Summer 8-6-2021

Incomplete Markets and Financial Innovation: Consequences for Risk-Sharing, Household Welfare, and Portfolio Choice

Nicolas Duvernois
University of New Orleans, duvernois.nicolas@gmail.com

Incomplete Markets and Financial Innovation: Consequences for Risk-Sharing, Household Welfare, and Portfolio Choice

A Dissertation

Submitted to the Graduate Faculty of the
University of New Orleans
in partial fulfillement of the
requirements for the degree of
Doctor of Philosophy in Financial Economics

By

Nicolas Duvernois

M.B.A. Nicholls State University, 2014
M.S. ESCE Business School, 2015
M.S. University of New Orleans, 2018
August, 2021

DEDICATION

To my parents.

ACKNOWLEDGMENTS

Firstly, I would like to thank my parents for their unwavering and unconditional support and love. I would not be where I am today without you. I would also like to thank my big brothers Christophe and Vincent for setting the right example for me to follow all these years. I will never be able to thank you enough.

I am deeply grateful to my advisor Neal Maroney for his invaluable mentorship, advice, and support. Special thanks to my dissertation committee, Dr. Turunen-Red, Dr. Zirek, and Dr. Pezzo for their support and comments. I would also like to thank Dr. John Lajaunie and Dr. Luke Cashen for setting me on this path.

I am grateful to my friends Dr. Blake Rayfield and Dr. Hasib Ahmed. Thank you for shaping me into the researcher I am today and everything else you've done. I would also like to recognize my peers Obed and Shin for being excellent study partners. I'd like to thank my friend Loïs for her always most excellent advice.

A huge thank you to Ms. Lisa for her hospitality, kindness, and generosity. Finally, thank you to Second Line Brewing where I clearly do my best thinking.

ABSTRACT

The dissertation consists of three chapters measuring the degree of risk-sharing in a panel of US households, and its impact on welfare and portfolio choice. Conventional wisdom suggests financial innovation improves risk-sharing by completing markets and lowering transaction costs–households engage in risk-sharing to insure against idiosyncratic income shocks to improve their own welfare. In the first chapter, using household level income and imputed consumption data, I find that households' ability to smooth permanent shocks has slightly increased while transitory insurance remained unchanged. However, I find that participating households have higher consumption insurance. Their ability to insure permanent shocks has improved while their ability to insure transitory shocks has decreased. I also document a change in the composition of risk where the variance of transitory shocks is increasing while the variance of permanent shocks is decreasing. I find significant heterogeneity among households. These results are robust to different income and consumption measurements. In the second chapter, I investigate the welfare and life satisfaction consequences of incomplete markets in a subset of US households. I use a set of parameters describing households' economic environments in terms of income growth, income risk, and transmission risk. I find that changes in risk-sharing have significant implications for household welfare. Cross-sectional differences in risk-sharing environment result in significantly different welfare criteria. I then use IV-regressions to separate the impact of permanent and transitory income and consumption shocks on life satisfaction. As suggested by consumption insurance theory, I find that transitory shocks have no effect on life satisfaction while permanent shocks do. This result suggest that risk-sharing environments have important consequences for households' well-being as well as a significant degree of insurance despite incomplete markets. In the third chapter, I consider the implications in consumption insurance for portfolio choice. Having documented a significant degree of risk sharing, I find that households experiencing positive labor income shocks invest more in risky assets. However permanent shocks are used to increase housing investment, evidenced by the increased allocation towards secured debts. Furthermore, transitory shocks are used to decrease households liabilities, evidenced by the decreased latent unsecured debt allocation.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

C h a p t e r 1

FINANCIAL INNOVATION AND RISK-SHARING

Abstract

What are the effects of the sharp acceleration of financial innovation over the last fifty years? Conventional wisdom suggests financial innovation improves risk-sharing by completing markets and lowering transaction costs–households engage in risk-sharing to insure against idiosyncratic income shocks to improve their own welfare. Using household level income and imputed consumption data, I find that households' ability to smooth permanent shocks has slightly increased while transitory insurance remained unchanged. However, I find that participating households have higher consumption insurance. Their ability to insure permanent shocks has improved while their ability to insure transitory shocks has decreased. I also document a change in the composition of risk where the variance of transitory shocks is increasing while the variance of permanent shocks is decreasing. I find significant heterogeneity among households. These results are robust to different income and consumption measurements.

JEL classification: E21,D12,D31,G52

Keywords: Risk-Sharing, Income Shocks, Financial Innovation

1.1 Introduction

What are the consumption insurance effects of financial innovation? Income inequality has dramatically increased over the past 50 years and is especially skewed in the United States compared to other developed nations (Saez and Zucman, 2016; Piketty, 2015; Wolff, 2016). Financial institutions and markets have undergone fundamental changes over the same period (Miller, 1986; Frame and White, 2004).

Since the turn of the second millennium, financial innovation has greatly accelerated, yet the empirical effects of these innovations remain largely unexamined. The standard view of financial innovation is that it facilitates the completion of markets, thereby allowing for the equalizing of marginal rates of substitution and improves welfare. Risk-sharing, or the cross-sectional alignment of marginal rates of substitution, should be more pervasive with successive financial market innovations under this traditional view. Cochrane (2009, p. 56) writes: "better risk sharing is much of the force behind financial innovation. Many successful new securities can be understood as devices to share risks more widely." However, recent theoretical papers have challenged the textbook view. Simsek (2013a) and Simsek (2013b) and Buss and Uppal (2017), in their model of financial innovation, consider heterogeneous information and beliefs which serve to increase speculative activity and volatility—thus attenuating welfare improvements.

There is extensive work on consumption insurance (Cochrane, 1991; Blundell and Preston, 1998; Krueger and Perri, 2006; Kaplan and Violante, 2010). However, much of the literature seeks to quantify the degree of consumption insurance and whether households are fully insured against idiosyncratic income shocks. While estimates of of consumption insurance differ, the literature agrees that evidence of full insurance cannot be found and markets are thus incomplete.

Some papers study specific mechanisms for risk-sharing (Olovsson, 2010; Ortigueira and Siassi, 2013; Blundell, Pistaferri, and Saporta-Eksten, 2016; Hryshko, Luengo-Prado, and Sørensen, 2010). However, to the best of my knowledge, I am the first to try to establish an empirical link between financial innovation and trends in risk-sharing. Given the development of exchange-traded funds (Iachan, Nenov, and Simsek, 2021; Huang, O'Hara, and Zhong, 2021) reducing transactions costs Turley (2012) and allowing households to form more optimal portfolios, I should see improving trends in risk-sharing, especially for market participants. Indeed, Guvenen (2007) argues that market participants might have lower insurance if they cannot optimally select a portfolio that minimizes their market risk exposure in addition to their labor income risk.

This paper is related to the consumption insurance literature by documenting trends in risk-sharing. It is also related to the literature on financial innovation by testing the claim that financial innovation improves risk-sharing. I use data that survey both households' income and consumption: the Panel Survey of Income Dynamics (PSID hereafter). To estimate a robust measure of consumption in the PSID, I follow the literature on consumption imputation (Johnson and Fisher, 2020; Attanasio and Pistaferri, 2014).

I first test a simple framework where idiosyncratic consumption responds to idiosyncratic income shocks. I confirm results from the literature. I find that a large portion of shocks are insured but I reject full insurance. I do not find improvements in households' ability to smooth idiosyncratic income shocks. There is a small increasing trend in the time-varying risk-sharing coefficients indicating a decline in consumption insurance. I do however find that entrepreneurs, wealthy, and participating households have lower risk-sharing coefficients suggesting insurance heterogeneity across household groups.

