Cyclical Dynamics in Idiosyncratic Consumption Risk

by

Allison Cole

B.A. Boston University (2013)
M.A. Boston University (2013)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Master of Science in Management Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2020

Author . . . .
Sloan School of Management
May 8, 2020

Certified by . . . .
Jonathan A. Parker
Robert C. Merton (1970) Professor of Finance
Thesis Supervisor

Accepted by . . . .
Catherine Tucker
Sloan Distinguished Professor of Management
Professor, Marketing
Faculty Chair, MIT Sloan PhD Program

Cyclical Dynamics in Idiosyncratic Consumption Risk

by

Allison Cole

Submitted to the Sloan School of Management
on May 8, 2020, in partial fulfillment of the
requirements for the degree of
Master of Science in Management Research

Abstract

This paper examines cyclical dynamics of idiosyncratic consumption risk using consumption data from the Nielsen Consumer Panel and the Panel Study of Dynamic Income. With GMM estimates and supplemental graphical analysis, I show that the idiosyncratic risk in consumption is i) highly persistent, with an autocorrelation coefficient near unity ii) strongly countercyclical, with the conditional variance rising by an average of 25 percent from peak to trough. Compared to previous findings on income dynamics, I show that the variance of idiosyncratic consumption risk is also countercyclical, but less so. Moreover, I do not find that consumption risk displays procyclical skewness, as has been shown with income risk. Furthermore, in a simple asset-pricing framework, the estimated countercyclical cross-sectional variance of consumption raises the equity premium by 4.1 percent from the representative-agent case, using a risk aversion of only 10-15.

Thesis Supervisor: Jonathan A. Parker
Title: Robert C. Merton (1970) Professor of Finance

Acknowledgments

I am grateful to my adviser, Jonathan Parker, for invaluable guidance and support throughout this project. I also thank Hui Chen, Dan Greenwald, Eben Lazarus, Debbie Lucas, Christopher Palmer, Larry Schmidt, and Adrien Verdelhan for helpful comments and suggestions. Pierre Jaffard, Lu Liu, Yura Olshanskiy, Pari Sastry, and participants at the Sloan Finance Workshop also provided helpful discussions. This paper presents the researcher's own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. All errors are my own.

I Introduction

Consumption outcomes are the main object of interest for the design of tax and transfer policies, benefits from insurance, and the welfare cost of business cycles. Moreover, consumption risk plays a key role in the pricing of assets. Several facts have been documented regarding the *income* process, with various (and sometimes conflicting) conclusions. Namely, that there appear to be differences in both the variance and the skewness of the distribution of income outcomes during economic contractions versus expansions. These findings verify the very natural conclusion that during a recession, the probability of a large income drop increases, while the probability of a large income increase does not. What is not yet well understood however, is how these income dynamics translate to consumption.

The primary reason why income has been more commonly studied than consumption is that better administrative data is available for income. However, the ultimate object of interest is consumption. Consumers have utility over *consumption*, not income; we price assets using *consumption* risk, not income risk. While we can draw some implications about consumption using income, in a world where not all consumers are hand-to-mouth, the conclusions are incomplete. In this paper, I circumvent this issue by bringing the methods used to study income to a rich source of consumption data. I use the Nielsen Consumer Panel (CP) dataset, which tracks consumption expenditures at the household level. This dataset has several advantages over other measures of household consumption. First, it is much larger than most other commonly used consumption datasets such as the Panel Study of Dynamic Income (PSID) and the Consumer Expenditure Survey (CEX), with about 40,000-60,000 observations per year. In addition, the CP is not subject to the same measurement error as survey based measures because its relies on scanner data, rather than participant recall. To corroborate the results found in the CP, I repeat all analyses on the PSID and find similar results.

My estimation results suggest that the variance of the persistent shock is strongly counteryclical. On average I find that the standard deviation rises by 25 percent from expansion to contraction. I also find that idiosyncratic consumption risk is highly persistent, with an annual autocorrelation coefficient close to unity. Moreover, I find no evidence that the the skewness of the persistent shock changes based on economic conditions, contrasting the results found for income, such as Guvenen, Ozkan and Song (2014). These results hold across both datasets, the CP and the PSID, as well as subsamples of the population that are more likely to be stockholders.

I begin the analysis by extending the parametric framework of Storesletten, Telmer and Yaron (2004b) (henceforth, STY). The key ingredient of this model is that it allows the variance of a persistent shock to income to vary based on the aggregate state of the economy. Following an extension developed by Busch and Ludwig (2018), I also assume that the skewness of the persistent shock is dependent on the aggregate state. The variance and skewness can be solved for in closed form and then estimated using the same Generalized Method of Moments (GMM) estimator developed by STY. Moreover, I extend the framework in a way which allows me to control for inequality, hence the trend increase in inequality over time does not confound the results on cyclicality. Both graphically and with the GMM results, I show that the volatility of the persistent shock increases significantly as the economy moves from expansion to contraction. However, there is insufficient power and precision in my sample to show that the skewness changes. This confirms the original finding of countercyclicality of STY, though I find a much smaller magnitude of countercyclicality in consumption than they did for income (a 25 percent increase from peak to trough versus a 75 percent increase). My dataset results in more precise estimation of the model's parameters and higher power in failing to reject the model.

The basic identification strategy relies on the fact that the persistent shock accumulates over the life-cycle such that the distribution of consumption observed for a given cohort widens as the cohort ages. The key is to take advantage of both year and cohort variation in consumption risk. As an example, suppose that we wish to examine the consumption outcomes of two groups of individuals aged 35. The first group was born in 1955 and the second was born in 1970. Both are subject to idiosyncratic consumption shocks, the variance of which is countercyclical. Hence, the former group should exhibit a larger within-group variance of consumption (at age 35, after conditioning on observables) because those born in 1955 have experienced more years of economic contractions than those born in 1975 (ie, there were more contractions from 1980-1990 than from 1995-2005). If consumption risk is systematically different between expansions and contractions, then the cohort variation will identify these differences. The same logic applies if we assume that the skewness varies across states as well.

As a next step, I use the modeling framework to derive a simple asset pricing relation. By calculating the implied growth rate under the specified consumption process, I derive the economy's stochastic discount factor. Then, I show that the risk premium is a function of the return's covariance with aggregate growth and the return's covariance with the variance of idiosyncratic components of consumption. As a result, this idiosyncratic risk amplifies the risk premium when its variance covaries negatively with the return. This is exactly what we would expect if the variance is in fact counterycyclical. Indeed, I observe this in the data: on average, the correlation of the real return with the model-implied variance of the idiosyncratic component of consumption is -.41. Combined with a risk aversion of only 12, this increases the equity premium by 4.1 percent during the post-war period. Over my sample period, the importance of this finding is evident during the financial crisis: as stock returns fell sharply, the variance of idiosyncratic risk rose sharply, thus producing a high required expected return during the financial crisis.

The contribution of this paper is three-fold. First, I document previously unknown facts about the cyclicality of the consumption distribution in the CP. This dataset has been widely used in many fields of research from Industrial Organization to Household Finance, hence understanding more precisely trends in the data and how it compares to income is valuable. Second, I highlight several important ways in which the consumption risk appears to differ from income risk. Lastly, I demonstrate how idiosyncratic risk can amplify the equity premium in a simple and commonly used framework for modeling shocks.

I.1 Related Literature

As mentioned above, the cyclicality of labor income risk has been extensively studied. STY conclude that idiosyncratic risk in labor income, is i) highly persistent, with an annual autocorrelation coefficient of approximately .95 ii) highly countercyclical, with a conditional standard deviation that increases by approximately 75 percent when the macroeconomy moves from contraction to expansion. Guvenen, Ozkan and Song (2014) refute this result, finding that the variance is not counteryclical but the the skewness is strongly procyclical. Busch and Ludwig (2019 WP) find that the variance is countercyclical and the skewness is procyclical. Krebs (2007) shows that job displacement risk (and hence income risk) has a large welfare cost. Schmidt (2016) shows that idiosyncratic labor risk in the tail is key driver of asset prices.

Many papers have also studied consumption risk. In a separate paper, Storesletten, Telmer, and Yaron (2004a) use PSID and Consumer Expenditure Survey (CEX) data to show that shocks to both income and consumption are highly persistent over the lifecycle. Gourinchas and Parker (2002) also document strong lifecycle dynamics in consumption using the CEX. Attanasio and Davis (1996) show the failure of consumption insurance and that there is large variation across groups based on education. Parker (2001) finds that medium-term consumption risk can in part explain the equity premium. Cogley (2002) constructs an equilibrium factor model based on consumption data from the CEX. While he documents the cross-sectional variation in consumption risk, the model finds that it does poorly in explaining the equity premium. Constantinides and Ghosh (2017) show that shocks to household consumption growth are negatively skewed, persistent, countercyclical, and drive asset prices. Despite the vast literature, an analogous study to STY or Guvenen, Ozkan and Song (2014) using complete consumption data, or any study of idiosyncratic consumption risk using the CP, has yet to be completed to the best of my knowledge[1].

This paper also contributes to the literature on consumption based asset pricing. Lucas (1994) shows that most idiosyncratic risk can be smoothed away via securities transactions. However, this assumes that all individuals are participating in such markets. Constantinides and Duffie (1996) suggest a model with labor income shocks and show that idiosyncratic risk lowers the risk free rate due to the precautionary saving effect. Heaton and Lucas (1996) study a model of idiosyncratic labor shocks with transactions costs and find that the direct effect of transactions costs on the ability to self-insure against these shocks can explain a significant proportion of the equity premium. Guvenen (2009) provides a model with both idiosyncratic labor-income risk and limited participation, showing that stockholders must be highly compensated for the risk they take when insuring against labor income shocks, which can explain the large equity premium. Duffee (2005) shows that consumption risk is priced and that consumption growth is closely tied to stock returns. Di Maggio, Kermani and Majlesi (2018) use administrative data from Sweden to document that household consumption responds strongly to stock returns, particularly at the lower end of the income distribution.

II Data

II.1 Nielsen Consumer Panel Data

The data underlying my estimation are from the Nielsen Consumer Panel (henceforth, CP). The CP is a longitudinal panel of approximately 40,000-60,000 households per year in the United States, members of which continually provide information to Nielsen about their household characteristics and purchases. Nielsen samples 48 states and their major markets and panelists are geographically dispersed and demograph-

[1]STY do replicate their analysis on *food* consumption, which is available in the PSID. They find similar, but less strong, results to what they find for income.

ically balanced. The data are available from 2004-2017 and the panel has steadily grown in size since 2004. The panel is unbalanced with some panelists staying in the sample for several years and others joining or dropping out in any given year. The CP is rich in demographic information including household size, age, education, race, occupation, income, and location. Demographic data on household members are collected at the time of entry into the panel, and are then updated annually through a written survey that takes place during the fourth quarter. Nielsen provides sample weights to make the sample representative of the U.S. population. I apply these weights in all of the analyses.

To record purchases, the Homescan panelists are given in-home scanners. Purchases include food and non-food items. The purchase information in the data includes the date of the trip, the retailer, the retailer's zip code, and total dollars spent. To get total consumption, I simply sum up the amount spent on all trips in a given household for the entire year. Note that because the CP only measures purchased products at the stores included by Nielsen, this measure of consumption does not account for services such as housing and utilities, health care, financial services, and most durable goods. This point bears more discussion, and I will return to it momentarily.

Nielsen offers the panelists a variety of incentives to join and remain in the panel. Some examples are monthly prize drawings, gift cards, and sweepstakes. Nielsen filters out households based on purchasing thresholds in order to remove households that are not active and may bias the results. My analysis includes only actively reporting households who have a least one trip per month over each year.[2]

There are several issues that arise when working with longitudinal data. Because the panel is unbalanced, limiting the analysis to only those who remain in the panel for the entire period would lead to a relatively small sample size. Additionally, such a panel would be subject to survivorship bias. Most problematic for the approach in this

[2]I refer interested readers to Kaplan and Menzio (2015), Broda and Weinstein (2010), and Broda and Parker (2014) for a more detailed description of the CP.

paper is that a longitudinal panel inherently features an average age that increases with time. Because the approach relies on variation in experience across household with macroeconomic conditions, this is problematic. As in STY, in order to avoid these issues I construct sequences of overlapping three-year subpanels consisting of households who remained in the sample for three consecutive years. This results in panels with a relatively stable average age as well as similar education and income characteristics. Table I shows summary statistics by year, while Table II shows the summary statistics for each subpanel. Overall, the average age and median income levels are slightly more stable for the subpanels, as desired.[3] I also drop cohorts with fewer than 1,000 observations in order to ensure that I have sufficient sample size within each cohort.

There are a few other selection criteria with which I limit the analysis. First, the head of the household must be aged between 25 and 60.[4] Additionally, they must have non-missing consumption, race, education, family size, and income data in order to be included. Table IV shows the summary statistics for the entire period relative to the Survey of Consumer Finance (SCF). As one can see by comparing columns two and three of Table IV, these restrictions do lead to some differences in the estimation sample versus the full sample. Participants in the estimation sample are on average younger, more college educated, and have higher income. While this may introduce some bias in the estimation, it is a necessary step for identification. Moreover, we know from the limited participation literature that wealthier, higher educated individuals are more likely to own stocks and thus this sample is exactly the relevant one when it comes to asset pricing implications.[5]

There is substantial evidence to support that the CP is a reliable source of consumption data. Returning to Table IV, one can see that Nielsen compares quite favorably to the SCF. The median income, mean age, and percentage of college ed-

[3]Note that the increasing age problem is less of an issue here than it was in STY, given the shorter range of the panel.

[4]When indicated, I further limit this to age 30-55.

[5]See for example Mankiw and Zeldes (1991).

ucated levels are similar. This is reassuring that the Nielsen data is a representative sample.[6,7]. Moreover, Einav, Leibtag and Nevo (2008) run a validation study of the CP by comparing the purchase data recorded by consumers to data recorded by the retailer. The most problematic field is the price variable. They show that errors are primarily due to two reasons: 1) standard inputting errors by panelists 2) Nielsen's imputation of prices using store averages. However, since I am primarily interested in cyclicality, any measurement error in the data should not affect my estimates as long as the errors do not vary over the business cycle.

Above, I raised the issue that the CP comprises primarily non-durable good consumption and does not include services or durable goods. This may raise concerns that it is difficult to make statements about cyclicality and asset prices without having a complete picture of consumption. Regarding cyclicality, however, having an estimate of non-durable consumption can only understate the true cyclicality in the data, as non-durable consumption is less cyclical than durable consumption. There is much empirical evidence to support this. For example, Berger and Vavra (2015) find that consumers adjust different types of consumption cyclically, particularly that durables are more sluggish to respond to shocks during recessions due to microeconomic frictions. Kaplan, Mitman and Violante (2016) show that the CP data can be used to replicate the qualitative results from Mian, Rao and Sufi (2013) who use housing supply elasticity as an instrument to measure the effect of the housing bust on household consumption. The Mian, Rao and Sufi (2013) analysis relies on automobile spending as a proxy for consumption. The fact that Kaplan, Mitman, and Violante (2016) are able to find similar results using the CP, which contains only information on non-durables, is reassuring.

As a further check on the validity of the CP data, I have completed a quick back of the envelope exercise to put the data in perspective. As of 2010, services

[6]While the mean income is much higher in the SCF, it is known that the SCF oversamples wealthy individuals, see for example Campbell (2006)

[7]The CP does not contain information about stock ownership. The estimates presented in this table are based on a probit model of stock ownership fitted from the PSID. I defer further discussion of this until Section 4.4.

made up approximately two-thirds of total expenditures in the CEX and the PCE.[8]. Moreover, durables make up about one-third of all goods consumption (about one-ninth of total consumption). So, in the CP, we expect to observe only about 20 percent of what we think of as true consumption. Indeed, the Bureau of Economic Analysis (BEA) estimates per capital PCE to be 63,563 dollars in 2010, implying about 12,000-13,000 dollars in consumption of non-durable goods.[9]. My sample has an average consumption of approximately 4,500 dollars over all the years, indicating that the CP is capturing only one-third of what we would expect. This is reasonable coverage given that we can't be certain that all members of the panel scan every purchase and that not all goods are covered by the Nielsen scanner system, as well as general measurement error.

One might also worry that the distribution of durables versus non-durable consumption is heterogeneous across income. As a result, my consumption estimates could be biased differently for those in different income groups. Based on the CEX, there is some evidence that the share of durables of total consumption is increasing in income (see Table V). Taking these moments from the 2010 CEX literally, one would conclude that the CP is missing 65 percent of consumption for those with an income between 50,000-69,999 dollars, but missing 70 percent of consumption for those with an income greater than 70,000 dollars. One can imagine an adjustment of the data that accounts for this, however that is beyond the scope of this paper. To alleviate concerns about this, I again make the point that durables are in fact *more* cyclical than non-durables. Hence, anything that I measure here represents a lower bound on the true implications of cyclicality for consumption. This may be even more relevant for those in higher income brackets. What is observed in the CP is interesting for the very fact that it represents items that are likely to be the last things to be cut in the face of recessions, such as a groceries and health and beauty products purchased at a drugstore.

---

[8]https://fred.stlouisfed.org/release/tables?rid=53&eid=43861&od=2010-01-01# https://www.bls.gov/news.release/cesan.nr0.htm

[9]https://fred.stlouisfed.org/series/A794RX0Q048SBEA#0

With regards to the asset pricing implications, the lack of information about durables may still be cause for concern. There is some evidence, such as Yogo (2006), that durable consumption is highly correlated with asset prices. However, this framework relies on extremely high risk aversion and non-separability of utility and thus may not be generalizable. Other evidence, such as Dunn and Singleton (1986), Eichenbaum and Hansen (1990), and Heaton (1993, 1995) consider the consumption Euler equation when utility depends on services from consumer durables. They show that adding consumer durables does not help explain the level of the equity premium. Piazezzi, Schneider and Tuzel (2007) develop a model that introduces "composition risk", that is the share of housing as total consumption, and that this risk induces low frequency movements in stock prices that are not driven by news about cash flows. Together, these pieces of evidence are mixed on the relevance of non-durable consumption for asset prices. Moreover, the standard in asset pricing is to use non-durable consumption, starting with Mehra and Prescott (1985). Applying the methodology used in this paper to a consumption data with more complete consumption is an interesting area for future work. The present paper aims to draw conclusions based on the best sources of consumption data that are publicly available. There is ample evidence to support that durable consumption and non-durable consumption co-move strongly[10], and thus there is still something to learn from looking at the available non-durable consumption.

II.2 The Panel Study of Dynamic Income

I supplement my analysis using data from the Panel Study of Dynamic Income (henceforth PSID) from 1999-2017 (biannually). This was the source of data for STY, however they used income data from 1968-1993. Each household in the PSID is interviewed annually from 1968-1997 and biannually after that. Sample sizes began at around 5,000, 3,000 of which were representative of the U.S. population and 2,000

[10] See for example Stock and Watson (1999).

were a low-income over sample. Households are asked detailed questions about their demographics, income, and spending, among other things.[11]

The PSID also contained food consumption data from 1968-1993. However, since 1999, the PSID has added more detailed consumption categories in an attempt to match the CEX. As such, I can use this data as another source of semi-complete household consumption data. In theory, the PSID should also contain information about durable and service spending, thus bypassing the issues discussed with the CP. However, the sample size is much smaller, thus limiting the power of inference.

I make the same restrictions on the PSID data as with the CP. That is, I construct subpanels in which the household must appear in two-consecutive (biannual) surveys. Moreover, the household must have non-missing education, consumption, and income data and must be of working age. Summary statistics for the PSID are shown in Table III. As expected, consumption in the PSID is much higher than in the CP. Table IV also shows how the PSID compares to the CP and the SCF. Overall, the PSID sample is younger, less white, less educated, and slightly lower-income.

II.3 Macroeconomic Data

I use four sources of macroeconomic data to classify years as expansions and contractions. The first uses the NBER business cycles. These are derived from the monthly NBER definitions of contractions and expansions. I convert the monthly definitions to yearly ones by classifying a year as contractionary if it has at least six months considered to be contractions. If contractions spanned more than six months over two calendar years, I designate the first year as contractionary. This results in a total of nine contractionary years from 1969[12]-2017. Appendix Table A1 shows how each year is classified using the different definitions. The shaded years in Figure II

[11] I refer readers to Blundell, Pistaferri, and Preston (2008) for a more detailed discussion of the PSID

[12] The first working year for a member of the sample.

indicate NBER recession years.

The second definition is real GDP growth from the National Income and Product Accounts (NIPA), maintained by the BEA. Using a similar definition as STY, I classify a year as a contraction if the growth is below the mean growth within the past five years an expansion if growth is above the mean growth for the past five years. This results in a total of 18 contractionary years over the same time period, hence this definition is more sensitive. The top left figure of Figure II shows GDP growth over time compared to the the five-year moving average.

The third definition uses unemployment data from the Bureau of Labor Statistics (BLS). I classify a year as a contraction if the national unemployment rate increases by more than one percent. This results in eight contractionary years over the same period. The top right figure of Figure II shows the unemployment rate over time.

Lastly, I use stock returns from CRSP. I calculate the equity premium by subtracting the annualized 30-day T-bill return from the annual market return. I then classify a year as a contractionary if the real return is less than five percent. This results in 16 contractionary years over the same period. The bottom figure of Figure II shows the excess return over time.

Table VI shows the correlation matrix for each measure of contractions. They are all positively correlated, with the NBER indicators and the unemployment indicators being particularly similar, with a correlation of .93. The equity premium indicator is the least similar to the other series. This becomes relevant when discussing the results. My findings are robust to each of these definitions, however they are stronger or weaker depending on the definition used.

III Estimation Approach

III.1 The Consumption Process

Following STY, I estimate the same time-series model for idiosyncratic risk, but for consumption rather than income.[13] This model has two key ingredients that make it ideal for the questions at hand. First, it accounts for the deterministic component of consumption that is not affected by idiosyncratic risk. Second, it allows for the moments of the distribution of shocks to change depending on the aggregate state of the economy. A third component that I add by extending upon the original framework, is controlling for inequality; this will be discussed below with assumption A.4. These three things together allow me to use this model to precisely estimate how cyclicality affects idiosyncratic consumption risk.

The first step is to decompose consumption into its observable components, time component, and an unexplained, or idiosyncratic component. Denote c_{it}^h as the logarithm of consumption for household i of age h at time t. Log consumption is specified as:

$$c_{i,t}^h = \theta_0 + \theta_1^T D(Y_t) + \theta_2^T x_{i,t}^h + u_{it}^h \tag{1}$$

where $D(Y_t)$ is a vector of year dummy variables for $t = 2004, ..., 2017$ and x_{it}^h is a vector of observable household characteristics. STY included the age of the male head of the household (plus age-squared and age-cubed), family size, and education level for the male head of the household in x_{it}^h. I also add race dummies and region-code dummies. The specification thus accounts for both aggregate variation in consumption and household observables. The main object of interest is the residual, u_{it}^h, which represents the random component of a household's consumption that is idiosyncratic to that household specifically. I specify that u_{it}^h follows a stochastic process, which is

[13] Similar models have been used in many other settings, such as Guvenen, Ozkan and Song (2014), Lettau and Ludvigson (2003), and Campbell and Cocco (2005).

ARMA(1,1) and has regime-switching conditional moments:

$$u_{it}^h = \alpha_i + z_{i,t}^h + \epsilon_{it}^h \tag{2}$$

$$z_{it}^h = \rho z_{it-1}^{h-1} + \eta_{it}^h \tag{3}$$

with $\alpha_i \sim_{iid} F_\alpha$, $\epsilon_{it}^h \sim_{iid} F_\epsilon$, $\eta_{it}^h \sim_{iid} F_\eta(s_t)$, $z_{it}^0 = 0$, $\forall i, t$.[14,15] The conditional variance of η_{it}^h is dependent on the aggregate state:

$$\mu_\eta^2 = \begin{cases} \mu_{\eta,E}^2 \text{ if t is a year of expansion} \\ \mu_{\eta,C}^2 \text{ if t is a year of contraction} \end{cases} \tag{4}$$

I extend the original STY model by also allowing the conditional skewness[16] of η_{it}^h to be dependent on the aggregate state:

$$\mu_\eta^3 = \begin{cases} \mu_{\eta,E}^3 \text{ if t is a year of expansion} \\ \mu_{\eta,C}^3 \text{ if t is a year of contraction} \end{cases} \tag{5}$$

The variables z_{it}^h and ϵ_{it}^h are the persistent and transitory shocks, respectively. α_i is a "fixed-effect" in the form of a shock at birth that is retained throughout life. Each component of the model serves a specific function. The fixed effects and transitory shocks allow us to better measure the primary objects of interest - μ_η^k and ρ. The way parameters will be identified involves examining how the cross-sectional moments of u_{it}^h change with age. However, much of the variation is common to households of all ages. Hence, if there were no fixed effect, the magnitude of μ_η^k would be overstated, therefore overestimating the amount of idiosyncratic risk that households face. Likewise, the transitory shock ϵ_{it}^h is included to capture measurement error.

