In particular,

$$\frac{P90 - P50}{P90 - P10} = 0.5 + \frac{\mathcal{S}_k}{2}, \tag{1.2}$$

which can be used to compute the fraction of overall dispersion (P90-P10) that is accounted for by the top tail (P90-50) and consequently by the bottom tail (P50-P10).

Armed with these definitions, we turn to the left panels of Figure 1.2. In all three countries, Kelley's skewness is procyclical. This pattern is particularly striking in Sweden and Germany, where movements in Kelley's skewness are almost perfectly synchronized with the business cycle as defined by ECRI. The notable exception is the fall in Kelley's skewness in 1996, but note that the cyclical component of GDP did indeed fall in 1996 as displayed in Figure 1.1.

[9] We have also analyzed the third standardized moment, and found very similar results.

Furthermore, Table II shows that the procyclicality of Kelley's skewness is (statistically) significant at the 1% level in all three countries. The coefficient is 1.67 for the U.S., 3.74 for Sweden, and 5.48 for Germany, showing more cyclicality when moving from the U.S. to Sweden and most for Germany. Thus, for example, if a typical recession in Sweden entails a drop in GDP growth of two standard deviations (from +1 to −1 sigmas, for a swing of $2 \times 0.0236 = 0.0472$), Kelley's skewness will fall by $0.0472 \times 3.74 = 0.177$. For the sake of discussion, suppose $\mathcal{S}_k^{\text{exp.}} = 0$ in an expansion, then $\mathcal{S}_k^{\text{rec.}} = -0.18$, which in turn implies from equation (1.2) that the upper tail to lower tail ratio, $(P90 - P50)/(P50 - P10)$ goes from 50/50 to 41/59 from an expansion to a recession. This is a large change in the relative size of each tail, especially for a country like Sweden, which might be thought of as displaying lower business cycle risk (due to the high unionization rate, among others).[10]

Inspecting the Tails

At the expense of some oversimplification, it might be useful to think about a shift towards more negative skewness as arising from either a compression of the right tail or an expansion of the left tail or both. Thus, a follow-up question is: which one of these changes is driving the cyclical changes in skewness for each country? Again, the pattern is particularly striking in Sweden, see the middle right panel of 1.2. It shows that the top tail is procyclical, whereas the bottom tail is countercyclical. The last two columns of Table II shows that this pattern is present and (statistically) significant in all three countries. This means that, in a recession, the positive half of the shock distribution compresses relative to the median, whereas the negative half expands. Thus, the shift towards negative skewness happens through both tails moving in unison during recessions.

Furthermore, notice that for all three countries it turns out that the magnitude of movement of each tail is similar to each other. For example, for the

[10] The corresponding changes in $\mathcal{S}_k$ for the U.S. and Germany are: 0.11 and 0.22 respectively.

U.S., the coefficient for L9050 is 0.57 and for L5010 is –0.68. The corresponding coefficients are 0.91 and –1.01 for Sweden, and 0.95 and –0.80 for Germany. Therefore, as log GDP growth fluctuates over the business cycle, the shrinking of one tail is matched closely by the expansion of the other tail, making the total dispersion, the L9010, move very little over the cycle. As a result, skewness becomes more negative in recessions without any significant change in the variance. This analysis shows that the behavior of higher-order risk is best understood by separately studying the top and bottom tails over the cycle, which can move together or independently. Focusing simply on a directionless moment, such as the L9010 or the variance, can miss important asymmetries that can matter for the nature of earnings risk. As we will see in a moment, whenever we observe cyclical dispersion, it is driven by *asymmetric* movements of the tails, and should not be thought of as a pure change in L9010 or the variance (which would imply an expansion/compression of *both* tails).

Survey vs. Administrative data

The earlier work on higher order income risk for male earnings, Guvenen *et al.* (2014), used administrative data from the U.S. Social Security Administration (SSA) records. As mentioned in the introduction, it lacks information on income components beyond earnings and one cannot link household members to each other. Similar restrictions apply to the administrative data we use for Germany (SIAB). This is why we use survey data (PSID for the U.S. and SOEP for Germany) to answer questions regarding insurance provided within households and by the government. These data sets however suffer from having fairly few observations, which may imply that higher moments are imprecisely estimated. Reassuringly however, the results for individual earnings are very similar in PSID and SSA data, and in SIAB and SOEP data respectively. Specifically, we have re-run regression 1.1 using moments from the SSA data, as reported in Guvenen et al. (2014), and from SOEP data. The resulting coefficients for

U.S. males using SSA data for each of the four moments are –0.07 , 2.31*** , 1.02*** , and –1.09**, respectively. These numbers are strikingly similar to those in the first row of the top panel in Table II The equivalent numbers using SOEP data are –1.33**, 1.76***, –0.21, and –1.12***. While these numbers differ somewhat from those in the first row of the bottom panel in Table II, they tell the same story. In particular, male earnings changes in both SOEP and SIAB is characterized by asymmetric movements of the tails rather than uniform expansions and contractions of both tails.[11] The bottom tail is countercyclical in both data sets while the top tail is procyclical in SIAB but acyclical in SOEP. As a result, the L9010 is acyclical in SIAB and countercyclical in SOEP, but in both data sets skewness is procyclical.

1.4.1 Differences by Gender

We now turn to the cyclicality of higher-order risk for female workers and examine how they compare to the patterns for males. Focusing on the second row of each panel in Table II, we see three main patterns. First, the L9010 of earnings changes is *procyclical* for U.S. and Swedish women but acyclical for German women. This is different from men, who displayed acyclical dispersion in all countries. Second, Kelley's measure of skewness is always procyclical–left-skewness is countercyclical–as indicated by the positive coefficient on log GDP growth, which is highly significant for Sweden (1% level), significant for Germany (5% level), and only slightly significant for the U.S. (10% level).

Third, inspecting the top and bottom tails separately (last two columns), we observe the expected pattern of cyclicality, whenever the coefficient is significant. In particular, L9050 is procyclical and significant for the U.S. and Sweden, whereas the L5010 is countercyclical and significant for Sweden and Germany.[12]

[11] We have also run 1.1 using the standard deviation of earnings changes as our measure for overall dispersion instead, and the coefficients are small (0.07 (SIAB), –0.12 (SOEP)) and insignificant (t-stat of 0.42 (SIAB), –054 (SOEP)) in both data sets.

[12] It is somewhat surprising that women in the U.S. seem to face less downside risk as measured by the L5010 differential compared with these two European countries.

Thus, just as for the case of male workers, the behavior of the variance is driven by an asymmetric movement of the two tails rather than a uniform expansion of both tails. In our view, this finding reiterates our earlier point that the L9010 or the variance are not ideal statistics to focus on when it comes to measuring higher-order earnings risk over the business cycle. Finally, it is worth noting that the magnitudes of the fluctuations in both Kelley's skewness and in the upper and lower tails separately are somewhat attenuated for women compared with men.

1.4.2 Differences across groups of workers

To shed light on the possible sources of cyclicality of higher-order income risk we now examine if it differs across observationally distinct groups. First we divide male and female workers into groups by education (college vs non-college graduates), or by private and public employment. These are two dimensions by which the three countries differ greatly. In Germany. 12 percent of men and 8 percent of women are college educated. In Sweden and the U.S. the equivalent numbers are 16 and 25 for males and 17 and 25 for females respectively. Differences in the size of public sector employment is even larger. Defining the public employment public administration, health care and education (sectors which in Germany and Sweden are dominated by public sector jobs or by jobs funded by the public),[13] the share of public sector employment in Sweden is more than twice as large as in Germany or the U.S.[14] Moreover, public sector jobs are often thought of as less risky, offering generous employment protection and less volatile compensation, so it is interesting to ask if this is borne out in the data.

[13] Formally we define a worker as working in the public sector, if he/she works in these sectors in both years t and $t+k$ (where $k = 1, 5$). Historically most workers in these sectors were employed by the public; this is less true today.

[14] In Sweden about 23% of men and 63% of women work in the public sector (these figures have been relatively stable over the considered time period). In Germany a stable 10% of men work in the public sector, while the share of women steadily increased from about 23% to about 36% over the considered time period. In the U.S. 13 percent of males and 18 percent of females are employed in the public sector.

For each of these groups we analyze higher order income risk by first computing average (standardized) moments across years and countries by quartiles of (standardized) log GDP change as shown in figures 1.3 and 1.4. The standardization of moments and log GDP change is performed independently for each country before pooling across countries, which implies that a deviation from zero indicates a standardized deviation from the country-specific mean of the moment. For each quartile the bars correspond to the average moment for (ordered from the left) the full sample (blue), college graduates (green), non-college graduates (red), private employment (cyan) and public employment (magenta), respectively. Figure 1.3 shows that earnings risk is very similar across all male subgroups; overall dispersion is acyclical (upper left panel), Kelley's skewness is procyclical (upper right panel), the top tail is procyclical and the bottom tail is countercyclical. Turning to females, Figure 1.4 shows a similar picture and, as noted above, that fluctuations in earnings risk is somewhat attenuated for women. For both males and females we see a strong asymmetric cyclical change of the distribution of earnings changes across groups.

Figure 1.3: Average moments by quartiles of log GDP change: Males

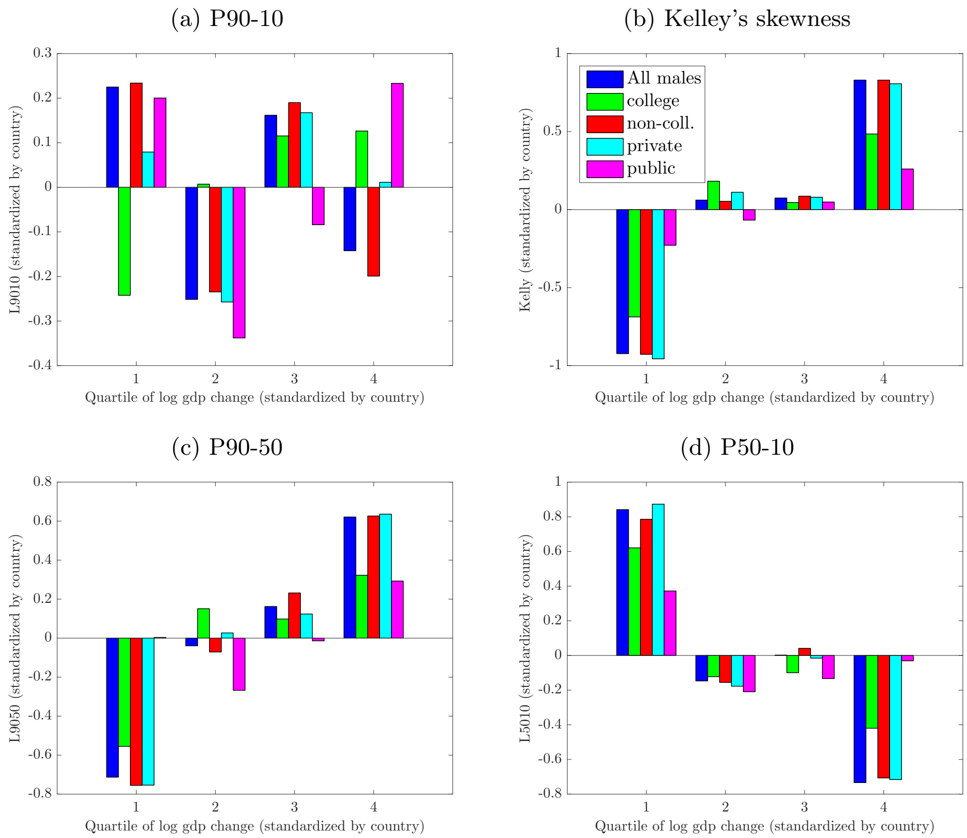

Note: For different samples, each bar shows the average moment across years and countries by quartiles of log GDP change. Both log GDP changes and moments are standardized by country.

Figure 1.4: Average moments by quartiles of log GDP change: Females

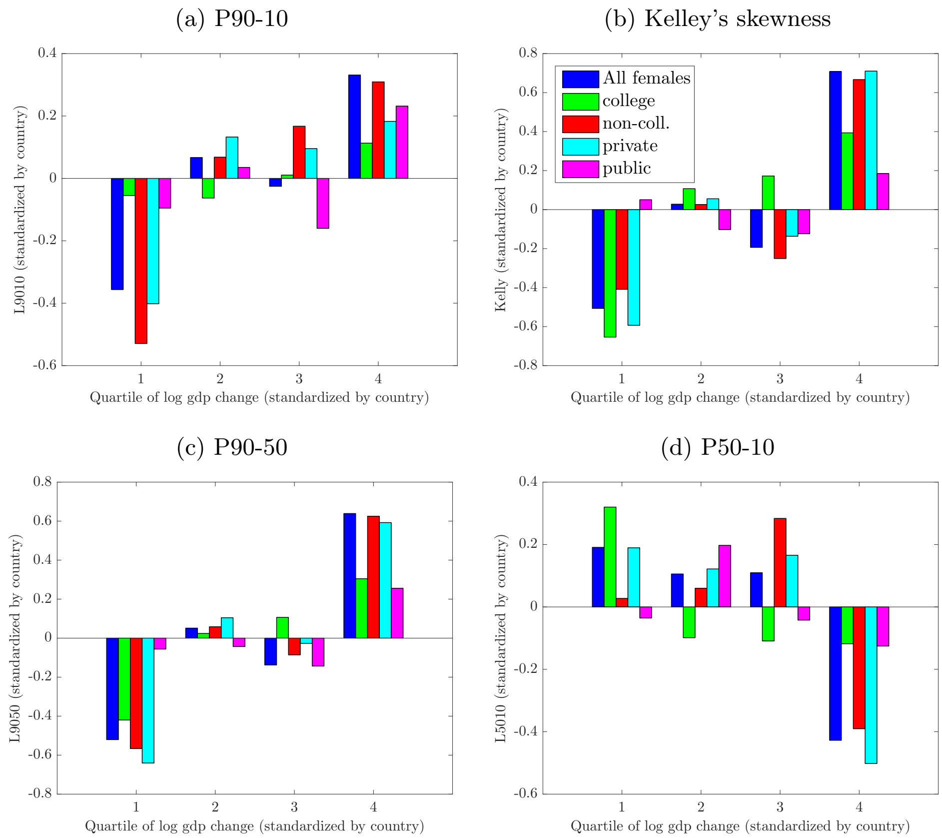

Note: See notes to figure 1.3.

For each group we have also computed correlations between the moments and log GDP change using equation 1.1 separately by country. These are displayed in Figure 1.5. Detailed results can be found in Appendix **??**. Each panel in the figure shows, starting from the left, the regression coefficients (from equation 1.1) with confidence intervals for males (solid) in the U.S.(blue), Sweden (red) and Germany (green), then followed by the equivalent regressions coefficients for females (dotted). Within each country-gender grouping, the regression coefficients are (ordered from the left) those from the full sample, college graduates, non-college graduates, private and public employment, respectively.