I then estimate the empirical framework designed by Blundell, Pistaferri, and Preston (2008). Unlike Cochrane (1991) who uses proxies to test whether households are insured against permanent or transitory shocks, Blundell, Pistaferri, and Preston (2008) uses the structure of the joint distribution of income and consumption autocovariances to describe the variance of permanent and transitory shocks as well as the proportion of these shocks that is transmitted to households' consumption (they refer to the latter parameters as partial insurance coefficients). The aim is to describe how the partial insurance coefficients evolve through time. I therefore allow these parameters to freely time vary throughout my sample period (1980-2016). I find that permanent insurance has increased in the whole sample and increased for most subgroups of households. However, I do not find an improvement in permanent insurance for entrepreneurs (which have the lowest permanent insurance parameters and thus the best insured), nonparticipating households, transfer recipients and low wealth households (defined as having below median wealth). Transitory insurance is found to be constant in the whole sample. However, I find that market participants, nonparticipating households, and low wealth individuals have lower insurance parameters. The heterogeneity in transitory insurance is smaller than the heterogeneity in permanent insurance, suggesting that the differences observed in the simple transmission framework are caused by permanent shocks.

I then estimate income shocks by removing different income streams such as financial, business, or transfer income to determine how subgroups smooth their income shocks. I find that financial income has little to no impact on the risk-sharing parameters estimated. Business income has a large impact on high wealth and participating households (and obviously entrepreneurs). This result suggests that entrepreneurs may have better formal insurance mechanisms. Alternatively, entrepreneurs may have superior information on their future income shocks which I could be classifying as insurance. Unfortunately I cannot test this assertion. Transfer income is also a significant source of risk-sharing. This is consistent amongst groups, although the impact is stronger for low-wealth households. Overall, I document small improvements in the risk-sharing environment compared to the sizeable development in financial markets; therefore casting doubt that financial innovation does improve risk-sharing.

The rest of the paper is organized as follows. Section 1.2 will connect this work to the related literature. Section 1.3 will describe the data I use, the sample selection, and the imputation technique used. Section 1.4 will show the results for static, time-varying risk-sharing, and the joint distribution of permanent and transitory income shocks. Section 1.5 will conclude.

1.2 Literature Review

The degree of market completeness has been widely researched and discussed in the literature. In a perfectly complete market, individual households' consumption would only react to aggregate risk but not their idiosyncratic risk. This assumption, common to many models, is flawed. As Cochrane (1991) points out, conventional wisdom would refute the existence of contingent claims insuring households on all idiosyncratic states of nature. Imperfect risk-sharing is seen as a potential solution to solve well-known asset pricing puzzles (Mehra and Prescott, 1985; Mankiw, 1986) and has given rise to a new class of heterogeneous agent models (Constantinides and Duffie, 1996; Constantinides and Ghosh, 2017).

A wide body of literature has studied the joint distribution of consumption and income. Krueger and Perri (2006) use the Consumer Expenditure Survey to document that despite the rise in income volatility, consumption has remained smooth providing evidence for some risk-sharing. The CEX is one the datasets that is widely used in this literature[1]. Mace (1991) finds that only one specification is consistent with full insurance. Nelson (1994) argues that some results might be partially driven by measurement error. More recently, Gervais and Klein (2010) also use CEX data but address the shortcomings of the data structure by estimating quarterly income from annual income. This allows them to estimate a consistent estimator of risk-sharing. They do find evidence of imperfect risk-sharing and reject the null hypothesis of perfect risk-sharing.

The Panel Survey of Income Dynamics is another dataset favored by researchers (Cochrane, 1991; Guvenen, 2007; Hryshko, Luengo-Prado, and Sørensen, 2010; Heathcote,

[1] Cutler et al., 1991; Blundell and Preston, 1998; Aguiar and Bils, 2015

Storesletten, and Violante, 2014). The longitudinal aspect of the PSID provides certain advantages; however consumption data is limited to food categories, rent and a few other items on a consistent basis. The PSID also surveys a wide variety of proxies related to income and consumption shocks. Income is better represented as a combination of transitory and permanent processes (Meghir and Pistaferri, 2004; Abowd and Card, 1989). However, survey data does not differentiate between the two income processes and researchers often study only the response of consumption to income shocks.

Cochrane (1991) uses a large selection of proxies (short unemployment for transitory shocks and disability for permanent shocks). He finds that illnesses less than 100 days or work loss due to strikes is fully insured. He cannot reject the null hypothesis of full insurance for these specifications. However, permanent shocks such as long-term illnesses and involuntary job loss are not fully insured. Blundell, Pistaferri, and Preston (2008), in their seminal paper, do not use proxies to determine the nature of shocks. They model permanent and transitory shocks of income[2]. Their income process allows their to measure the variance of permanent and transitory shocks as well as the proportion of these shocks that translate into consumption shocks. Blundell, Pistaferri, and Preston (2008) refer to these parameters as partial insurance parameters. Their results are consistent with Cochrane (1991). They find that transitory shocks are fully insured (95% and in some specifications insignificantly different from 0) while permanent shocks are not fully insured (36% of permanent income shocks get transmitted to consumption.).

Kaplan and Violante (2010) argues that Blundell, Pistaferri, and Preston (2008)'s insurance coefficients are central to macroeconomic research and are the benchmark to measure household insurance. Their methodology has been applied to multiple countries that have datasets similar to the PSID. Kubota (2020) applies the partial insurance framework to Japanese data and finds that transitory shocks are almost fully insured despite increasing over time. On the other hand, permanent shocks are insured by half and remain constant throughout the sample period. Casado (2011) uses the Spanish Household Budget Continuous Survey and find similar results. He also finds that home-owners, high-wealth households, and college educated households have higher degree of permanent insurance. Santaeulalia-Llopis and Zheng (2018) use a longitudinal panel dataset of Chinese households. They find that Chinese households experienced a decline in permanent insurance combined with increased levels of income risk accompanying the rapid economic growth during the past 30 years. Furthermore, they show that the welfare effects of growth can be as large as the welfare effects of risk and insurance.

[2]See section 1.4.2 for a more detailed description of their methodology

While the papers discussed above attempt to quantify the amount of consumption insurance in various economies, another strand of literature discusses the specific mechanisms through which households insure the income risk. Cochrane (1991) does not discuss the mechanisms for risk-sharing but does outline several formal and informal mechanisms for consumption insurance. Such mechanisms include unemployment, disability Michaud and Wiczer (2018) or medical insurance, social security (Olovsson, 2010), charities and other family (Ortigueira and Siassi, 2013; Blundell, Pistaferri, and Saporta-Eksten, 2016) and community safety net, wealth or home-ownership (Hryshko, Luengo-Prado, and Sørensen, 2010).

However, a fundamental mechanism for risk-sharing remains largely unexplored empirically: financial innovation. As Cochrane (2009) states, the fundamental drive of financial innovation is risk-sharing. Cochrane (2009, p.56) writes: "better risk sharing is much of the force behind financial innovation. Many successful new securities can be understood as devices to share risks more widely". Furthermore the past 60 years have seen a sizeable development in financial innovations (Iachan, Nenov, and Simsek, 2021; Miller, 1986; Frame and White, 2004). Zero coupon bonds, options, financial futures, options on futures, options on indexes, collateralized mortgages, home equity loans, currency swaps are but a few examples of financial innovation in the past 50 years (see Miller, 1986). The notion that financial innovation improves risk-sharing is well anchored in the literature (Allen and Gale, 1988; Allen, Gale, et al., 1994; Weil, 1992; Elul, 1997).

Simsek (2013a) differentiates between two types of financial innovation: product and process. Product innovation creates new assets while process innovation reduces transaction costs. Consistent with increases in exchange traded products (Iachan, Nenov, and Simsek, 2021), Turley (2012) finds that trading costs have dramatically declined. Similarly Huang, O'Hara, and Zhong (2021) show that ETFs are used by informed traders to hedge. They conclude that market efficiency is improved by this financial innovation. Bai, Philippon, and Savov (2016) find that stock prices have become more informative Dynan, Elmendorf, and Sichel (2006) attribute the economic stability of the 1980s to financial innovation. Calvet, Gonzalez-Eiras, and Sodini (2004) develop a model in which new instruments are introduced, and allow participants to share more risk. Under this framework, market participation increases, the equity premium decreases, and real rates increase as observed in the data.