[14] In the original STY paper, normality was also imposed on the distribution of α_i, ϵ_{it}^h and η_{it}^h. I relax that in order to estimate higher order moments; this will be discussed further in section 3.2

[15] In what follows, μ_y^x represents the xth moment of parameter y.

[16] Note that the model does not rule out regime switching for moments higher than the third. In this paper, I focus only on the second and third moment, but it is possible that higher order moments are regime switching as well.

The i.i.d. assumptions imply that α_i is not correlated with residual consumption nor the persistent or transitory shock. This is a strong assumption. To understand this, note that it rules out the possibility that α_i is time varying or that later-in-life realizations of shocks could affect it. For example, this means that the fixed effect received at birth is drawn from the same distribution for everyone, regardless of when they are born. In other words, it assumes that there are no changes in inequality over time. However, this assumption, as well as the initial condition that $z_{it}^0 = 0$, are essential to the estimation approach. These conditions allow us to interpret (2) and (3) as a collection of finite processes and thus we can condition on age. I now move on to discuss these assumptions in more detail.

Discussion of assumptions

As with any parametric estimation, the model requires some assumptions in order to provide structure. In this section, I'll provide a brief discussion of the key assumptions underlying the estimation.

A.1 Functional form of the aggregate and the deterministic component

As shown in equation (1) The aggregate component is modeled using year dummies. The deterministic component is:

$$f(x_{it}^h) = \hat{\beta}[1, h, h^2, h^3, education_{it}, familysize_{it}, region_{it}] + \varepsilon \tag{6}$$

This specification of the aggregate and deterministic components is essential for identifying the residual, but beyond that it is not central. One potential issue that could arise out of this specification is that the deterministic component is invariant across time and individuals. This may inflate the residual, as we know that individuals from different socioeconomic backgrounds may have different deterministic profiles and that the profile may change over time[17]. Moreover,

[17]See for example, Guvenen (2007) and Carroll and Summers (1991).

modeling education as a deterministic component likely results in underestimating the overall risk that agents face. However, estimating schooling decisions is beyond the scope of this paper. That being said, this specification is standard in the earnings literature (Hubbard, Skinner and Zeldes (1994)). As a robustness check, I estimate the model on different subsets of the population based on education, income, and time period in Appendix Tables A4, A5, and A6, respectively. Some of the coefficients are significantly different from each other, particularly the race coefficients in the education and income regressions. The regressions split by time period also exhibit quite a bit of dispersion across the race and age coefficients. However, the order of magnitude and direction of each coefficient is preserved across each specification. More importantly, the concavity in age is maintained among each of the samples. In section 4.2, I will show how the results of my first-stage estimation compare to others in the literature; the results are reassuring. Moreover, I will show in section 4.4 that my findings are robust to using specific subsets of the population based on income and education.

A.2 Functional form of the residual

The residual, shown in equation (2) is modeled a a function of a fixed effect, a persistent shock, and a transitory shock. This is, again, standard in the literature and there is ample empirical evidence to support its validity (Abowd and Card (1989), Hubbard, Skinner and Zeldes (1994), Heaton and Lucas (1996) Blundell, Pistaferri, and Preston (2008), Guvenen, Ozkan, and Song (2014), and Schmidt (2016)).

A.3 The fixed effect α_i is time invariant, drawn once at birth, and i.i.d.

This is a more substantial assumption, as it rules out time variation in α_i and dependence between α_i, for a given i, and subsequent realizations of z_{it} and ϵ^h_{it}. In other words, the fixed effect is drawn from the same distribution for everyone, regardless of the year of birth. This rules out an increase in inequality (due to the fixed effect) over time. Time invariance in α_i is actually not essential, but it being uncorrelated with subsequent realizations of the other shocks is. Imposing time invariance is the most straightforward and tractable way of achieving this. The payoffs for making this assumption, along with assumption 4, which I discuss next, are large: they me to interpret the process for z_{it}^h in equation (3) as a collection of finite processes, and thus condition on age.

A.4 The persistent component follows an AR(1) process

- The initial condition: $z_{it}^0 = 0 \quad \forall \quad i, t$
- The innovations η_{it}^h are i.i.d.
- The second and third moments of the AR(1) innovations are regime switching

The assumption that $z_{it}^0 = 0$ for all i and t is perhaps the strongest assumption imposed. However, it is very important for the estimation strategy: if z_{it}^0 varied across individuals then I would not be able to identify the parameters of interest in equation (3). In light of more recent empirical evidence of rising inequality, there is reason to believe that this assumption may bias my estimates.[18]. Indeed, looking at Figure I, one can clearly see that the residual variance of both consumption and income has increased over time.[19] Even just in the period from 2004-2017 of the CP data, it has increased by about 75 percent on average. In the PSID, it has almost doubled since 1968. This means that it is difficult to disentangle the effects of business cycles and inequality. For example, consider two groups of 30-year olds. One enters the workforce in 1980 and the other in 2007. Looking at Appendix Table A1, one can see that both face three contractions of GDP growth in their first five years of working (between ages 25-30). Taking the model literally, this would imply that they should exhibit, on average, the same cross-sectional variance (and skewness) at age 30 (in 1985 and 2012, respectively). However, if one believes that inequality has increased

[18]See Blundell, Pistaferri, and Preston (2008), Aguiar and Bils (2015), Carrol, Tokuoka, and White (2017), and Patterson (2018)

[19]These figures present only raw data, so some of the increase may also be due to change in the sampling procedures and composition of Nielsen and the PSID, respectively, over time.

from 1980-2012, then this may not be the case. Again looking at Figure I, we can see that residual variance has increased over time even when controlling for cohort. As a result, when observing higher variance among younger cohorts, it is possible to falsely attribute this to cyclicality when it may truly be driven by inequality.

However, there is a way to maintain this assumption while still accounting for the fact that the initial distribution of shocks has likely changed over time. I do so by controlling for the initial residual variance of consumption in the year that a cohort enters the workforce using food consumption data in the PSID going back to 1968. To achieve this, I add the initial residual variance of food consumption in year zero to x_{it}^h in equation (1).[20,21,22] While using only food consumption likely understates the true variance of consumption, it is the best proxy available going back far enough to provide controls for the oldest individuals in my sample.[23] Adding this control, we can still assume that $z_{it}^0 = 0$ for all i and t, but it should now be thought of as the initial shock net of the observed initial dispersion, which is plausibly unchanging over time.[24] To see this, recall that u_{it}^h in equation (1) represents the *idiosyncratic*, or unobserved component of log consumption. I specify that it follows an ARMA (1,1) process:

$$
\begin{aligned}
u_{it}^h &= \alpha_i + \rho z_{it-1}^{h-1} + \eta_{it}^h + \epsilon_{it}^h && (7)\\
&= \alpha_i + \rho^h z_{it-h}^0 + \sum_{k=0}^{h-1} \rho^k \eta_{it-k}^{h-k} + \epsilon_{it}^h && (8)
\end{aligned}
$$

[20]That is, I control for the orange line in the bottom panel of Figure I

[21]Note that adding this control does not substantially change the results of the first stage regression. The first stage regression results without the control are shown in Appendix Table (A3).

[22]In an alternative specification, I also run this regression controlling for both the variance and the skewness of initial food consumption. The coefficient on skewness is not significantly different from zero and makes no difference in the estimation results using the residuals. Using a larger dataset than the PSID, in which skewness is more precisely estimated, would be a natural next step, were such a dataset available going back to the 1960s.

[23]I also repeat the analysis using the variance of income in the PSID as a control and the results are similar.

[24]This is difficult to test as I observe very few individuals at age 25, when the initial variance should be the same across all cohorts, conditional on the current state of the economy, if this assumption is valid. In Appendix A, I compare the initial residual variances for the cohorts that I do observe in year 0 and show that they are similar.

Hence, for somebody in their first year of the workforce, the residual is:

$$u_{it}^1 = \alpha_i + \rho z_{it-1}^0 + \eta_{it}^1 + \epsilon_{it}^1 \tag{9}$$

Note that the fixed effect, α_i, the transitory shock, ϵ_{it}^h, and the AR(1) innovation, η_{it}^h, are drawn from the same distribution for everyone. Thus, when aggregating across cohorts, the sum of these components is the same on average. Hence if there is variation in u_{it}^1 across cohorts, as observed in Figure I, it must be the case that z_{it}^0 is not the same across cohorts. Moreover, z_{it}^0 represents the initial starting point of the permanent shock. We do not know exactly what determines this starting point, but it is specified to be *unobservable*. However, we can in fact observe something about where cohorts start off - i.e. the level and variance of their residual consumption or income when they enter the workforce. Hence, adding this control moves a portion of the determining factors of z_{it}^0 from the unobservable part of equation (1) to the observable part.[25]

On the regime switching moments, this is simply a way to add structure to the model and identify cyclical effects. While this is a simplification for the sake of tractability, it is a useful starting point.[26]

A.5 The transitory component, ϵ_{it}^h, is i.i.d.

Again, this is a standard assumption that provides tractability.

III.2 GMM

I follow the same approach as STY in estimation. First, I run the regression in equation (1) to obtain residuals. The second step fits the stochastic process (2) to the cross-sectional moments of the distribution of residual log consumption. The system

[25] An alternative route would be to estimate z_{it}^0 as a free parameter. This is an interesting avenue for future research, but beyond the scope of the current paper.

[26] One can imagine a model with moments that are continuous, rather than binary, and scale proportionally to the severity of the contraction/expansion. This is an interesting avenue for future work, but beyond the scope of the current paper

in (2) and (3) implies the following moments of this distribution[27]:

$$\textbf{Variance: } \mu^2(u_{it}^h;\theta) = \mu_\alpha^2 + \mu_\epsilon^2 + \sum_{j=0}^{h-1} \rho^{2j}\mu_\eta^2(s_{t-j}) \tag{10a}$$

$$\textbf{Covariance: } \mu^{11}(u_{it}^h, u_{it+1}^{h+1};\theta) = \mu_\alpha^2 + \rho\sum_{j=0}^{h-1} \rho^{2j}\mu_\eta^2(s_{t-j}) \tag{10b}$$

$$\textbf{Skewness: } \mu^3(u_{it}^h;\theta) = \mu_\alpha^3 + \mu_\epsilon^3 + \sum_{j=0}^{h-1} \rho^{3j}\mu_\eta^3(s_{t-j}) \tag{10c}$$

$$\textbf{Coskewness: } \mu^{21}(u_{it}^h, u_{it+1}^{h+1};\theta) = \mu_\alpha^3 + \rho\sum_{j=0}^{h-1} \rho^{3j}\mu_\eta^3(s_{t-j}) \tag{10d}$$

where $\theta = (\rho, \mu_\alpha^2, \mu_\epsilon^2, \mu_{\eta,E}^2, \mu_{\eta,C}^2, \mu_\alpha^3, \mu_\epsilon^3, \mu_{\eta,E}^3, \mu_{\eta,C}^3)$ is the vector collecting all of the parameters to be estimated. $\mu^2(u_{it}^h;\theta)$ and $\mu^3(u_{it}^h;\theta)$ denote the second and third central moment; $\mu^{11}(u_{it}^h, u_{it+1}^{h+1};\theta)$ and $\mu^{21}(u_{it}^h, u_{it+1}^{h+1};\theta)$ denote the covariance and a measure of co-skewness between u_{it}^h and u_{it+1}^{h+1}. The covariance and coskewness terms enable me to separately identify the moments of the fixed effect α_i and the transitory shock ϵ_{it}^h.

As described in the previous section, the conditional second and third moments of the AR(1) innovation, η_{it}^h, are allowed to be state dependent. The aggregate state can either be an expansion or a contraction, and the variance (skewness) takes on a separate value in each state. Define an indicator variable $\mathbb{1}_{s(t)=E} = 1$ if the economy is an an expansion (denoted by E) at time t. Then we have:

$$\mu_\eta^2(s(t)) = \mathbb{1}_{s(t)=E}\mu_{\eta,E}^2 + (1 - \mathbb{1}_{s(t)=E})\mu_{\eta,C}^2 \tag{11}$$

$$\mu_\eta^3(s(t)) = \mathbb{1}_{s(t)=E}\mu_{\eta,E}^3 + (1 - \mathbb{1}_{s(t)=E})\mu_{\eta,C}^3 \tag{12}$$

In order to ensure that the central moments are precisely estimated, I collapse each observation into a five-year age groups, beginning at age 22. Hence, I have eight total age groups, denoted by h_g: 22-26, 27-31, 32-36, 37-41, 42-46, 47-51, 52-56, 57-

[27]See Appendix B for derivation.

61.[28]. This results in an average of 3,990 observations in each age by year cell. I can thus calculate the empirical moments :

$$\mu^k(u_{it}^{h_g}) = \frac{1}{N_{h_g t}} \sum_{h \in h_g} (\mu^k(u_{it}^h)) \text{ for } k \in (2, 11, 3, 12) \tag{13}$$

This represents the mean of the given moment (variance, covariance, skewness, or co-skewness) within an age group, h_g, at time t. The theoretical counterpart is:

$$\mu^k(u_{it}^{h_g}; \theta) = \frac{1}{N_{h_g t}} \sum_{h \in h_g} N_{h_g t}(\mu^k(u_{it}^h; \theta)) \text{ for } k \in (2, 11, 3, 12) \tag{14}$$

for a specific combination of the parameters, θ. And so the moment conditions are

$$E[\mu^2(u_{it}^{h_g}) - \mu^2(u_{it}^{h_g}; \theta)] = 0 \tag{15a}$$

$$E[\mu^{11}(u_{it}^{h_g}) - \mu^{11}(u_{it}^{h_g}; \theta)] = 0 \tag{15b}$$

$$E[\mu^3(u_{it}^{h_g}) - \mu^3(u_{it}^{h_g}; \theta)] = 0 \tag{15c}$$

$$E[\mu^{21}(u_{it}^{h_g}) - \mu^{21}(u_{it}^{h_g}; \theta)] = 0 \tag{15d}$$

for each year t and age group h_g. All in all, I have $h_g \times t$ cross-sectional measures of variance and skewness and $h_g \times (t-1)$ measures of covariance and co-skewness. This results in a total of $2 \times h_g \times t + 2 \times h_g \times (t-1)$ empirical moments.[29] Given the classification of years as expansions of contractions, the system in 15 will uniquely identify the nine parameters in θ.

III.3 Identification

The key to the identification strategy is taking advantage of variation by cohort in macroeconomic history. The fact that the persistent shocks accumulate over time means that the methodology incorporates more business cycles into the analysis than

[28] Ages below 25 and above 60 are still dropped, as discussed in section 2.1

[29] So, there are $2\times8\times14+2\times8\times13 = 432$ moments conditions in the CP and $2\times8\times10+2\times8\times9 = 304$ moment conditions in the PSID.

are covered by the sample. In particular, examining equation (10a), shows us that the variances of the persistent shock innovation, η, accumulate as the cohort ages. Additionally, if the innovation variance is higher in contractionary years, then a cohort that has lived through more contractions will have a higher residual income variance at a given age when compared to cohorts that have lived through fewer contractions.

Revisiting the example from the introduction, recall our two groups of 35-year olds. One is born in 1955 and the other is born in 1970. The first enters the workforce at age 25, in 1980. Over the subsequent decade, half of the ten years are classified as a NIPA contraction (see Appendix Table A1 and Figure II). The second enters the workforce at age 25 in 1995. Only two of the following ten years are classified as NIPA contractions. So, the cohort born in 1955 will accumulate more high-variance persistent shocks while aged 25-35. This will show up when comparing the cross-sectional residual variances of cohorts at age 35.[30]

More concretely, consider two cohorts which we observe beginning their working life during the sample period, from 2004-2017. Those born from 1975-1979 enter the workforce between 2000-2004. Those born from 1980-1984 enter the workforce between 2005-2009. In the top panel of Figure III, I plot the residual variance of each of these two cohorts over the number of years that they have been in the workforce (i.e. the x-axis is their age minus 25). One can immediately see that the 1980 cohort has much higher residual variance after only three years in the workforce. This perfectly demonstrates the point that those who have lived through more contractions will show higher variance, as those born from 1980-1985 entered the workforce right before or during the financial crisis. The bottom panel of Figure III shows how the residual variance evolves for all of the cohorts observed in the sample. As we can see, the residual variance always increases as the cohort ages. However, there is dispersion between cohorts at the same age. This is exactly how the parameters in the model are identified via the GMM estimation.

[30]Recall that this is *residual* variance, hence I have already controlled for age, family size, region, year effects, and the initial distribution of consumption for each cohort.

The same logic applies to the skewness, as can be seen in equation (10c). If the probability of a large negative income shock is higher (and the probability of a large positive income shocks is lower) during contractions, then skewness during contractionary periods would be more negative. Comparing two cohorts at same age, we should see the cumulative effects of this in a measure of cross-sectional skewness. that is, cohorts that have worked through more contractions will exhibit more negative skewness.

Looking again at (10a), we see that the sum of the variances of the transitory shock and the fixed effect, $(\mu_\alpha^2 + \mu_\epsilon^2)$ is identified as the intercept of the variance profile over age. The same holds for $(\mu_\alpha^3 + \mu_\epsilon^3)$. Moreover, the magnitude of the increase in the cross-sectional variance (skewness) over age identifies the variance (skewness) of the persistent shocks. The difference between $\mu_{\eta,C}^k$ and $\mu_{\eta,E}^k$ is identified by the difference of the kth moment of different cohorts of the same age. Roughly speaking, the average of the two represents the rate of increase over the age profile. Additionally, the shape of the profile over age identifies ρ; a linear shape implies an autocorrelation coefficient close to one. In summary, variation across age drives the estimates of the autocorrelation coefficient, ρ while variation across time drives the estimates of μ_η^k. Lastly, we can see from equations (10b) and (10d) that the covariance and co-skewness allows me to separately identify μ_α^k and μ_ϵ^k.

The identification strategy is robust as long as ρ is not close to zero. As ρ approaches zero, it becomes impossible to identify the moments of the AR(1) innovation, as the shocks would no longer accumulate. There is ample evidence that there is persistence in these types of shocks, thus I am not concerned that the identification is threatened by a ρ close to zero.[31]

[31] See for example, McCurdy (1982) and Abowd and Card (1989) Carroll and Summers (1991), Carroll (1992) Gourinchas and Parker (2002).

IV Results

IV.1 GMM Preview: Graphical Analysis

Many of the main results can be understood by simply looking at the data. From the system that describes the dynamics of u_{it}^h, (2) and (3), one can see that the estimation of autocorrelation ρ is mostly driven by variation in age. On the other hand, the estimation of the regime switching variance, $\mu_{\eta,C}^2$ and $\mu_{\eta,E}^2$, of the innovation η_{it}^h to the persistent shock z_{it}^h, is driven by variation in time. This motivates a reduced-form representation of the data. Following Deaton and Paxson (1994), I decompose cross-sectional variances (10a) into cohort and age effects:

$$\tilde{Var}(u_{it}^h) = a_c + b_h + e_{h,t} \tag{16}$$

$c = t - h$ is a cohort (birth year) and the parameters a_c and b_h are cohort and age effects, respectively.[32]

I first plot the age coefficients, b_h from the regression in (16) for consumption in the CP in the top panel of Figure IV in blue. This demonstrates that there is much dispersion when people are young, and that it increases substantially over the lifetime. The residual variance increases from about .15 to .4, or approximately 2.7 times, from age 25 to 60. Moreover, the increase is roughly linear, indicating an autocorrelation coefficient close to one.

The bottom panel of Figure IV plots the cohort coefficients, a_c but differently. The x-axis is the fraction of contractions that a cohort has lived through.[33] If the variance is indeed counteryclical, then an implication of the model is that a cohort which has worked through more contractions than another at the same age should, on average, exhibit higher cross-sectional variance. Figure IV shows that this indeed

[32]Regression results are available upon request.

[33]I.e. $\frac{\text{Number of contractions at time t}}{Age_t - 25}$

the case. Plotting the cohort coefficients against the fraction of NBER contractions that a cohort has worked through shows an obvious positive relationship. The OLS slope coefficient is .41. In Appendix Figure A2, I show that this positive relationship is robust to each of the three other definitions of contractions.

Next, I turn to the third-moment. I run the same regression as above, but with skewness on the left hand side

$$\tilde{Skew}(u_{it}^{h}) = a_c + b_h + e_{h,t} \tag{17}$$

I plot the coefficients in Figure V. From the top panel, we can see that the residual skewness also increases with age. It increases from about -1.5 at age 25-30 to -.2 at age 60. From equation (10b), it's not obvious that this is what we would expect, as the skewness can take on negative or positive values. If ρ is close to one, then we expect the skenwnesses to accumulate, however many years of varying positive or negative skewness might cancel each other out. What this figure indicates is that, on average, the skewness of the AR(1) innovation is negative, but , on net, it is positive. This makes sense given that the average fraction of NBER contractions in my sample is less than 20 percent, but only if skewness in expansionary years is positive. If we expect the consumption results to match closely those found in income, then the skewness should be more negative during contractions. Hence, with mostly non-contractionary years, the skewness could be more positive and hence accumulate over the lifecycle.

The bottom panel of Figure V shows the cohort coefficients from the regression in (17) against the fraction of NBER contractions that a cohort has worked through. The positive correlation is apparent - The OLS slope coefficient is 5.3. This is a surprising finding. Most of the income literature has found that skewness tends to be *procyclical*, and thus we might expect skewness to be more negative if a cohort has worked through more contractions. The fact that this is not the case indicates that consumption may not exhibit the same pro-cyclical skewness that income does. In Appendix Figure A3, I show that this positive relationship is robust to each of the three other definitions of contractions. I'll return to this discussion in Section 4.2, when I discuss the GMM estimation results, which verify this conjecture.

Turning now to the PSID data, the same figures show consumption data in the PSID in red. Looking again at the top panel of Figure IV, we see that the cross-sectional variance also increases in age, going from about .5 at age 30 to .8 at at 60, or about 1.6 times.[34] This is smaller than the increase observed in the CP data, but the shape of the increase is also close to linear, indicating unity in the autocorrelation coefficient. Note that the variances are also higher on average in the PSID, which makes sense given that the PSID consumption covers many more products. The bottom panel of Figure IV shows the cohort coefficients against the fraction of contractions the cohort has worked though. Again, we observe a positive relationship, with an OLS coefficient of .48.

The top panel of Figure V shows how the cross-sectional skewness evolves with age. Unlike with the CP data, it is relatively flat. This indicates that the persistent shocks to consumption are not, on net, positive in this data. Rather, they must be negative on average, with some positive and negative realizations canceling each other out as an individual accumulates shocks over the lifecycle. The bottom panel shows a negative relationship between the cross-sectional skewness and the fraction of contractions worked: the OLS coefficient is -.88. This contrasts the very large positive coefficient that is found when looking at consumption in the CP, but indicates that consumption, as measured by the PSID, may exhibit similar procyclical skewness as does income. However, the standard error on the OLS coefficient is .67, indicating that is it very imprecisely estimated and not significantly different from zero. I will revisit this discussion following the estimation results in the next section.

[34] Those younger than 30 were dropped from this exercise due to small sample sizes.

IV.2 Estimation Results

The estimation begins with the first stage regression specified in equation (1) and repeated here for clarity:

$$c_{i,t}^h = \theta_0 + \theta_1^T D(Y_t) + \theta_2^T x_{i,t}^h + u_{it}^h \tag{18}$$

Recall that the year dummies $D(Y_t)$ are used to capture the aggregate component of consumption and the vector x_{it}^h captures the deterministic component of consumption. x_{it}^h contains a cubic in age for the head of household, family size, education level of the head of household, race dummies, state dummies, and an inequality control in the form of the initial residual variance of food consumption for the year in which the head of household enters the workforce. Table VII shows the results of this regression. The estimates are quite similar to previous studies, including STY and others, such as Hubbard, Skinner, and Zeldes (2004). Consumption is concave in age and increasing in family size. Also note that the coefficient on the initial standard deviation of food consumption is large and highly significant. This confirms evidence that inequality has increased over time, hence this initial dispersion affects the level of consumption later in life.[35]

A few differences from the regression in STY are worth noting. First, the education coefficient is negative. It is also economically small and barely significant, with a t-statistic of less than two. In STY, the coefficient is positive and highly significant, indicating a strong positive relationship between education and earnings. Another important difference is the R-squared values. In the regression presented here, these variables explain 12.8 percent of the variation in log of consumption. In STY, the R-squared is 23 percent. Additionally, STY do not include race or region dummies, hence they have fewer explanatory variables but more explanatory power. This is

[35]In Appendix Table A3, I show this regression without the inequality control. Note that the coefficients on other observables do not substantially change, nor does the R-squared. This validates the use of this proxy for inequality as a control, as without it, its impact would be captured by the residual, which is exactly what I want to avoid.

relevant, as the object that I am primarily interested in is the *residual* of this regression. A lower R-squared indicates that there may be more unobserved factors that influence residual consumption than income.