Figure 1.5: Cyclicality of Individual Earnings and Wages: United States, Sweden, and Germany (SIAB)

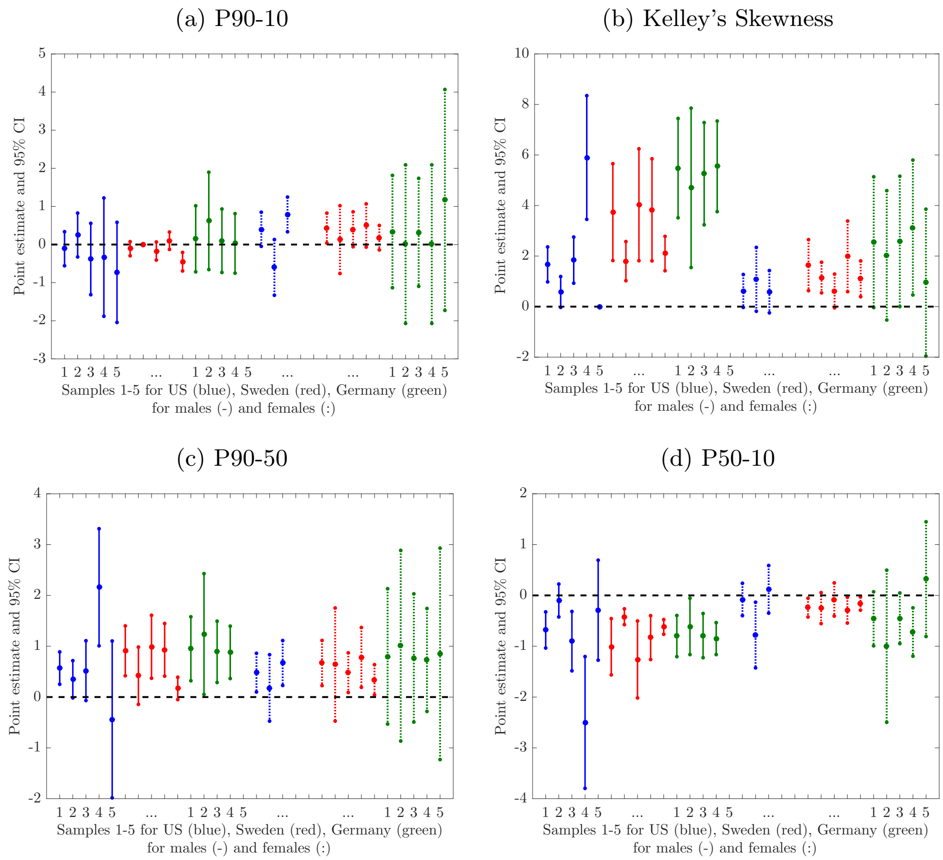

Note: The samples are (1) Earnings: full sample , (2) Earnings: college graduates, (3) Earnings: non-college graduates, (4) Earnings: private sector, (5) Earnings: public sector. For details of samples see text.
For the regressions, see note to table II. Each • reports reports the coefficient on log GDP change.

Figure 1.5 confirms the picture that emerged in Figures 1.3 and 1.4; higher order earnings risk is similar across groups. There are however some noteworthy differences. The magnitude of cyclicality is stronger for non-college graduates as compared to college graduates—particularly in the U.S. and Sweden, where it is about three times stronger. Moreover the magnitude of cyclicality for public sector workers is weaker in all countries. For example, in Sweden, the procyclicality of Kelley's measure of earnings is lower for the public sector (2.10*** for males and 1.10*** for females) compared with the private sector (3.83*** for males and 1.99*** for females). For males this is due to differences in the top tail; it compresses strongly for private sector employees, whereas it is acyclical in the public sector. The L5010 gap on the other hand fluctuates by comparable magnitudes for both groups. For women, the reduced cyclicality is due to both tails fluctuating less.

Overall, it is somewhat surprising that, even for workers in the public sector in a country like Sweden with a reputation for high levels of public insurance, there is robust evidence of higher downside risk in recessions—compression of the top and expansion of the bottom—even if the magnitudes are somewhat smaller than in the private sector. This finding further strengthens the conclusion of this section that increasing downside earnings risk appears to be a robust feature of business cycles in developed countries, even with very different labor market institutions.

Differences Across Occupations

We now turn to occupations and explore the heterogeneity of cyclical earnings changes along this dimension. We are able to conduct this analysis for Germany; the SIAB provides time-consistent occupational codes based on the KldB-88, the 1988 version of the classification of occupations by the German Federal Employment Agency. We now run the cyclicality regressions separately for each occupation, where a worker contributes to the earnings changes of occupation j from t to $t+1$ if in year t he or she works in that occupation.

We first consider the highest level of disaggregation in the KldB-88, which defines five broad *occupational areas*; (1) farming, gardening, animal breeding, fishing, and similar occupations, (2) mining and mineral extraction, (3) manufacturing and fabrication, (4) technical occupations like engineering or laboratory work, and, (5) service occupations. For each occupation we have computed correlations between moments and log GDP change using equation 1.1. Detailed results can be found in Appendix **??**. The results are quite similar as compared with those for the full sample; in particular for male workers in manufacturing occupations, technical occupations, and service occupations.

We also consider a more disaggregated analysis and re-run the regressions for 30 *occupational segments*. While there are then more variation across occupations in terms of earnings cyclicality, the general pattern seen in the full sample of male and female remain; the lower tail is countercyclical for most occupational segments and the upper tail is mostly procyclical. For both males and females the tail movements translate into the cyclicality procyclicality of Kelley's skewness in the by now familiar way.

Summing up, we find that broad occupational groups experience similar cyclicality with farming and mining related occupations being less cyclical. Regressions at finer level of disaggregation point towards interesting heterogeneity of earnings cyclicality across occupations.

1.4.3 Cyclicality of Earnings vs. Wages

A natural question that is raised by these results is whether the observed cyclicality of earnings changes can be attributed mainly to changes in wages or to increased risk of unemployment in economic downturns. The SIAB contains detailed information on the duration of each employment spell and on whether it is a part-time or full-time job. Focusing on full-time workers, we analyze the cyclicality of the distribution of wage changes and compare the results to the ones on earnings changes. We define a worker as full time if his or her full-time spells add up to at least 50 weeks of employment in a given year. (A less strict definition of full-time workers as 45 weeks of employment does not change the results.) The wage variable is the average daily wage rate, where the average is taken over all full-time spells. The same measure has also been used in Dustmann *et al.* (2009); Card *et al.* (2013).[15]

[15] In Germany, a full-time worker is entitled to an annual vacation time of 4 to 6 weeks, which is counted as part of the employment spell.

Table III: Cyclicality of Individual Earnings vs. Wages; Germany (SIAB)

	L9010	Kelley	L9050	L5010
		Males		
Earnings	0.15	5.48***	0.95***	–0.80***
	(0.36)	(5.80)	(3.14)	(–4.11)
Full-Time Wages	–0.09	4.73***	0.30***	–0.39***
	(–0.54)	(6.31)	(3.77)	(–3.20)
Full-Time Wages	–0.12	4.98***	0.28***	–0.40***
(Firm Stayers)	(–0.81)	(5.78)	(3.29)	(–3.20)
		Females		
Earnings	0.34	2.55**	0.80	–0.46*
	(0.48)	(2.05)	(1.25)	(–1.80)
Full-Time Wages	0.03	2.12***	0.17**	–0.14
	(0.18)	(5.11)	(2.61)	(–1.58)
Full-Time Wages	0.02	2.28***	0.16***	–0.14
(Firm Stayers)	(0.13)	(4.84)	(3.17)	(–1.61)

Note: See notes for Table II.

In Table III, rows 1 and 4 reproduce the results from Tables II for completeness. The first set of new results are in rows 2 and 5: these report the cyclicality regressions using average daily wages instead of annual earnings. The main finding for both males and females is that the cyclicality of wages for full-time workers are remarkably similar to the cyclicality of earnings. Specifically, both measures of dispersion of wages are acyclical as was the case for earnings, and the point estimates for both skewness measures are very close for wages and earnings.[16] Naturally, the dispersion of earnings changes is wider than the distribution of wage changes, which is reflected by the point estimates on the tails (last two columns), which are about half as big for wage changes.

A question that remains is what happens to the wages of workers that stay at the same firm. We therefore further restrict the sample to those workers that work at least 50 weeks for the same employer in both year t and t+1.[17] The

[16] The sample of full-time female workers contains about 73% of women (who make for only 54% of the observations) that contribute to the measures of earnings change for women. The corresponding figures were 88% of individuals and 82% of observations for males. This implies that part-time employment plays a more important role for the female sample.

[17] The sample of full-time female workers that do not switch firms contains about 61% of

second set of new results is in rows 3 and 6: the cyclicality regressions for average daily wages for those workers who work at the same firm. The remarkable result is that even for those we observe the same qualitative pattern of cyclicality of wage changes. By and large, these results strongly indicate that the cyclicality results are driven by changes in wages even for full time workers and not by hours.

1.5 Introducing Insurance

We now turn to various sources of insurance available in modern economies and gauge the extent to which they are able to mitigate such downside risk over the business cycle.

1.5.1 Within-Family Insurance

In the previous section, we have shown that higher-order moments drive individual earnings risk over the business cycle. While it is important to understand the underlying nature of labor income risk and the systematic differences across groups, most of our samples are composed by individuals in cohabitation.[18] Assuming pooling of resources within the household, the relevant income measure for many economic decisions is the joint labor income in the household, not individual income. We therefore shift our attention to joint labor earnings at the household level in order to shed light on the role of informal insurance mechanisms within the household. As mentioned earlier, it is not possible to link individuals in SIAB, so we rely on SOEP data instead.

women (who make for about 40% of the observations) that contribute to the measures of earnings change for women. The corresponding figures were 80% of individuals and 65% of observations for males.

[18]Only 12% of our benchmark individual sample in the United States lives in a singe-person household, for example.

Mixed Evidence of Within-Family Insurance

The first row of each panel in Table IV displays the cyclicality of each moment of household earnings changes. In order to get a feeling for the decrease (or increase) of exposure to business cycle fluctuations, we compare these results to the corresponding measures for individual earnings from Table II and in particular male earnings as these on average constitute 71, 60, and 62 percent of household earnings in the United States, Sweden, and Germany, respectively. Additional evidence comes from the graphical analysis of the dispersion, skewness, and the tails, of male earnings changes and household earnings changes in Figures 1.6 and 1.7, respectively.

Table IV: Cyclicality of Household Earnings

	L9010	Kelley	L9050	L5010
	United States			
Earnings	0.23	1.97***	0.93***	–0.71***
	(0.74)	(6.17)	(4.96)	(–3.20)
Post-Gov	0.59**	1.17***	0.72***	–0.14
	(2.44)	(3.13)	(3.42)	(–0.86)
Disposable	0.63*	1.13***	0.74***	–0.12
	(1.90)	(4.83)	(3.75)	(–0.65)
	Sweden			
Earnings	–0.02	2.24***	0.50***	–0.52*
	(–0.08)	(3.33)	(4.94)	(–2.00)
Post-Gov	–0.41*	0.94**	–0.03	–0.38**
	(–2.00)	(2.38)	(–0.44)	(–2.33)
Disposable	–0.43	1.50***	0.06	–0.49**
	(–1.64)	(3.89)	(0.61)	(–2.67)
	Germany (SOEP)			
Earnings	–1.31***	1.88**	–0.05	–1.26***
	(–3.60)	(2.68)	(–0.18)	(–4.26)
Post Gov	–0.18	0.66	0.07	–0.25
	(–1.09)	(0.85)	(0.32)	(–1.28)
Disposable	–0.16	0.56	0.05	–0.22
	(–1.11)	(0.67)	(0.21)	(–1.19)

Note: See notes for Table II.

Considering cyclicality of dispersion, the patterns and magnitudes for household earnings line up with the ones described for individual male earnings for all countries: household earnings changes display no cyclicality of dispersion. This is true especially for Sweden and the United States. The countercyclical measure of dispersion (as measured by L9010) for Germany is driven by the lower tail and thus the overall pattern here mirrors the one of male earnings dispersion in SOEP (see Section 1.4).

The analysis of Kelley's skewness–and the inspection of the tails–yields very interesting results when comparing the three countries. In Sweden, intra-family insurance plays an important role in reducing downside risk over the business cycle: The estimated coefficients on household earnings are smaller than those on male earnings and quite close to those on female earning. For example, the coefficient on Kelley's skewness is about 2.2 as compared to 3.7 and 1.6 on male and female earnings respectively. The difference is primarily driven by both tails reacting less than those for male earnings; the lower tail by about half and the upper tail by almost as much as compared to male earnings. Repeating the illustrative calculation from above, this would imply a move from an upper tail to lower tail ratio of 50/50 in a typical expansion to 45/55 in a recession—much smaller compared to the change to a ratio of 41/59 for male earnings and very similar to a ratio of 46/54 for females.

Evidence of within-family insurance is weaker for the United States and Germany. In both economies, the results suggest somewhat higher downside risk in recessions for household earnings that that for male earnings, and much higher risk than that for female earnings. Considering the tails separately for the two countries, the slightly stronger reaction of Kelley's skewness is primarily driven by larger movements in the upper tail in United States, whereas it is the lower tails that widens more in Germany.

In order to shed further light on the insurance within households, we consider the cyclicality of income for actual households in comparison to income changes for randomly formed couples. This way we want to see if there is anything special about households visible in the data, or if the dynamics of household income just represent the dynamics of male and female income. We therefore randomly pair heads and spouses for each t to t+1 change. For each random couple, we make sure that artificial income is above the lower income. The first set of results in each country panel of table **??** shows the bootstrapped mean, standard deviation and 10-90 confidence band of the regression coefficients. In both the US and Germany, we find the random couples to experience lower downside risk than actual households as measured by the cyclicality of L5010. For Sweden, the random couples' L5010 shows the same cyclicality as actual households. The next rows show the same results when not randomly pooling all heads and spouses, but controlling for some observables on the side of the head. When we control for age, we group heads into 7 age groups and in the pool of spouses for each age group are all spouses of heads in the actual data. Finally, we do the random coupling by age and education groups. As expected, the cyclicality experienced by random couples is more and more similar to actual households. Still, for the US and Germany we find actual households experiencing slightly higher cyclicality of earnings changes than their artificial counterparts. This suggests that the correlation between head's and spouse's labor market income is higher than for a random counterpart and uncontrolled characteristics play some role - like, e.g., most heads and spouses working in the same local labor markets.