However, Calvet, Gonzalez-Eiras, and Sodini (2004) assumes that all of the financial innovation increases agents' ability to insure, in line with the traditional view. Recent papers contradict this view. Simsek (2013b) argues that investors' belief disagreements on the value of new assets have implications for portfolio risk. Specifically, he argues that disagreement will lead to speculation which increases portfolio risk. In his model, he separates the risk-sharing motive for trading and the speculative one. He defines portfolio risk as a trader's variance of net worth. The variance is decomposed into uninsurable variance (i.e., variance under no belief disagreement) and the speculative variance. He finds that financial innovation (modelled as the introduction of new assets) increases the possibilities for risk-sharing, thus decreasing the uninsurable variance. However, when belief disagreements are sufficiently high enough which increases the speculative variance, the average variance increases. In this case, financial innovation can increase portfolio risk and decrease the overall risk-sharing.

Simsek (2013a) models the different impact of different types of financial innovation (product and process) on risk-sharing. His model yields the following results. When belief disagreement is high, both types of innovation will increase portfolio risk. Surprisingly, he also finds that product innovation can lead to increase portfolio risk in the absence of disagreement. Indeed, when new assets are introduced, traders have an increased ability to hedge, which in turns increases the size of their speculative bets. Simsek (2013a) and Simsek (2013b) refer to this effect as the *hedge-more/bet-more* effect.

Buss and Uppal (2017) also challenges the assumption of homogeneous beliefs. They argue that the presence of informed and uninformed traders (relative to new securities) can have consequences for asset prices. They show that even in cases of improved risk-sharing through financial innovation, it can increase return volatility and asset premium. In their model, they introduce two classes of risky assets: a traditional asset which two groups of investors have access to and an innovative asset which only experienced agents initially have access to. They define financial innovation as granting inexperienced investors access to the innovative asset. They show that inexperienced investors overweight the risk-free asset as a consequence of having to learn the payoff structure of the new asset. They show that despite the increased risk-sharing and the volatility of experienced investors' stochastic discount factor (SDF) declining, inexperienced investors' SDF increases significantly more, leading to increased asset premium.

Economic theory is divided. On the one hand financial innovation should increase risk-sharing. On the other, risk-sharing will depend on price agreement; and even if risk-sharing is improved through process innovation, investors' risk might be increased. These implications have not been tested empirically. Although Krueger and Perri (2006) find that an increase in income inequality has not led to an increase in consumption inequality. They conclude that the development of credit markets is a likely explanation for their result and leave the matter for future research. Furthermore, their data spans from 1980 to 1998, omitting the rapid growth of mutual funds and exchange traded derivatives observed after

2000 (Iachan, Nenov, and Simsek, 2021).

Guvenen (2007) does investigate the risk-sharing capabilities of stockholders and non-stockholders to assess the importance of market incompleteness. While he is able to reject perfect insurance for stockholders he cannot for non-stockholders. He argues that stockholders might be exposed to certain types of risk that non-stockholders are exposed to (consistent with Simsek, 2013a; Simsek, 2013b). Stockholders tend to be wealthier and face more entrepreneurial risk which is harder to insure. He offers an alternative explanation that is market based. Should investors face information or trading constraints (high costs, prohibited or limited short-selling), they will not be able to form an optimal portfolio, thus bearing more risk.

Given the recent development in the literature, my contribution is two-fold. First, I will estimate the joint distribution of idiosyncratic income and consumption risk through time. Indeed, given the tremendous transformation of the financial sector, I can test whether the conventional wisdom is correct. If so, I should see downward trends of risk-sharing (i.e., an appreciation in households' insurance). Second, I test whether financial innovation is responsible for the observed trends in risk-sharing. I can estimate risk-sharing for households participating and not participating in the financial market. Given the lower transaction costs and the availability of exchange-traded products, investors should form more optimal portfolios, thereby reducing their risk. With a longer horizon than Guvenen (2007), I can test whether stockholders have higher entrepreneurial risk and face additional market risk. Overall, my contribution simply lies in the question: has financial innovation improved risk-sharing?

1.3 Data

The empirics of consumption insurance are somewhat rendered difficult by the availability (or lack thereof) data. Indeed, the identification strategy requires a large panel dataset with a detailed distribution of household income and consumption. The Panel Survey of Income Dynamics (PSID) tracks households through time and surveys their income; however, only food consumption is consistently available in the survey. The Consumer Expenditure Survey (CEX) provides both detailed income and consumption. As a repeated cross-section it is unsuited to this analysis. I therefore derive a consumption series using imputation technique proposed by Attanasio and Pistaferri (2014) using PSID only.

The PSID is a longitudinal survey. It began in 1968 with 5,000 households. It follows households interviewed in 1968 and the households formed by their descendants, increasing the number of households being tracked over time. The PSID reports various household characteristics, food spending, and income. Data are collected in the survey year[3], however the timing of reporting is not always clear. While income is always reported for the prior year, consumption and wealth can be reported for the current or prior year. I therefore treat consumption and wealth to be reported for the prior year to coincide with income.

To impute consumption, I exploit the fact that food expenditure is consistently available[4]. The temporary measure of consumption in the PSID is the annual expenditure on food at home and food away from home. Several adjustments must be made to ensure comparability across waves. I add food delivered (including food stamps) to food at home when available. I also add food away from home purchased with food stamps to the food away from home category when available. When no distinction between the use of food stamps is available, I add the monetary value of food stamps to the food at home category. I remove top-coded values for each category of food consumption. Starting with the 1994 wave, the PSID provides food consumption as well as the time units used by the household in their reporting. I therefore adjust the categories into annual data points. These adjustments sometimes result in obvious outliers which are removed from the sample. I then sum all food consumption categories for each household-year: $f_{i,t}$.

In 1999, the PSID started to include consumption categories, which now corresponds to 70% of the categories covered by the CEX (Blundell, Pistaferri, and Saporta-Eksten, 2016). These categories include health expenditures, utilities, gasoline, car maintenance, transportation, education, and child care[5]. I remove top coded observations for each category. I calculate a measure of consumption net of food consumption for each household-year for the 1999 to 2017 PSID waves: $n_{i,t}$[6].

To match Attanasio and Pistaferri (2014)'s sample selection[7], I remove households from the SEO sample. The SRC sample was designed to be representative of the US population; the exclusion of the SEO sample allows me to proceed without using sample weights. I also remove households from the Latino and Immigrant subsamples. The households were included in 1999 and would therefore not have any imputed values. Furthermore, their inclusion could change the regression's projection. I exclude households whose reference person is younger than 25 or older than 65. Finally, I remove households whose hourly wages are below

[3]Households were interviewed each year until 1999. Post-1999 interviews are conducted on a biennial basis.

[4]Only the monetary value of food stamps is available for the 1988, and 1989 waves. I therefore do not impute values for these waves.

[5]Clothing and entertainment were added in 2005 and are not used in the imputation for consistency.

[6]This measure also includes rent-equivalent expenditure calculated as 6% of house value (Attanasio and Pistaferri, 2014; Flavin and Yamashita, 2002) for homeowners; rent is a expenditure category that is consistently measured in the PSID.

[7]This is also commonly done in the literature.

half of the federal minimum wages, or have missing data on state of residence[8], employment, or marital status.

I estimate the following approximation of a demand system using pooled OLS regression:

$$ln(N_{i,t}) = Z'_{i,t}\beta + p'_t\gamma + g(f_{i,t}, \theta) + u_{i,t} \quad , \tag{1.1}$$

where Z is a vector of socioeconomic variables[9], p is a vector of prices[10], $g(\cdot)$ is a third-degree polynomial of food consumption, and u is the error term. Imputed consumption is then:

$$\widehat{C_{i,t}} = f_{i,t} + exp\{Z_{i,t}\hat{\beta} + p'_t\hat{\gamma} + g(f_{i,t}, \hat{\theta})\} \tag{1.2}$$

I depart from Attanasio and Pistaferri (2014) in that I use the multiple imputation methodology proposed by Rubin (2004) and used by Fisher and Johnson (2020). I produce 100 estimates of $\widehat{C_{i,t}}$ by adding random noise to each coefficient. This random noise has mean 0 and standard deviation equal to the standard error of the coefficient. This technique allows me to account for potential measurement error in the data. The final measure of consumption is the average of all the estimated $\widehat{C_{i,t}}$. To check the validity of the imputation, I plot the first two moments of my imputation and the measure of total consumption provided by the PSID in figure (1.1).