Running the regression in (18) allows me to identify the residual, u_{it}^h which follows the process in equation (2). I estimate the parameters specified via the systems in (10) and (15) using GMM. The results are shown in Table IX. The first three columns shows the results for the CP data for three of the contraction definitions.[36] Looking first at the results for the autocorrelation coefficient, ρ, we see that it is indeed close to unity and quite precisely estimated.[37] On average, the sum of the variances of the fixed effect and the transitory shock, shown in the 2nd and 3rd row respectively, is .1587. This matches well with the intercept in the top panel of Figure IV, which is approximately .13.

Now turning to the results for the second moment of the AR(1) innovation, the countercyclical volatility is immediately apparent. The standard deviation rises by an average of 23 percent from expansion to contraction.[38]

The last four rows in Table IX show the results for the parameters of the third moment. Again, the sum of the skewness of the transitory shock and the fixed effect identifies the intercept in the top panel of Figure V. The average sum of the two is -1.56, compared to the intercept of -1.54 in the figure. Additionally, the fact that the skewness of the transitory shock, μ_ϵ^3 is highly negative confirms that transitory shocks to consumption are negatively skewed: negative shock realizations have more weight than positive ones.

Moving now to the skewness of the AR(1) innovation, we see that $\mu_{\eta,C}^3$ and $\mu_{\eta,E}^3$

[36] I exclude the results for the unemployment contractions as they are almost identical to those of the NBER indicators. The NBER and unemployment indicators are highly correlated (see Table VI).

[37] The 95 percent confidence interval for the estimate using the NBER indicators is [.9764, .9988].

[38] The differences between $\mu_{\eta,C}^2$ and $\mu_{\eta,E}^2$ are significantly different at the 1 percent level for the NIPA and Negative returns definition; they are significantly different at the 10% level for the NBER definition.

are not significantly different from each other. Additionally, they are very imprecisely estimated with large confidence bands. Although the estimation results with each measure of contractions do indicate that there is slightly procyclical skewness (i.e. $\mu^3_{\eta,C}$ is more negative that $\mu^3_{\eta,E}$ in all cases), there is not enough power to reject that the two estimates are equal.

Next, I discuss the complementary estimation results using the PSID data. The sample sizes are much smaller in this data, with between 3,000-6,000 observations per year (see Table III). Additionally, the survey is only biannual from 1999-2017, thus I have only 10 years of observations, as opposed to the 14 in the CP data. Even so, the results, shown in the last three columns of Table IX, are comparable. The autocorrelation coefficient is close to one for each of the three recession definitions. The volatility exhibits even more countercyclicality here than in the CP, rising by 68, 30, and 66 percent for each of three recession types, respectively.[39] Moreover, the skewness still does not change substantially between the two states, and is imprecisely estimated.

IV.3 Comparison to Income Results

In order to compare the dynamics of consumption to those of income, I first repeat the analysis on the PSID over the same sample period, 1999-2017. Beginning again with the graphical analysis, the top panel of Figure VI shows the age coefficients from the regression in (16) with the variance of residual income on the y-axis. As with consumption, the variance increases with age, rising by about 1.5 times from age 25 to 60. The magnitude of the increase is similar to that found in PSID consumption. The initial dispersion, or the sum of the variance of the fixed effect and the transitory shock, is higher in the income data, about .55 versus .38. The bottom panel of Figure VI plots the cohort coefficients from the regression in (16) against the fraction of

[39]For each of the recession types, the conditional variances for expansions and contractions are significantly different at the 1 percent level.

contractions that a cohort has worked through. As with consumption, the relationship is positive, with an OLS slope coefficient of .57, which is slightly higher than the coefficient for consumption (.48). Hence, it seems that income and consumption in the PSID both exhibit strong autocorrelation and countercyclicality.

Looking now at Figure VII, we see that the dynamics of the third moments of income in the PSID are similar to those of consumption in the PSID as well. As with consumption, the top panel shows that the cross-sectional skewness is not increasing over the lifecycle, but rather stays relatively flat and negative. This indicates that the persistent shock to income is not positive, on net, as it was with consumption. In the PSID, negatively skewed shocks dominate as the innovations accumulate. Looking at the bottom panel, we see that the OLS slope coefficient is -.11, with a very large standard deviation of 1.15.[40], compared to an imprecisely estimated measure of -.87 in the PSID consumption data. Although this slope is small in magnitude and not significantly different from zero, it provides some weak evidence of possible procyclicality of skewness in the PSID income data. It's important to interpret these results with caution, though, due to the small sample size in the PSID and imprecision of measurement with higher order moments.

In Table X, I show the estimation results on PSID income from 1999-2017. As with consumption, the autocorrelation coefficient, ρ, is close to unity. The conditional standard deviation, μ_η^2 rises by approximately 34 percent from expansion to contraction. This is very similar to increase for consumption, which was 30 percent for the full sample in the PSID using NBER indicators. The difference between the state-conditional variances for income is significantly different at the 95 percent level. In the skewness here, however, we see that it does substantially increase in magnitude from expansion to contraction. The values are significantly different from each other at the 95th percent.[41] Hence, there is evidence that the skewness of income is

[40]When using the definitions for contractions that rely on the equity premium and GDP growth, the OLS coefficients are -1.12 with a standard deviation of .97 and -.66 with a standard deviation of .65, respectively.

[41]The values for the skewness during expansions, -.0007, is not significantly different from zero.

procyclical; this evidence is absent in the consumption data.

In Table X, I also duplicate the main estimation results using NBER indicators from STY Table 2, for comparison. Besides the fact that STY used data on income while I am using data on consumption, their estimation was also for the time period from 1968-1993, while mine is from 2004-2017 in the CP and 1999-2017 in the PSID. Hence, there may be differences due to the time period, rather than due to differences in income and consumption. Comparing these results to my results for income, a few differences are immediately apparent. First, the variance of the fixed effect and the transitory shock are higher in the more recent data. Similarly, the conditional variances of the persistent shock are larger. On the other hand, the increase from expansion to contraction is much *smaller* in the updated data. In my estimation from 1999-2017, it rises by about 33 percent from peak to trough, versus 69 percent in the STY estimation.

To summarize, residual income and consumption both exhibit strong autocorrelation and countercyclicality in the variance of the persistent shock using data from the PSID from 1999-2017. This is also observed by STY using income data from the PSID from 1968-1993. However, the magnitude of the increase of the variance from peak to trough in STY is much larger than I find in the more recent data for both consumption and income. Moreover, I do observe procyclical skewness in the PSID income data, with the skewness being almost seven times more negative during contractions. However, I do not observe a similar pattern in the consumption data; the conditional skewness in each state are not statistically different from each other.

IV.4 Stockholders

One of the primary issues with consumption-based asset pricing is *limited participation*. Namely, when not all consumers participate in the stock market, the SDF is attenuated and therefore the equity premium is overestimated.[42] For this reason, it is difficult to make asset pricing implications when looking at the consumption of individuals who may not participate in the stock market.

To deal with this, I re-estimate the parameters on specific subsamples of the population that are most likely to own stocks. The CP does not contain data on stock ownership, however there are several proxies available in the demographic information that is provided. First, I drop those who have less than a college education, leaving about 40 percent of the sample. Second, I drop those who have an income of fewer than 100,000 dollars, leaving about 20 percent of the sample. Lastly, I fit a model based on the PSID data, which does include stock ownership information, to predict stock ownership in the CP. Following Malloy, Moskowitz and and Vissing-Jorgenson (2009), I fit a probit model of stock ownership on age, income, education, and year effects. I then apply these estimates to the CP data to predict stock ownership. Appendix Table A7 shows the results of the probit regression. Table IV shows the estimated stock ownership in the CP data. Following this method, approximately 16 percent of the estimation sample in the CP is predicted to own stocks.[43]

The results for the estimation on the high-education, high-income, and stock owning subsamples are shown in Table XI. This table uses the NBER recession indicators, and thus is comparable with the first column of Table IX. The results are qualitatively similar across each of the subsamples and the full sample. However, there are some notable differences. First of all, these subsamples exhibit slightly *more* countercyclicality in the estimates of the conditional variance of the persistent shock, particularly the high income and stock holding samples. The standard deviation rises by 28 percent, 33 percent and 34 percent from economic expansion to contraction for each subsample, respectively, compared to 23 percent for the full sample. Second, the variances of the transitory shock and the fixed effect (the sum of which identifies the initial dispersion), are lower for the high education and stock-

[42] See Campbell (2018) and Brav, Constantinides and Geczy (2002) for a more thorough discussion.
[43] Summary statistics for each of these subpanels are shown in Appendix Tables A8-A10, respectively.

holding groups. Comparing the second and third rows of Table XI, we can see that the difference is primarily driven by the variance of the fixed effect This implies that, conditional on being someone who ultimately obtains a higher education or becomes a stockholder, the fixed effect drawn at birth exhibits less variation. Looking now at the estimates for the third moments, the skewness does not appear to differ substantially conditional on the aggregate state, but the estimates are imprecise. This is similar to the findings in the full sample. The skewness of the fixed effect and the transitory shock are also smaller in magnitude within these subsamples than in the full sample.

IV.5 Growth rates

I now discuss the implications of my findings for growth rates. Simply first differencing log consumption in the model gives:

$$\begin{aligned}\Delta c_{it+1} = c_{it+1} - c_{it} =& (\theta_1^T (D(Y_{t+1} - D(Y_t)) + \theta_2^T (x_{it+1} - x_{it}) \\ &+ n_{it+1} + \sum_{k=0}^{h-1} ((\rho^{k+1} - \rho^k)\eta_{it-k}^{h-k}) + (\epsilon_{it+1}^{h+1} - \epsilon_{it}^h)\end{aligned} \tag{19}$$

Thus, as the autocorrelation coefficient, ρ, approaches one, the past shocks becomes less and less significant. That is, the idiosyncratic growth rate is primarily driven by the current AR(1) innovation and the change in the transitory shock.[44]

The residual is obtained by running an equivalent first-stage regression with growth rates on the left-hand side:

$$\Delta c_{i,t+1}^h = \theta_0 + \theta_1^T D(Y_t) + \theta_2^T \Delta x_{i,t}^h + \Delta u_{it}^h \tag{20}$$

Where Δx_{it}^h is the change in observables from time $t-1$ to t and Δu_{it}^h is the residual.

[44] This is consistent with the finding in Kandel and Stambaugh (1990) that the variance of consumption growth is countercyclical.

From equation (19), the cross-sectional variance of the residual growth rate is:

$$var(\Delta u_{it}^h) = \sigma_{\eta,t}^2 + \sum_{k=1}^{h-1} (\rho^k - \rho^{k-1})^2 \sigma_{\eta,t-k}^2 + 2\sigma_\epsilon^2 \tag{21}$$

Thus, we expect the variance of the residual growth rate to increase both with age and as a function of past contractions.

The residual variance of the growth rate should also be lower in magnitude than the residual variance of the level. To see this, note that

$$var(\Delta u_{it}^h) - var(u_{it}^h) = \sigma_\epsilon^2 - \sigma_\alpha^2 + \sum_{k=1}^{h-1} [(\rho^k - \rho^{k-1})^2 - \rho^{2k}] \sigma_{\eta,t-k}^2 \tag{22}$$

Now, take the case where $\rho = 1$, and we have

$$var(\Delta u_{it}^h) - var(u_{it}^h) = \sigma_\epsilon^2 - \sigma_\alpha^2 - \sum_{k=1}^{h-1} \sigma_{\eta,t-k}^2 \tag{23}$$

Thus, the residual variance of the growth rate is almost surely smaller than the residual variance of the level. Based on the parameter estimates from the GMM estimation in Table IX, we can see that the difference is always negative, given that σ_α^2 is about two times σ_ϵ^2. Note that this may not hold if i) the variance of the transitory component were extremely large ii) the variance of the AR(1) innovations were extremely small iii) if the variance of the transitory component were large and the individual i were young (i.e. h is low). This is entirely intuitive, as if the transitory shock had a much higher variance, we would expect growth rates to exhibit a higher variance than levels. Moreover, if the variance of the persistent shocks were low, the effect of the transitory shock would dominate, again making the variance of the growth rate higher than that of the levels. However, we can see from the parameter estimates that this is not the case; as long as the persistent shocks have a similar volatility to the transitory shocks, the effect of the persistent shock will always dominate.

Following the procedure above, I first plot the age coefficients from the Deaton and Paxon (1994) regression in equation (16) in the top panel of Figure VIII. Not surprisingly, the residual variance of consumption growth is increasing in age. That is, holding cohort constant, the residual variance increases in age, as the model would predict. Moreover, when comparing this figure to Figure IV, we can see that the level of variance is much lower (about .07 on average versus .3).

However, looking simply at the residual variance by age without controlling for cohort, we see that it *decreases* with age, as shown in the blue line in Figure VIII. The fact that we see the variance decreasing in age when not controlling for cohort could tell us several things. First, notice that as ρ approaches one, we would not be able to estimate (21) using the GMM strategy described above for estimating levels. This is precisely because the identification of the model comes from the fact that the shocks are persistent and there is variation in macroeconomic history by cohort. This vanishes as ρ approaches one. Thus, its likely that what we are observing is the result of the rise in inequality. Younger cohorts are more exposed to the inequality increase that has occurred over time, as documented in Blundell, Pistaferri, and Preston (2008), Deaton and Paxon (1994) and several others. Hence, the younger people in my sample exhibit more cross-sectional dispersion. When I control for the starting point (i.e. cohort) of each individual, we see the dispersion increase with age, as expected. Moreover, this speaks to the possible mis-specification in the model. Examining equation (21), one can see that if the variance of the growth decreases over the lifecycle, it must be that either the variance of the AR(1) innovation, $\sigma^2_{\eta,t}$, or the variance of the transitory shock, σ^2_ϵ, are decreasing as individuals age.[45] Clearly this is not possible when the shocks are modeled as i.i.d. However, this again makes the point that there are likely secular trends happening over the lifecycle of the members of my sample which are not captured in this model.

Next, I plot the regression coefficients on cohort from equation (16) in the bottom panel of Figure VIII. A clear increasing trend is apparent; cohorts which have worked

[45]This is in fact what we would expect, as consumption tends to stabilize as people age, see Deaton and Paxson (1994).

through more contractions, controlling for age, exhibit higher cross-sectional variation in the residual of growth rates. The OLS coefficient is .43. This is consistent with the findings on levels. This also holds when not controlling for age; even between cohorts, the residual variance is increasing with more contractions (see the bottom panel of Figure VIII). This indicates that the result is not driven by inequality.

Consumption growth is the primary object of interest in this analysis when it comes to asset pricing. In equation (19), we see that the variance of residual consumption growth is driven by the variance of the current AR(1) innovation. The evidence I have presented on growth rates here gives a few insights. First, the past fraction of contractions that a cohort has worked through is positively correlated with the residual variance. Thus, the evidence is consistent with the finding of counteryclical variance. Second, the effects of inequality seem to play a large role in determining the distribution of growth rates. However, when I discuss the asset pricing implications in the next section, we will see evidence that the incidence of recessions in indeed an important driver of the equity premium. Moreover, in my analysis on *levels*, I have controlled for inequality, as explained in section 3.1. Thus, my GMM estimates are still useful in telling us how idiosyncratic consumption risk interacts with asset prices, controlling for inequality.

V Implications for Asset Pricing

V.1 Asset pricing framework

With some standard assumptions, I can calculate the required excess return implied by the model.[46] The Euler equation of consumer i for security j is given by:

$$E_t[R^j_{t+1}\beta\frac{u_i'(c_{i,t+1},x_{i,t+1})}{u_i'(c_{i,t},x_{i,t})}]=1 \tag{24}$$

[46]See Appendix D for more details on the derivation that follows.

where $u_i'(c_{i,t}, x_{it})$ is the marginal utility for consumer i with deterministic characteristics x_{it}, of consumption at time t, R_{t+1}^j is the return of security j at time $t+1$, and β is a discount factor. $E_t[X]$ is the conditional expectation of variable X at time t. $\beta\frac{u'(c_{i,t+1})}{u'(c_{i,t})}$ represents the stochastic discount factor (SDF). If we further assume CRRA utility with risk aversion γ, and plug in for the consumption process[47], we have

$$E_t[R_{t+1}^j\beta(\frac{exp(\theta_1^T D(Y_{t+1}) + \ \alpha_i + z_{it+1}^{h+1} + \epsilon_{it+1}^{h+1}])}{exp(\theta_1^T D(Y_t) + \alpha_i + z_{it}^h + \epsilon_{it}^h])})^{-\gamma}] = 1 \tag{25}$$

$$\Longrightarrow E_t[R_{t+1}^j\beta exp[-\gamma(\ \underbrace{\theta_1^T(D(Y_{t+1}) - D(Y_t)}_{\text{Growth of aggregate component}}\ + \underbrace{(z_{it+1}^{h+1} - z_{it}^h) + (\epsilon_{it+1}^{h+1} - \epsilon_{it}^h)}_{\text{Growth of idiosyncratic component}}))]] = 1 \tag{26}$$

Thus, the SDF can be written as:

$$M_{t+1} = \beta E_{t+1}^*[exp[-\gamma(\theta_1^T(D(Y_{t+1} - D(Y_t)) + (z_{it+1} - z_{it}) + (\epsilon_{it+1}^{h+1} - \epsilon_{it}^h)]] \tag{27}$$

Where E_t^* denotes the cross-sectional expectation.[48] This follows from the fact that any investor's marginal rate of substitution is a valid SDF; therefore the cross-sectional average of all investors' marginal rates of substitution is also a valid SDF. I take the cross-sectional average because the methodology relies on using information from repeated cross-sections of households. Given the data limitations, I do not have a long time series for any given individual or cohort. By taking the cross-sectional average, I can learn more from the macroeconomic history.

Assuming that cross-sectional consumption growth is log-normal[49,50], taking logs

[47] Note that I am modeling the deterministic components in equation (6) as utility shifters, thus they no longer enter directly into consumption growth. Information about these observables (age, family size, education, and region) is contained in the individual's utility function $u_i(c_{it}, x_{it})$, which is allowed to vary across consumers. I further show in Appendix D that excluding the deterministic component is a harmless assumption, as it has no effect on the asset pricing relation.

[48] That is, $E_t^* X_{kt} = (1/K)\sum_{k=1}^K X_{kt}$

[49] In the case of regime switching skewness, the normality assumption is not valid. However, given that I find no significant difference in the skewness between states in my GMM estimates, I will assume normality for the sake of the asset pricing relationships. Alternatively, I could add a skewness correction as in Martin (2013). This would be quantitatively small and would not substantially impact my results.

[50] This assumption also relies on the fact that the sum of two normal distributions is also a normal

of equation (27) gives:

$$\begin{aligned} m_{t+1} =& log(\beta) - \gamma E^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] \\ &+ \frac{\gamma^2}{2} var^*_{t+1}((z_{it+1} - z_{it}) + (\epsilon^{h+1}_{it+1} - \epsilon^h_{it})) \end{aligned} \tag{28}$$

Where var^*_{t+1} denotes the cross-sectional variance.[51]

Now plugging in for the variances of the persistent and transitory shocks, we have

$$\begin{aligned} m_{t+1} =& log(\beta) - \gamma E^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] \\ &+ \frac{\gamma^2}{2}[\sigma^2_{\eta,t+1} + E^*_{t+1}[\sum_{k=0}^{h-1}(\rho^{k+1} - \rho^k)^2 \sigma^2_{\eta,t-k}] + 2\sigma^2_\epsilon] \end{aligned} \tag{29}$$

Where $E^*_{t+1}[\sum_{k=0}^{h-1}(\rho^{k+1} - \rho^k)^2 \sigma^2_{\eta,t-k}]$ is the cross-sectional average of the weighted sum of past variances for all cohorts at time $t+1$.[52]

Returning now to the the basic Euler equation we have that

$$E_t[R^j_{t+1} M_{t+1}] = 1 \tag{30}$$

By the law of iterated expectations[53], the following unconditional Euler equation also holds:

$$E[R^j_{t+1} M_{t+1}] = 1 \tag{31}$$

$$\Longrightarrow E[r^j_{t+1}] + E[m_{t+1}] + \frac{1}{2}[var(r^j_{t+1}) + var(m_{t+1}) + 2cov(r^j_{t+1}, m_{t+1})] = 0 \tag{32}$$

This holds for all assets, including the risk-free asset. Thus, the risk premium is given

distribution.

[51] That is, $var^*_t X_{kt} = (1/K)\sum_{k=1}^K (X_{kt} - E^*_t X_{kt})^2$

[52] All of the covariances terms are zero, as the transitory and persistent shocks are i.i.d. and the residual is orthogonal to the aggregate component, by assumption.

[53] See Lettau and Ludvigson (2009)

by:

$$E[log(R^m_{t+1})] - log(R^f_{t+1}) = -cov(r^m_{t+1}, m_{t+1}) \tag{33}$$

Where m_{t+1} is specified by (29).[54] So the risk premium is a function of the market return's covariance with the changes in the aggregate component and the variance of changes in the idiosyncratic components.

Expanding out equation (33), while mathematically tedious, yields some insights:

$$\begin{aligned} E[log(R^j_{t+1})] - log(R^f_{t+1}) = {} & \gamma cov(r^j_{t+1}, E^*_{t+1}[\theta_1(D(Y_{t+1})) - D(Y_t))]) \\ & - \frac{\gamma^2}{2} cov(r^m_{t+1}, \sigma^2_{\eta,t+1}) \\ & - \frac{\gamma^2}{2} cov(r^m_{t+1}, E^*_{t+1}[\sum_{k=0}^{h-1} (\rho^{k+1} - \rho^k)^2 \sigma^2_{\eta,t-k}]) \\ & - \gamma^2 cov(r^m_{t+1}, \sigma^2_\epsilon) \end{aligned} \tag{34}$$

Note that the final line of equation (34), $cov(r^m_{t+1}, \sigma^2_\epsilon)$ is equal to zero, as the variance of the transitory shock does not change over time (by assumption). The third line is approaching zero, as the difference between realizations of the cross-sectional average of past shock variances and the sample average of the same object will equal zero in large samples.[55] Thus, the equity premium can be written more succinctly as:

$$\begin{aligned} E[logR^j_{t+1})] - log(R^f_{t+1}) = {} & \gamma cov(r^j_{t+1}, E^*_{t+1}[\Delta c_{t+1}]) \\ & - \frac{\gamma^2}{2} cov(r^m_{t+1}, \sigma^2_{\eta,t+1}) \end{aligned} \tag{35}$$

The first line, representing the covariance of the return with aggregate growth is what is represented in the typical representative Euler equation. Thus the line of interest is the second, representing the covariance of the return with the cross-sectional

[54] $cov(X, Y)$ represents the unconditional covariance between X and Y.

[55] Note that I could also assume a unit root and reach the same conclusion. If using my GMM estimates of approximately $\rho = .98$, this term becomes trivially small.

variance of the persistent shock's contemporaneous component. When this covariance is negative, the risk premium is larger. Hence, idiosyncratic risk has the potential to amplify risk premia in this model of consumption.[56]

Specifically, the second line tells us the expected rate of return is larger if the market return covaries negatively with the shocks' variance. This is exactly what we expect when this variance is indeed counterycyclical. To put it differently, if the cross-sectional variance of idiosyncratic shocks is larger when the return is low, investors require a higher expected return going forward. When the economy moves from a period of expansion to a period contraction, if idiosyncratic variances increases, we expect to see the required rate of return increase as well.

V.2 Results

Using my GMM estimates for predicted stockholders in section 4,4, I can now show exactly how the risk premium is affected by idiosyncratic risk. Additionally, I will estimate what is the required risk aversion in order for this model to explain the data. Figure IX shows the times series of the log equity premium, log aggregate consumption growth, and the variance of the persistent shock to consumption's growth.[57] The covariances of these series are what governs the equity premium in equation (35). The top figure shows these series for the entire post-war period, 1950-2017. I calculate the variance of idiosyncratic growth by imputing my GMM estimates for the AR(1) innovation variance for stockholders from Table XI. That is, the variance of idiosyncratic growth takes on a value of .0411 if the year is an expansion and a value of .0741 if the year is a contraction. The bottom figure shows the same three lines, zoomed in over my sample period. There is also an additional series during the sample period, which is a reduced-form way of calculating idiosyncratic growth variance. I do so by

[56] Note the same is true in any model of idiosyncratic consumption/income risk such as Constantinides and Duffie (1996), Heaton and Lucas (1996), Guvenen (2009) and many others.

[57] The equity premium is calculated as the annual market return (including dividends) of the value-weighted NYSE/AMEX index minus the annualized 30-day T-bill return. Aggregate consumption growth is from the CEX.

taking the cross-sectional variance of the residual from the first-stage regression in equation (1). To understand this, note that if the model were exactly correct, the "GMM Estimated" and "Regression Estimated" lines would be equivalent.