We conclude that the responses of gross household earnings are heterogeneous across countries, with Sweden being the only economy where the family plays a clear insurance role against aggregate fluctuations. However, it is hard to extract further conclusions in disconnection to taxes and transfers payed and received by the household. In order to shed light on this issue, we move on to considering the role of social insurance policy over the business cycle.

Figure 1.6: Standard Deviation (Left) and Skewness (Right) of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden

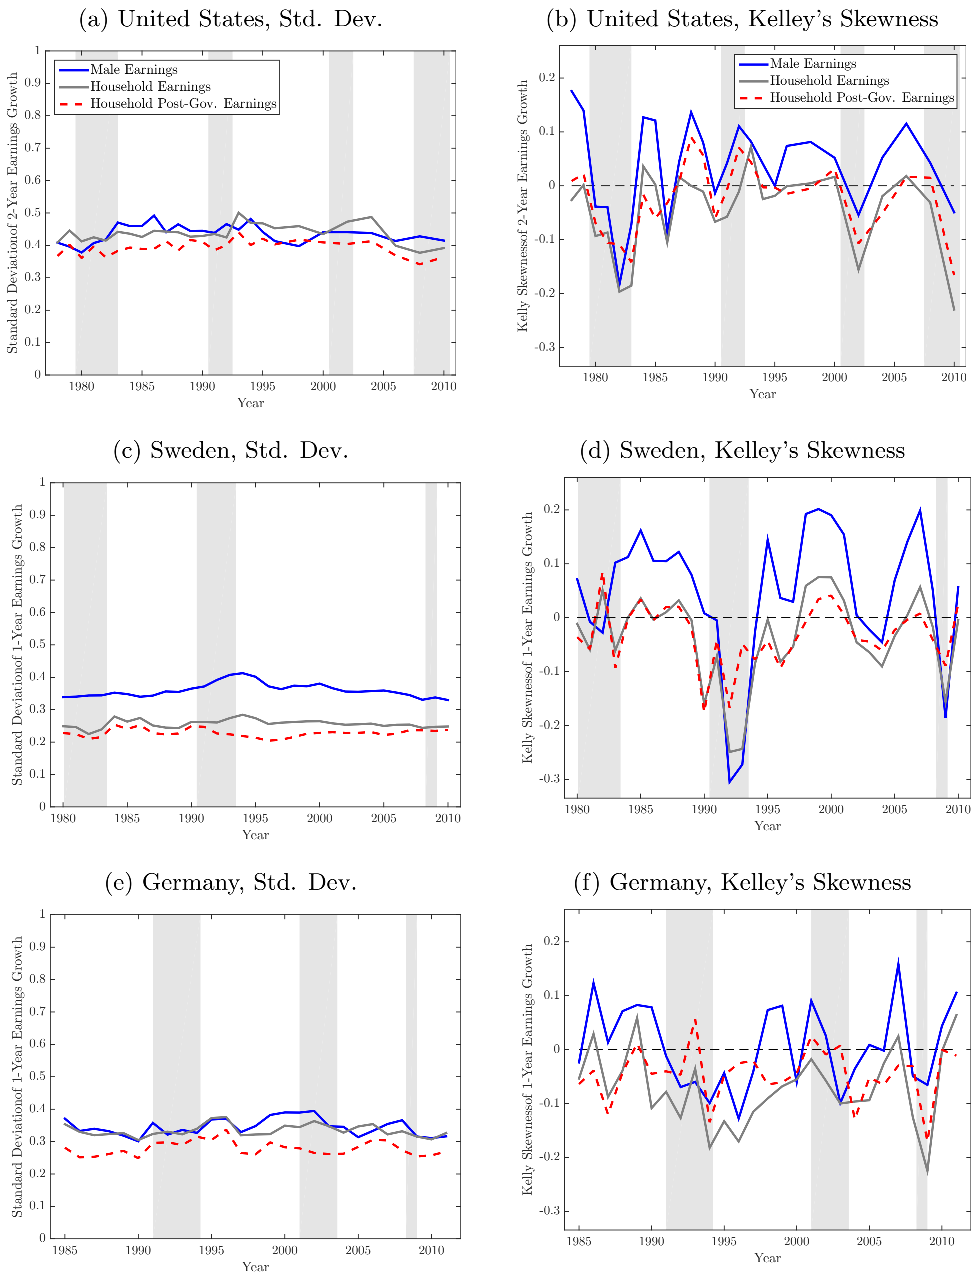

Note: Linear trend removed, centered at sample average.

Figure 1.7: Tails of Short-Run Earnings Growth: United States, Germany (SOEP), and Sweden

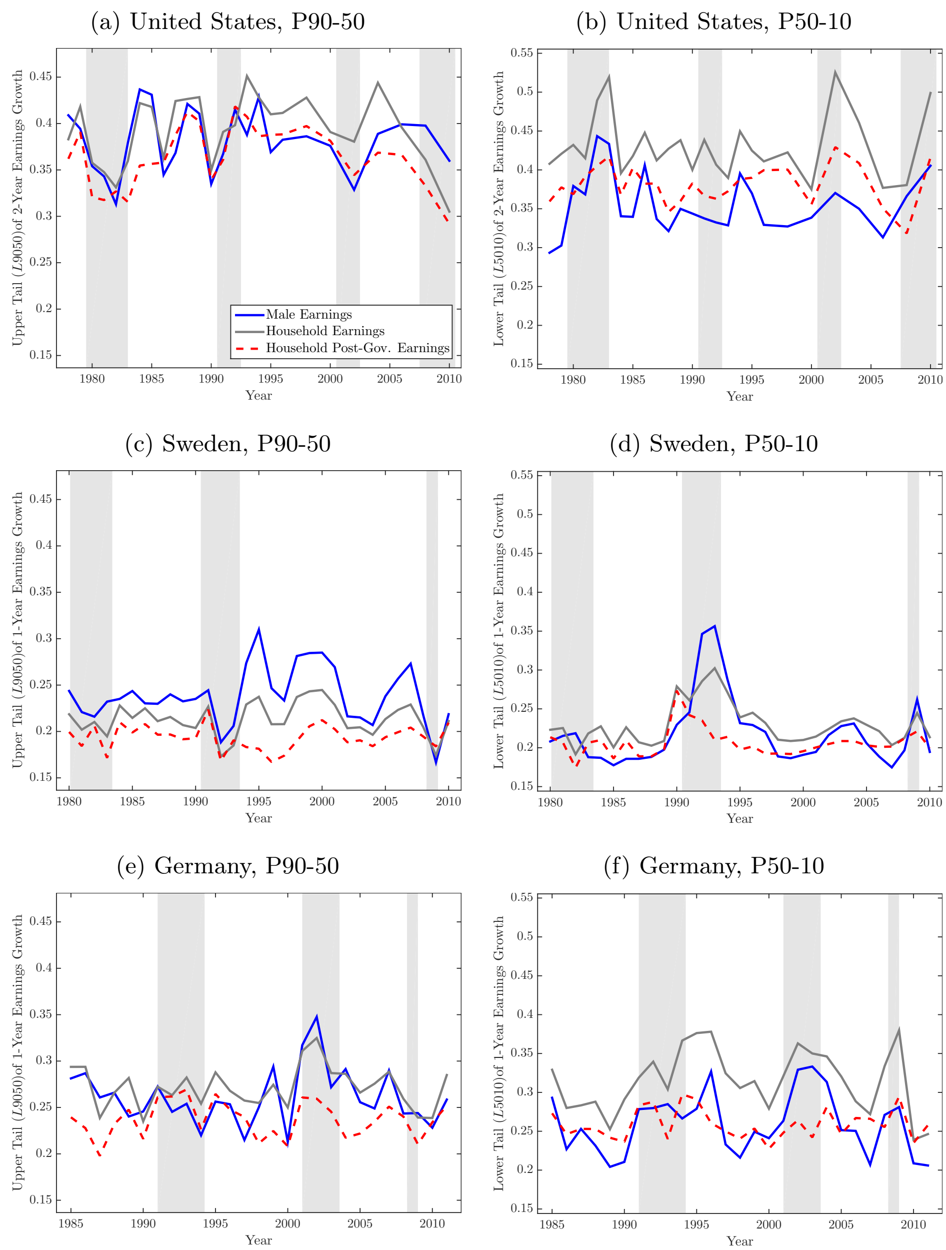

Note: Linear trend removed, centered at sample average.

1.5.2 Government and Social Insurance Policy

Focusing on the household as the relevant unit, we analyze the effectiveness of social policy in mitigating business cycle risk in addition to any insurance arrangements made within households. We evaluate the total insurance effect of the tax and transfer system by analyzing the cyclicality of post-government earnings as compared to household gross earnings. In order to gain insights on the effectiveness of different policies, we then evaluate the relative importance of several subcomponents of transfers using the empirical tools employed in the previous analysis on income measures that in turn add certain transfers to household gross earnings.

For the analysis of subcomponents, we consider three main groups of transfers that are comparable across countries and for each country are consistently measured over time. The groups are (1) labor-market-related policies, (2) aid to low-income families, and (3) "pensions," and are listed in Table V. Labor-market-related policies mainly consist of unemployment benefit payments—this component of social insurance policy is of particular importance for the mitigation of increased downside household earnings risk in recessions, if the nature of downside risk is (temporary) job loss of household head or spouse.

Table V: Components of Social Policy

	LINDA	SOEP	PSID
1. Labor Market Transfers:	Unemployment benefits; Labor market programs	Unemployment benefits	Unemployment benefits; Workers' compensation
2. Aid to Low-Income Families:	Family support; Housing support; Cash transfers from the public; (no private transfers)	Subsistence allowance; Unemployment assistance (up to 2004); Unemployment benefits II (since 2005)	Supplemental Security Income; Aid to Families with Dependent Children (AFDC); Food Stamps; Other Welfare
3. Social Security and Pensions:	(Old Age) Pensions	Combined old-age, disability, civil service, and company pensions	Combined (Old Age) Social Security and Disability (OASI)

Note: Table lists the measures used in the three data sets to construct subcomponents of transfers.

The second component considered, "aid to low-income families," consists of several measures of social insurance policies specifically aimed at at-risk households. The relevance of this type of transfer can therefore be expected to matter most for low-income households who have a higher likelihood of falling down to fulfilling 'at-risk' criteria in the course of a recession. The third component, pension payments, is not directly connected to business cycle considerations. It can still play a relevant role for household members near or at retirement age, who may take up pension payments instead of unemployment payments if they decide to leave the labor market upon job loss.

The Overall Effect of the Tax and Transfer System

We begin with a brief discussion on the overall effect of the government, comparing the cyclicality of pre- and post-government measures of household earnings listed in rows 1 and 2 of Table IV. Again, Figures 1.6 and 1.7 visualize the findings. We find that social policy is an important source of insurance against aggregate fluctuations in all three economies, with very similar overall effects. Motivated by the considerations from above sections, we directly consider the reactions of the upper and lower tails of income changes. In all three economies, downside risk is mitigated successfully by the tax and transfer system. In both the United States and Germany, the lower tail of post-government earnings changes is unresponsive to the business cycle–while significantly countercyclical for pre-government earnings. In Sweden, lower tail counter-cyclicality is dampened but still statistically significant (from a point estimate of –0.52 to –0.38).

Considering the cyclicality of the upper tail reveals differences between the countries. In Germany, it is unresponsive to the cycle for both pre-and post-government earnings. While both the U.S. and Sweden reveal procyclicality of L9050 of pre-government earnings changes, the L9050 of post-government earnings changes is acyclical in Sweden, but still procyclical in the United States. The different reactions of the tails translates into procyclical overall dispersion of post-government earnings changes in the U.S., and countercyclical dispersion in Sweden. Summarizing the reaction of overall dispersion and tails results in procyclicality of Kelley's skewness measure for both countries; though the procyclicality is much smaller for post- than for pre-government earnings.

To sum up, the analysis suggests that downside risk in recessions is mitigated by taxes and transfers. In Sweden, an additional effect are lowered upside chances in expansions. This lines up with considerations of Sweden as a country with a high degree of redistribution.

The Role of Subcomponents of Social Policy

The measure of post-government earnings used so far lumps a lot of very different transfers received and taxes paid by households. While this measure is appropriate for assessing the overall effect of the tax and transfers system, it is not as well suited for understanding the success of different social policies that specifically aim at mitigating downside risk or that aims at aiding low-income families, who can be expected to be especially vulnerable in recessionary periods. Therefore, we now consider different types of transfers separately. The results of the cyclicality analysis are listed in Table VI. As for for the estimates of total taxes and transfers, we compare the coefficients to the ones from the household gross earnings analysis in row 1 of Table IV. Recall that in order to be in the year t base sample for the analysis, the lowest considered income measure of a household needs to be above the income threshold for that year. This way, we ensure that the sample is stable at the lower end of the distribution and results are not driven by low-income households entering the sample for a certain type of transfer but are not in the sample when considering another.

Table VI: Cyclicality of Household Earnings - transfers added separately

	L9010	Kelley	L9050	L5010
	United States			
+ Labor transfers	0.60	1.59***	0.92***	–0.33
	(1.54)	(5.20)	(4.20)	(–1.34)
+ Aid to low-income	0.21	1.90***	0.89***	–0.69***
	(0.77)	(6.13)	(5.16)	(–3.33)
+ Pensions	0.22	1.82***	0.86***	–0.64***
	(0.80)	(5.61)	(4.79)	(–3.06)
	Sweden			
+ Labor transfers	–0.22	1.14***	0.13*	–0.35**
	(–1.23)	(4.23)	(2.04)	(–2.58)
+ Aid to low-income	–0.07	2.11***	0.42***	–0.49**
	(–0.38)	(3.72)	(4.51)	(–2.47)
+ Pensions	–0.07	2.34***	0.48***	–0.55**
	(–0.43)	(3.55)	(4.50)	(–2.68)
	Germany (SOEP)			
+ Labor transfers	–1.09***	1.34**	–0.13	–0.96***
	(–2.96)	(2.50)	(–0.60)	(–3.65)
+ Aid to low-income	–1.32***	1.66**	–0.11	–1.21***
	(–3.82)	(2.40)	(–0.47)	(–4.08)
+ Pensions	–1.21***	1.80***	–0.04	–1.17***
	(–3.30)	(3.10)	(–0.18)	(–4.58)

Note: See notes for Table II.

The results in Table VI show that out of the three transfer components, labor market related transfers (which have unemployment benefits as the main component) accounts for most of the reduction in downside. The other two components of transfers do not have any impact on cyclicality as measured by our cyclicality regressions. For all three economies, the point estimates when adding aid to low-income families or pensions are almost identical to the ones for gross earnings.