In computing the first two moments, I scale both consumption measures by the OECD adult equivalent scale. However, in the following sections, I use the raw imputed consumption. I also scale the standard deviation of consumption to be indexed at 0 when the real data starts. Figure (1.1) shows that my imputation fits the data relatively well. The correlation between the imputed and real series is 80%.

Although data can be imputed backwards to 1968, I will restrict my analysis to the 1980 to the 2017 waves[11]. This is done for several reasons. First, the demand system is approximated with non-concurrent data (as opposed to imputation techniques that rely on the CEX). Consumption heterogeneity introduced by preference shifts would not be captured by the imputation, thus the quality of the imputation may decline in the PSID's earlier waves. Furthermore, my analysis will rely on the PSID wealth module that was first introduced in the 1984 wave. While I can reasonably assume that households' situation in 1984 was similar to their situation in 1980, it is less likely further back in time.

[8]mostly households residing abroad

[9]Reference person's hours worked, family size, dummy variables for reference person's age, race, employment status, marital status, whether self-employed, whether disabled, home-ownership and state of residence

[10]CPI for all items, rent, food at home, and food away from home.

[11]Meaning my data spans 1979 to 2016. 1979 is kept to have the first differenced data start in 1980.

Figure 1.1: Imputed Consumption Vs. PSID Consumption

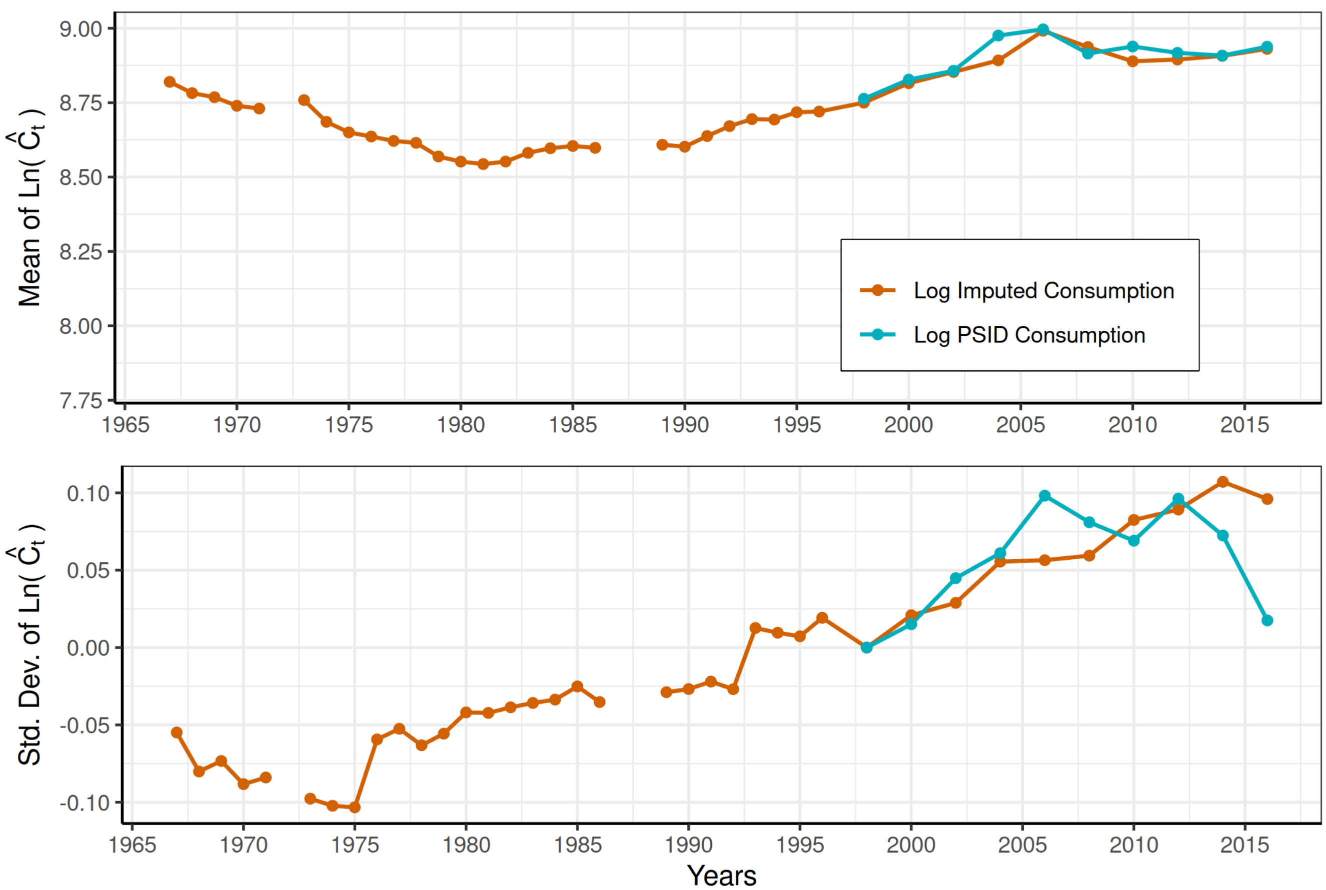

The second aspect of risk-sharing is income. I use disposable income as the main measure of income. I calculate disposable income as the sum of a household taxable income (including reference person, spouse, and other family members), household transfers, social security income, and financial income[12] to which I subtract federal taxes. Note that prior to the 1992 wave, the PSID provided a tax liability variable. I therefore make use of NBER Taxsim and guidelines from Feenberg and Coutts (1993) and Kimberlin, Kim, and Shaefer (2014) to estimate household taxes for the 1992 wave onward.

I then exclude households with top-coded values for each category of income. I also remove income outliers defined as households whose income grows by more than 500%, falls by more than 80%, or is less than \$100. I then require each household to have at least 4 consecutive data points for income. Although this filter is arbitrary, 4 observations are required to identify the parameters in the income process.

A note must be made on my definition of household. In order to identify households

[12]Financial income is measured as the sum of income from dividends, interest, trust funds, and royalties.

through each wave, I track the reference person. If the reference person changes or there is a major family composition change[13], I drop the wave where the change occurs and consider the household to be a newly formed in the subsequent waves. Blundell, Pistaferri, and Preston (2008) requires households to have no changes to the reference person or their marital status. Indeed, they consider only continuously married couples. However, Blundell, Pistaferri, and Saporta-Eksten (2016) show that female labor supply is an important risk-sharing mechanism, thus their sample selection may overstate the degree of consumption insurance. I impose no such requirement. Households resulting from a change in reference person must however be present in at least 4 consecutive waves to be included in the final sample.

Starting in 1984 (and every five years until 1999, when it becomes consistently available), the PSID included questions designed to assess wealth. I define wealth as the sum of assets (farm and business, checking and savings, other real estate, stocks, vehicles, other assets, annuities/IRA, and home equity) minus debt (debt on farms, businesses, real estate, credit card debt, student loans, medical and legal bills, and loans from relatives). From these data, I determine which households participated in the stock market. Participating households are defined as having nonzero values for amounts invested in stocks, mutual funds, investment trusts, or IRAs. As these variables are not available in all waves, I follow Guvenen (2007) in identifying participating households. I first assume that the first recorded status has not changed in the prior years. That is a household entering the sample prior to a "wealth wave" will be assign the first non-missing participation status. Then, in order to determine participation between waves, I require households to have the same status between waves. As, I cannot identify when the household would exit or enter the market, I simply treat those observations as missing. Although, an imperfect identification, this is only an issue for waves prior to 1999. Figure (1.2) plots the average market participation and the average proportion of risky assets held in households' portfolio. I define the risky share of a portfolio as the sum of assets invested in financial markets divided by the sum of financial assets[14].