From this figure, three important facts emerge. First, the correlation between the return and the variance of idiosyncratic growth is strongly negative. This is exactly what we would expect with countercyclical variance. This holds both with my GMM estimates and with the regression estimates. The correlation between the GMM estimates of idiosyncratic growth variance and the log return is -.41 in the post-war period and -.45 during the sample period. The correlation between the regression estimated idiosyncratic growth variance and the log return during the sample period is -.09. Second, the correlation between the return and the GMM estimated variance of idiosyncratic growth has much larger magnitude than the covariance with aggregate growth. The correlation of the return with aggregate growth is .07 for the post-war period and .16 for the sample period. An extensive literature has shown that the covariance of returns with aggregate consumption growth is small and cannot explain the size of the equity premium.[58] Thus, the covariance of the return with idiosyncratic growth must be large in order to generate a sizable equity premium: this is exactly what I observe. Looking at the reduced-form regression estimates, while the magnitude of the correlation is smaller, it is still of comparable magnitude to the covariance with aggregate growth, meaning that it will at least double the risk premia. Lastly, the regression and model estimated idiosyncratic growth variances are positively correlated. While they are not equivalent, as would be the case with a perfectly specified model, the positive correlation indicates the model is qualitatively capturing movement in the consumption residual.

Next, I use the unconditional covariances of the series shown in Figure IX to calculate the implied equity premium from equation (35). The results are shown in Table XII. Again, I show the results for both the entire post-war period and the sample period. In addition to the GMM and Regression estimated calculation described

[58]See Campbell (2018) for a review of this evidence.

above, I also show the results in the case of no idiosyncratic risk, that is eliminating the second term from (35). The power of idiosyncratic risk is immediately evident. Looking first at the post-war period, using my GMM estimates for idiosyncratic growth variance raises the equity premium by 4.11 percent at a risk aversion of only 10. Over this time period, the average log equity premium was 6.7 percent. Hence, the model matches the level of the equity premium at a risk aversion between 10-15. This is a substantial improvement over the textbook CRRA case, which requires risk aversion of greater than 200.

Now looking at the in-sample period, the same results holds when using the GMM estimates of idiosyncratic risk variance: the model can explain the observed level of the equity premium at risk aversion between 10-15. The regression estimated idiosyncratic growth variance also does remarkably well. While this covariance is lower in magnitude than the GMM estimated covariance, it matches the level of the equity premium at a risk aversion of only 20. Hence, idiosyncratic risk, as estimated in the model and using a reduced-form regression, can significantly amplify risk premium in this model of consumption.

VI Conclusion

In this paper, I use a detailed data set on consumption expenditures to study the cyclicality of consumption risk. After controlling for aggregate, deterministic, and inequality trends, I show that idiosyncratic consumption risk is both highly persistent and countercyclical. That is, the variance of the persistent shock to residual consumption increases by about 25-30 percent from economic contractions to expansions. The results are robust to several different definitions of contractions and across both the CP and the PSID data.

When examining the skewness of the consumption data, I do not find that it is procyclical, as has been shown with income. The GMM estimates on third moments are imprecise and show no significant differences between skewness in varying states of the economy. In the PSID data, there is some weak graphical evidence of procyclical skewness, but this does not bear out in the GMM estimation. I do confirm STY's previous finding, however, of countercyclicality of the variance in income on a more recent sample of PSID data. I also find evidence that *income* data does exhibit procyclical skewness, as in Guvenen, Ozkan, and Song (2014). Hence, my findings indicate that consumption may not exhibit as strong of procyclical skewness that income does.

There are several possible explanations for this discrepancy between income and consumption dynamics. Perhaps the most obvious is consumption smoothing: if consumer's smooth, then large increases or decreases in income do not necessarily translate directly into consumption. Moreover, it could simply be that these data sets are not capturing the part of consumption that is most responsive to macroeconomic conditions. Indeed, the CP captures only a small percentage of total consumption and it is likely to be goods for which demand is fairly inelastic such as groceries and health and beauty supplies purchased at a drug store. In the PSID data, while consumption is more complete, it is likely that there is substantial measurement error due to the interview/survey nature of the data collection. The sample sizes in the PSID are also very small. Lastly, I exclude housing costs from the measure of PSID consumption, hence it is still not providing the full picture of consumption.

My results also reaffirm previous findings in the literature that idiosyncratic risk can amplify risk premium. Indeed, without idiosyncratic risk, my data imply that a risk aversion of greater than 50 is needed to match the observed level of the equity premium. When I account for idiosyncratic risk, risk aversion need only be around 10-20, depending on how idiosyncratic risk is measured. Moreover, I have presented clear evidence that idiosyncratic risk covaries negatively with the market return; the variance of the idiosyncratic risk rises sharply when the return falls, as was the case in 2008 during the financial crisis.

An important next step in this research would be to apply the analysis techniques on a more complete consumption data set. Moreover, there are several interesting extensions that could be made to the model. First, the assumption that the regime-switching moments are binary is simplistic: one could extend this to allow for a continuous various/skewness that scales proportionally to the severity of the expansion/contraction. Second, the persistent shocks need not follow an AR(1) - one could consider other specifications that introduce other potentially interesting dynamics, such as a persistent shock with an MA(1) process, and thus a slow adjustment. Lastly, the relationship between income and consumption could be formalized. Then one could use both income and consumption data to see how income shocks translate into consumption expenditures, as in Blundell, Preston and Pistaferri (2008).

Bibliography

[1] Abowd, John and Card, David, (1989), On the Covariance Structure of Earnings and Hours Changes, Econometrica, 57, issue 2, p. 411-45.

[2] Aguiar, Mark, and Mark Bils. (2015). "Has Consumption Inequality Mirrored Income Inequality?" American Economic Review, 105 (9): 2725-56.

[3] Attanasio, O., & Davis, S. (1996). Relative Wage Movements and the Distribution of Consumption. Journal of Political Economy, 104(6), 1227-1262.

[4] Berger, D., & Vavra, J. (2015). Consumption Dynamics During Recessions. Econometrica, 83(1), 101-154.

[5] Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." American Economic Review, 98 (5): 1887-1921.

[6] Brav, A., Constantinides, G. M., & Geczy, C. C. (2002). Asset Pricing With Heterogeneous Consumers and Limited Participation: Empirical Evidence. Journal of Political Economy, 110 (4)

[7] Broda, Christian & Parker, Jonathan A., (2014). "The Economic Stimulus Payments of 2008 and the aggregate demand for consumption," Journal of Monetary Economics, Elsevier, vol. 68(S), pages S20-S36.

[8] Broda, Christian, and David E. Weinstein.(2010). "Product Creation and Destruction: Evidence and Price Implications." American Economic Review, 100 (3): 691-723.

[9] Busch, Christopher and Alexander Ludwig. “Higher-Order Income Risk over the Business Cycle: A Parametric Approach.” November 30, 2017. Working Paper.

[10] Campbell, John Y. “Household Finance.” The Journal of Finance, vol. 61, no. 4, 2006, pp. 1553–1604.

[11] Campbell JY. Financial Decisions and Markets: A Course in Asset Pricing. (2018).

[12] Campbell, John Y. & Cocco, Joao F., 2007. "How do house prices affect consumption? Evidence from micro data," Journal of Monetary Economics, Elsevier, vol. 54(3), pages 591-621, April.

[13] Carroll, Christopher D. (1992). "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," Brookings Papers on Economic Activity, Economic Studies Program, The Brookings Institution, vol. 23(2), pages 61-156.

[14] Carroll, Christopher, (1997), Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis, The Quarterly Journal of Economics, 112, issue 1, p. 1-55.

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

[16] Carroll, Christopher D., and Lawrence H. Summers. 1991. "Consumption Growth Parallels Income Growth: Some New Evidence." In National Saving and Economic Performance, ed. B. Douglas Bernheim and John B. Shoven, 305-43. Chicago: University of Chicago Press.

[17] Cogley, T. (2002). Idiosyncratic risk and the equity premium: evidence from the consumer expenditure survey. Journal of Monetary Economics, 49(2), 309-334.

[18] Constantinides, G. & Duffie, D. (1996). Asset Pricing with Heterogeneous Consumers, Journal of Political Economy, 104(2), 219-40.

[19] Constantinides, G. M. and Ghosh, A. (2017), Asset Pricing with Countercyclical Household Consumption Risk. The Journal of Finance, 72: 415-460

[20] Deaton, A., & Paxson, C. (1994). Intertemporal Choice and Inequality. Journal of Political Economy, 102(3), 437-467.

[21] Di Maggio, Marco and Kermani, Amir and Majlesi, Kaveh. (2018). "Stock Market Returns and Consumption" Forthcoming in the Journal of Finance.

[22] Duffee, G. R. (2005), Time Variation in the Covariance between Stock Returns and Consumption Growth. The Journal of Finance, 60: 1673-1712.

[23] Dunn, K., Singleton, K., 1986. Modeling the term structure of interest rates under non-separable utility and durability of goods. Journal of Financial Economics 17, 27–55. E

[24] Eichenbaum, M., Hansen, L.P., 1990. Estimating models with intertemporal substitution using aggregate time series data. Journal of Business and Economic Statistics 8, 53–69.

[25] Gourinchas, P., & Parker, J. (2002). Consumption over the Life Cycle. Econometrica, 70(1), 47-89.

[26] Green, D, Melzer, B., Parker, J., & Rojas, A. (2018). Accelerator or Brake? Cash for Clunkers, Household Liquidity, and Aggregate Demand." NBER Working Paper 22878.

[27] Guvenen, Fatih. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" The American Economic Review, vol. 97, no. 3, 2007, pp. 687–712.

[28] Guvenen, F. (2009). A Parsimonious Macroeconomic Model for Asset Pricing. Econometrica, 77(6), 1711-1750.

[29] Guvenen, F., Ozkan, S., & Song, J. (2014). The Nature of Countercyclical Income Risk. Journal of Political Economy, 122(3), 621-660. doi:10.1086/675535

[30] Heaton, J., 1993. The interaction between time non-separable preferences and time aggregation. Econometrica 61, 353–385.

[31] Heaton, J., 1995. An empirical investigation of asset pricing with temporally dependent preference specifications. Econometrica 63, 681–717.

[32] Heaton, J. & Lucas, D. (1996). Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing. Journal of Political Economy, 104(3).

[33] Hubbard, R., Skinner, J., & Zeldes, S. (1995). Precautionary Saving and Social Insurance. Journal of Political Economy, 103(2), 360-399.

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

[35] Kandel, Shmuel, and Robert F. Stambaugh. "Expectations and Volatility of Consumption and Asset Returns." The Review of Financial Studies, vol. 3, no. 2, 1990, pp. 207–232.

[36] Kaplan, Greg & Menzio, Guido. (2015). "The Morphology Of Price Dispersion," International Economic Review, Department of Economics, University of Pennsylvania and Osaka University Institute of Social and Economic Research Association, vol. 56, pages 1165-1206, November

[37] Kaplan, Greg, Mitman, Curt & Violante, Giovanni L. (2016). "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data," NBER Working Papers 22232, National Bureau of Economic Research, Inc.

[38] Krebs, Tom. (2007.) "Job Displacement Risk and the Cost of Business Cycles." American Economic Review, 97 (3): 664-686.

[39] Lettau, Martin (2002) "Idiosyncratic Risk and Volatility Bounds, or, Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle?" Review of Economics and Statistics 84, 376-380.

[40] Lettau, Martin and Ludvigson, Sydney. (2004). "Understanding Trend And Cycle In Asset Values: Reevaluating The Wealth Effect On Consumption," American Economic Review, v94(1,Mar), 276-299.

[41] Lettau, Martin and Ludvigson, Sydney. (2009). "Euler Equation Errors," Review of Economic Dynamics, Elsevier for the Society for Economic Dynamics, vol. 12(2), pages 255-283, April.

[42] Lucas, Deborah, J. (1994). "Asset pricing with undiversifiable income risk and short sales constraints: Deepening the equity premium puzzle", Journal of Monetary Economics, Volume 34, Issue 3, Pages 325-341,

[43] MaCurdy, Thomas E., (1982), The use of time series processes to model the error structure of earnings in a longitudinal data analysis, Journal of Econometrics, 18, issue 1, p. 83-114.

[44] Malloy, C.J., Moskowitz, T.J. and Vissing-Jorgensen, A. (2009), Long-Run Stockholder Consumption Risk and Asset Returns. The Journal of Finance, 64: 2427-2479

[45] Mankiw, N. Gregory and Zeldes, Stephen, (1991), The consumption of stockholders and nonstockholders, Journal of Financial Economics, 29, issue 1, p. 97-112.

[46] Martin, Ian W.R. (2013). Consumption-Based Asset Pricing with Higher Cumulants. The Review of Economic Studies. Volume 80. Issue 2. Pages 745–773,

[47] Mehra, Rajnish and Prescott, Edward C. (1985). The equity premium: A puzzle. Journal of Monetary Economics,Volume 15, Issue 2, Pages 145-161.

[48] Mian, A., K. Rao, and A. Sufi (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” The Quarterly Journal of Economics, 128, 1687–1726.

[49] Mian, A. & Sufi, A. (2010). Household Leverage and the Recession of 2007–09. IMF Economic Review, 58(1), 74–117.

[50] Ogaki, M. (1993). Generalized Method of Moments: Econometric Applications. Econometrics, Handbook of Statistics. 11. Edited by G. S Maddala, C. R. Rao, and H. D. Vinod. Amsterdam: Elsevier Sci.

[51] Parker, J. (2001). The Consumption Risk of the Stock Market. Brookings Papers Econ. Activity, 2, 279–333.

[52] Patterson, Christina. 2019. "The Matching Multiplier and the Amplification of Recessions," 2019 Meeting Papers 95, Society for Economic Dynamics.

[53] Piazzesi, Monika, Martin Schneider and Selale Tuzel. "Housing, Consumption, and Asset Pricing." Journal of Financial Economics 83 (March 2007): 531-569

[54] Schmidt, Lawrence, Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk (March 3, 2016).

[55] Stock J. and Watson, M. Business Cycle Fluctuations in U.S. Macroeconomic Time Series. In: Taylor J, Woodford M Handbook of Macroeconomics. Amsterdam: Elsevier ; 1999. pp. 3-64.

[56] Storesletten, K. Telmer, C., & Yaron, A. (2004). Consumption And Risk Sharing Over The Life Cycle. Journal of Monetary Economics, 51(3), 609-633.

[57] Storesletten, K., Telmer, C., & Yaron, A. (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. Journal of Political Economy, 112(3), 695-717.

Appendix A: Initial Residual Variance

In the discussion of Assumption A.4, I note that the initial AR(1) shock z_{it}^0 is plausibly the same for all cohorts after I have controlled for the inequality by including the variance of initial food consumption in the first stage regression. Observing equation (10a), we can see that the variance for someone in their first year of working is simply:

$$\mu^2(u_{it}^1;\theta) = \mu_\alpha^2 + \mu_\epsilon^2 + \mu_\eta^2(s_t) \tag{36}$$

However, if $z_{it}^0 \neq 0$, then this becomes:

$$\begin{aligned} \mu^2(u_{it}^1;\theta) &= \mu_\alpha^2 + \mu_\epsilon^2 + var(\rho z_{it-1}^0 + \eta_{it}) &(37)\\ &= \mu_\alpha^2 + \mu_\epsilon^2 + \rho^2 var(z_{it-1}^0) + \mu_\eta^2(s_t) &(38)\end{aligned}$$

Thus, if I could see each cohort at the time in which it enters the workforce and show that the residual variance (after controlling for inequality) is similar in year zero, conditional on whether year zero is a contraction or an expansion, this would be strong evidence that the assumption is valid. To see this, note that the fixed effect and the transitory shock are drawn from the same distribution for everyone, so no variation across cohorts comes from those terms. Moreover, the variance of the AR(1) innovation should be the same for all individuals which are in the same economic state in a given year. So, if the residual variance differs across cohorts in year one, conditional on year one being the same macroeconomic state, it must be that $var(z_{it-1}^0) > 0$ and hence all of the z_{it}^0 initial shocks are not the same.

Unfortunately, after I put the restrictions on my sample about consecutive data and eliminating cohorts with less than 1,000 observations, I only have observations from the cohorts 1979-1983 and each has very few observations in the year that they are aged 25.

In order to provide some evidence on this assumption, I relax the restrictions put on my estimation sample and look at all available data. This means that I have observations for cohorts 1979-1992 at age 25. I then collapse these into 3-year cohorts and groups based on whether or not the year in which they are 25 is an expansion or a contraction. The number of observations in each cohort in the total sample and at age 25 is shown in Table A2. So, I am still dealing with very small sample sizes and thus any inference must be made with extreme caution. Figure A1 shows the residual variance of each of the 3-year cohorts beginning at age 25. The red bands are 95 percent confidence intervals. The top panel shows the cohorts which are aged 25 during an expansion and the bottom shows those which are aged 25 during a contraction. With the exception of those born from 1987-1989, each of the groups that enter during an expansion are very close to each other with nearly overlapping confidence bands. Those that enter during a contraction are much more dispersed, however it is important to note that only those born in 1982, 1983, or 1984 enter the workforce during a contraction so the sample size is extremely small.

Overall, it is difficult to draw conclusions from this evidence without larger samples. Aside from noise due to small sample size, there are several reasons why the observed initial residual variances may differ across cohorts entering in the same economic conditions that do not undermine the assumption made in the model. For example, the model simplifies the effect of macroeconomic conditions on the variance by assuming that it is binary; in reality, it maybe be continuous and the variance scales depending on the severity of the expansion/contraction. In the bottom panel of Figure A1, we see that the 1984 cohort, which enters the workforce in 2009 has much lower variance than the 1982 and 1983 cohorts, which enter the workforce in 2007 and 2008. It could be that initial impact of the recession had a more of an impact on consumption than there was in 2009. Similarly, the cohorts observed in the top panel entered in any year from 2004-2017 that was not during the financial crisis. Even though these are not classified as contractions in this binary framework, economic conditions varied widely between all of those years.

Appendix B: Derivation of empirical moments, equations (10)

Begin by rewriting the stochastic process (3) by recursive substitution:

$$u_{it}^h = \alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j} \tag{39}$$

Variance

Variance is given by:

$$\begin{aligned}
\mu^2(u_{it}^h) =& E[(u_{it}^h - E[u_{it}^h])^2] && (40)\\
=& E[(\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j} - E[\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j}])^2] && (41)\\
=& E[\alpha_i - E[\alpha_i] + \epsilon_{it}^h - E[\epsilon_{it}^h] + \sum_{j=]}^{h-1} (\rho^j \eta_{it-j}^{h-j} - E[\rho^j \eta_{it-j}^{h-j}]))^2] && (42)
\end{aligned}$$

Due to the assumptions of independence of each of the random variables, this becomes

$$\mu^2(u_{it}^h) = \mu_\alpha^2 + \mu_\epsilon^2 + \sum_{j=0}^{h-1} \rho^{2j} \mu_\eta^2(s_{t-j}) \tag{43}$$

Skewness

Skewness is given by:

$$\mu^3(u_{it}^h) = E[(u_{it}^h - E[u_{it}^h])^3] \tag{44}$$

$$= E[(\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j} - E[\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j}])^3] \tag{45}$$

$$= E[\alpha_i - E[\alpha_i] + \epsilon_{it}^h - E[\epsilon_{it}^h] + \sum_{j=]}^{h-1} (\rho^j \eta_{it-j}^{h-j} - E[\rho^j \eta_{it-j}^{h-j}]))^3] \tag{46}$$

Due to the assumptions of independence of each of the random variables, this becomes

$$\mu^3(u_{it}^h) = \mu_\alpha^3 + \mu_\epsilon^3 + \rho \sum_{j=0}^{h-1} \rho^{3j} \mu_\eta^3(s_{t-j}) \tag{47}$$

Covariance

Covariance is given by:

$$\mu^{11}(u_{it}^h, u_{it+1}^{h+1}) = E[u_{it}^h - E[u_{it}^h](u_{it+1}^{h+1} - E[u_{it+1}^{h+1}])] \tag{48}$$

Using (39) above, this becomes

$$\begin{aligned} \mu^{11}(u_{it}^h, u_{it+1}^{h+1}) = & E[(\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j} - E[\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j}]) \\ & \times (\alpha_i + \epsilon_{it+1}^{h+1} + \sum_{j=0}^{h} \rho^j \eta_{it+1-j}^{h+1-j} - E[\alpha_i + \epsilon_{it+1}^{h+1} + \sum_{j=0}^{h} \rho^j \eta_{it+1-j}^{h+1-j}]) \end{aligned} \tag{49}$$

$$= \mu_\alpha^2 + \rho \sum_{j=0}^{h-1} \rho^{2j} \mu_\eta^2(s_{t-j}) \tag{50}$$

When using the PSID data, which is biannual, I also need the covariance from time t to $t+2$:

$$cov(u_{it}^h, u_{it+2}^{h+2}) = \mu_\alpha^2 + \rho^2 \sum_{j=0}^{h-1} \rho^{2j} \mu_\eta^2(s_{t-j}) \tag{51}$$

Coskewness

Covariance is given by:

$$\mu^{21}(u_{it}^h, u_{it+1}^{h+1}) = E[(u_{it}^h - E[u_{it}^h])^2 (u_{it+1}^{h+1} - E[u_{it+1}^{h+1}])] \tag{52}$$

Using (39) above, this becomes

$$\begin{aligned} \mu^{21}(u_{it}^h, u_{it+1}^{h+1}) =& E[(\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j} - E[\alpha_i + \epsilon_{it}^h + \sum_{j=0}^{h-1} \rho^j \eta_{it-j}^{h-j}])^2 \\ & \times (\alpha_i + \epsilon_{it+1}^{h+1} + \sum_{j=0}^{h} \rho^j \eta_{it+1-j}^{h+1-j} - E[\alpha_i + \epsilon_{it+1}^{h+1} + \sum_{j=0}^{h} \rho^j \eta_{it+1-j}^{h+1-j}]) \end{aligned} \tag{53}$$

Multiplying out each of these terms and factoring in the expectations, all of the terms that involve cross products with random variables to the first power cancel out due to the independence assumptions. Thus we have

$$\mu^{21}(u_{it}^h, u_{it+1}^{h+1}) = \mu_\alpha^3 + \rho \sum_{j=0}^{h-1} \rho^{3j} \mu_\eta^3(s_{t-j}) \tag{54}$$

When using the PSID data, which is biannual, I also need the coskewness from time t to $t+2$:

$$\mu^{21}(u_{it}^h, u_{it+2}^{h+2}) = \mu_\alpha^3 + \rho^2 \sum_{j=0}^{h-1} \rho^{3j} \mu_\eta^3(s_{t-j}) \tag{55}$$

Appendix C: GMM Estimation

The GMM estimator solves the following minimization problem

$$\hat{\Theta} = argmin_{\Theta \in \Xi} \hat{g}(\Theta)' \hat{W} \hat{g}(\Theta) \tag{56}$$

Under standard GMM assumptions, we have

$$\sqrt{N}(\hat{\Theta} - \Theta) \rightarrow N(0, \Sigma) \tag{57}$$

with

$$\Sigma = (G'W_0G)^{-1}G'W_0\Omega W_0G(G'W_0G)^{-1} \tag{58}$$
$$G = E[\frac{\partial g(i,h,t;\Theta)}{\partial \Theta'}] \tag{59}$$
$$\Omega = V[g(i,h,t;\Theta)] \tag{60}$$

I follow the two-step GMM procedure, originally described in Ogaki (1993). In the first step, I set $\hat{W} = I$. The second step uses the optimal weighting matrix, given by:

$$\hat{S} = [g(i,h,t;\hat{\Theta})g(i,h,t;\hat{\Theta})'] \tag{61}$$

where $\hat{\Theta}$ are the parameter values estimated in the first step.

Each estimation is overidentified and uses covariances and coskewness. Due to the unbalanced panel and the overlapping nature of the subpanels, there are a few adjustments to the typical GMM methodology required, Denote $N = min(N_{ht})$ as the minimum number of observations in a given h/t cell among all such cells and $\kappa_{ht} = \frac{N_{ht}}{N}$. I use the matrix

$$\Lambda = diag[\kappa_{ht}, \kappa_{ht}] \tag{62}$$

to scale the block diagonal matrix for the asymptotic covariance. With this adjustment, the estimator of the asymptotic variance is

$$\frac{1}{N}(\hat{G}'\hat{W}\hat{G})^{-1}\hat{G}'\hat{W}(\hat{S}\Lambda)\hat{W}\hat{G}(\hat{G}'\hat{W}\hat{G})^{-1} \tag{63}$$

Recall the systems described in (10) and (15). There are four moment conditions for each of eight age groups over the 14 years of the CP from 2004-2017. Our parameter vector, $\Theta = (\rho, \mu^2_{\alpha}, \mu^2_{\epsilon}, \mu^2_{\eta,E}, \mu^2_{\eta,C}, \mu^3_{\alpha}, \mu^3_{\epsilon}, \mu^3_{\eta,E}, \mu^3_{\eta,C})$, contains nine parameters to be estimated.