A closer look at the estimated coefficients reveal some interesting differences between the countries. In Sweden, labor market transfers account for almost the whole difference between pre- and post-government earnings cyclicality in the tails (compare rows 1 and 2 in Table V with row 1 in Table VI). Thus only a tiny amount is accounted for by the Swedish tax system or other transfers.

In the U.S., labor market transfers similarly accounts for the entire reduction in lower tail cyclicality. But the upper tail is unaffected by labor market transfer, and it is barely affected by aid to low-income families or pensions. This suggest that the lower procyclicality of the upper tail of U.S. post-government earnings changes is accounted for by U.S. tax system (or some interaction between taxes and transfers).

Finally, in Germany labor market transfers also mitigates downside risk, but it does so to a lesser extent than in Sweden and the United States. Rows 1 and 2 in Table V and row 1 in Table VI, shows that households earnings plus labor market transfers display significant down-side risk, (smaller but quite similar coefficients that those on household earnings), whereas post-government earnings changes are acyclical. The former finding is corroborated on individual earnings changes using our larger sample based on the SIAB data base. Besides individual earnings SIAB also contains information on unemployment benefits at the individual level. Table VII shows results for individual level regressions for male and female earnings separately, when unemployment benefits are excluded (rows 1 and 3) and included (2 and 4). These individual level results line up well with the household level analysis conducted using SOEP data; labor market transfers mitigate the cyclicality of the tails but there is still significant higher order income risk even when unemployment benefits are included in the income measure. This suggest that the German tax system (or interaction terms between taxes and transfers) is the primary reason for post-government earnings being acyclical.

Table VII: Cyclicality of Individual Earnings including unemployment benefits in Germany (SIAB)

	L9010	Kelley	L9050	L5010
Male Earnings	0.11	5.71***	0.97***	–0.86***
	(0.26)	(5.32)	(2.93)	(–4.40)
+Unempl. benefits	0.15	5.12***	0.84**	–0.70***
	(0.34)	(5.24)	(2.61)	(–4.01)
Female Earnings	0.46	2.69*	0.89	–0.44*
	(0.60)	(1.92)	(1.26)	(–1.74)
+ Unempl. benefits	0.50	2.43*	0.82	–0.32
	(0.67)	(1.82)	(1.22)	(–1.43)

Note: See notes for Table II. Difference to estimates in II are due to the fact that regressions start in 1981 instead of 1976.

1.5.3 Sensitivity of results to choice of lag length

All results reported in the text refer to the distribution of what we label transitory, i.e., one-year changes of several income measures.[19] Given the focus of Storesletten *et al.* (2004) or Guvenen *et al.* (2014), to which we relate our results, on persistent income changes this choice needs to be discussed. The main reason for us to focus on one-year changes is that we choose a regression framework as our main tool of analysis. We make this choice, because we compare the cyclicality of income risk across countries. While for the US it is widely accepted to base the dating of business cycles on NBER recession dates, this dating is less clear cut for both Germany and Sweden. More generally, it is not clear that in a cross-country comparison the dating of business cycles is of the same quality in terms of capturing actual economic conditions. Our regression framework allows a very clear interpretation and comparison of cyclicality of income changes.

Moving to five-year changes—which are closer to capturing persistent changes—would imply problems with the regression analysis for two reasons. One option would

[19]Recall that for the US we define two-year changes as transitory in order to account for the biannual nature of the PSID since 1997.

be to use non-overlapping five-year changes of income and GDP, another would be to use overlapping changes. The first option would give too few data points for a regression analysis, while the second would open the door to usual problems of overlapping data.

The time-series of five-year changes is shown in figures **??** to **??** in Appendix **??**. Comparison to the one-year changes suggest the same qualitative patterns.

1.6 Conclusion

This paper has studied how higher-order income risk varies over the business cycle, as well as the extent to which such risks can be smoothed within households or with government social insurance policies. To provide a broad perspective on these questions, we studied panel data on individuals and households from the United States, Germany, and Sweden, covering more than three decades of data for each country.

We find that the underlying variation in higher-order risk is remarkably similar across these countries that differ in many details of their labor markets. In particular, in all three countries, the variance of earnings shocks is almost entirely constant over the business cycle, whereas the skewness of these shocks becomes much more negative in recessions.

Government provided insurance, in the form of unemployment insurance, welfare benefits, aid to low income households, and the like, plays a more important role reducing downside risk in all three countries; the effectiveness is weakest in the United States, and most pronounced in Germany. For Sweden we find that insurance provided within households plays a similar role.

Overall, we have provided evidence of the important role played by government policy in insuring households against aggregate fluctuations that originate asymmetric changes in the earnings distribution. Furthermore, we have shown how the effects vary by the characteristics of the individuals and the specific public instrument. A quantitative analysis is out of the scope of this paper.

However, our results call for further research on public policy design that accounts for the asymmetric response of idiosyncratic earnings risk to business-cycle fluctuations.

Chapter 2

Employment Protection Legislation and Earnings Risk

2.1 Introduction

What are the implications of employment protection legislation for individual income risk? In this paper, we characterize the dynamics of earnings uncertainty in dual labor markets. The purpose of this study is two-fold: On the one hand, we seek to shed light on the institutional determinants of earnings risk over the workers' lifetime. In other words, how different exposure to stringent or flexible labor markets can influence income uncertainty and, hence, individual behavior. On the other hand, we want to investigate the particular case of income dynamics in dual labor markets.

Dual labor markets are limited by two forms of contracting: highly protected open-ended contracts and unprotected temporary contracts. The latter were introduced with the goal of lowering unemployment and acting as stepping stones towards better job matches. Nevertheless, this labor market regime has been found to favor various detrimental labor market outcomes, namely lower human capital accumulation, lower productivity, and higher inequality Bonhomme and Hospido (2013) than comparable fully rigid or fully flexible arrangements. Yet, to the best of our knowledge, the literature has been silent about its implications on earnings dynamics.

There is evidence that employment protection legislation is failing at decreasing the exposure of workers to income volatility. Cross-country empirical studies in Europe (OECD (2011)) have found a significant positive correlation between employment protection and earnings volatility. Given the high firing costs and the low churning rate of regular forms of employment in most of these countries, this relation is puzzling. We will show that our results suggest that this relation can be explained by the fact that the countries with the highest firing costs also have a high fraction of the active population in temporary contracts.

A key element of our analysis is the use of high-quality administrative panel data from Spain. We combine the detailed information on employment histories from the Spanish Social Security Administration (MCVL hereafter, following its acronym in Spanish) and the linked income data from the tax registers. The former contains very detailed information on all the jobs a worker has held over her career, including employer, type of contract, number of days, and salary. The Tax data complements the salary reported in the MCVL with other forms of income and, importantly, with non-top-coded earnings. The combination of these two linked sources of data is a useful and unique resource for our study, but also presents challenges. The main concern we face is the representativeness of the sample. While the MCVL selects a 4% representative sample of the population, the tax supplement excludes some groups. In Section 2.2 we explore the importance of these issues.

In our analysis, we pose a model of income dynamics that captures the salient features of the life-cycle. We estimate the model using a method of moments and allow variation by the different exposure to temporary forms of employment. In particular, we estimate the age profile of the persistence and variance of labor income shocks, separately for workers that spent most of their career in fixed-term contracts – job-unstable – and the rest – job-stable –. We use college-educated males as the benchmark group in order to isolate from observed heterogeneity. We exploit the panel dimension of our data and include individual fixed effects to control for unobserved heterogeneity.

We find that the job-unstable workers face larger uncertainty, as represented by the variance of permanent shocks, at all ages after the age of 25. This uncertainty is U-shaped for the job-unstable group, while it is mostly decreasing in age for the job-stable group. The persistence of shocks is hump-shaped for both groups, peaking around the mid-career years at 0.95, but it is on average higher for the job-unstable group (0.85 as compared to 0.75). We perform the same analysis for females and less educated groups, with varying results. Finally, we illustrate the importance of capturing the differential dynamics of temporary and permanent workers in a simple application.

The rest of the paper proceeds as follows. Section 2.2 describes the data and the sample in detail. Section 2.3 presents the empirical strategy and estimation. Section 2.4 discusses the results. Section 2.5 concludes.

2.2 Data and Sample

2.2.1 The Data: Social Security and Tax Records

We use administrative data from the Continuous Sample of Working Histories on earnings and working histories of Spanish workers. The data is provided by the Spanish Social Security Administration in cooperation with the IRS counterpart in Spain.

The MCVL consists of a 4% representative random sample of all workers affiliated with the social security administration within a given year between 2004 and 2015. Besides, starting in 2005, the sample has a panel design: all individuals present in each wave subsequently remain in the sample. Retroactive information on the whole working history is provided as early as 1962 for work variables and 1980 for earnings. Bonhomme and Hospido (2013) show that the sample is representative at least since the late 1980s. The main drawback is that earnings data is bottom- and top-coded. We, therefore, complement the earnings data with the tax supplement provided by the IRS and matched to the Social Security records. The tax supplement contains non-top-coded information on annual earnings but does not include such detailed job characteristics or earnings a higher-frequency levels.

2.2.2 Sample Selection

We select individuals at least 19 and at most 60 years old. For the case of the college graduates, we further restrict the sample to individuals at least 22 years old. That is, we drop those who graduated before the age of 22.

2.2.3 Main Variables

Earnings. The earnings data are extracted from the "Annual summary of retentions and payments for the personal income tax on earnings, economic activities, awards, and income imputations" (known as *Modelo 190*). All employers are required to fill out Modelo 190 with the total compensation paid to each of their employees during the year, independently of whether or not they pay labor income taxes. To obtain a measure of total annual labor earnings, we add all the incomes that correspond to each worker during the reference year.

Educational Attainment. We define four education groups:

1. Less than high school graduate
2. Finished first two years of high-school, and possibly a lower-ranked vocational training.
3. High school graduates, and possibly an advanced vocational training, but not college.
4. College and above.

Annual Employment Status. Given that our period of observation is one year, it is not uncommon to find workers that hold different simultaneous jobs or that change jobs within the same year. In some cases, some of those contracts are temporary and some permanent. This poses a challenge when defining employment spells at the annual level. To be able to exploit the daily employment information back to the 1980s, we define employment status in terms of share of annual time spent in each kind of job: permanent, temporary, or none. Workers who have zero annual earnings or earn less that the corresponding amount to a month minimum-wage salary are considered unemployed.

Young **Employment Status.** We define young permanent workers as those who spent most of their 20s (20 to 30 years old) under a permanent contract. We define young temporary workers and young unemployed accordingly.

Prime **Employment Status.** We define prime permanent workers as those who spent most of their peak working years (35 to 60 years old) under a permanent contract. We define prime temporary workers and prime unemployed accordingly.

Lifetime **Employment Status.** We define lifetime permanent workers as those who spent most of their working years under a permanent contract. We define lifetime temporary workers and lifetime unemployed accordingly.

2.3 Empirical Analysis

2.3.1 A Framework to Measure Life-Cycle Earnings Risk

In this section, we pose and estimate a stochastic income process to measure the uncertainty faced by different workers at each age. We group workers based of their exposure to temporary forms of employment during the first 10 years of their working lives. The statistical framework follows Karahan and Ozkan (2013).

More specifically, let Y_{iat} be Annual earnings for individual i at age a in year t. We assume $\log Y_{iat}$ is given by

$$\log Y_{iat} = \beta X_{iat} + y_{iat}, \tag{2.1}$$

where X_{iat} is a vector of observable characteristics that includes a quartic polynomial in age, year dummies, and region dummies. β is assumed constant in time.

The residual income from the first stage y_{iat} is decomposed into three components: (1) an individual deterministic component $\alpha_i + \gamma_i a$, formed by a fixed effect and a linear trend; (2) a stochastic persistent component z_{iat}, modeled as an AR(1) with persistence $\rho < 1$; and (3) a stochastic transitory component u_{iat}, represented by a MA(1). The specific structure is given below by equations

(2.2)-(2.5):

$$y_{iat} = \alpha_i + \gamma_i a + u_{iat} + z_{iat} \qquad \alpha_i \sim N(0, \sigma^2_{\alpha,a}), \gamma_i \sim N(0, \sigma^2_{\gamma,a}) \tag{2.2}$$

$$u_{iat} = \varepsilon_{iat} + \theta \varepsilon_{i,a-1,t-1} \qquad \varepsilon_{iat} \sim N(0, \sigma^2_{\varepsilon,a}) \tag{2.3}$$

$$z_{iat} = \rho z_{i,a-1,t-1} + \eta_{iat} \qquad \eta_{iat} \sim N(0, \sigma^2_{\eta,a}) \tag{2.4}$$

$$z_{i0t} = 0, \quad \varepsilon_{i0t} = 0. \tag{2.5}$$

To capture the evolution of uncertainty over life, parameters ρ, σ_ε, and σ_η are functiones of age. The exact form of age dependence will be discussed next.

2.3.2 Estimation

Our baseline model assumes that $\sigma^2_{\varepsilon,a}$, $\sigma^2_{\eta,a}$, and ρ are all cubic functions of age, and θ, σ_α, and σ_γ are fixed across ages.

We minimize the distance between the empirical and the model-implied closed-form covariance matrix to estimated the parameters, using Generalized Method of Moments with efficient weighting matrix, to estimate the parameters.

$$(\theta, \sigma_\alpha, \sigma_\gamma, \rho(\cdot), \sigma_\varepsilon(\cdot), \sigma_\eta(\cdot)) \tag{2.6}$$

2.3.3 Grouping Workers

Based on the definition of *young* employment status, we define two groups of workers: those that have spent 50% or more of their days before the age of 30 in a temporary contract, which we call job-unstable, and the rest, denoted job-stable. We estimate the process separately for each group.

2.4 Results

2.4.1 Life-Cycle Earnings Dynamics of Temporary and Permanent Workers

This section reports the estimates at every age. Overall, the job-unstable group experiences a larger degree of uncertainty and shock persistence. However, the more interesting results concern the dynamics over the life-cycle that we discuss next.

Notice how workers that have been more exposed to temporary contracts face more uncertainty over life. Moreover, looking at the variance of permanent shocks, at all ages above 25 uncertainty is U-shaped for the job-unstable group, while it is mostly decreasing in age for the job-stable group. The persistence of shocks is hump-shaped for both groups, peaking around 0.95, but it is on average higher for the job-unstable group and it peaks much later in life for those more exposed to temporary contracts.

Looking at the dynamics of transitory shocks, while those in the job-stable group feature a clear life-cycle profile, decreasing in the middle of the career and then increasing as retirement approaches, workers in the job-unstable group barely see a change in their variance.