From figure (1.2), the average risky share has increased in the first half of the sample and stabilizes in the second. Market participation is generally increasing, although it reached a peak in 2000 and declined slightly afterwards. It goes from 27% in 1983 to 40% in 2016. Note the dashed line for market participation. This represents the fact that data are missing between these waves. While the interpolation would suggest that market participation is

[13]defined as a change of reference person or spouse.

[14]Financial assets are defined as assets invested in the stock market plus the value of checking and savings accounts, money market funds, certificates of deposits, government bonds, or treasury bills.

Figure 1.2: Average Financial Market Participation in PSID Households

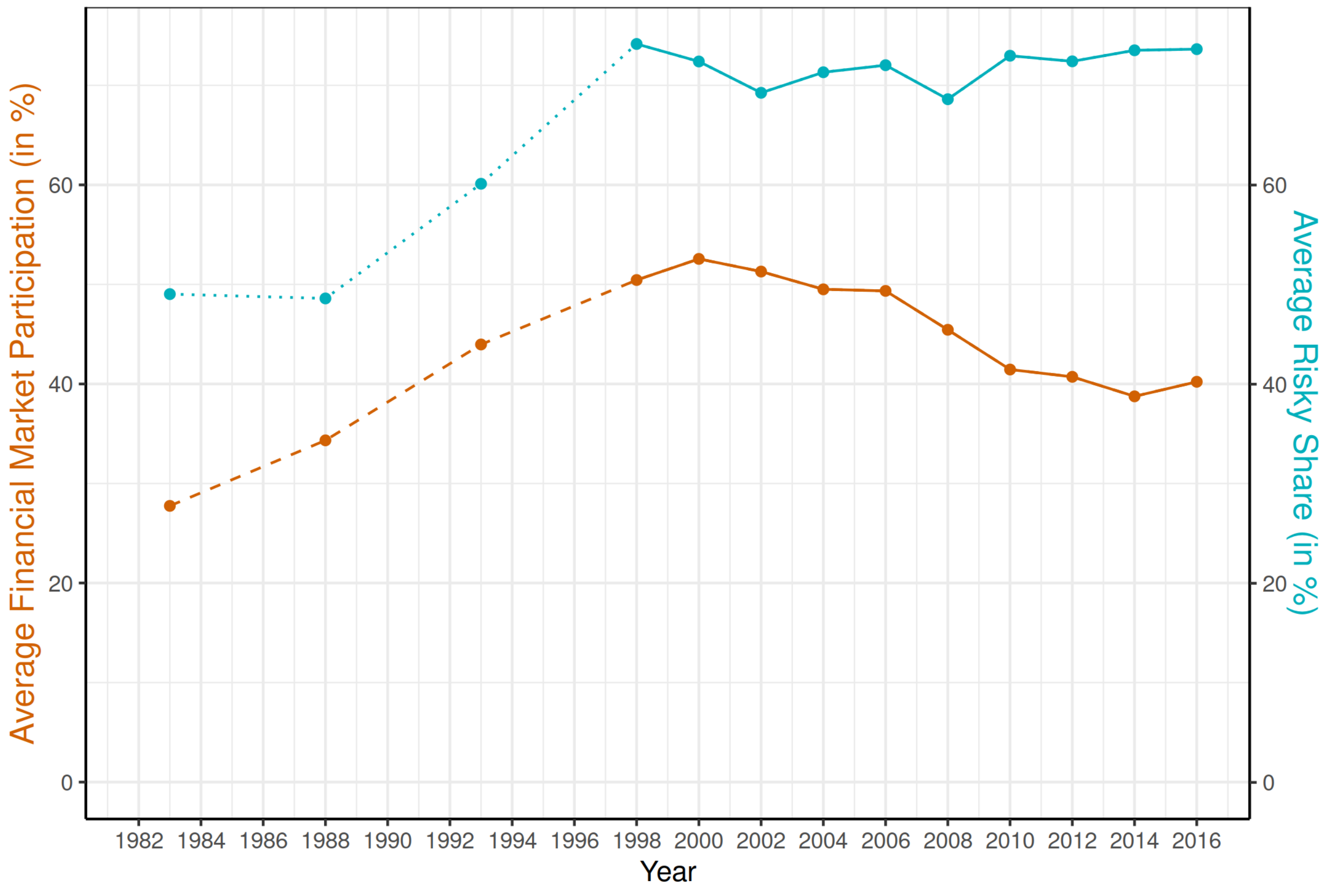

steadily increasing between waves, my identification technique (not represented in this graph) yields a different pattern. Since I am keeping households' status constant, the average participation remains constant[15] between waves and spikes to its actual level when data is available.

I also identify households' positions within the wealth distribution using a similar identification scheme. Total wealth is carried backwards from the most recent non-missing value (limited to 5 years and only prior to the 1999 wave). I then identify households in the top and bottom 50% of the total wealth distribution. This is again an imperfect identification scheme. However, Fisher and Johnson (2020) find that there is little wealth mobility. I differentiate between entrepreneurs and employees where entrepreneur households are those where at least one household member reports non-zero business income. Finally, I identify transfer recipients as those with non-zero transfer income (including family members other than reference person and spouse).

[15]There is very small decrease caused by household attrition between wealth waves.

The final sample is composed of 6,291 unique households and 65,591 household-years observations. Table 1.1 reports summary statistics for the sample. Panel A reports summary statistics for all households in the final sample. Income and consumption data are deflated to 1982 dollars. The median income is $30,303 and the median consumption is $13,312. 61% of households' reference person has some college education or higher, 87% are employed, and 16% of households report owning a business or having a financial interest in a business enterprise. The median wealth is $43,222. Unsurprisingly, the distribution is highly skewed and is dominated by a few highly wealthy individuals. In panels B and C, I report summary statistics for market participants and not participating households respectively. Participants have higher median levels of income ($41,855 vs $23,844) and consumption ($16,705 vs $11,542). Participating households tend to be more educated (79% vs 49%), more likely to be employed (90% vs 84%) or have an interest in a business (23% vs 11.5%). Market wealth (the sum of stocks and IRAs) median level is $23,303 and similarly to wealth, is highly skewed.

1.4 Results

1.4.1 Idiosyncratic Income and Consumption

To capture the idiosyncratic components of income and consumption, I use OLS regressions to remove the deterministic effects of observable household characteristics. The regression is similar to Blundell, Pistaferri, and Preston (2008). I regress log income and log consumption on dummy variables for education, race, employment status, and marital status of the reference person. These characteristics are allowed to vary with time. I include year dummies, and year-of-birth dummies to control for cohort effects. I also include dummy variables for the presence of outside dependents, region of residence, whether the spouse earns income, and the presence of income earners other than the reference person and spouse. As noted before, I do not scale the dependent variables into adult equivalent shares. I do however control for family size and the number of children. The residuals $y_{i,t}$ and $c_{i,t}$ are used in the rest of the analysis. The adjusted R^2 for the income regression is 49% while the adjusted R^2 for the consumption regression is 70%, indicating that a larger proportion of households' consumption is directly observable. I then calculate the growth rates of $y_{i,t}$ and $c_{i,t}$ labelled $\Delta y_{i,t}$ and $\Delta c_{i,t}$. Both growth rates are trimmed at the top and bottom 1%.

1.4.2 Static and Time-Varying risk-sharing

In its simplest form, risk-sharing can be understood as the degree of transmission of idiosyncratic income shocks to idiosyncratic consumption shocks. Empirically, β is measured

Table 1.1: Summary Statistics

This table reports summary statistics for the final sample. The final sample is comprised of 6,291 unique households and 65,591 household-year observations. Panel A reports mean, median, standard deviation,25^{th}, and 75^{th} percentile for all households. Disposable income is the sum of labor, business, transfer, social security, and financial income minus federal taxes. Consumption is imputed following Attanasio and Pistaferri (2014). College takes the value of 1 if the reference person reports having at least attended college. Employed takes the value of 1 if the reference person reports being employed. Entrepreneur takes the value of 1 if any household member reports owning a business or having a financial interest in a business enterprise. Market participation takes the value of 1 for households with non-zero stock or IRA wealth. Guvenen (2007)'s methodology is used to imputed missing data points. Market wealth is defined as the sum of stocks and IRA wealth. Total Wealth is defined as the sum of assets minus debt. Data are deflated to 1982 dollars.