In the case with no covariances or coskewness, we have

$$\hat{G} = \frac{1}{N_{ht}} \sum_{i=1}^{N_{ht}} \frac{g(i,h,t;\hat{\Theta})}{\partial\Theta'}|_{224\times 9} \tag{64}$$

$$\hat{S} = \frac{1}{N_{ht}} \sum_{i=1}^{N_{ht}} g(i,h,t;\hat{\Theta})g(i,h,t;\hat{\Theta})'|_{2HT\times 2HT} \tag{65}$$

$$\implies \hat{G}_1 = \begin{bmatrix} \frac{\partial\hat{g}^1(h,t)}{\partial\rho} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^2_{\alpha}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^2_{\epsilon}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^2_{\eta,E}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^2_{\eta,C}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^3_{\alpha}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^3_{\epsilon}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^3_{\eta,E}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^3_{\eta,C}} \\ \frac{\partial\hat{g}^2(h,t)}{\partial\rho} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^2_{\alpha}} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^2_{\epsilon}} & \frac{\partial\hat{g}^1(h,t)}{\partial\mu^2_{\eta,E}} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^2_{\eta,C}} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^3_{\alpha}} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^3_{\epsilon}} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^3_{\eta,E}} & \frac{\partial\hat{g}^2(h,t)}{\partial\mu^3_{\eta,C}} \end{bmatrix}_{2HT\times 9} \tag{66}$$

As there are eight age groups and 14 years, $8 \times 14 = 112 \times 2 = 224$. So, the moment conditions for the variance and skewness for each year/age groups are stacked on top of one another.

When I include covariances and coskewness, we add the moment conditions corresponding to (15b) and (15d).

$$\Longrightarrow \hat{G}_2 = \begin{bmatrix} \frac{\partial \hat{g}^3(h,t)}{\partial \rho} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^2_\alpha} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^2_\epsilon} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^2_{\eta,E}} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^2_{\eta,C}} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^3_\alpha} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^3_\epsilon} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^3_{\eta,E}} & \frac{\partial \hat{g}^3(h,t)}{\partial \mu^3_{\eta,C}} \\ \frac{\partial \hat{g}^4(h,t)}{\partial \rho} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^2_\alpha} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^2_\epsilon} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^2_{\eta,E}} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^2_{\eta,C}} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^3_\alpha} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^3_\epsilon} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^3_{\eta,E}} & \frac{\partial \hat{g}^4(h,t)}{\partial \mu^3_{\eta,C}} \end{bmatrix}_{2H(T-1)\times 9} \tag{67}$$

Where here there are eight age groups for 13 years, $8 \times 13 = 104 \times 2 = 208$. Stacking $\hat{G}_1$ and $\hat{G}_2$, we have our complete set of moment conditions with $2HT + 2H(T-1) = 2 \times 8 \times 14 + 2 \times 8 \times 13 = 432$ moment conditions. $\hat{S}$ is then given by (61) with

$$g = \begin{bmatrix} g^1(h,t) \\ g^2(h,t) \\ g^3(h,t) \\ g^4(h,t) \end{bmatrix} \tag{68}$$

$$= \begin{bmatrix} E[\mu^2(u^{h_g}_{it}) - \mu^2(u^{h_g}_{it};\theta)] \\ E[\mu^3(u^{h_g}_{it}) - \mu^3(u^{h_g}_{it};\theta)] \\ E[\mu^{11}(u^{h_g}_{it}) - \mu^{11}(u^{h_g}_{it};\theta)] \\ E[\mu^{21}(u^{h_g}_{it}) - \mu^{21}(u^{h_g}_{it};\theta)] \end{bmatrix}_{(2HT+2H(T-1))\times 1} \tag{69}$$

Lastly, I test the overidentifying restriction using the “J-statistic"

$$J = N \times \hat{g}'(\hat{S}\Lambda)^{-1}\hat{g} \sim \chi^2(\#\text{ of moments} - \#\text{ of parameters}) \tag{70}$$

So, we have 423 $(432 - 9)$ degrees of freedom.

When using the PSID data, the same procedure is followed, but there are $2HT + 2H(T-1) = 2 \times 8 \times 10 + 2 \times 8 \times 9 = 304$ moment conditions. Moreover, since the PSID is completed biannually after 1999, the covariance and coskewness are lagged by two years, as shown in Appendix C, equations (51) and (55).

Appendix D: Derivation of Equity Premium

Suppose log consumption follows the process:

$$c_{i,t}^h = \underbrace{\theta_1^T D(Y_t)}_{\text{Aggregate component}} + \underbrace{\theta_2^T x_{it}}_{\text{Deterministic component}} \underbrace{+ \alpha_i + [\rho^h z_{i,t-h}^0 + \sum_{k=0}^{h-1} \rho^k \eta_{it-k}^{h-k}] + \epsilon_{it}^h}_{\text{Idiosyncractic component}} \quad (71)$$

with $\alpha_i \sim iid(0, \sigma_\alpha^2)$, $\epsilon_{it}^h \sim iid(0, \sigma_\epsilon^2)$, $\eta_{it}^h \sim iid(0, \sigma_{\eta,t}^2)$.l[59] The conditional variance of η_{it}^h is regime switching:

$$\sigma_{\eta,t}^2 = \begin{cases} \sigma_E^2 \text{ if t is a year of expansion} \\ \sigma_C^2 \text{ if t is a year of contraction} \end{cases}$$

With CRRA utility and risk aversion γ, the Euler equation of consumption of consumer i for security j is given by

$$E_t[R_{t+1}^j \beta (\frac{C_{i,t+1}}{C_{i,t}})^{-\gamma}] = 1 \quad (72)$$

where $E_t[X]$ represents the conditional expectation at time t of variable X. Plugging in for the consumption process, we have

$$E_t[R_{t+1}^j \beta (\frac{exp(\theta_1^T D(Y_{t+1}) + \theta_2^T x_{it+1} + \ \alpha_i + z_{it+1} + \epsilon_{it+1}^{h+1}])}{exp(\theta_1^T D(Y_t) + \theta_2^T x_{it} + \alpha_i + z_{it} + \epsilon_{it}^h])})^{-\gamma}] = 1 \quad (73)$$

$$\Longrightarrow E_t[R_{t+1}^j \beta exp[-\gamma(\underbrace{\theta_1^T (D(Y_{t+1} - D(Y_t)}_{\text{Growth of aggregate component}} + \underbrace{\theta_2^T (x_{it+1} - x_{it})}_{\text{Growth of deterministic component}}$$
$$+ \underbrace{(z_{it+1} - z_{it}) + (\epsilon_{it+1}^{h+1} - \epsilon_{it}^h)}_{\text{Growth of idiosyncratic component}}))]] = 1 \quad (74)$$

[59]Note that I have modified the assumptions from the original consumption process and put more structure on the distributions of the stochastic components. Because I find no significant difference in skewness in my GMM estimation, I assume normality going forward, for simplicity.

Note that the individual's consumption growth $\frac{C_{i,t+1}}{C_{i,t}}$ is given by

$$\frac{C_{i,t+1}}{C_{i,t}} = exp(\theta_1^T(D(Y_{t+1}) - D(Y_t)) + \theta_2^T(x_{it+1} - x_{it}) + (z_{it+1} - z_{it}) + (\epsilon_{it+1}^{h+1} - \epsilon_{it}^{h})) \tag{75}$$

Whereas aggregate consumption growth is simply

$$\frac{C_{t+1}}{C_t} = exp(\theta_1^T(D(Y_{t+1}) - D(Y_t)) \tag{76}$$

To see this, we only need to recall that each of the variables α_i, ϵ_{it}^h, and η_{it}^h are mean zero and hence have zero net effect when consumption is aggregated and we apply the law of large numbers:

$$\begin{aligned} lim_{n\to\infty}(1/N)\sum_{i=1}^{N} c_{it} =& lim_{n\to\infty}(1/N)[\sum_{i=1}^{N}\theta_1^T D(Y_t) + \sum_{i=1}^{N}\alpha_i + \sum_{i=1}^{N} z_{it} + \sum_{i=0}^{N}\epsilon_{it}^h] & (77) \\ =& \theta_1^T D(Y_t) & (78) \\ \Longrightarrow & \frac{C_{t+1}}{C_t} = \frac{exp(\theta_1^T(D(Y_{t+1}))}{exp(\theta_1^T(D(Y_t)))} = exp(\theta_1^T(D(Y_{t+1}) - D(Y_t)) & (79) \end{aligned}$$

So, the SDF can be written as:

$$\begin{aligned} M_{t+1} =& \beta E_{t+1}^*[exp[-\gamma(\underbrace{\theta_1^T(D(Y_{t+1}) - D(Y_t)) + \theta_2^T(x_{it+1} - x_{it})}_{\text{Representative agent terms}} \\ &+ \underbrace{(z_{it+1} - z_{it}) + (\epsilon_{it+1}^{h+1} - \epsilon_{it}^{h})}_{\text{Additional term due to idiosyncratic risk}}]] \end{aligned} \tag{80}$$

Where E_t^* denotes the cross-sectional expectation, $E_t^* X_{kt} = (1/K)\sum_{k=1}^{K} X_{kt}$. This follows from the fact that any investor's marginal rate of substitution is a valid SDF; therefore the cross-sectional average of all investors' marginal rates of substitution is also a valid SDF.

Assuming that cross-sectional consumption growth is log-normal[60,61], taking logs of equation (80) gives:

$$\begin{aligned} m_{t+1} =& log(\beta) - \gamma E^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] - \gamma E^*_{t+1}[\theta_2^T(x_{it+1} - x_{it})] \\ &+ \frac{\gamma^2}{2}[var^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] + var^*_{t+1}[\theta_2^T(x_{it+1} - x_{it})] \\ &+ var^*_{t+1}[(z_{it+1} - z_{it}) + (\epsilon^{h+1}_{it+1} - \epsilon^h_{it})]] \end{aligned} \tag{81}$$

Where var^*_{t+1} denotes the cross-sectional variance, $var^*_t X_{kt} = (1/K)\sum_{k=1}^K (X_{kt} - E^*_t X_{kt})^2$. I am ignoring the covariance terms for now and will revisit them shortly.

Now, we can see that several terms drop out. Firstly, $E^*_{t+1}[\theta_2^T(x_{it+1} - x_{it})] = 0$. Similarly, $var^*_{t+1}[\theta_2^T(x_{it+1} - x_{it})] = 0$. Lastly, $var^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] = 0$, as this is the same for all consumers at time t. Note that I have now shown that we can ignore the deterministic components, as I did in the main body of the paper. Thus, we are left with:

$$\begin{aligned} m_{t+1} =& log(\beta) - \gamma E^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] \\ &+ \frac{\gamma^2}{2} var^*_{t+1}((z_{it+1} - z_{it}) + (\epsilon^{h+1}_{it+1} - \epsilon^h_{it}))] \end{aligned} \tag{82}$$

Now plugging in for the variances of the persistent and transitory shocks, we have

$$\begin{aligned} m_{t+1} =& log(\beta) - \gamma E^*_{t+1}[\theta_1(D(Y_{t+1}) - D(Y_t))] \\ &+ \frac{\gamma^2}{2}[\sigma^2_{\eta,t+1} + E^*_{t+1}[\sum_{k=0}^{h-1}(\rho^{k+1} - \rho^k)^2\sigma^2_{\eta,t-k}] + 2\sigma^2_\epsilon] \end{aligned} \tag{83}$$

Where $E^*_{t+1}[\sum_{k=0}^{h-1}(\rho^{k+1} - \rho^k)^2\sigma^2_{\eta,t-k}]$ is the cross-sectional average of the weighted sum

[60] In the case of regime switching skewness, the normality assumption is not valid. However, given that I find no significant difference in the skewness between states in my GMM estimates, I will assume normality for the sake of the asset pricing relationships. Alternatively, I could add a skewness correction as in Martin (2013). This would be quantitatively small and would not substantially impact my empirical results.

[61] This assumption also relies on the fact that the sum of two normal distributions is also a normal distribution.

of past variances for all cohorts at time $t+1$. All of the covariance terms are zero, as the transitory and persistent shocks are i.i.d. and the residual is orthogonal to the aggregate component, by assumption. Furthermore, the deterministic component carries no covariance terms, either.

The variances of the shocks follow from the fact that z_{it} is an AR(1) process, so

$$z_{it+1} - z_{it} = z^0_{it}(\rho^{h+1} - \rho^h) + \eta^h_{it+1} + \sum_{k=0}^{h-1} \eta^{h-k}_{it-k}(\rho^{k+1} - \rho^k) \tag{84}$$

and that $z^0_{it} = 0$ for all i and t, by assumption.

Returning now to the the basic Euler equation we have that

$$E_t[R^j_{t+1} M_{t+1}] = 1 \tag{85}$$

By the law of iterated expectations[62], the following unconditional Euler equation also holds:

$$E[R^j_{t+1} M_{t+1}] = 1 \tag{86}$$

$$\Longrightarrow E[r^j_{t+1}] + E[m_{t+1}] + \frac{1}{2}[var(r^j_{t+1}) + var(m_{t+1}) + 2cov(r^j_{t+1}, m_{t+1})] = 0 \tag{87}$$

Note that *var* and *cov* represent the unconditional moments. This holds for any asset, including the risk free asset:

$$r^f = -E[m_{t+1}] - \frac{1}{2}var(m_{t+1}) \tag{88}$$

assuming that the risk free rate has zero variance and zero covariance with the SDF.

[62] See Lettau and Ludvigson (2009)

Thus, the log risk premium is given by:

$$E[r^m_{t+1}] - r^f_{t+1} = -\frac{1}{2}[var(r^j_{t+1}) + 2cov(r^m_{t+1}, m_{t+1})] \tag{89}$$

$$\implies E[logR^m_{t+1})] - log(R^f_{t+1}) = -cov(r^m_{t+1}, m_{t+1}) \tag{90}$$

Where m_{t+1} is specified by (83). So the risk premium is a function of the market return's covariance with the changes in the aggregate component and the variances of the idiosyncratic components.

In an economy without idiosyncratic risk, the SDF implied by equation (87) is:

$$m_{t+1} = log(\beta) - \gamma E[\Delta c_{t+1}] + \frac{\gamma^2}{2} var(\Delta c_{t+1})) \tag{91}$$

Thus, the risk premium is a function of the return's covariance only with aggregate growth and the variance of aggregate growth. Hence, idiosyncratic risk has the potential to amplify risk premium in this model of consumption.

Expanding out equation (90), while mathematically tedious, yields some insights:

$$\begin{aligned} E[logR^j_{t+1})] - log(R^f_{t+1}) = {} & \gamma cov(r^j_{t+1}, E^*_{t+1}[\theta_1(D(Y_{t+1})) - D(Y_t))]) \\ & - \frac{\gamma^2}{2} cov(r^m_{t+1}, \sigma^2_{,\eta,t+1}) \\ & - \frac{\gamma^2}{2} cov(r^m_{t+1}, E^*_{t+1}[\sum_{k=0}^{h-1}(\rho^{k+1} - \rho^k)^2 \sigma^2_{\eta,t-k}]) \\ & - \gamma^2 cov(r^m_{t+1}, \sigma^2_\epsilon) \end{aligned} \tag{92}$$

Note that the final line of equation (34), $cov(r^m_{t+1}, \sigma^2_\epsilon)$ is equal to zero, as the variance of the transitory shock does not change over time (by assumption). The third line is approaching zero, as the difference between realizations of the cross-sectional average of past shock variances and the sample average of the same object will equal zero in large samples.[63] Thus, the equity premium can be written more

[63]Note that I could also assume a unit root and reach the same conclusion. If using my GMM

succinctly as:

$$
\begin{aligned}
E[logR^j_{t+1})] - log(R^f_{t+1}) = \gamma cov(r^j_{t+1}, E^*_{t+1}[\Delta c_{t+1}]) \\
- \frac{\gamma^2}{2} cov(r^m_{t+1}, \sigma^2_{,\eta,t+1})
\end{aligned}
\tag{93}
$$

The first line, representing the covariance of the return with aggregate growth is what is represented in the typical representative Euler equation. Thus the line of interest is the second, representing the covariance of the return with the cross-sectional variance of the persistent shock's contemporaneous component. When this covariance is negative, the risk premium is larger. Hence, idiosyncratic risk has the potential to amplify risk premia in this model of consumption.[64]

estimates of approximately $\rho = .98$, this term becomes trivially small.

[64]Note the same is true in any model of idiosyncratic consumption/income risk such as Constantinides and Duffie (1996), Heaton and Lucas (1996), Guvenen (2009) and many others.

Figures

Figure I: Residual variance over time, CP consumption (top), CP consumption by cohort (middle) and PSID Income and Consumption (bottom)

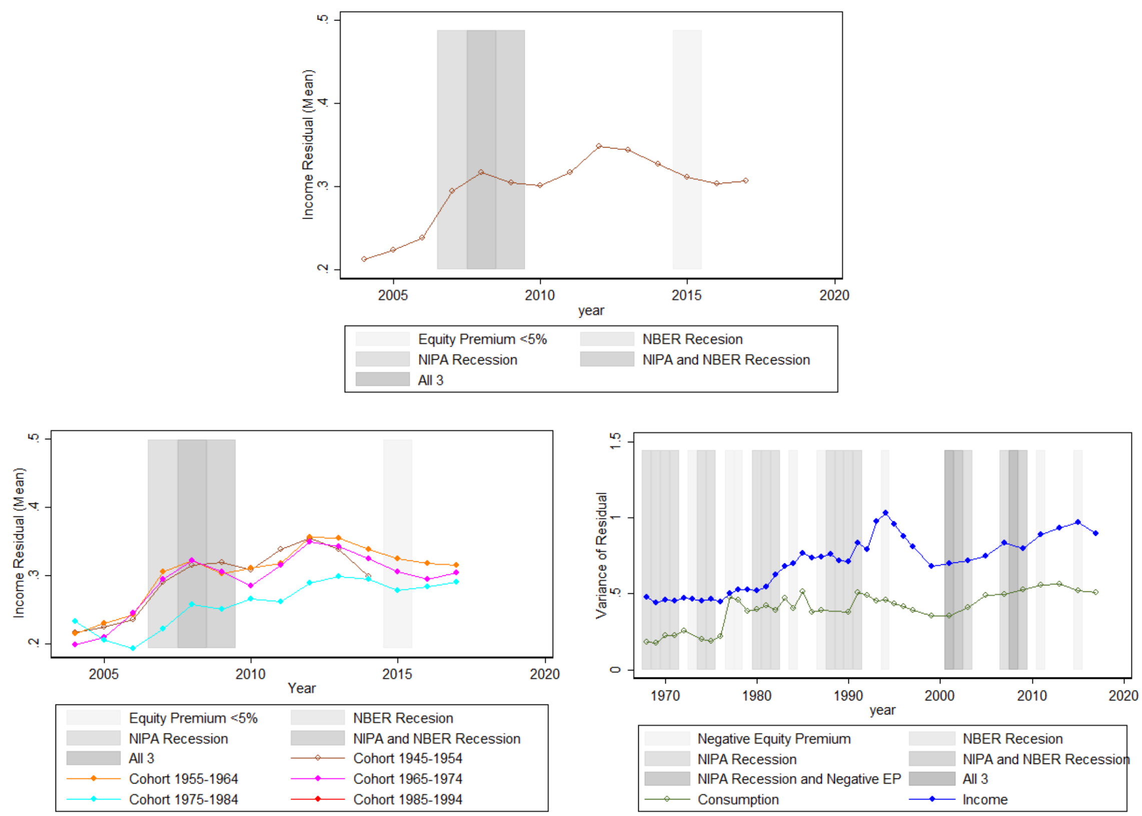

Note: This figure shows the variance of the residual from the regression in equation (1) for consumption in the CP (top), consumption in the CP by cohort (middle) and food consumption and income in the PSID (bottom).

Figure II: GDP Growth, Unemployment Rate, and Excess Return, 1968-2017

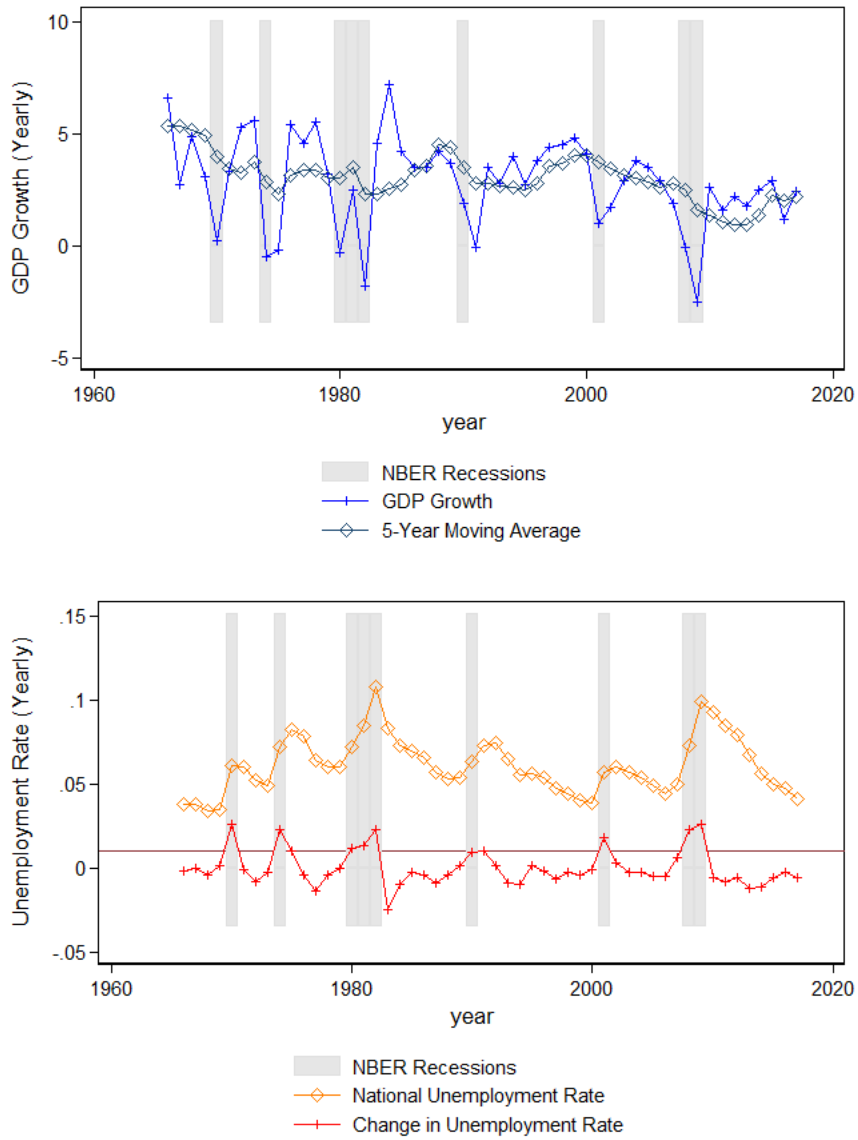

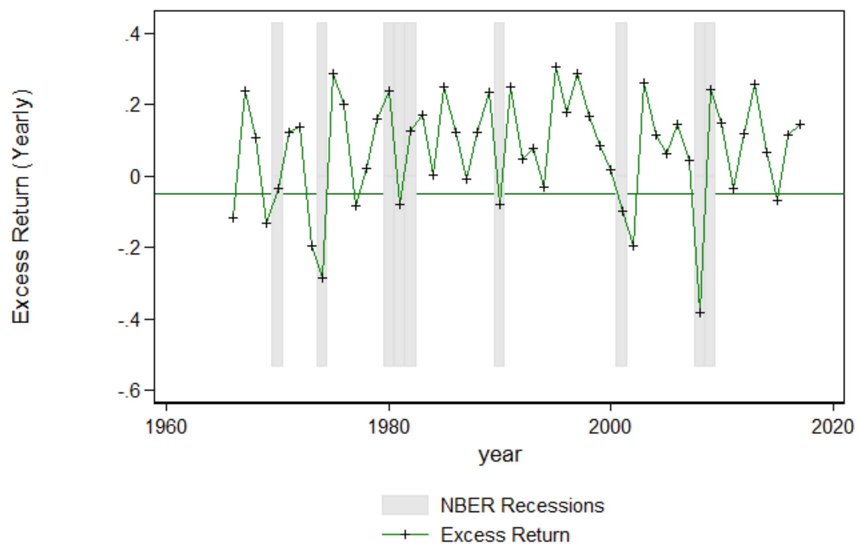

Note: These figures show each of the definitions used to classify years as contraction. The top panel shows GDP growth from 1968-2017 as well as the 5-year moving average. If growth in a given year is below the moving average, it is classified as a contraction. The middle panel shows the unemployment rate over the same period. The year is classified as a contraction if the unemployment rate falls by more than one percent. The bottom panel shows the excess return over the same period. The excess return is calculated as the annual market return (including dividends) of the value-weighted NYSE/AMEX index minus the annualized 30-day T-bill return. A year is classified a contraction if the excess return is less than five percent. Each graph also shows the NBER recession years, shaded in gray.
Sources:

- GDP Growth - Bureau of Economic Analysis: `https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2#reqid=19&step=2&isuri=1&1921=survey`
- Unemployment Rate - Bureau of Labor Statistics: urlhttps://data.bls.gov/timeseries/LNS14000000
- Excess Return: The Center for Research in Security Prices (CRSP)
- NBER Recessions: `https://www.nber.org/cycles.html`

Figure III: Cross-sectional variance of residual consumption by cohort

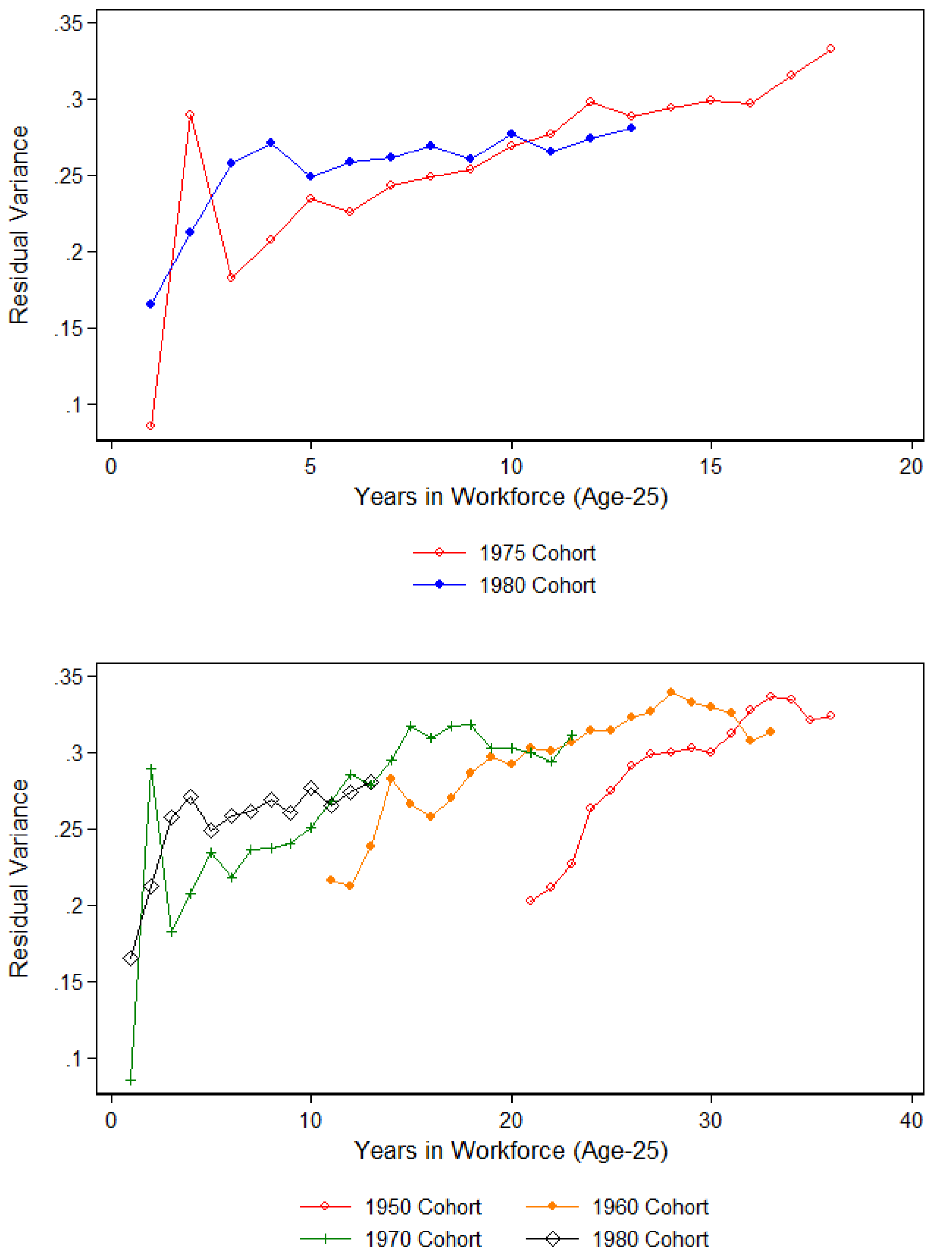

Note: This figure shows the variance of the residual from the regression in equation (1) for different cohorts. Age is normalized such that the x-axis point of zero corresponds to the cohort's first year in the workforce (at age 25) .

Figure IV: Cross-sectional variance of residual consumption by age (top), cross-sectional variance by macroeconomic history (bottom): regression coefficients

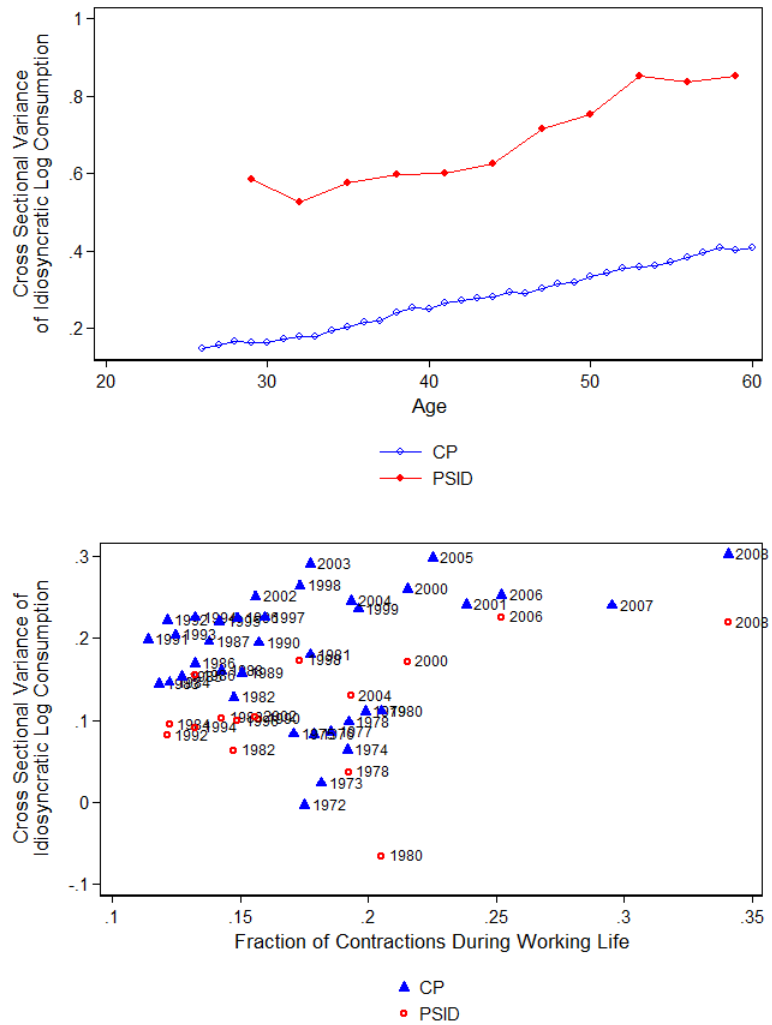

Note: these figures plot the age and cohort coefficients from the regressions specified in (16) for consumption in the CP (blue) and the PSID (red). The top figure controls for “cohort effect" by regressing the cross-sectional residual variance on cohort and age dummies. The coefficients are rescaled to match the level of variance at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels of the data points indicate the year in which the cohort entered the workforce, or when its members were aged 25.

Figure V: Cross-sectional skewness of residual consumption by age (top), cross-sectional skewness by macroeconomic history (bottom): regression coefficients

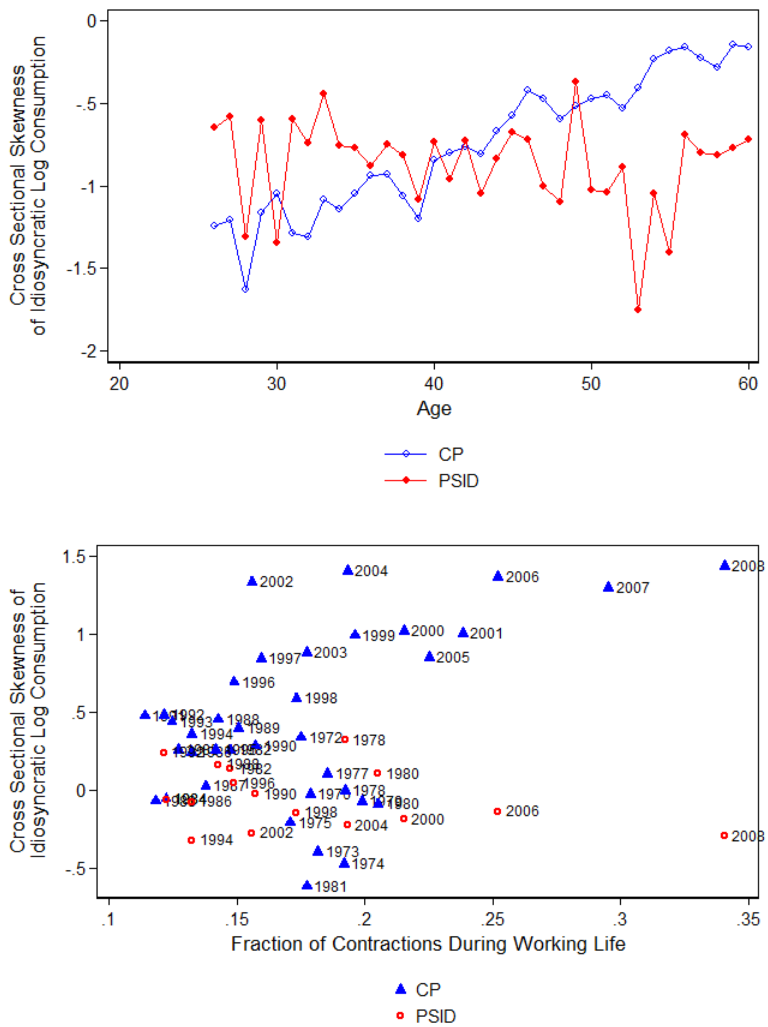

Note: these figures plot the age and cohort coefficients from the regressions specified in (17) for consumption in the CP (blue) and the PSID (red). The top figure controls for “cohort effect" by regressing the cross-sectional residual skewness on cohort and age dummies. The coefficients are rescaled to match the level of skewness at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels of the data points indicate the year in which the cohort entered the workforce, or when its members were aged 25.

Figure VI: Cross-sectional variance of residual income by age (top), cross-sectional variance by macroeconomic history (bottom): regression coefficients

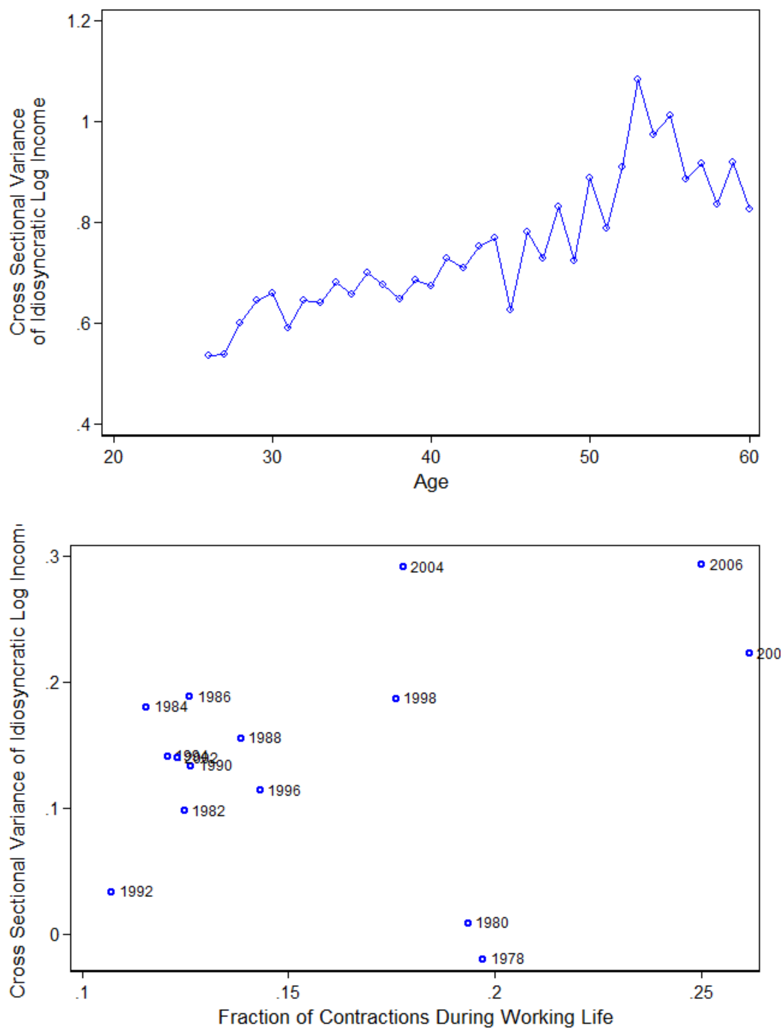

Note: these figures plot the age and cohort coefficients from the regressions specified in (16), with residual income variance in the PSID from 1999-2017 on the left-hand side. The top figure controls for "cohort effect" by regressing the cross-sectional residual skewness on cohort and age dummies. The coefficients are rescaled to match the level of skewness at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions.

Figure VII: Cross-sectional skewness of residual income by age (top), cross-sectional skewness by macroeconomic history (bottom): regression coefficients

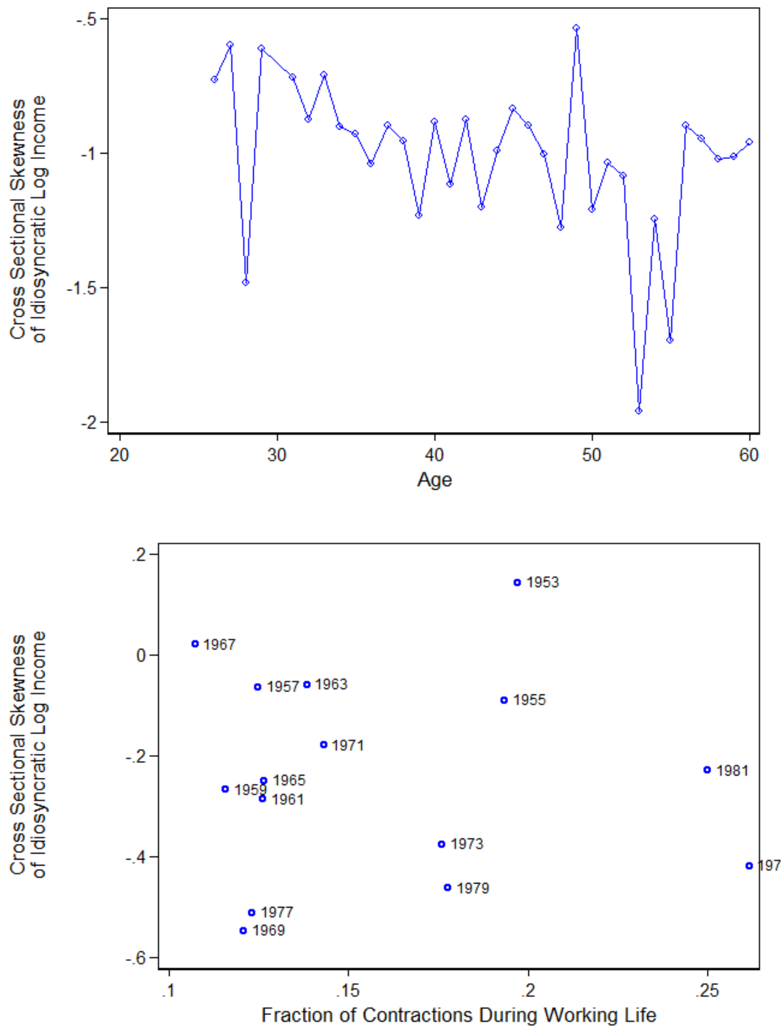

Note: these figures plot the age and cohort coefficients from the regressions specified in (17), with residual income skewness in the PSID from 1999-2017 on the left-hand side. The top figure controls for “cohort effect" by regressing the cross-sectional residual skewness on cohort and age dummies. The coefficients are rescaled to match the level of skewness at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions.

Figure VIII: Cross-sectional variance of residual consumption growth by age (top), Cross-sectional variance by macroeconomic history (bottom)

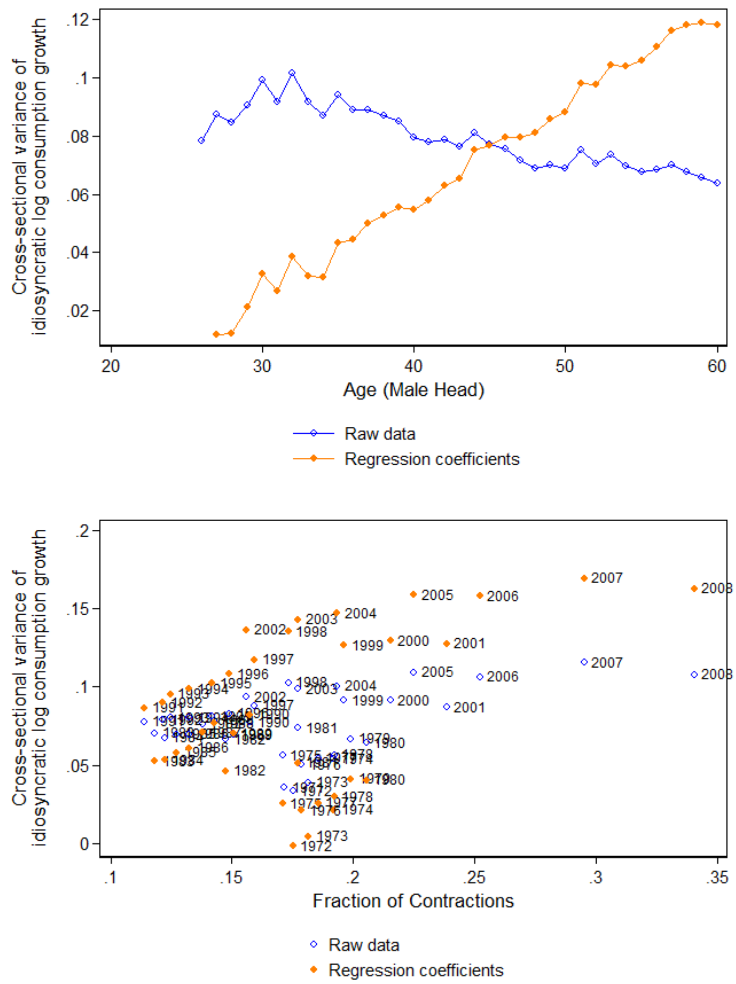

Note: these figures plot the cross-sectional variance by age and macroeconomic history, based on raw data (blue) and the regression specified in (16) (orange) for consumption growth in the CP (weighted using sample weights provided by Nielsen). The blue line in the top figure shows the average of the residual variance by age. The orange line controls for “cohort effect" by regressing the cross-sectional residual variance on cohort and age dummies. The coefficients are rescaled to match the level of variance at age 45. The blue points in the bottom panel show the average of the residual variance against the fraction of contractionary years during which the members of the cohort have worked. The orange points show the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked. Both are based on the NBER definition of contractions. The labels of the data points indicate the year in which the cohort entered the workforce, or when its members were aged 25.

Figure IX: Components of the SDF: Log returns, log aggregate consumption growth, and variance of log idiosyncratic consumption growth (post-war period (left) and in-sample (right))

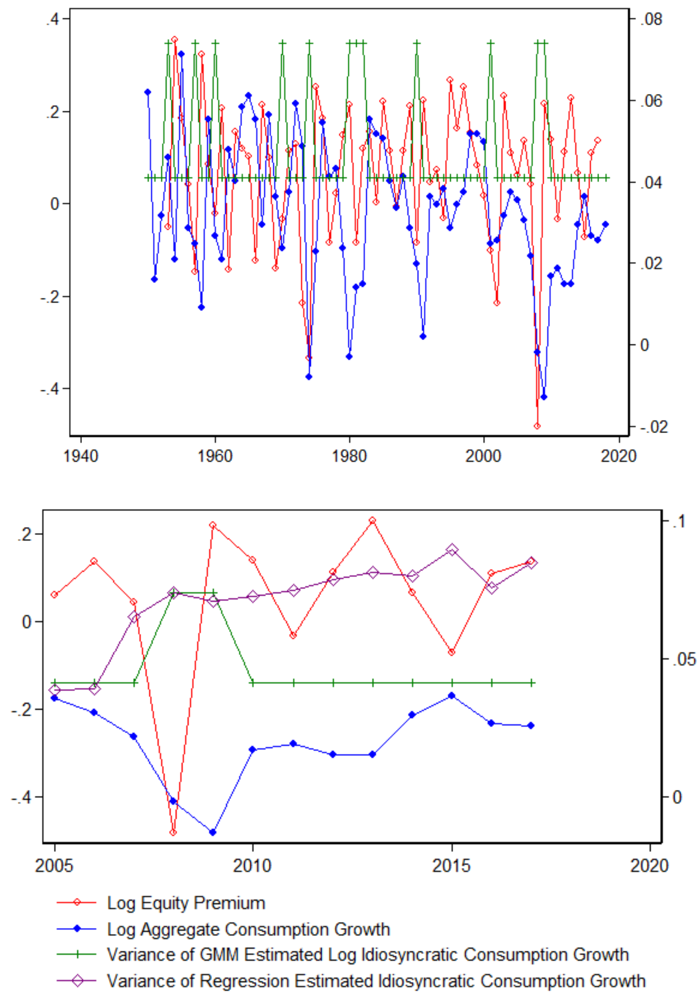

Note: This figure shows the time series of the observed log equity premium, log aggregate consumption growth (based on the CEX), and the variance of log idiosyncratic consumption growth. The top figure shows the entire post-ward period (1950-2017) and the bottom figures shows the series during my sample period (2005-2017). The covariances of these objects are what drives the risk premia in equation (35). The variance of log idiosyncratic consumption growth is calculated in two distinct ways: 1) imputing my GMM estimates for stockholders from Table XI for the variance of the AR(1) innovation (see equation (29)) 2) a more reduced form approach in which I simply take the variance of the residual from the first stage regression in equation 1. See text for details. The equity premium is calculated as the annual market return (including dividends) of the value-weighted NYSE/AMEX index minus the annualized 30-day T-bill return

Tables

Table I: Summary Statistics by Year - Nielsen Consumer Panel

Year	Number of Observations	Age (Mean)	Education Level (Median)	Income Level (Mean $)	Consumption (Mean $)
2004	10,921	46	Some College	60,000-69,999	4,257
2005	13,009	45	Some College	45,000-49,999	4,151
2006	15,207	46	Some College	60,000-69,999	4,050
2007	22,949	45	Some College	60,000-69,999	4,357
2008	23,712	45	Some College	60,000-69,999	4,435
2009	24,236	45	Some College	60,000-69,999	4,355
2010	23,538	46	Some College	60,000-69,999	4,362
2011	23,582	46	Some College	60,000-69,999	4,520
2012	23,001	46	Some College	60,000-69,999	4,778
2013	22,795	47	Some College	60,000-69,999	4,707
2014	22,339	47	Some College	60,000-69,999	4,756
2015	22,043	47	Some College	70,000-99,999	4,653
2016	19,092	48	Some College	70,000-99,999	4,752
2017	16,125	49	Some College	70,000-99,999	4,547

This table shows summary statistics by year in the estimation sample of the Nielsen Consumer Panel, the construction of which is described in Section 2.1. Education levels and income levels are given as the median bucket for the Nielsen categories. Consumption is calculated as the sum of all purchases across the given year. All data is based on the head of household. Data are weighted using the Nielsen projection weights. All dollar figures are in 2009 dollars.

Table II: Summary Statistics by Subpanel - Nielsen Consumer Panel

	Number of Participants (Number of Observations ÷3)	Age (Mean)	Education Level (Median)	Income Level (Mean $)	Consumption (Mean $)
2004	10,922	46	Some College	50,000-59,999	4,206
2005	11,354	46	Some College	50,000-59,999	4,129
2006	11,493	46	Some College	50,000-59,999	3,999
2007	18,738	45	Some College	50,000-59,999	4,389
2008	17,302	46	Some College	50,000-59,999	4,451
2009	16,167	46	Some College	50,000-59,999	4,514
2010	16,702	46	Some College	50,000-59,999	4,645
2011	17,631	47	Some College	50,000-59,999	4,779
2012	16,807	47	Some College	50,000-59,999	4,833
2013	16,455	47	Some College	50,000-59,999	4,811
2014	16,314	48	Some College	50,000-59,999	4,829
2015	16,125	48	Some College	50,000-59,999	4,769

This table shows summary statistics for the three-year subpanels, the construction of which is described in Section 2.1. Each row corresponds with the subpanel beginning in the year indicated. Education levels and income levels are given as the median bucket for the Nielsen categories. Consumption is calculated as the sum of all purchases across the given year. All data is based on the head of household. Data are weighted using the Nielsen projection weights. All dollar figures are in 2009 dollars.