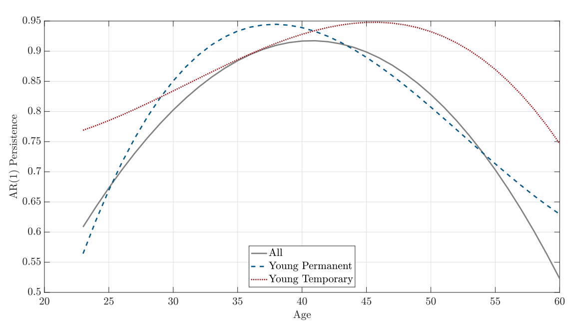

Figure 2.1: Persistence

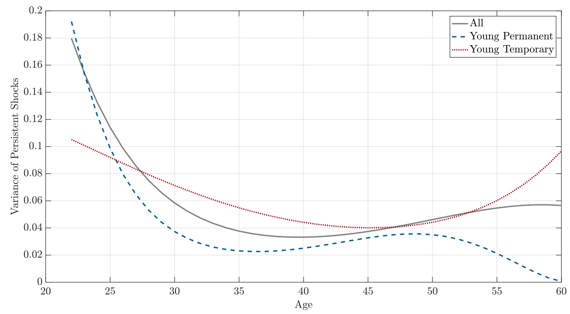

Figure 2.2: Variance of Persistent Shocks

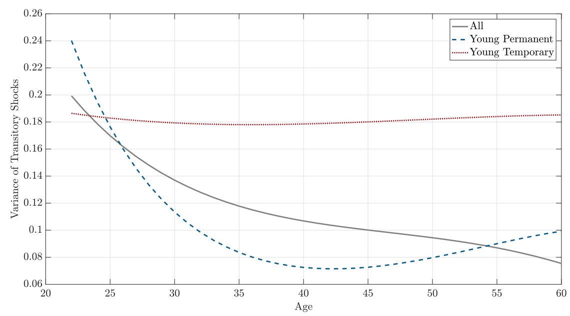

Figure 2.3: Variance of Trasitory Shocks

2.5 Conclusion

This chapter studied the role of labor market institutions in shaping individuals' exposure to earnings uncertainty over their life cycle. Using administrative data from Spain and a method of moments, we estimate the labor income profiles for different workers, depending on their exposure to temporary forms of employment. In particular, we estimate the age profile of the persistence and variance of labor income shocks, separately for workers that spent most of their career in fixed-term contracts – job-unstable – and the rest – job-stable –. We use college-educated males as the benchmark group, in order to isolate forms of observed heterogeneity. We exploit the panel dimension of our data and include individual fixed effects.

We find that the job-unstable workers face larger uncertainty, as represented by the variance of permanent shocks, at all ages after the age of 25. This uncertainty is U-shaped for the job-unstable group, while it is mostly decreasing in age for the job-stable group. The persistence of shocks is hump-shaped for both groups, peaking around the mid-career years at 0.95, but it is on average higher for the job-unstable group (0.85 as compared to 0.75).

These results have implications for the design of public policy. In particular, Cabrales et al (2017) show that allowing for the specific dynamics for each type of worker is crucial to evaluate the effects of the introduction of income-contingent public loans to finance higher education.

Chapter 3

Consumption and Self-Insurance with Higher-Order Earnings Risk

3.1 Introduction

To what extent are a household's labor income fluctuations uninsurable? More broadly, what characterizes the joint distribution of income and consumption changes? These questions are critical for the design of a social insurance policy that is effective for lessening the burden of idiosyncratic income uncertainty on individuals. They are also at the core of understanding the changes in consumption inequality and welfare as a result of shifts in the distribution of labor income. Nevertheless, these questions remain open, largely due to the scarcity of nonparametric studies of the income distribution and consumption panels.

Few studies have inspected changes in income growth from an agnostic, nonparametric perspective. Those who have (Geweke and Keane, 2000; Bonhomme and Robin, 2009; Guvenen *et al.*, 2016), however, reach a remarkably common conclusion: earnings changes are not normally distributed and, in particular, 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; a few experience extreme events. Yet, with few notable exceptions, studies that look at joint changes in income, consumption, and marginal responses 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 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, 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.

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

[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.* (2014).

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 non-durable 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.

3.1.1 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, 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.

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 and 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 3.2 empirically inspects the marginal and joint distribution of earnings and consumption changes, the baseline model and its calibration are described in Section 3.3, Section 3.4 explains the main results and implications. Section 3.5 concludes.

3.2 The Tails of Earnings and Consumption Changes in The Data

This section presents an empirical characterization of tail events in income and consumption changes, as well as their relation. After presenting the data and sample selection, I show the moments of the distribution of the variables of interest, with a focus on the higher-order moments. Next, I examine the joint behavior of consumption and earnings changes. When the marginal distributions are fat-tailed, the tail area dependence might be quite different to that suggested by the correlation. I then define a concept of tail dependence between consumption and earnings changes, separately for durable and nondurable goods, that is novel in this literature.

3.2.1 Data and Sample

This section describes the data used both in the empirical analysis and to construct the moments used in the calibration.

The *new* PSID

The PSID is a longitudinal study of a representative sample of U.S. households, tracking a wide variety of socioeconomic variables from 1968 to 2013. It is also one of the most commonly used micro-data sources for the study of income dynamics. Therefore, in this section, I give a brief overview and focus on the recent 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 PSID originally started as a survey for the study of poverty, hence its focus on socioeconomic covariates and income. Before the 1999 wave, the only information on expenditures was food and rent[3]. Starting in 1999, a wide set of consumption categories was added, comparable to the CEX in the aggregates—although with less detail in the subcategories. The *new* PSID spans 14 years (1998–2012) and contains information on income, consumption, and wealth. In addition, for some exercises, I include the previous five waves (1993–1997) in order to construct controls that concern past family and income characteristics.

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 2012. I keep both the original representative SRC sample and the SEO sample from the PSID, with the appropriate weighting.

For the main analysis, I use data for the period 1998–2012, which corresponds to the waves containing detailed consumption information. In addition to the availability of consumption data, this shorter panel has several advantages over the longer horizon. First, the level of attrition is higher in the initial years of the survey, and the number of households remaining in the sample 30 years later – by the time expenditures are recorded – is very small. Second, the 1993 PSID wave – corresponding to information for the 1992 year – underwent a major revision in main variables concerning labor income; starting in that year eliminates spurious variation and the need to make assumptions to homogenize those variables over time.

The final sample with information on income and consumption is comprised of around 20000 observations, corresponding to approximately 5000 households over 15 years. More details, including the number of observations left at each step of the sample selection, are given in Appendix B.1.

[3]Some 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.

Definitions 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 toÂ 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, 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 B.1.

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.

3.2.2 Tail Changes at the Household Level

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 $\mathrm{arc}\Delta^s(x)$ 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. 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 valueÂ from last year, no adjustment is recorded. For the rest of the cases, I define different situations that are explained in Appendix B.1 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.

Higher-order Moments and Deviations from Normality

To provide a definition of tail changes, I first look at the empirical distribution of labor income changes. The first panel of Figure 3.1 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 3.1.

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,

$$\mathrm{P}k \equiv k\text{th percentile}$$

$$\mathrm{P}k\ell \equiv \mathrm{P}k - \mathrm{P}\ell$$

$$\text{Kelley Skewness} \equiv \frac{P9050 - P5010}{P9010} \tag{3.1}$$

$$\text{Crow-Siddiqui Kurtosis} \equiv \frac{P97.5 - P2.5}{P7525}. \tag{3.2}$$

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.2. 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 B.1 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	
$\log \Delta^2 y_t^{ind}$	**0.654**	**-0.045**	**9.899**
Y	0.660	-0.027	8.701
O	0.649	-0.074	11.059
$\log \Delta^2 y_t^{hh}$	**0.571**	**-0.043**	**7.018**
Y	0.584	-0.048	6.542
O	0.559	-0.040	7.444
$\log \Delta^2 y_t^{post}$	**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	
$\log \Delta^4 y_t^{ind}$	**0.725**	**-0.126**	**8.151**
Y	0.746	-0.121	7.797
O	0.703	-0.164	8.755
$\log \Delta^4 y_t^{hh}$	**0.632**	**-0.095**	**5.838**
Y	0.654	-0.090	5.854
O	0.611	-0.089	5.782
$\log \Delta^4 y_t^{post}$	**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 (3.1) and (3.2) 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 3.2.1 for detailed definitions. Y: Age group 25-44, O: Age group 45-60. See Table A.2 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 ($\log\Delta$)	Average Size ($)
$0 \leq \text{abs}\Delta < 1SD$	82.79	0.15	6671.80
$1SD \leq \text{abs}\Delta < 2SD$	10.77	0.69	24536.16
$2SD \leq \text{abs}\Delta < 3SD$	2.90	1.21	36703.40
$3SD \leq \text{abs}\Delta$	3.53	2.80	43867.68
N	18,524		

Note: $abs\Delta$ 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

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 3.1 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 B.1, and (2) Quantile-Quantile plots (QQ plots hereafter).

The lower panel in Figure 3.1 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[4]. 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

[4] Appendix B.1.3 includes a stylized example of QQ plots for usual distributions.

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.

Figure 3.1: Empirical Distributions: Histograms and Deviations from Normality

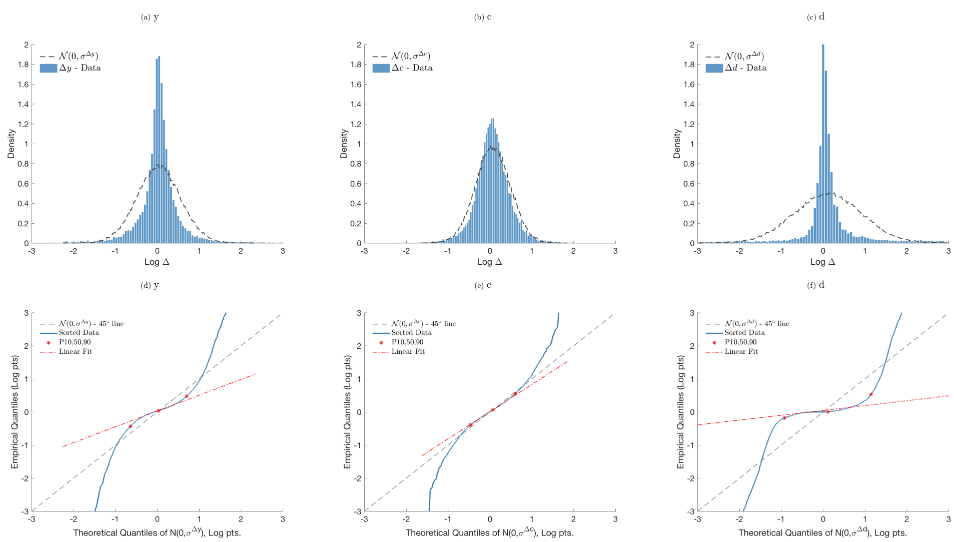

3.2.3 Tail Dependence and the Joint Distribution of Earnings and Consumption

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 is defined 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:

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

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.

Figure 3.2 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.3. 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.

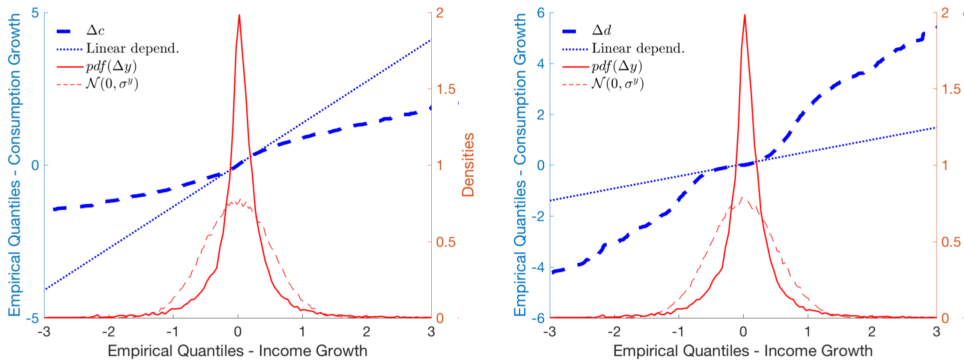

Figure 3.2: The Joint Distribution of Income and Consumption: nondurables (left) vs. durables (right)

3.3 Model and Calibration

3.3.1 Model

In this section I present a life-cycle consumption-savings model with income uncertainty and incomplete markets, with two additions to the standard[5] 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.

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.

In the remaining of this section, I first describe the specifics of the idiosyncratic shocks and durable choices. Next, I go over the remaining elements of the model and formally state the household problem and its solution algorithm. Finally, I proceed to the calibration and estimation.

[5] Aiyagari (1994); Kaplan and Violante (2010).

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 distributed[6] with mean 0 and standard deviation σ_ε. Equations (3.3)-(3.6) formally summarize these relations.

Finally, equation (3.7) 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 int he data.

$$\log Y_{it} = g_t^a + y_{it} \tag{3.3}$$

$$y_{it} = z_{it} + \varepsilon_{it} \tag{3.4}$$

$$z_{it} = \rho z_{it-1} + \eta_{it} \tag{3.5}$$

$$\varepsilon_{it} \sim \mathcal{N}(0, \sigma_\varepsilon) \tag{3.6}$$

[6]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.

$$\eta_{it} \sim \begin{cases} \mathcal{N}(\mu_1, \sigma_1) & \text{with prob.} \quad p \\ \mathcal{N}(\mu_2, \sigma_2) & \text{with prob.} \quad 1-p \end{cases} \tag{3.7}$$

Durable Consumption Choice and Adjustment Costs

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 $\mathcal{D} = \{d_0, \cdots, d_{n_d}\}$. 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:

$$A(d_{t+1}, d_t) = \begin{cases} 0 & \text{if } d_{t+1} = (1-\delta)d_t \\ \frac{\chi}{2}((1-\delta)d_t + d_{t+1}) & \text{otherwise} \end{cases} \tag{3.8}$$

Notice that equation (3.8) can be rewritten as $A(d_{t+1}, d_t) = \chi(d_t + e_t^d)$, where $e_t^d \equiv d_{t+1} - (1-\delta)d_t$ denotes durable expenditures at age t. This alternative definition clearly exposes the presence of a fixed component, χd_t, independent on the size of the adjustment, making the adjustment costs non-convex.