Variable	Mean	Median	Std. Dev	25^{th} P.	75^{th} P.
Panel A. All Households					
Disposable Income	35,874	30,303	32,262	19,840	43,745
Imputed Consumption	14,618	13,312	7,249	9,641	18,132
College	0.612	1	0.487	0	1
Employed	0.872	1	0.334	1	1
Entrepreneur	0.163	0	0.370	0	0
Total Wealth	144,254	43,222	550,287	8,231	130,590
Market Participation	0.391	0	0.488	0	1
Market Wealth	39,081	0	281,805	0	14,686
Panel B. Market Participants					
Disposable Income	50,289	41,855	43,885	29,948	58,307
Imputed Consumption	17,980	16,705	7,975	12,558	21,908
College	0.788	1	0.408	1	1
Employed	0.897	1	0.303	1	1
Entrepreneur	0.229	0	0.420	0	0
Total Wealth	267,018	120,041	779,559	47,538	264,535
Market Wealth	91,779	23,303	428,923	5,435	76,253
Panel C. nonparticipating Households					
Disposable Income	26,608	23,844	20,736	15,421	34,071
Imputed Consumption	12,730	11,542	6,310.570	8,330	15,698
College	0.489	0	0.500	0	1
Employed	0.841	1	0.365	1	1
Entrepreneur	0.115	0	0.319	0	0
Total Wealth	49,083	15,513	241,920	1,392	48,958

as the coefficient in the following regression:

$$\Delta c_{i,t} = \alpha + \beta \Delta y_{i,t} + \varepsilon_{i,t} \tag{1.3}$$

The coefficient β in equation (1.3) is the share of the variance of idiosyncratic income that gets transmitted to consumption. β is theoretically bounded between 0 and 1, 0 indicating perfect risk-sharing, 1 indicating no risk-sharing. In other words, a lower coefficient indicates a higher degree of consumption insurance. To test whether different groups of households share risk differently, I interact idiosyncratic income with an indicator variable ($I_{i,t}$) and estimate the following regression:

$$\Delta c_{i,t} = \alpha + \beta \Delta y_{i,t} + \delta \Delta y_{i,t} * I_{i,t} + \gamma I_{i,t} + \varepsilon_{i,t} \tag{1.4}$$

I consider the following groups: participating households (vs. nonparticipating), entrepreneurs (vs. employees), high wealth (vs. low wealth), no transfers (vs. households receiving at least $1 of transfer income), and stock owners (vs. non-stock owners). Note that I has a t subscript. Indeed, households can be part of a group one wave and the other group for the next wave. Furthermore, groups are not mutually exclusive. I estimate equations (1.3) and (1.4) with 3 different methods. First, I use OLS with the imputed consumption series. Second, I use least absolute deviation (LAD) to ensure the estimation is robust to outliers (i.e., large positive or negative shocks likely caused by measurement errors). Third, I use OLS with non-imputed consumption (using 2000 to 2016 data only). Table 1.2 reports the results[16].

In specification (1), I use all households in the sample. The results are consistent with the literature (Gervais and Klein, 2010) in that I cannot reject complete risk-sharing; there is however a significant amount of it. Indeed, I find that a 10% shock to idiosyncratic income results in a 1.1% change in consumption. The estimation is robust to outliers. In panel B, I find a coefficient of 0.1155. In panel C, I find a slightly higher coefficient of 0.1508 significant at the 1% level, implying a lower degree of risk-sharing. This could be indicative of two things. One, the imputation technique could produce a smoother consumption series thus reducing shocks. This would overstate the degree of consumption insurance in Panel A. Two, since I use only the 2000 to 2016 waves[17], consumption insurance could have deteriorated in the second half of the sample. This implication is considered later in the paper.

I then turn to specifications (2) through (6). Note that the groups in the header of table (1.2) have $I_{i,t} = 0$. Since the coefficient δ measures the change in slope, the sum of

[16] Equations (1.3) and (1.4) are estimated with the constant term α although omitted from the table.
[17] 1998 has real data but is first differenced so the estimation starts with the 2000 wave.

Table 1.2: Static Risk-Sharing Estimation

*This table reports the estimated risk-sharing coefficients using equations (1.3) and (1.4). Standard errors are reported below the coefficient in parentheses. Households are participating if they have non-zero ownership in stocks, mutual funds or IRAs. Households are entrepreneurs if they report non-zero business or farm income. Households are high wealth if they are in the top 50% of the wealth distribution. Household are transfer recipients if they report positive transfer receipts. Households are stockowners if they have non-zero wealth in stocks. Panel A reports OLS estimates using imputed data; panel B reports LAD estimates using imputated data; and panel C reports OLS estimates using PSID measured data. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

$I_{i,t} = 0$	All (1)	Participation (2)	Entrepreneurs (3)	High wealth (4)	No transfer (5)	Stockowners (6)
	Panel A. OLS with Imputed Data					
β	0.1123***	0.0787***	0.0560***	0.0649***	0.1000***	0.0820***
	(0.0029)	(0.0057)	(0.0065)	(0.0043)	(0.0048)	(0.0070)
δ		0.0433***	0.0720***	0.0672***	0.0185***	0.0382***
		(0.0071)	(0.0073)	(0.0060)	(0.0067)	(0.0079)
γ		0.0015	0.0035	0.0050***	-0.0057***	0.0009
		(0.0021)	(0.0027)	(0.0018)	(0.0021)	(0.0023)
	Panel B. LAD with Imputed Data					
β	0.1155***	0.0751***	0.0530***	0.0700***	0.1056***	0.0815***
	(0.0030)	(0.0057)	(0.0069)	(0.0042)	(0.0046)	(0.0068)
δ		0.0491***	0.0776***	0.0640***	0.0173**	0.0392***
		(0.0073)	(0.0078)	(0.0065)	(0.0070)	(0.0080)
γ		0.0020	-0.0002	0.0041**	-0.0021	0.0006
		(0.0022)	(0.0029)	(0.0020)	(0.0022)	(0.0023)
	Panel C. OLS with Measured Data					
β	0.1508***	0.0991***	0.0636***	0.0860***	0.1361***	0.1119***
	(0.0069)	(0.0120)	(0.0160)	(0.0108)	(0.0117)	(0.0164)
δ		0.0908***	0.1035***	0.0941***	0.0365**	0.0475**
		(0.0157)	(0.0180)	(0.0148)	(0.0162)	(0.0185)
γ		-0.0048	-0.0019	-0.0134***	-0.0043	-0.0111*
		(0.0049)	(0.0069)	(0.0046)	(0.0052)	(0.0058)

β and δ gives the degree of risk-sharing for groups with $I_{i,t} = 1$. I find that participating households tend to smooth consumption more than nonparticipating households. Indeed, I find that a 10% shock to participating households' idiosyncratic income causes a 0.8% shock to consumption, while a similar shock to nonparticipating households results in a 1.22% change in idiosyncratic consumption. This result is in contradiction with Guvenen (2007) who find that participating households share risk less completely. However, I use a larger panel and use a consumption measure that is broader the food consumption measure he uses. Furthermore, I find similar results in panel C. This result strengthens my claim that participants share risk more completely. Indeed, I do not have to rely on imputed data, nor do I rely on an imperfect identification of participation. The magnitude of the difference in risk-sharing is also larger in panel C. The portion of income shocks transmitted to consumption for shareholders is half of that for nonparticipating households. Results are similar when considering direct participation through stock ownership only. I also find that entrepreneurs and high-wealth households smooth consumption to a higher degree than their counterparts. This is also in contradiction of Guvenen (2007) who claim that high wealth households (and participating households) take on more entrepreneurial and market risk. However, high wealth individuals have increased savings capacity and thus can more easily weather shocks. Furthermore, it is possible that entrepreneurs have a better information on their future income and are able to preemptively adjust their consumption (see Kaufmann and Pistaferri (2009) for a discussion on insurance vs information). Regardless of the channel, these results show that entrepreneurs and high-wealth households seem to be less impacted by income shocks. Finally, in specification (5), I find that transfer recipients smooth consumption to a less degree than households who do not receive any government transfers. However, while statistically different, the difference in slope is 0.0185 which is lower than the δ estimated for other groups. This result indicates that formal government mechanisms are imperfect yet provide a significant amount of consumption insurance.