Table III: Summary Statistics by Year - PSID

Year	Number of Observations	Age (Mean)	Education Level (Median Years)	Income (Mean $)	Consumption (Mean $)	Food Consumption (Mean $)
1999	3,299	37	12	57,948	21.721	7,482
2001	3,898	38	12	59,297	22,724	7,690
2003	4,554	39	12	57,132	23,043	7,486
2005	4,908	40	12	57,232	23,024	7,576
2007	5,271	41	12	56,909	23,772	8,348
2009	5,635	41	13	58,020	21,395	6,948
2011	5,616	43	13	52,457	19,752	6,799
2013	5,297	43	13	52,176	20,332	6,841
2015	4,490	45	13	54,717	21,196	7,048
2017	3,852	46	13	58,641	21,528	7,785

This table shows summary statistics by year in the estimation sample of the Panel Study of Dynamic Income, the construction of which is described in Section 2.2. Education levels are based on the number of years of education of the head of household. Consumption is the sum of all consumption expenditure categories in the PSID, excluding housing expenditures. All data is based on the head of household. Data are weighted using the PSID weights. All dollar figures are in 2009 dollars.

Table IV: Comparison of demographics across datasets

	SCF	CP - Full Sample	CP - Estimation Sample	PSID - Full Sample	PSID - Estimation Sample
Income (Mean)	92,200	59,909[2]	72,752	66,047	71,302
Income (Median)	54,100	54,000	64,000	47,269	53,247
Age	51	53	49	45	40
% College Educated	28.4	30.2	40.6	25.7	26.0
% White	72.2	82.9	81.7	58.8	56.4
% Holding Stock	51.2	16.2[3]	15.8	17.5	15.6

[1] SCF data is averaged over 2006-2016. Nielsen data is averaged over 2004-2017. PSID data is averaged over 2001-2017 (biannually).

[2] Nielsen income values are based on midpoints of income bins provided by Nielsen.

[3] Nielsen stock ownership data is based on a probit model fitted to the PSID, then applied to the Nielsen data. I predict stock ownership as function of age, income, education, and year effects. The results of this regression are shown in Appendix Table A7.

Table V: Distribution of Consumption Types by Income, CEX

Income	Percent distribution of consumer units	Goods	Durables	Non-durables	Services
Less than $5,000	4.0	0.45	0.10	0.35	0.55
$5,000 to $9,999	4.4	0.52	0.10	0.41	0.47
$10,000 to $14,999	6.7	0.50	0.10	0.40	0.47
$15,000 to 19,999	6.8	0.50	0.12	0.38	0.46
$20,000 to $29,999	12.2	0.53	0.13	0.40	0.47
$30,000 to $39,999	10.8	0.53	0.14	0.38	0.47
$40,000 to $49,999	9.5	0.51	0.14	0.37	0.47
$50,000 to $69,999	14.3	0.48	0.13	0.35	0.51
$70,000 and more	31.5	0.44	0.14	0.30	0.56

This table shows the distribution of consumption types across income groups. Based on CEX data from 2010.

Table VI: Correlation Matrix Between Types of Contractions

	NIPA	NBER	Unemployment Increase	Excess Return $< 5\%$
NIPA	1.0000			
NBER	.5533	1.0000		
Unemployment Increase	.5155	.9318	1.0000	
Excess Return $< 5\%$	.1324	. 4115	.3584	1.0000

Table VII: First Stage Regression Results: Log Consumption

	Log of Consumption
Black	-0.2044
	(0.0062)
Asian	-0.2105
	(0.0086)
Other	-0.1020
	(0.0069)
Age (Male Head)	-0.0696
	(0.0169)
Age Squared	0.0017
	(0.0004)
Age Cubed	-0.0000
	(0.0000)
Education	-0.0007
	(0.0004)
Household Size	0.1354
	(0.0013)
Initial residual variance of food consumption	0.1773
	(0.0224)
Constant	8.5850
	(0.2457)
Observations	282,549
Adjusted R^2	0.128

This table shows the results of the regression specified in (1) The regression also includes year dummies and state dummies, which are omitted here for brevity. Standard errors are in parentheses.

Table VIII: First Stage Regression Results: Consumption Growth

	Log consumption growth
Δ age	0.0000
	(0.0000)
Δ age-squared	-0.0003
	(0.0007)
Δ age-cubed	0.0000
	(0.0000)
Δ education	0.0013
	(0.0005)
Δ family size	0.0125
	(0.0016)
Constant	0.0520
	(0.0410)
Observations	234,286
Adjusted R^2	0.012

This table shows the results of the regression specified in (20) The regression also includes year dummies and a dummy equal to one if the individual has moved states from time t to $t+1$, which are omitted here for brevity. Standard errors are in parentheses.

Table IX: GMM Results: CP & PSID Consumption

	CP			PSID		
	NBER Contractions	NIPA Contractions	Negative Returns	NBER Contractions	NIPA Contractions	Negative Returns
$\hat{\rho}$	.9876	.9846	.9825	.9835	.9670	.9861
	(.0057)	(.0040)	(.0051)	(.0212)	(.0233)	(.0211)
$\hat{\mu_\alpha^2}$	.1203	.0966	.1152	.2566	.2001	.2131
	(.0032)	(.0057)	(.0036)	(.0091)	(.0103)	(.0095)
$\hat{\mu_\epsilon^2}$	.0503	.0411	.0527	.1203	.1102	.1008
	(.0038)	(.0050)	(.0042)	(.0152)	(.0202)	(.0189)
$\hat{\mu_{\eta,E}^2}$	.0668	.0631	.0589	.0378	.0593	.0601
	(.0041)	(.0026)	(.0052)	(.0124)	(.0110)	(.0092)
$\hat{\mu_{\eta,C}^2}$	.0853	.1012	.0973	.0641	.1670	.1666
	(.0049)	(.0072)	(.0064)	(.0182)	(.0095)	(.0134)
$\hat{\mu_\alpha^3}$	-.8998	-.8950	-.9621	-.7904	-.8122	-.8306
	(.0017)	(.0013)	(.0015)	(.0043)	(.0049)	(.0059)
$\hat{\mu_\epsilon^3}$	-.6997	-.5432	-.6852	-.1500	-.1497	-.1512
	(.0071)	(.0073)	(.0071)	(.0292)	(.0284)	(.0233)
$\hat{\mu_{\eta,E}^3}$	-.0038	.0032	-.0034	-.0048	-.0049	-.0048
	(.0049)	(.0044)	(.0046)	(.0092)	(.0130)	(.0098)
$\hat{\mu_{\eta,C}^3}$	-.0043	.0047	-.0047	-.0050	-.0051	-.0053
	(.0010)	(.0025)	(.0009)	(.0041)	(.0050)	(.0042)
p-value	.9999	.9999	.9999	.9120	.8974	.9054

This table shows the GMM estimation results from the process in (3) and (2) using the CP data on consumption from 2004-2017 (first three columns) and PSID data on consumption from 1999-2017 (last three columns). The parameters $\mu_{\eta,C}^2$ and $\mu_{\eta,E}^2$ denote the conditional standard deviation of the AR(1) innovation of the persistent component, conditional on the aggregate state, E (expansion) or C (contraction). $\mu_{\eta,S}^3$ are the same, but for conditional skewness. Expansions and contractions are defined based on NBER recession indicators, NIPA GDP growth indicators, or the level of real stock market returns, as indicated by the column headers. Standard errors, shown in parentheses, are calculated using the Newey-West method. The moments used are defined in (15) and (10), collapsed across eight age groups for each of the 14 years of available data in the CP and 10 years of available data in the PSID. The p-value is from the overidentifying test of moment conditions. For further details on the GMM estimation, see Appendix C.

Table X: GMM Results: PSID Income

	1999-2017	1968-1993
$\hat{\rho}$	.9639	.943
	(.0110)	-
$\hat{\mu_\alpha^2}$	.3502	.150
	(.0089)	-
$\hat{\mu_\epsilon^2}$	.2003	.064
	(.0143)	-
$\hat{\mu_{\eta,E}^2}$	.0362	.014
	(.0064)	-
$\hat{\mu_{\eta,C}^2}$	.0669	.040
	(.0182)	-
$\hat{\mu_\alpha^3}$	-.5503	-
	(.0025)	-
$\hat{\mu_\epsilon^3}$	-.2533	-
	(0137)	-
$\hat{\mu_{\eta,E}^3}$	-.0007	-
	(.0051)	-
$\hat{\mu_{\eta,C}^3}$	-.0053	-
	(.0010)	-

This table shows the GMM estimation results from the process in (2) and (3) using PSID data on income. The first column shows my estimation results using data from 1999-2017. The second column duplicates the results from Row D, using NBER indicators, of Table 2 in Storesletten, Telmer, and Yaron (2004b). STY report standard deviations, thus I square their estimates in order to get variances and omit the standard errors. The parameters $\mu_{\eta,C}^2$ and $\mu_{\eta,E}^2$ denote the conditional standard deviation of the AR(1) innovation of the persistent component, conditional on the aggregate state, E (expansion) or C (contraction). $\mu_{\eta,S}^3$ are the same, but for conditional skewness. Expansions and contractions are defined based on NBER recession indicators. Standard errors, shown in parentheses, are calculated using the Newey-West method. The moments used are defined in (15) and (10), collapsed across eight age groups for each of the 10 years of available data in the PSID. The p-value is from the overidentifying test of moment conditions. For further details on the GMM estimation, see Appendix C.

Table XI: GMM Results: Subsamples

	CP		
	High Education	High Income	Stockholders
$\hat{\rho}$	.9826	.9702	.9687
	(.0058)	(.0083)	(.0084)
$\hat{\mu_{\alpha}^2}$	.0752	.1001	.0466
	(.0030)	(.0012)	(.0014)
$\hat{\mu_{\epsilon}^2}$	.0302	.0344	.0216
	(.0039)	(.0022)	(.0020)
$\hat{\mu_{\eta,E}^2}$	.0590	.0430	.0411
	(.0061)	(.0061)	(.0059)
$\hat{\mu_{\eta,C}^2}$	.0971	.0762	.0741
	(.0056)	(.0021)	(.0019)
$\hat{\mu_{\alpha}^3}$	-.7211	-.7522	-.5291
	(.0017)	(.0026)	(.0027)
$\hat{\mu_{\epsilon}^3}$	-.3902	-.4311	-.2502
	(.0066)	(.0059)	(.0061)
$\hat{\mu_{\eta,E}^3}$	-.0040	-.0045	-.0037
	(.0048)	(.0063)	(.0066)
$\hat{\mu_{\eta,C}^3}$	-.0044	-.0049	-.0041
	(.0010)	(.0014)	(.0014)

This table shows the GMM estimation results from the process in (3) and (2) using Nielsen data on consumption from 2004-2017. The sample is limited to those who have at least a college education, at least 100,000 dollars of income, or are predicted to be stockholders. The parameters $\mu_{\eta,C}^2$ and $\mu_{\eta,E}^2$ denote the conditional standard deviation of the AR(1) innovation of the persistent component, conditional on the aggregate state, E (expansion) or C (contraction). $\mu_{\eta,S}^3$ are the same, but for conditional skewness. Expansions and contractions are defined based on NBER recession indicators. Standard errors, shown in parentheses, are calculated using the Newey-West method. The moments used are defined in (15) and (10), collapsed across eight age groups for each of the 14 years of available data in the CP. The p-value is from the overidentifying test of moment conditions. For further details on the GMM estimation, see Appendix C.

Table XII: Implied Equity Premium for Various Levels of Risk Aversion

	Post-War: 1950-2017		In-Sample: 2005-2017		
Risk Aversion	GMM Estimated	No Idiosyncratic Risk	GMM Estimated	Regression Estimated	No Idiosyncratic Risk
2	0.20	0.04	0.28	0.13	0.08
5	1.12	0.09	1.46	0.52	0.20
10	4.29	0.18	5.43	1.68	0.40
15	9.52	0.27	11.90	3.47	0.60
20	16.8	0.36	20.90	5.91	0.81

Note: This table shows the equity premium implied by equation (34):

$$E[logR^j_{t+1})] - log(R^f_{t+1}) = \gamma cov(r^j_{t+1}, E^*_{t+1}[\Delta c_{t+1}]) - \frac{\gamma^2}{2} cov(r^m_{t+1}, \sigma^2_{,\eta,t+1}) \tag{94}$$

for various levels of risk aversion, γ. The calculations labeled as "GMM Estimated" are calculated by imputing my GMM estimates for stockholders in Table XI for $\sigma^2_{,\eta,t+1}$ in the equation above. The calculations labeled as "Regression Estimated" are calculated by taking the cross-sectional variance of the residual from the first-stage regression in equation (1) and imputing this value for $\sigma^2_{,\eta,t+1}$. The calculations labeled as "No Idiosyncratic Risk" are calculated by eliminating the second term from the above equation and simply taking the covariance of the return with aggregate consumption growth. See text for more details.

Appendix Tables and Figure

Figure A1: Cross-sectional variance of residual consumption by cohort, Cohorts Observed at Age 25 Entering the Workforce in an Expansion (top) or a Contraction (bottom)

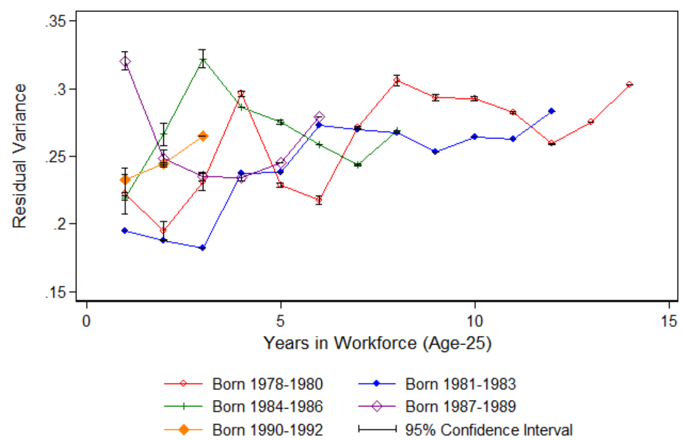

Note: This figure shows the cross-sectional variance of residual consumption from the regression in (1) conditional on cohort, for cohorts that I observe at age 25 during my sample period. This includes all individuals born from 1978-1992. They are then bucketed into 3-year cohort groups. The x-axis is the number of years that a cohort has been working, or their age minus 25. The top panel shows cohorts which entered the workforce during an expansion (2004, 2005, 2006, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017). The bottom panel shows those who entered the workforce during a contraction (2007, 2008, 2009).

Figure A2: Cross-sectional variance by macroeconomic history: regression coefficients

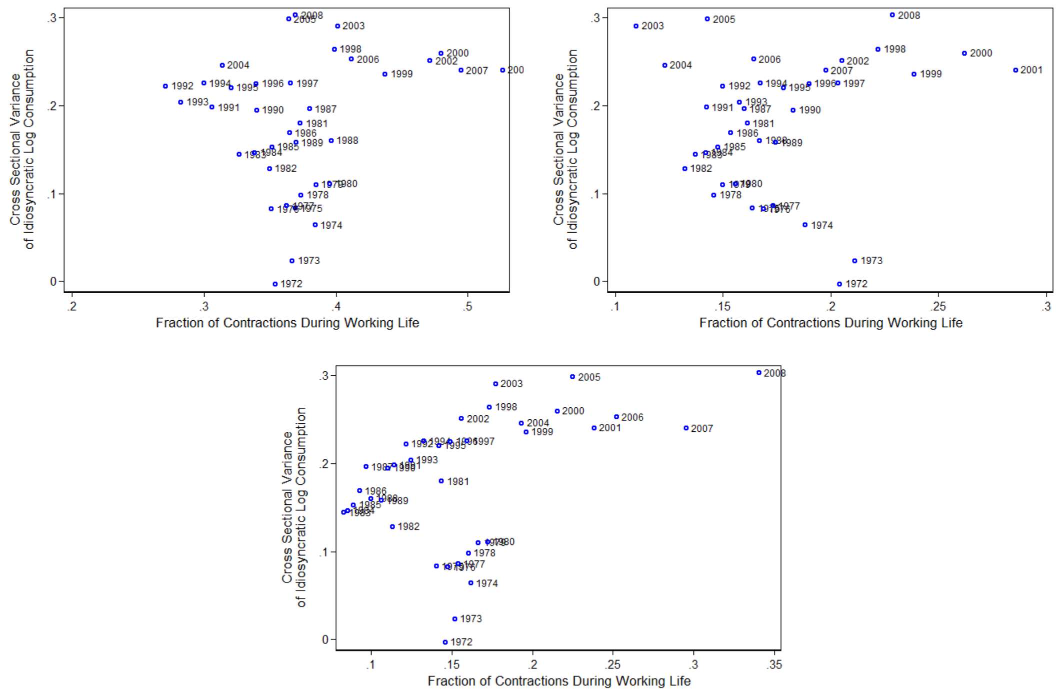

Note: these figures plot the cohort coefficients from the regressions specified in (16) against the fraction of contractionary years during which the members of the cohort have worked. The labels indicate the year in which the cohort entered the workforce, at age 25. All data are from the CP, 2004-2017. The top figure is based on the NIPA definition of contractions. The middle is based on the equity premium being less than five percent. The bottom is based on the unemployment rate rising by more than one percent.

Figure A3: Cross-sectional skewness by macroeconomic history: regression coefficients

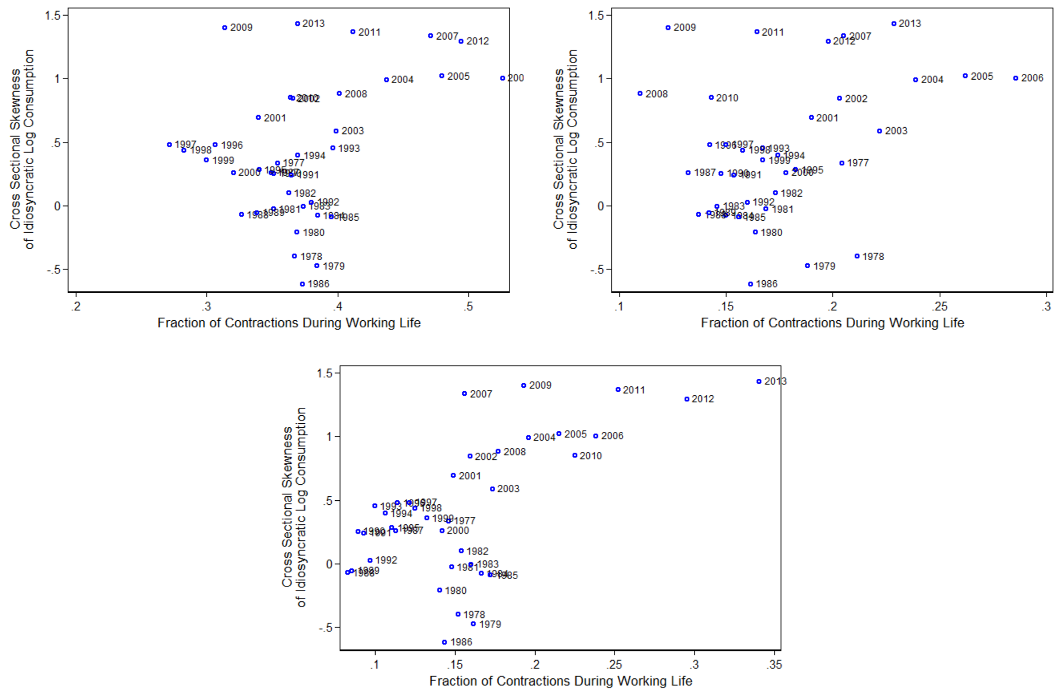

Note: These figures plot the cohort coefficients from the regressions specified in (17) against the fraction of contractionary years during which the members of the cohort have worked. The labels indicate the year in which the cohort entered the workforce, at age 25. All data are from the CP, 2004-2017. The top figure is based on the NIPA definition of contractions. The middle is based on the equity premium being less than five percent. The bottom is based on the unemployment rate rising by more than one percent.

Figure A4: High Education Sample: Cross-sectional variance of residual consumption by age (top), cross-sectional variance by macroeconomic history (bottom): regression coefficients

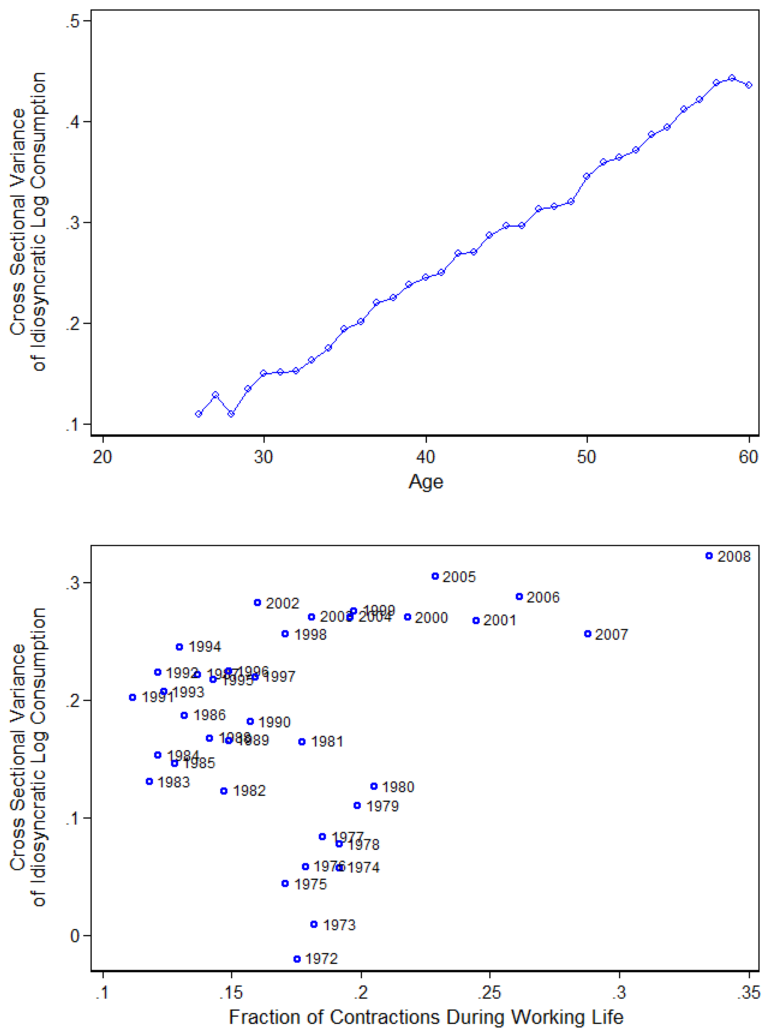

Note: These figures plot the age and cohort coefficients from the regressions specified in (16), run on only those with a college education or higher in the CP, 2004-2017. The top figure controls for “cohort effect" by regressing the cross-sectional residual variance on cohort and age dummies. The coefficients are rescaled to match the level of variance at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels indicate the year in which the cohort entered the workforce, at age 25.