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:

$$s_t^d = \kappa d_t, \kappa > 1 \tag{3.9}$$

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 c and d, respectively. Both types of goods conform the consumption aggregate following a Cobb-Douglas specification[7]. 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,

$$E_0 \sum_{t=1}^{T} \beta^{t-1} \xi_t \frac{\mathcal{C}\left(c_t, s_t\right)^{1-\gamma}}{1-\gamma} \tag{3.10}$$

$$\mathcal{C}(c_t, s_t) \equiv c_t^{\alpha} (s_t)^{1-\alpha}, \tag{3.11}$$

where $\mathbb{E}_0$ is the expectation operator with respect to the stochastic processes introduced in subsection 3.3.1, 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 downpayment requirement in the case of adjustment:

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

$$a_{t+1} \quad \geq \quad -\lambda^{y} \underline{\mathrm{Y}}_{t} - \lambda^{d} d_{t+1} \tag{3.12}$$

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

Pensions. Income at retirement mimics the US system. Pensions are a function of lifetime average gross earnings[8]. 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:

$$P\left(\bar{Y}_i^R\right) = \begin{cases} 0.9\bar{Y}_i^R & \text{if} \quad \bar{Y}_i^R \leq 0.3\bar{Y} \\ 0.27 + 0.32\left(\bar{Y}_i^R - 0.3\right) & \text{if} \quad 0.3\bar{Y} < \bar{Y}_i^R \leq 2\bar{Y} \\ 0.81 + 0.15\left(\bar{Y}_i^R - 2\right) & \text{if} \quad 2\bar{Y} < \bar{Y}_i^R \leq 4.1\bar{Y} \\ 1.13\bar{Y} & \text{if} \quad 4.1\bar{Y} < \bar{Y}_i^R \end{cases} , \tag{3.13}$$

where

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

[8] 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 $t = 1, \cdots, T_r - 1$

$$\begin{aligned}
V_t(a_t, d_t, z_t) \quad &= \quad \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} \\
s.t. \qquad & c_t + a_{t+1} + d_{t+1} + A(d_t, d_{t+1}) = Y_t + (1+r)a_t + (1-\delta)d_t \\
& Y_t \text{ given by equations (3.3)} - \text{(3.7)} \\
& a_{t+1} \geq -\lambda^y \underline{Y}_t - \lambda^d d_{t+1}, \quad c_t \geq 0
\end{aligned}$$

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

$$\begin{aligned}
V_t(a_t, d_t, z_t) \quad &= \quad \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \xi_t \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} \\
s.t. \qquad & c_t + \frac{\zeta_t}{\zeta_{t+1}} a_{t+1} + d_{t+1} + A(d_t, d_{t+1}) = P(\bar{Y}) + (1+r)a_t + (1-\delta)d_t \\
& P(\bar{Y}) \text{ given by equation (3.13)} \\
& a_{t+1} \geq -\lambda^y \underline{Y}_t - \lambda^d d_{t+1}, \quad c_t \geq 0 \\
& V_{T+1} = 0
\end{aligned}$$

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.2.

3.3.2 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 (3.3)-(3.7) 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 IV. 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.

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 IV: Income Process Estimates

$\boldsymbol{\sigma_\varepsilon}$	(Variance 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
$\boldsymbol{\sigma_1}$	(Variance of 1 persistent shock)	0.075
$\boldsymbol{\sigma_2}$	(Variance of 2 persistent shock)	1.189
$\boldsymbol{\sigma_{z_0}}$	(Variance of initial distribution)	0.753

.

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)[9]. 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 V includes the different values for typically used definitions of nondurable consumption.

[9] 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).

	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 V: 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.

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.

Internally Calibrated Parameters

To calibrate the remaining parameters behind the accumulation of durables and liquid assets, namely the share of durables that can be used as a collateral and the adjustment cost parameter, I do a second SMM to target moments of consumption changes.

Table VI: Internally Calibrated Parameters

		Value	Target	Data	Model
β	(Discount factor)	0.976	Median W/Y	1.35	1.48
λ^d	(Collateralized borrowing)	0.720			
χ	(Adjustment costs parameter)	0.078			

.

3.4 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.

Higher-order moments in the model

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
$var(y_t^{post})$	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 3.3 shows the model counterpart of Figure 3.1. 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 3.3: Deviations from Normality: Durables vs. Nondurables

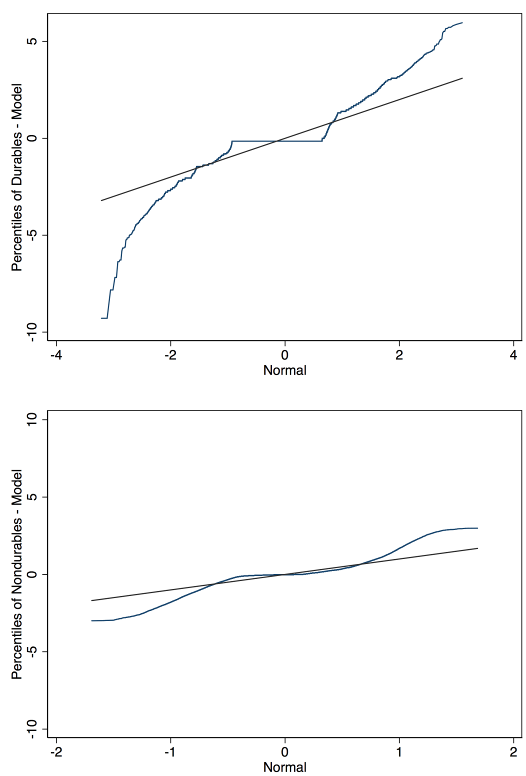

3.4.1 Tail risk and Partial Insurance

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

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

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:

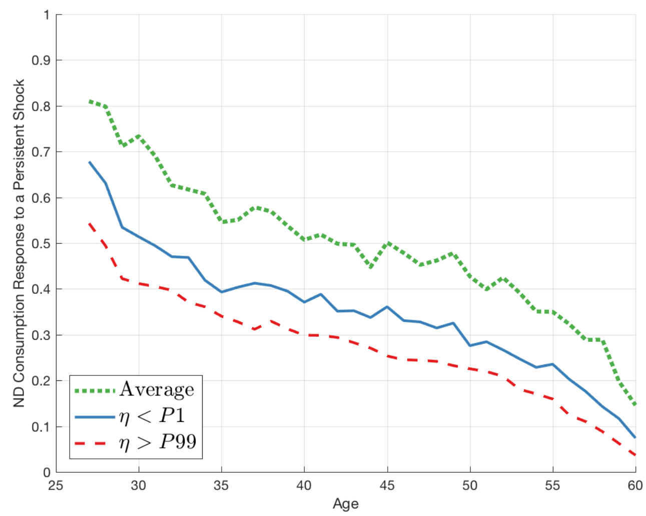

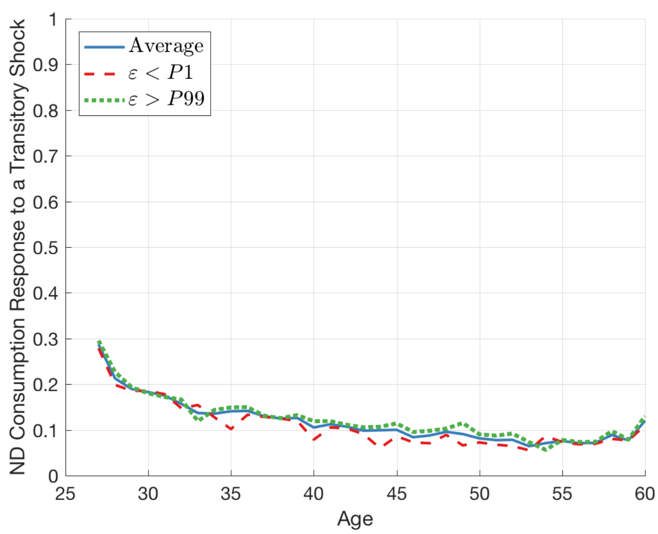

I will comment on the estimates for the persistent component, as the transitory component is highly insurable. It is interesting to notice that the non-Gaussian case exhibits a larger degree of self-insurance than the Gaussian case. Moreover, considering durable adjustments reduces the amount of partial 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 3.4: Response to Income Shocks of Different Size

(a) Nondurables (b) Durables

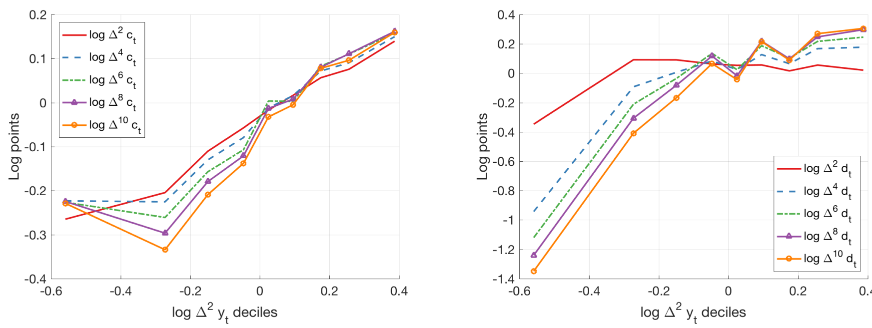

The x-axis shows 20 percentiles of the income changes distribution. The y-axis contains the average response to this size of shock. The different lines correspond to different horizons, from 1 to 10 years ahead.

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.

3.5 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 join 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.

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-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**, 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 Hospido, L. (2013). Earnings inequality in spain: new evidence using tax data. *Applied Economics*, **45** (30), 4212–4225.

— 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.

Busch, C., Guvenen, F., Domeij, D. and Madera, R. (2016). *Asymmetric Business Cycle Risk and Government Insurance*. Working paper.

Card, D., Heining, J. and Kline, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics*, **128** (3), 967–1015.

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.

Constantinides, G. M. and Ghosh, A. (2014). *Asset Pricing with Countercyclical Household Consumption Risk*. Working paper, University of Chicago.

Daly, M., Hryshko, D. and Manovskii, I. (2014). Reconciling estimates of earnings processes in growth rates and levels, manuscript.

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.

Dustmann, C., Ludsteck, J. and Schönberg, U. (2009). Revisiting the german wage structure. *Quarterly Journal of Economics*, **124** (2), 843–881.

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

Fitzenberger, B., Osikominu, A. and Völter, R. (2006). Imputation rules to improve the education variable in the IAB employment subsample. *Schmollers Jahrbuch. Zeitschrift für Wirtschafts- und Sozialwissenschaften*, **Jg. 126**.

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. (2016). *What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?* Working paper.

—, Ozkan, S. and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, **122** (3), 621–660.

— 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, 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.

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

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

Mankiw, N. G. (1986). The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics*, **17** (1), 211–219.

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.

Skedinger, P. (2007). *The Design and Effects of Collectively Agreed Minimum*

Wages: Evidence from Sweden. Working Paper Series 700, Research Institute of Industrial Economics.

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

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.

vom Berge, P., Burghardt, A. and Trenkle, S. (2013). *Sample-of-Integrated-Labour-Market-Biographies Regional-File 1975-2010 (SIAB-R 7510)*. Tech. rep., Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB).

Appendix A

Appendix to Chapter 1

A.1 Data Appendix

This appendix briefly describes the variables used for each of the data sets and lists the numbers of observations after the sample selection steps.

A.1.1 PSID

Variables

Demographic and Socioeconomic

Head and Relationship to Head. We 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.

Age. The age variable recorded in the PSID survey does not necessarily increase by 1 from one year to the next. This may be perfectly correct, since the survey date changes every year. For example, an individual can report being 20 years old in 1990, 20 in 1991, and 22 in 1992. We thus create a consistent age variable by taking the age reported in the first year that the individual appears in the survey and add 1 to this variable in each subsequent year.

Education Level. In the PSID, the education variable is not reported every year and it is sometimes inconsistent. To deal with this problem, we use the highest education level that an individual ever reports as the education variable for each year. Since our sample contains only individuals that are at least 25 years old, this procedure does not affect our education variable in a major way.

Income

Individual Male Wages and Salaries. This is the variable used for individual income in the benchmark case. It is the answer to: *How much did (Head) earn altogether from wages or salaries in year t-1, that is, before anything was deducted for taxes or other things?* This is the most consistent earnings variable over time reported in the PSID, as it has not suffered any redefinitions or change in subcomponents[1].

Individual Male Labor Earnings. `Annual Total Labor Income` includes all income from wages and salaries, commissions, bonuses, overtime and the labor part of self-employment (farm and business income). Self-employment in PSID is split into asset and labor parts using a 50-50 rule in most cases. Because this last component has been inconsistent over time[2], we subtract the labor part of business and farm income before 1993.

Individual Female Labor Earnings. There is no corresponding `Wages and Salaries` variable for spouses. We use `Wife Total Labor Income` and follow a similar procedure as in the case of heads.

[1]See Shin and Solon (2011) for a comparison of PSID male earnings variables in inequality analyses.

[2]In particular, total labor earnings included the labor parts of farm and business income up to the 1993 survey but not in subsequent waves.

Annual Hours. For heads and wives, it is defined as the 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.

Pre-Government Household Labor Earnings. Head and wife labor earnings.

Post-Government Household Labor Earnings. Pre-government household earnings *minus* taxes *plus* public transfers, as defined below.

Taxes. The PSID reports own estimates for total taxes until 1991. For the remaining years, we estimate taxes using TAXSIM.

Public Transfers. Transfers are considered at the family unit level, when possible. We group social and welfare programs in three broad categories. Due to changes in the PSID design, the specific definition of each program is different every year. We give an overview below and leave the specific replication details for the online Data Appendix.

Household Disposable Income. We construct this variable from Household Taxable Income (Head's and wife's income from assets, earnings, and net profit from self-employment) *minus* taxes *plus* public transfers.

Transfers

We refer to Table V in the main text for a description of the three groups of programs considered, as well as their subcomponents. In the PSID, obtaining an annual amount of each type of benefits is almost wave-specific. Every few survey years, the level of aggregation within the family unit and across welfare programs is different for at least one of our groups. To impose some common structure, we establish the following rules.

For survey years 1970-1993[3] and 2005-2011, the total annual amount of each program is reported for the head, spouse and others in the family unit. In occasions, the amount appears combined for several or all members.[4] Because in those cases it is impossible to identify separate recipiency of each member, we consider the benefit amount of the whole family. This is, we add up all available information for all family members, whether combined or separately reported.

In survey years 1994-2003, most benefits (except Food Stamps and OASDI) are reported separately for the head and the spouse only. The way amounts are reported changes as well. First, the reported amount ($\$X$) received is asked. Second, the frequency of that amount ($\$X$ per year, per month, per week, etc) is specified. We convert all amounts to a common frequency by constructing a monthly amount $\$x$ using these time values. Finally, the head and spouse are asked during which months the benefit was received. The final annual recipiency of transfers is then obtained multiplying $\$x$ by the number of months this benefit was received. For Food Stamps and OASDI, we follow the rules described for the other waves.