The grouping of households in table (1.2) is fairly coarse. In table (1.3), I separate households based on their position within the income distribution. Households are sorted into quartiles (Q_1 to Q_4). I separate households that have 0 values into Q_0. I know from table (1.2) that households with at least \$1 of business income tend to have a lower risk-sharing coefficient, so I will sort employees into Q_0 and sort households with non-zero business income into the Q_1 to Q_4. I consider the following variables and ratios when sorting households: after-tax income[18], wave, transfer, business, financial income to total income. I estimate the following regression:

[18]Note that Q_0 after-tax income are households who are below the poverty line according to the Census tables. This variable is provided directly by the PSID.

$$\Delta c_{i,t} = \sum_{j=1}^{4} \beta_j \Delta y_{i,t} * I_{i,t;j} + \sum_{j=1}^{4} \gamma_{i,t;j} + \varepsilon_{i,t} \tag{1.5}$$

Equation (1.5) is estimated without a constant term with Q_1 as the base group. The coefficients in Q_1 represents the baseline slope and coefficients Q_2 to Q_4 represent the change in slope and statistical significance. γ_j are not reported. Q_0 is estimated with equation (1.3). Panel A reports the results when estimated with imputed data, while panel B reports results estimated with PSID measured data.

In sorting households by after-tax income, I first divide their disposable income by the OECD equivalent scale. Indeed, a household with more income earners will naturally be more likely to be in a higher quartile[19]. I find that households bottom 25% of the after-tax income distribution smooth consumption to a lesser degree than households in the top 25%. However, the relationship is not monotonic as I do not find a statistical difference between Q_1 and Q_3 households. Almost 94% of shocks affecting households below the poverty line are smoothed away. I offer two possible explanations for this result. Consumption of households below the poverty line is likely to be a baseline of necessities and could hardly get any lower, thus unaffected by shocks. Households below the poverty line are also likely to receive a significant portion of income from government transfers and would thus experience fewer and small shocks to income. Note that these results are consistent when estimating equation (1.5) with measured data. The relationship between risk-sharing and after-tax income appears to be monotonically decreasing suggesting that high-income households are better able to smooth out shocks.

I then consider several sources of income as a ratio of total income. I scale these income streams by total income instead of disposable income to insure that all ratios are between 0 and 1[20]. Secondly, these different streams of income are likely to contribute differently to households' overall tax liability. Scaling by total income does not require assumptions on each stream's tax contribution.

I find that households who receive more transfer income tend to share risk more completely than those who receive a lower proportion of their income in the form of government transfers. This implies that formal risk-sharing mechanisms are efficiently insuring households against income shocks. This result also indicates why the magnitude of the δ

[19]Note that without scaling, I find that households with higher after-tax income share less risk relative to Q_1 households.

[20]Although business income can be negative indicating a loss. The households are sorted in Q_0. This only affects are very small number of households.

Table 1.3: Conditional Risk-Sharing - Income Distribution

This table reports the estimated risk-sharing coefficients based on household's position within the income distribution. Households are sorted into quartile and zero-values households are placed into Q_0. *The* β *coefficients from equation (1.5) are reported with standard errors in parentheses. Panel A is estimated with imputed data. Panel B is estimated with PSID measured data. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Q0	Q1	Q2	Q3	Q4
Panel A. Imputed Data					
After-tax Income	0.0673***	0.1262***	-0.0141*	-0.0061	-0.0230***
	(0.0154)	(0.0060)	(0.0085)	(0.0084)	(0.0082)
Wage / Total Income	0.0536***	0.0976***	0.0254***	0.0323***	0.0256***
	(0.0176)	(0.0048)	(0.0078)	(0.0082)	(0.0079)
Transfer / Total Income	0.1000***	0.1647***	-0.0285	-0.0422**	-0.0739***
	(0.0045)	(0.0142)	(0.0183)	(0.0174)	(0.0165)
Business / Total Income	0.1286***	0.1093***	-0.0443	-0.0556**	-0.0685**
	(0.0035)	(0.0225)	(0.0287)	(0.0276)	(0.0271)
Financial / Total Income	0.1147***	0.1404***	-0.0055	-0.0331***	-0.0670***
	(0.0039)	(0.0093)	(0.0130)	(0.0127)	(0.0118)
Panel B. Measured Data					
After-tax Income	0.0880**	0.1923***	-0.0332	-0.0408**	-0.0783***
	(0.0391)	(0.0140)	(0.0204)	(0.0198)	(0.0195)
Wage / Total Income	0.1005**	0.1338***	0.0282	0.0279	0.0355*
	(0.0412)	(0.0114)	(0.0185)	(0.0196)	(0.0191)
Transfer / Total Income	0.1361***	0.2386***	-0.0339	-0.0676*	-0.1069***
	(0.0109)	(0.0329)	(0.0421)	(0.0405)	(0.0388)
Business / Total Income	0.1671***	0.0997	0.0205	-0.0794	-0.0775
	(0.0082)	(0.0651)	(0.0810)	(0.0768)	(0.0740)
Financial / Total Income	0.1532***	0.2043***	-0.0691**	-0.0517	-0.0937***
	(0.0087)	(0.0251)	(0.0344)	(0.0333)	(0.0315)

coefficient in table (1.2) is small. Q_1 coefficient is 0.1647 while Q_4 is 0.0908. Households who receive most of their income (or its entirety) in the form of transfers share as much risk as households who receive no transfers.

Households who receive a small portion of the income from owning a business appear to share risk more completely than those who receive no business income (0.1093 vs 0.1286). However, the large difference in risk-sharing observed in table (1.2) appears to come from households who receive a significant portion (above median) of business income. Indeed, I find that almost 95% and 96% of income shocks are insured for Q_3 and Q_4 households. Q_1 and Q_2 are not statistically different from each other, though appear larger than Q_0.

I also consider financial income. Recall that financial income is defined as the sum of income from dividends, interest, trust funds, and royalties and thus may be a good proxy for market participation and financial wealth. I find that households with a small proportion of their income in the form of financial income have a risk coefficient of 0.1404 (.14 for Q_1 compared to 0.1147 for Q_0 households. This result is consistent with Guvenen (2007)'s claim that participating in financial markets increases households' risk. Indeed, it can be argued that the increased risk outweighs the risk-sharing benefits from households that cannot sufficiently diversify. However, the relation is monotonically decreasing and households who receive most of their income in the form of financial income have a risk-sharing coefficient of 0.0734 (consistent with the 0.0787 of market participants).

These results are confirmed when looking at the distribution of wages to total income. Households that receive most of their income from wages tend to risk-share less than those who receive little from wages (implying more income from transfer, business, or financial income). Note that results are similar in magnitude in panel B of table (1.3).

In table (1.4), I estimate equation (1.5) with households sorted based on wealth variables instead of income variables. Note that prior to 1998, I record a household's position within the distribution when data are available (1983,1988,1993) and carry their position backwards. Although imperfect, the low wealth mobility observed in the literature allows me to do such an operation with a limited bias. I consider the following variables and ratios: total wealth, financial, market, and stock to total wealth.