Figure A5: High Education Sample: Cross-sectional skewness of residual consumption by age (top), cross-sectional skewness by macroeconomic history (bottom): regression coefficients

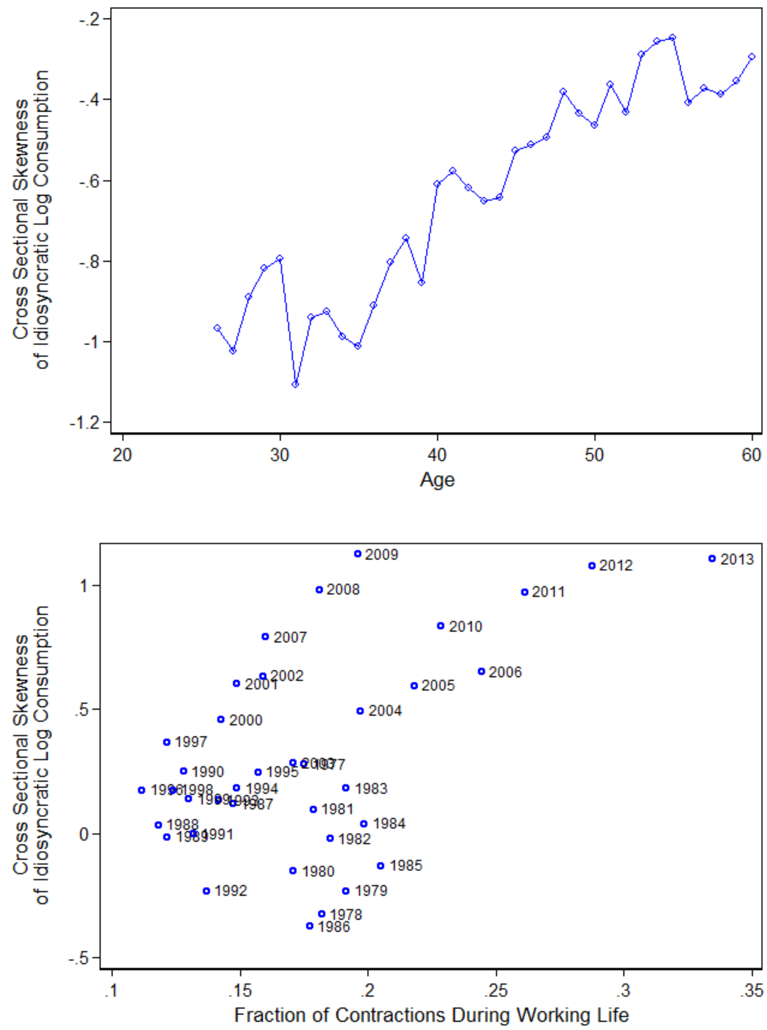

Note: These figures plot the age and cohort coefficients from the regressions specified in (17), run on only those with a college education or greater in the CP, 2004-2017. The top figure controls for “cohort effect" by regressing the cross-sectional residual skewness on cohort and age dummies. The coefficients are rescaled to match the level of skewness at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels indicate the year in which the cohort entered the workforce, at age 25.

Figure A6: High Income Sample: Cross-sectional variance of residual consumption by age (top), cross-sectional variance by macroeconomic history (bottom): regression coefficients

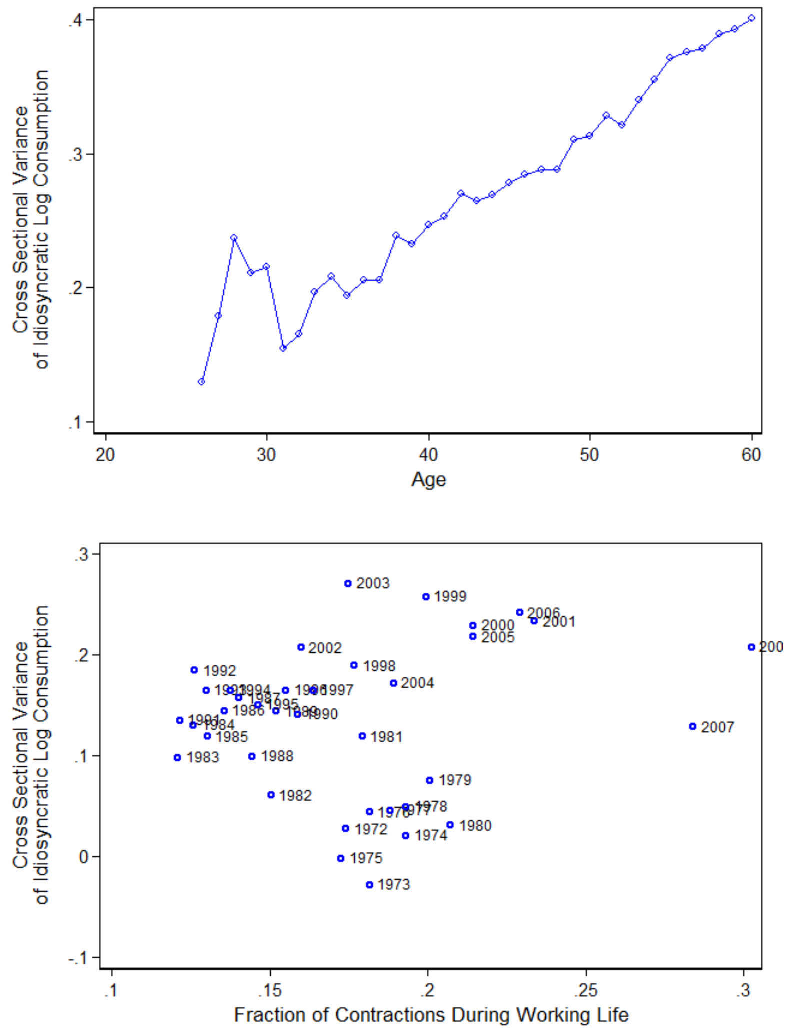

Note: these figures plot the age and cohort coefficients from the regressions specified in (16), run on only those with income greater than or equal to 100,000 dollars in the CP, 2004-2017. The top figure controls for "cohort effect" by regressing the cross-sectional residual variance on cohort and age dummies. The coefficients are rescaled to match the level of variance at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels indicate the year in which the cohort entered the workforce, at age 25.

Figure A7: High Income Sample: Cross-sectional skewness of residual consumption by age (top), cross-sectional skewness by macroeconomic history (bottom): regression coefficients

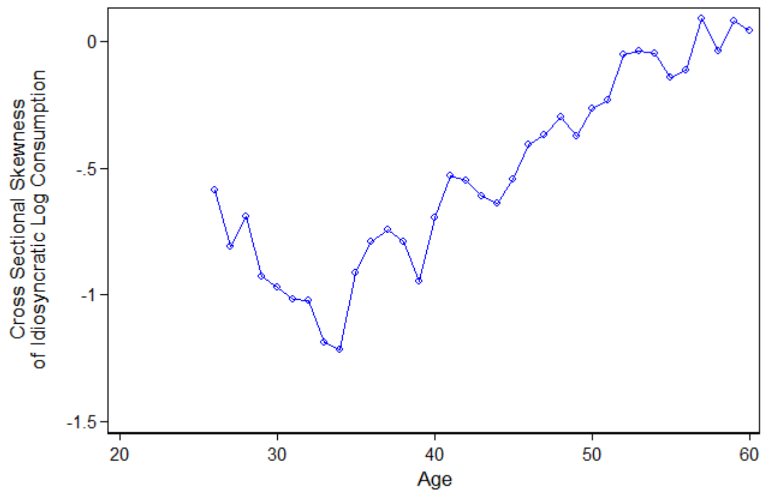

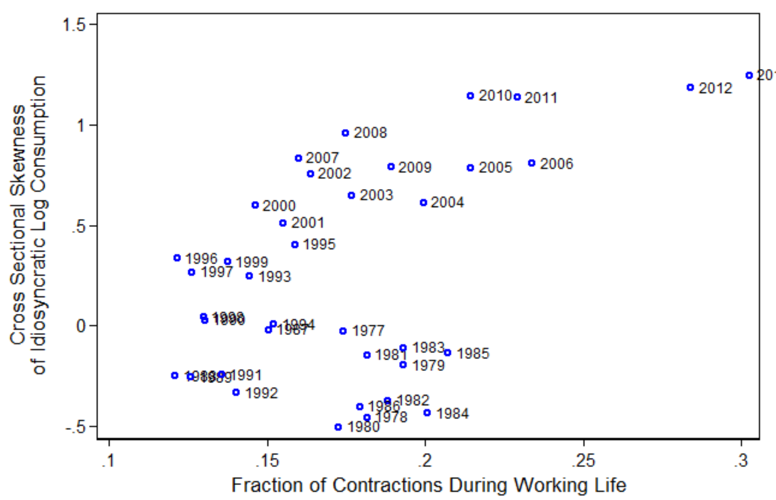

Note: these figures plot the age and cohort coefficients from the regressions specified in (17), run on only those with an income greater than or equal to 100,000 dollars in the CP, 2004-2017. The top figure controls for "cohort effect" by regressing the cross-sectional residual skewness on cohort and age dummies. The coefficients are rescaled to match the level of skewness at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels indicate the year in which the cohort entered the workforce, at age 25.

Figure A8: Predicted Stockholder Sample: Cross-sectional variance of residual consumption by age (top), cross-sectional variance by macroeconomic history (bottom): regression coefficients

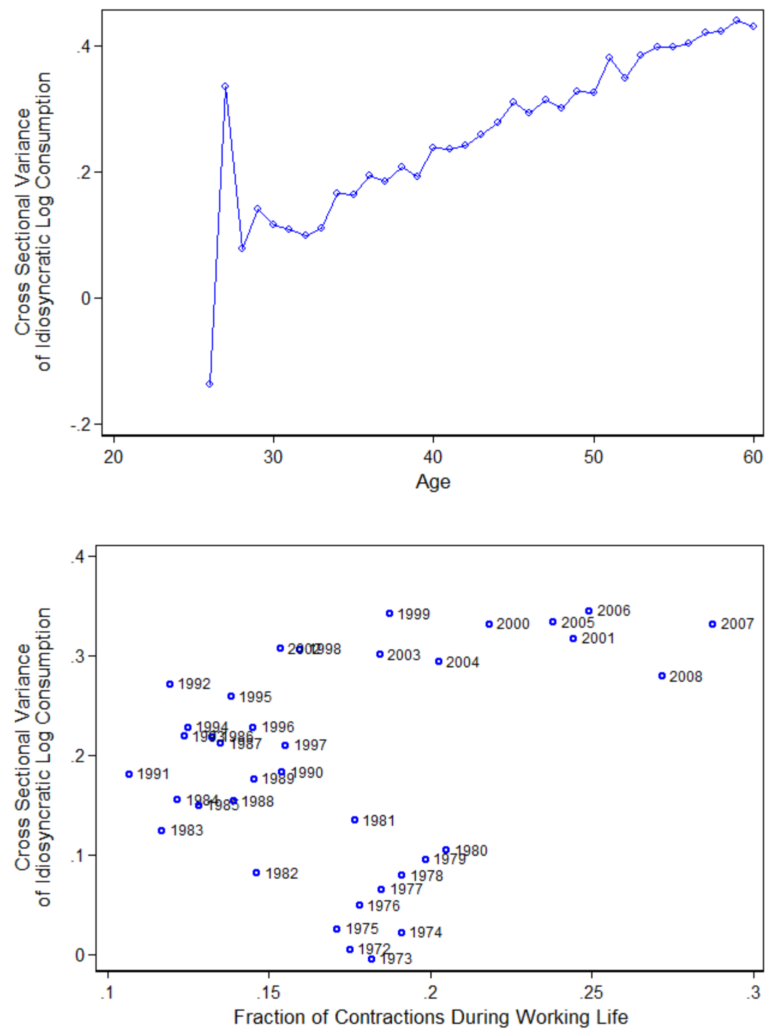

Note: these figures plot the age and cohort coefficients from the regressions specified in (16), run on only those who are predicted to be stockholders in the CP, 2004-2017. The top figure controls for “cohort effect" by regressing the cross-sectional residual variance on cohort and age dummies. The coefficients are rescaled to match the level of variance at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels indicate the year in which the cohort entered the workforce, at age 25.

Figure A9: Predicted Stockholder Sample: Cross-sectional skewness of residual consumption by age (top), cross-sectional skewness by macroeconomic history (bottom): regression coefficients

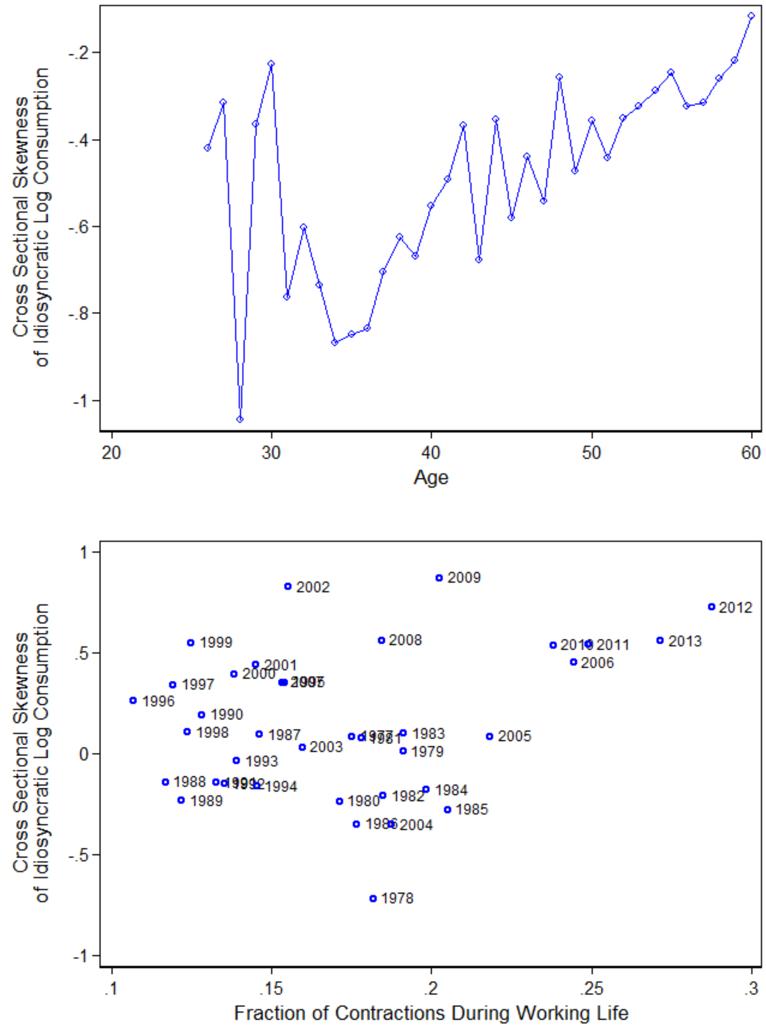

Note: these figures plot the age and cohort coefficients from the regressions specified in (17), run on only those who are predicted to be stockholders in the CP, 2004-2017. The top figure controls for "cohort effect" by regressing the cross-sectional residual skewness on cohort and age dummies. The coefficients are rescaled to match the level of skewness at age 45. The bottom figure plots the cohort coefficients from the same regression against the fraction of contractionary years during which the members of the cohort have worked, based on the NBER definition of contractions. The labels indicate the year in which the cohort entered the workforce, at age 25.

Table A1: Classification of Contractions by Year

Year	NIPA	NBER	Unemployment Increase	Excess Return $< -5\%$
1969	✓			✓
1970	✓	✓	✓	
1971	✓			
1972				
1973				✓
1974	✓	✓	✓	✓
1975	✓			
1976				
1977				✓
1978				
1979				
1980	✓	✓	✓	
1981	✓	✓	✓	✓
1982	✓	✓	✓	
1983				
1984				
1985				
1986				
1987				
1988	✓			
1989	✓			
1990	✓	✓		✓
1991	✓			
1992				✓
1993				
1994				
1995				
1996				
1997				
1998				
1999				
2000				
2001	✓	✓	✓	✓
2002				✓
2003	✓			
2004				
2005				
2006				
2007	✓			

Year	NIPA 5 year	NBER	Unemployment Increase	Excess Return $< -5\%$
2008	✓	✓	✓	✓
2009	✓	✓	✓	
2010				
2011				
2012				
2013				
2014				
2015				✓
2016				
2017				

Table A2: Cohorts Observed at Age 25

Cohort	Total Observations	Total Observations at Age 25: All Years	Expansions	Contractions
1978-1980	12,956	53	53	0
1981-1983	8,920	152	23	129
1984-1986	4,881	125	83	42
1987-1989	2,29	132	132	0
1990-1992	544	193	193	0

This table shows the number of observations in each cohort that I observe at age 25 during my sample period in the CP, from 2004-2017. The number of observations at age 25 is much less than the total observations because many people do not enter the CP until they are older than 25. The last two columns split the observations at age 25 by whether or not the year they entered the workforce was an expansion or a contraction.

Table A3: First Stage Regression Results, Without Controlling for the Variance of Initial Food Consumption: Log Consumption

	Log of Consumption
Black	-0.2061
	(0.0062)
Asian	-0.2111
	(0.0087)
Other	-0.1007
	(0.0069)
Age (Male Head)	-0.0881
	(0.0168)
Age Squared	0.0022
	(0.0004)
Age Cubed	-0.0000
	(0.0000)
Education	-0.0008
	(0.0004)
Household Size	0.1354
	(0.0013)
Constant	8.9377
	(0.2440)
Observations	282,562
Adjusted R^2	0.127

This table shows the results of the regression specified in (1), without the control for initial food consumption variance. The regression also includes year dummies and state dummies, which are omitted here for brevity. Standard errors are in parentheses.

Table A4: First Stage Regression Results, by Education Groups: Log Consumption

	(1) Less than High School Diploma	(2) High School Diploma or Greater
Black	-0.2320	-0.1983
	(0.0172)	(0.0067)
Asian	-0.2666	-0.1954
	(0.0416)	(0.0088)
Other	-0.0892	-0.1025
	(0.0186)	(0.0075)
Age (Male Head)	-0.0960	-0.0846
	(0.0449)	(0.0181)
Age Squared	0.0024	0.0021
	(0.0010)	(0.0004)
Age Cubed	-0.0000	-0.0000
	(0.0000)	(0.0000)
Education	0.0075	-0.0105
	(0.0011)	(0.0007)
Household Size	0.0922	0.1388
	(0.0040)	(0.0014)
Constant	9.1305	9.0309
	(0.6452)	(0.2629)
Observations	24547	258015
Adjusted R^2	0.097	0.135

This table shows the results of the regression specified in (1) for different education groups. The regression also includes year dummies and state dummies, which are omitted here for brevity. Standard errors are in parentheses.

Table A5: First Stage Regression Results, by Income Groups: Log Consumption

	(1) Income < $35,000	(2) Income $ 35,000-99,999	(3) Income >= $100,000
Black	-0.1812	-0.2078	-0.2073
	(0.0151)	(0.0078)	(0.0135)
Asian	-0.2100	-0.1912	-0.2719
	(0.0384)	(0.0124)	(0.0126)
Other	-0.0984	-0.0983	-0.0752
	(0.0170)	(0.0088)	(0.0144)
Age (Male Head)	-0.1231	-0.0648	-0.1268
	(0.0384)	(0.0211)	(0.0398)
Age Squared	0.0027	0.0016	0.0030
	(0.0009)	(0.0005)	(0.0009)
Age Cubed	-0.0000	-0.0000	-0.0000
	(0.0000)	(0.0000)	(0.0000)
Education	-0.0095	-0.0041	-0.0056
	(0.0008)	(0.0005)	(0.0009)
Household Size	0.1490	0.1252	0.1093
	(0.0037)	(0.0016)	(0.0026)
Constant	9.5430	8.6951	9.7657
	(0.5475)	(0.3057)	(0.5860)
Observations	37552	184723	60287
Adjusted R^2	0.150	0.119	0.102

This table shows the results of the regression specified in (1) for different income groups. The regression also includes year dummies and state dummies, which are omitted here for brevity. Standard errors are in parentheses.

Table A6: First Stage Regression Results, by Time Period: Log Consumption

	(1) 2004-2007	(2) 2008-2012	(3) 2013-2017
Black	-0.2544	-0.2317	-0.1413
	(0.0122)	(0.0109)	(0.0093)
Asian	-0.2225	-0.2395	-0.1730
	(0.0189)	(0.0147)	(0.0127)
Other	-0.1186	-0.1322	-0.0527
	(0.0126)	(0.0117)	(0.0114)
Age (Male Head)	-0.0234	-0.0915	-0.0629
	(0.0303)	(0.0268)	(0.0376)
Age Squared	0.0005	0.0023	0.0018
	(0.0007)	(0.0006)	(0.0008)
Age Cubed	-0.0000	-0.0000	-0.0000
	(0.0000)	(0.0000)	(0.0000)
Education	-0.0008	-0.0001	-0.0011
	(0.0008)	(0.0006)	(0.0006)
Household Size	0.1454	0.1376	0.1251
	(0.0027)	(0.0021)	(0.0021)
Constant	8.1400	8.9411	8.4321
	(0.4286)	(0.3831)	(0.5642)
Observations	62094	118074	102394
Adjusted R^2	0.169	0.131	0.101

This table shows the results of the regression specified in (1) for different time periods in the sample period. The regression also includes year dummies and state dummies, which are omitted here for brevity. Standard errors are in parentheses.

Table A7: Probability of Stockownership: PSID 1999-2017

	Stockholder
Education	0.1628
	(0.0022)
Log income	0.5070
	(0.0065)
Age	-0.0203
	(0.0017)
Age squared	0.0004
	(0.0000)
Constant	-8.4439
	(0.0719)
Observations	68,846
Pseudo R^2	.2227

This table shows the results of a probit regression to predict stock ownership based on the PSID data from 1999-2017. The regression also includes year dummies, which I omit here for brevity. Standard errors are in parentheses.

Table A8: Summary Statistics by Year: High Education Subsample

Year	Number of Observations	Age (Mean)	Education Level (Median)	Income Level (Mean $)	Consumption (Mean $)
2004	4,418	45.31	Graduated College	70,000-99,999	4,080
2005	5,239	44.66	Graduated College	60,000-69,999	3,954
2006	6,198	44.95	Graduated College	70,000-99,999	3,866
2007	9,354	44.51	Graduated College	70,000-99,999	4,138
2008	9,714	44.06	Graduated College	70,000-99,999	4,224
2009	9,890	44.03	Graduated College	70,000-99,999	4,136
2010	9,704	44.23	Graduated College	70,000-99,999	4,115
2011	9,880	44.27	Graduated College	70,000-99,999	4,222
2012	9,575	44.72	Graduated College	70,000-99,999	4,526
2013	9,597	45.27	Graduated College	70,000-99,999	4,433
2014	8,766	45.92	Graduated College	>100,000	4,525
2015	8,647	46.26	Graduated College	>100,000	4,453
2016	7,474	47.61	Graduated College	>100,000	4,569
2017	6,382	48.38	Graduated College	>100,000	4,423

This table shows summary statistics by year in the subsample of highly educated members of the Nielsen Consumer Panel, the construction of which is described in Section 2.1 and Section 4.3. Education levels and income levels are given as the median bucket for the Nielsen categories. Consumption is calculated as the sum of all purchases across the given year. All data is based on the head of household. Data are weighted using the Nielsen projection weights. All dollar figures are in 2009 dollars.

Table A9: Summary Statistics by Year: High Income Subsample

Year	Number of Observations	Age (Mean)	Education Level (Median)	Income Level (Mean $)	Consumption (Mean $)
2004	1,614	46.54	Graduated College	>100,000	4,939
2005	1,981	46.00	Graduated College	>100,000	4,826
2006	2,659	46.16	Graduated College	>100,000	4,554
2007	4,201	46.11	Graduated College	>100,000	4,799
2008	4,704	45.93	Graduated College	>100,000	4,872
2009	49,27	46.21	Graduated College	>100,000	4,871
2010	4,698	46.70	Graduated College	>100,000	4,807
2011	4,972	47.03	Graduated College	>100,000	4,914
2012	5,057	47.19	Graduated College	>100,000	5,238
2013	5,129	47.74	Graduated College	>100,000	5,112
2014	5,264	47.92	Graduated College	>100,000	5,209
2015	5,412	47.71	Graduated College	>100,000	5,135
2016	5,034	48.44	Graduated College	>100,000	5,278
2017	4,635	49.01	Some College	>100,000	5,058

This table shows summary statistics by year in the subsample of high income members of the Nielsen Consumer Panel, the construction of which is described in Section 2.1 and Section 4.3. Education levels and income levels are given as the median bucket for the Nielsen categories. Consumption is calculated as the sum of all purchases across the given year. All data is based on the head of household. Data are weighted using the Nielsen projection weights. All dollar figures are in 2009 dollars.

Table A10: Summary Statistics by Year: Predicted Stockholder Subsample

Year	Number of Observations	Age (Mean)	Education Level (Median)	Income Level (Mean $)	Consumption (Mean $)
2004	2,266	47.94	Post College Grad	>100,000	4,406
2005	3,331	47.55	Graduated College	>100,000	4,332
2006	2,794	47.89	Post College Grad	>100,000	4,092
2007	5,612	46.76	Graduated College	>100,000	4,395
2008	4057	47.20	Post College Grad	>100,000	4,377
2009	4,767	47.36	Post College Grad	>100,000	4,356
2010	3,423	48.30	Post College Grad	>100,000	4,307
2011	3,830	47.83	Post College Grad	>100,000	4,479
2012	2,520	47.38	Post College Grad	>100,000	4,765
2013	3,376	48.66	Post College Grad	>100,000	4,748
2014	2,058	47.94	Post College Grad	>100,000	4,795
2015	2,901	48.89	Post College Grad	>100,000	4,707
2016	1,709	48.65	Post College Grad	>100,000	4,814
2017	2,252	50.57	Post College Grad	>100,000	4,666

This table shows summary statistics by year in the subsample of members of the Nielsen Consumer Panel who are predicted to own stock, the construction of which is described in Sections 2.1 and Section 4.3. Education levels and income levels are given as the median bucket for the Nielsen categories. Consumption is calculated as the sum of all purchases across the given year. All data is based on the head of household. Data are weighted using the Nielsen projection weights. All dollar figures are in 2009 dollars.