Detailed Sample Selection

We start with an initial sample of 584,392 SRC individuals interviewed between 1976 and 2011. 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, we have cleaned the raw data and corrected duplicates and inconsistencies (for example, zero working hours with positive labor income). We also require that the individuals have non top-coded observations in income.

[3]Our main sample refers to survey years 1977-2011, but complementary results are provided for the annual subsample of the PSID. This is, for 1970-1997. We drop the first two waves in all cases, since benefits such as OASDI, UI and WC are only reported for the family head; and benefits such as SSI are not reported at all.

[4]This is always the case for Food Stamps.

1. The individual must be from the original main PSID sample (not from the Survey of Economic Opportunities or Latino subsamples).

2. In the benchmark individual sample, we select male heads of family. In the reference household sample, we require at least two adult members in the unit and that individuals had no significant changes in family composition. More specifically, we require that they responded either "no change" or "change in family members other than the head or wife" to the question about family composition changes.

3. The household must not have missing variables for the head or wife labor income, or for education of the head. The individuals must not have missing income or education themselves.

4. The individual must not have income observations that are outliers. An outlier is defined as being in the top 1% of the corresponding year.

5. We require the income variable of analysis to be positive.

6. Household heads must be between 25 and 65 years old.

Table A.1: Number of Observations Kept in Each Step

	Male Heads	Households	All Females
SRC	586,187	586,187	586,187
Family Composition	90,106	75,202	110,711
Non-Missing y or College	83,039	69,443	97,990
Positive Income	63,875	58,551	54,214
Outliers	63,065	57,262	53,257
Age Selection	54,593	50,102	45,330
Final #Obs for transitory changes	42,623	38,171	33,687
Final #Obs for persistent changes	34,985	30,985	27,269

Note: Table lists number of person-year, or household-year, observations in the three panels for the sample from PSID.

A.1.2 LINDA

Variables

Demographic and Socioeconomic

Head and Relationship to Head. LINDA is compiled from the Income Register based on filed tax reports and other registers. Statistics Sweden samples individuals and then adds information for all family members, where family is defined for tax purposes. This implies that there is no information about 'head of households'. We therefore define the head of a household as the sampled male.

Age. As defined by Statistics Sweden

Education Level. LINDA contains information about eduction from 1991 and onwards. An individual is assigned "college" education if it has at least 3 years of university education.

agraphPrivate / Public employment

An individual is defined as as working in the public sector, if he/she works in public administration, health care or education. Linda contains consistent comparable information for the years 1991 and onwards. For the years 1991-92 the public sector employment is defined as those we use SNI90 codes 72000-72003, 90000-93999 and >=96000. For 1993-2006 we use SNI92 codes 64110-64202, 73000-74110, 75000-92000, 92500-92530 and >=96000. For 2007 we use SNI2007 codes 64110-64202, 73000-74110, 75000-92000, 92500-92530 and >=96000.

Income

For the years 1985-2010 we use the measures suggested by Statistics Sweden to be comparable between years in LINDA. We construct comparable measure for the years 1979-1984.

Individual labor earnings. Labor earnings consist of wages and salaries, the part of business income reported as labor income, and taxable compensation for sick leave and parental leave.

Pre-Government Household Labor Earnings. Defined as the sum of individual labor income within the family.

Post-Government Household Labor Earnings. Post-government earnings is calculated as pre-government earnings *minus* taxes *plus* public transfers.

Household Disposable Income. Disposable income consists of the sum of factor income and minus taxes and plus public transfers.

Taxes. LINDA provides observations of total taxes paid by the individual. Since taxed paid on capital income constitute a small part of total tax payments, and since we cannot separate taxes on capital income from those on labor income, we assume that all taxes are labor income taxes.

Public Transfers. LINDA provides observations of total public transfers at the individual level (Statistics Sweden has individualized transfers given to families) and at the household level. We also consider three subcategories of transfer as listed below.

Transfers

Transfers in subcategory 1 and 3 are individual level transfers. Transfers in subcategory 2 are family level transfers but have been individualized by Statistics Sweden. For each subcategory, we take all transfers received by all members of the households.

- *HH-level transfers subcategory 1 (labor market transfers)*: sum of unemployment benefits received by all members of household.

- *HH-level transfers subcategory 2 (family aid)*: sum of transfers to support families received by all members of household.
- *HH-level transfers subcategory 3 (pensions)*: sum of old-age pensions received by all members of household.

Detailed Sample Selection

To be included in the individual sample the individual has to be sampled and between 25 and 60 years old. A family is included in the household sample if the sampled individual is a man between 25 and 60 years old and there are at least two members aged 25-60 in the family.

A.1.3 SIAB

We use the scientific use file SIAB-R7510 provided by the Institute for Employment Research (IAB). The SIAB data from which the scientific use file is constructed is a 2% random sample of all individuals covered by a dataset called IEB. This data set is from four different sources, which can be identified in the data. For construction of our sample we use earnings data stemming from BeH (employee history) and transfer data from LeH (benefit recipient history). Records in BeH are based on mandatory social security notifications from employers and hence cover individuals working in employment subject to social security, which excludes civil servants, students and self-employed. A new spell starts whenever there is a new notification, which happens when either a new employment relationship changes, an ongoing contract is changed, or with the start of a calendar year. BeH covers all workers subject to social security contributions, which excludes civil servants, self-employed and students. For details on the data set see vom Berge *et al.* (2013).

Variables

Demographic and Socioeconomic

Head and Relationship to Head. SIAB does not contain information on households. We use only individual level data.

Age. Birth year is reported consistently in SIAB data.

Education Level. Each individual spell in SIAB contains information on the highest degree of formal education as reported by the employer. In order to construct a consistent measure of education we apply imputation rules proposed by Fitzenberger *et al.* (2006).

agraphPrivate / Public employment

An individual is defined as as working in the public sector, if he/she works in public administration, health care or education. SIAB contains consistent comparable information for all years of the sample. We use the classification WZ93 as provided in the data, which aggregates 3 digit codes of the original WZ93 classification into 14 categories. The industry of an employer is registered once a year and assigned to the worker spells of that year. This implies that for some individual spells there is no information on the industry. For each year a worker is assigned the industry from the longest spell in that year. We classify as public employment those in sectors 13 (3-digit WZ93 801-804, 851-853: Education, social and health-care facilities) and 14 (751-753, 990: public administration, social security).

Income

Individual labor earnings. We calculate annual earnings as the sum of total earning from all valid spells for each individual. As marginal employment spells were not reported before 1999, we drop marginal employment in the years where they are reported to obtain a time consistent measure. For the same reason we drop spells with reported average daily wage rate below the highest marginal employment threshold in the sample period, which is 14.15 Euros (in 2003 Euros). There are two drawbacks in the available data: structural break of the wage measure in 1984 and top-coding.

Structural break in wage measure Since 1984 the reported average daily wage rate from an employment spell includes one-time payments. We correct for this structural break following a procedure based on Dustmann *et al.* (2009): we rank individuals from 1976 to 1983 into 50 quintiles of the annual full-time wage distributions. Then we fit locally weighted regressions of the wage growth rate from 1982-1983 on the quintiles in 1983 and the same for 1983-1984. We then define as the correction factor the difference between the quintile-specific smoothed value of wage growth between 1984 and 1983. The underlying assumption is that wage growth should be higher from 1983-1984 because the wage measure includes one-time payments. In order to control for overall wage growth differences we subtract the average of the correction factor of the second to 20th quintiles. The resulting percentile-specific correction factor is then applied to wages in 1976-1983.

Imputation of top-coded wages Before aggregating earnings from all spells we correct full-time wage spells for the top-coding. We therefore follow Daly *et al.* (2014) and fit a Pareto tail to the cross-sectional wage distribution. The Pareto distribution is estimated separately for each year by age-group and sex. We define seven age groups: 25-29,30-34,...,55-60. As starting point for the Pareto we choose the 60th percentile of the subgroup-specific distribution. As in Daly *et al.* (2014), we draw one random number by individual which we then apply to the annual specific distributions when assigning a wage to the top-coded workers. We apply the imputation method to the annual distribution of average full-time wages and hence an individual can be below the cutoff limit if, e.g., from two full-time spells in a year only one is top-coded. We therefore define as top-coding limit the annual specific limit minus 3 DM (1995 DM) as in Dustmann *et al.* (2009).

Transfers

In SIAB we observe consistently over time unemployment benefits at the individual level.

Detailed Sample Selection

To be included in the sample the individual has to be between 25 and 60 years old and earn a gross income above 520*0.5*minimum wage. We drop all workers which have at least one spell reported in East Germany.

SOEP

Variables

Demographic and Socioeconomic

Head and Relationship to Head. For each individual in the sample, SOEP reports the relationship to the head of household in any given wave. Whenever there is a non-couple household, i.e., no spouse is reported, the reported head is classified as head. Whenever we observe a couple household and the reported head is a male we keep this; when the reported head is a female and the reported spouse is a male, we reclassify the male to be head and the female to be spouse.

Age. The age is measured by subtracting year of birth from the current year.

Education Level. The education variable used categorizes the obtained maximum education level by ISCED 1997. An individual with category 6 is assigned "college" education, an individual with categories 1-5 is assigned "non-college". Category 6 includes a degree obtained from university, from technical college, from a university abroad, and a PhD. An individual still in school (category 0) is assigned a missing. For a small number of individuals the described procedure yields inconsistencies in the sense that for some year t the assignment is "college" and some later year t+s the assignment is "non-college"; in these cases we assign "college" to the later year.

Income and Hours

Individual labor income. Labor earnings are calculated from individual labor income components and includes income from first job, secondary job, 13th and 14th salary, christmas bonus, holiday bonus, profit sharing. For consistency with the PSID measure we assign 50% of income from self-employment to labor income.

Household level labor income. Defined as the sum of individual labor income of head and spouse.

Annual Hours. SOEP measures the average actual weekly hours worked and the numbers of months an individual worked. From these measures SOEP provides a constructed measure of annual hours worked of an individual.

Pre-Government Household Labor Earnings. Head and spouse labor earnings.

Post-Government Household Labor Earnings. Pre-government household earnings *minus* taxes *plus* public transfers, as defined below.

Taxes. SOEP provides estimates of total taxes at the household level.

Public Transfers. Transfers are considered at the family unit level and at the individual level. We group social and welfare programs in three broad categories as listed below.

Household Disposable Income. We construct this variable from Household Taxable Income (Head's and wife's income from assets, earnings, and net profit from self-employment) *minus* taxes *plus* public transfers. SOEP provides a measure of household asset flows, which is calculated as income from renting minus operating costs, plus dividend income.

Transfers

Transfers are partly observed at the individual level and partly at the household level. For each subcategory, we take all transfers received by all members of the households.

- *HH-level transfers*: we use transfers received by all individual household members in order to calculate measures that are consistent over time. For each individual, total transfers are the sum of the following components: old-age pensions, widow's pensions, maternity benefit, student grants, unemployment benefits, subsistence allowance, unemployment assistance (up to 2004); at the hh-level we measure received child allowances and the total unemployment benefits II received by all household members (since 2005 replacing unemployment assistance).

- *HH-level transfers subcategory 1 (labor market transfers)*: sum of unemployment benefits received by all members of household.

- *HH-level transfers subcategory 2 (family aid)*: sum of subsistence allowance of all members, + sum of unemployment assistance received by all members (up to 2004), + hh-level measure of unemployment benefits II (since 2005).

- *HH-level transfers subcategory 3 (pensions)*: sum of old-age pensions received by all members of household.

Sample Selection

In order to be in the initial sample for a year, the individual or household head must be between ages 25 and 60 and live in West Germany. In order to have a consistent sample, we drop the immigrant subsample and the high income subsample. This gives initial sample sizes of 87,582 individual-year observations for the male sample, 76,249 individual-year observations for the female sample, and 76,051 household-year observations for the household sample. The sample selection then follows the steps listed below for each sample. All cross-sectional statistics are calculated using appropriate cross-sectional individual or household weights, respectively.

1. drop if no info on education or if no degree obtained yet
2. drop if currently working in military
3. drop if no info on income
4. drop if no info on hours worked
5. keep if income $>$ 0 and hours $>$ 520
6. drop if in highest percentile (sample outliers)
7. drop if below 520*0.5**minimum wage*, where *minimum wage* is set to be 6€ in year 2000 Euros
8. for transitory change measure: keep if in sample in t and t-1
9. for permanent change measure: keep if in sample in t and t-5

Table A.2: Number of observations in the three panels after each selection step

selection step	Male Heads	Households	All Females
initial	87,582	76,051	76,249
drop if no coll. info	86,737	75,310	75,270
drop if in military	86,712	75,293	75,268
drop if no obs on ymin	79,547	75,070	50,374
drop if no obs on hours	79,547	75,070	50,374
keep if >=520 hrs and ymin>0	77,265	71,389	42,245
drop top 1% of ymin per year	76,404	70,627	41,830
drop if ymin<.5*520*min wage	76,268	70,097	41,434
Final #Obs for transitory changes	64,572	59,209	31,612
Final #Obs for persistent changes	38,399	34,792	16,792

Note: Table lists number of person-year, or household-year, observations in the three panels for the sample from SOEP.

Appendix B

Appendix to Chapter 3

B.1 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.

B.1.1 The PSID

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.

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 Busch *et al.* (2016), 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)[1].

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

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

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.

B.1.2 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 $\widetilde{Y}$	22373
Not enough consequtive obs	20954
Final # Observations	20954
Final # Households	4697

Table A.1: Detailed Sample Selection

B.1.3 Detailed Summary Statistics and Extra Moments

Table A.2: 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: non-durable consumption, d: durable consumption. See Section 3.2.1 for detailed definitions.

Table A.3: 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					

Comparison with SSA Data for Male Earnings

Figure A.1: U.S. Males Annual Earnings 1-Year Log-Changes: SSA vs. PSID

(a) Density (b) Log-Density

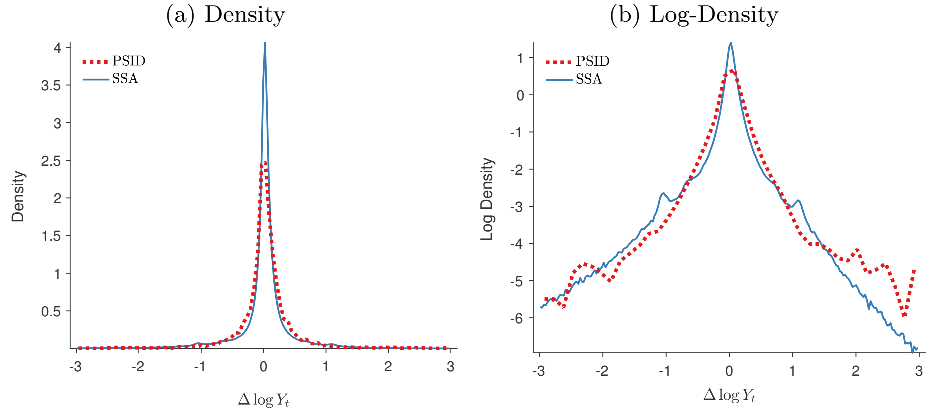

Figure A.2: U.S. Males Annual Earnings 5-Year Log-Changes: SSA vs. PSID

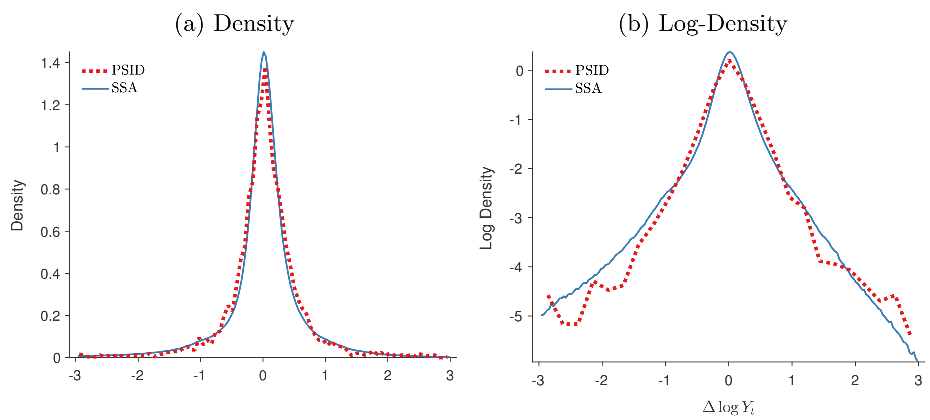

Understanding QQ Plots

Figure A.3: Understanding QQ Plots

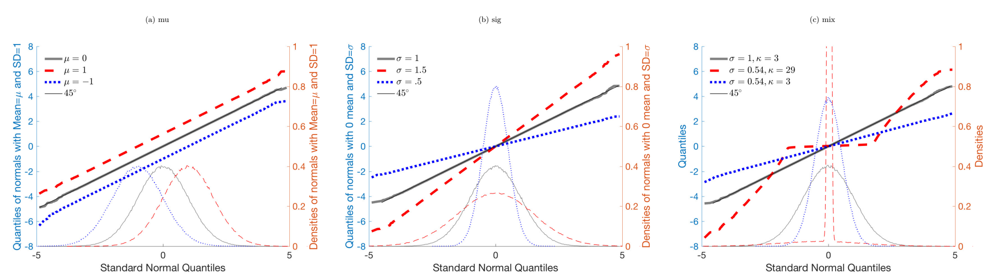

Log-Densities

Graphs are truncated at +/- 3.

2-Year Changes

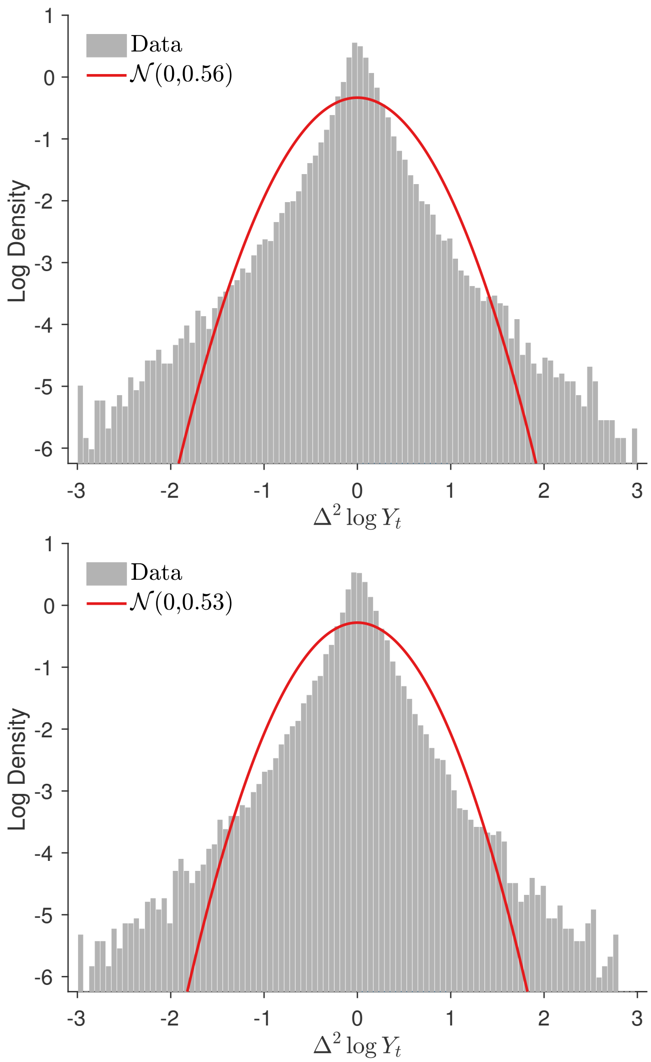

Figure A.4: Household Income (Pre-Gov and Post-Gov)

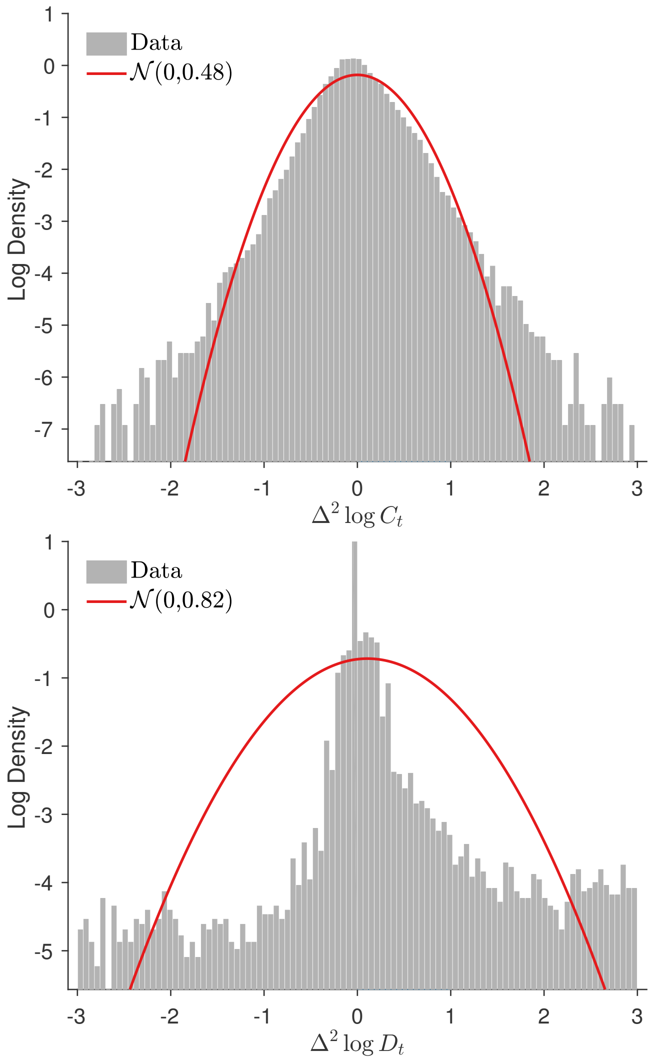

Figure A.5: Consumption (Nondurable and Durable)

4-Year Changes

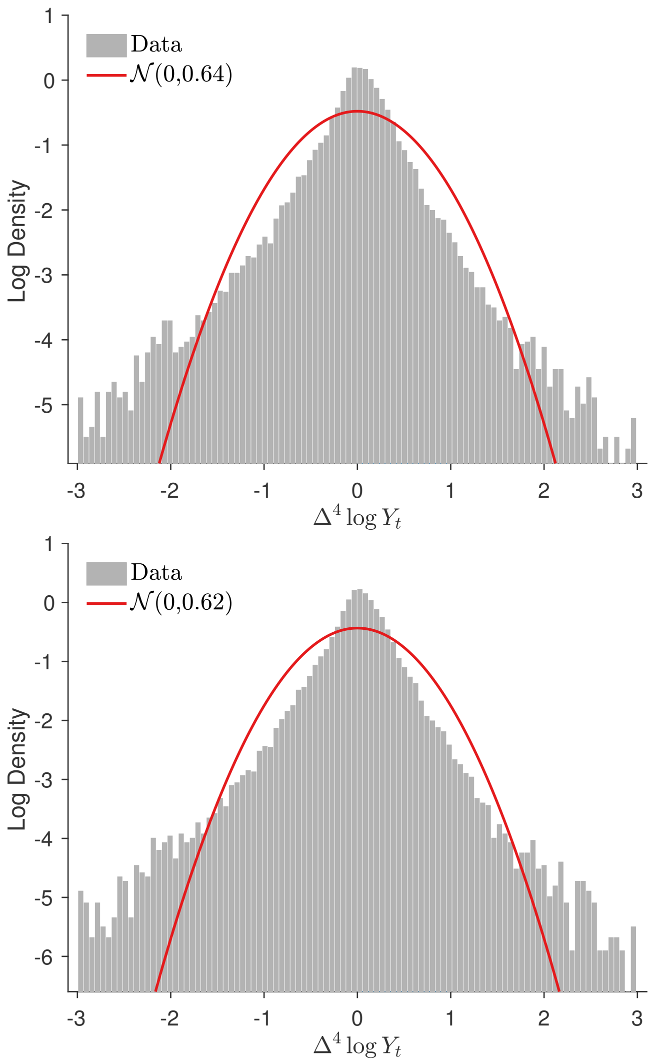

Figure A.6: Household Income (Pre-Gov and Post-Gov)

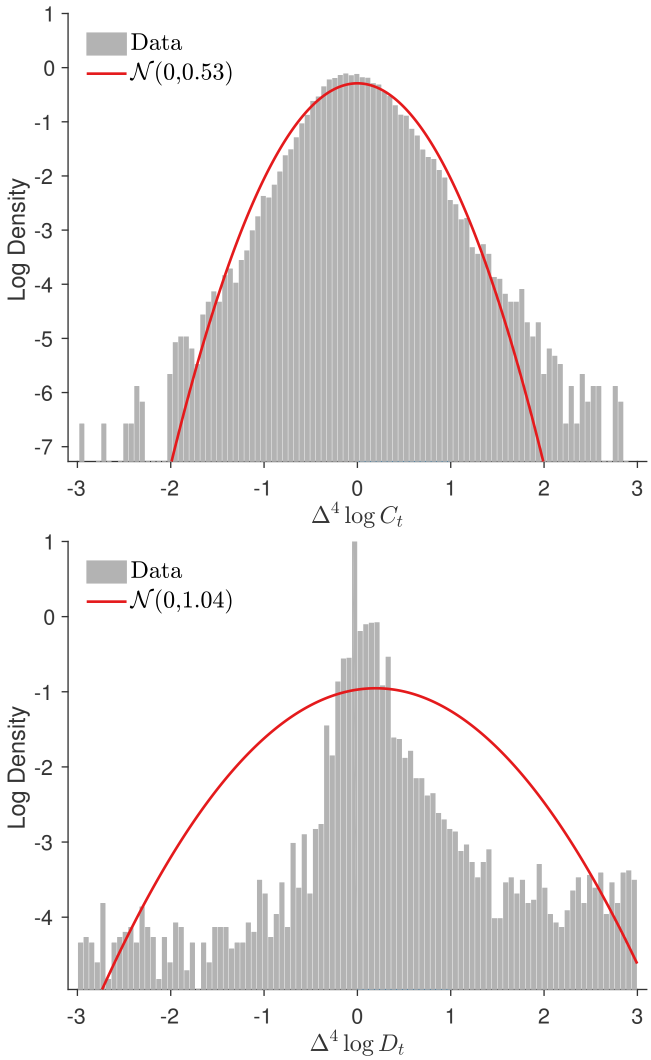

Figure A.7: Consumption (Nondurable and Durable)

B.2 Numerical Appendix

B.2.1 Solution Details

Optimality Conditions (Euler Equation)

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

$$V_t(a_t, d_t, z_t) = \max_{c_t, d_{t+1}, a_{t+1}} \{u(c_t, s_t) + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, d_{t+1}; z_t)\} \tag{B.1}$$

$$s.t. \quad c_t + a_{t+1} = Y_t + (1+r)a_t + (1-\delta)d_t - d_{t+1} - A(d_t, d_{t+1}) \tag{B.2}$$

$$Y_t \text{ given by equations (3.3) – (3.7)} \tag{B.3}$$

$$a_{t+1} \geq -\lambda^y \underline{y}_{t+1} - \lambda^d d_{t+1}, \quad c_t \geq 0, d_{t+1} \in D \tag{B.4}$$

To facilitate the solution of (B.1), I implement a set of transformations described in the next paragraphs.

First, let

$$b_{t+1} \equiv a_{t+1} + \lambda^y \underline{y}_{t+1} + \lambda^d d_{t+1}. \tag{B.5}$$

Then, the borrowing constraint can be rewritten as

$$c_t + (b_{t+1} - \lambda^y \underline{y}_{t+1} - \lambda^d d_{t+1}) = Y_t + (1+r)(b_t - \lambda^y \underline{y}_t - \lambda^d d_t) + (1-\delta)d_t - d_{t+1} - A(d_t, d_{t+1})$$
$$c_t + b_{t+1} = \underbrace{Y_t + (1+r)b_t + (1-\delta)d_t - (1-\lambda^d)d_{t+1} - A(d_t, d_{t+1}) - \lambda^d(1+r)d_t + \lambda^y(\underline{y}_{t+1} - (1+r)\underline{y}_t)}_{m(d_{t+1})} \tag{B.6}$$

Notice that (B.6) 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

$$\begin{aligned} \tilde{V}_t(b_t, d_t, z_t) &= \max_{d_{t+1}, b_{t+1}} \left\{ \mathcal{C}(m(d_{t+1}) - b_{t+1}, s_t)^{1-\gamma} + \beta \mathbb{E}_t \tilde{V}_{t+1}^{1-\gamma}(b_{t+1}, d_{t+1}; z_t) \right\}^{\frac{1}{1-\gamma}} \quad \text{(B.7)} \\ s.t. \quad & (b_{t+1}, d_{t+1}) \in \{b_{t+1}, d_{t+1} : b_{t+1} \in [0, m(d_{t+1}; b_t, a_t, z_t)], d_{t+1} \in D\}. \end{aligned}$$

The FOC therefore are given by

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

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

And the envelope condition

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

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

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

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

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:

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

B.2.2 Numerical Details 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 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