The relationship between risk-sharing and wealth is less striking than income. Indeed, I find less statistical differences between quartiles. I do nonetheless find interesting patterns. I find that the difference between households above and below the median level of wealth is caused by households in the top 25% of the wealth distribution. Furthermore, I find that the differences between participating and nonparticipating households are mainly driven by

Table 1.4: Conditional Risk-Sharing - Wealth Distribution

*This table reports the estimated risk-sharing coefficients based on household's position within the income distribution. Households are sorted into quartile and zero-values households are placed into Q_0. The β coefficients from equation (1.5) are reported with standard errors in parentheses. Panel A is estimated with imputed data. Panel B is estimated with PSID measured data. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Q0	Q1	Q2	Q3	Q4
Panel A. Imputed Data					
Total Wealth	0.1167***	0.1244***	-0.0007	-0.0058	-0.0369***
	(0.0099)	(0.0060)	(0.0088)	(0.0088)	(0.0082)
Financial / Total Wealth	0.1139***	0.0999***	0.0115	0.0218**	0.0154*
	(0.0075)	(0.0060)	(0.0088)	(0.0087)	(0.0087)
Market / Total Wealth	0.1213***	0.0532***	0.0165	0.0287*	0.0422***
	(0.0044)	(0.0106)	(0.0149)	(0.0148)	(0.0145)
Stock / Total Wealth	0.1197***	0.0652***	0.0144	0.0184	0.0282
	(0.0039)	(0.0127)	(0.0177)	(0.0179)	(0.0172)
Panel B. Real Data					
Total Wealth	0.1632***	0.1498***	0.0270	-0.0191	-0.0076
	(0.0209)	(0.0146)	(0.0212)	(0.0211)	(0.0196)
Financial / Total Wealth	0.1548***	0.1468***	-0.0091	0.0139	0.0046
	(0.0175)	(0.0147)	(0.0213)	(0.0209)	(0.0205)
Market / Total Wealth	0.1899***	0.0959***	-0.0242	0.0492	-0.0127
	(0.0110)	(0.0227)	(0.0312)	(0.0307)	(0.0307)
Stock / Total Wealth	0.1594***	0.1308***	-0.0010	-0.0160	-0.0489
	(0.0088)	(0.0317)	(0.0431)	(0.0443)	(0.0423)

households with a relatively small proportion of their wealth invested in financial markets. Indeed, the risk-sharing coefficient for Q_1 households is 0.0532 relative to Q_0's coefficient of 0.1213. Furthermore, risk-sharing is monotonically increasing (though I do not find a statistical difference between Q_1 and Q_2 households). Households with most of their wealth invested in the stock market (directly or indirectly) have a risk-sharing coefficient of 0.0954. When considering only stock wealth (i.e. direct financial market participation), I do not find a statistically significant relationship.

In panel B, I repeat the analysis with the PSID consumption series after 2000. This subsample is also characterized by continuously measured wealth data. However, Q_1 to Q_4 households do not appear to be statistically different for any ratio. It is to be noted that segmenting households into quartiles (plus Q_0) results in each quartile being relatively small. Thus, all tests will have lower power. While not significant, the signs indicate that households with more of their wealth invested in the stock market (directly and indirectly) appear to share more risk.

I then ask: has risk-sharing changed over time? Indeed, given the transformation of financial markets in recent decades (Iachan, Nenov, and Simsek, 2021), I should see clear improvement in risk-sharing. To establish a trend in risk-sharing, I estimate equation (1.4) with interacting $\Delta y_{i,t}$ with year-dummies. I plot the results in figure (1.3)

Figure 1.3: Time varying risk-sharing - All households

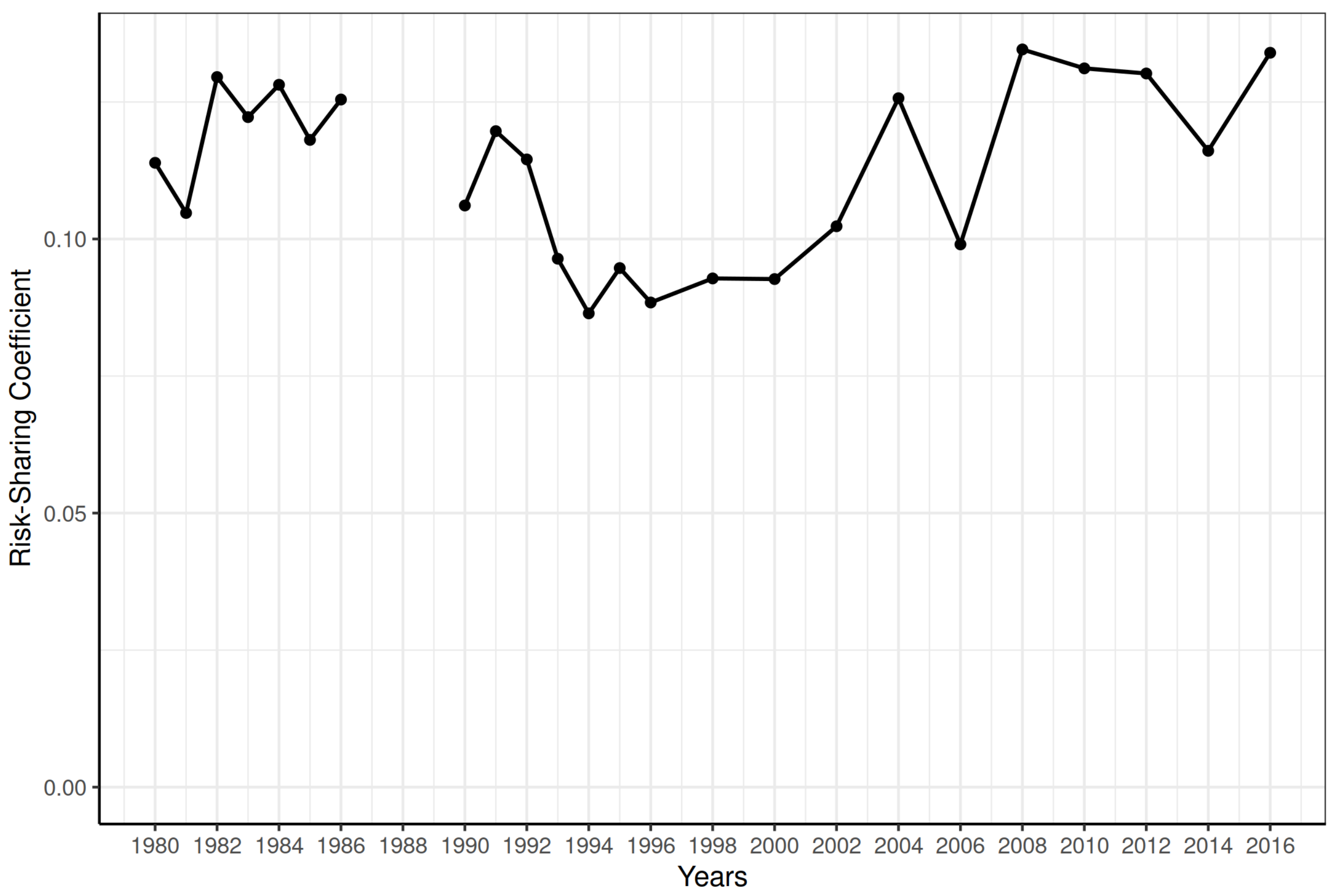

The results of figure (1.3) are surprising. Indeed, the risk-sharing coefficients display an upward trend. The coefficients go from 0.1138 in 1980 to 0.1339 in 2016. The coefficients remain relatively stable in the 1980s, fall in the 1990s, and steadily increase in the 2000s. Although increasing, the economic significance of the change is relatively small. A 10% to income would cause a 1.1% shift in consumption in 1980 and a 1.3% shift in consumption in 2016. Furthermore, the risk-sharing coefficient is the ratio of the covariance of idiosyncratic income and consumption shocks over the variance of shocks. Looking at the β coefficients might obscure the transmission process. In figure (1.4), I plot the cross-sectional covariance between income and consumption shocks. This plot shows a clearer upward trend. Although I cannot strictly interpret these covariances as risk-sharing, the impact of income shocks on consumption shocks seems to be intensifying.

Figure 1.4: Time varying covariances - All households

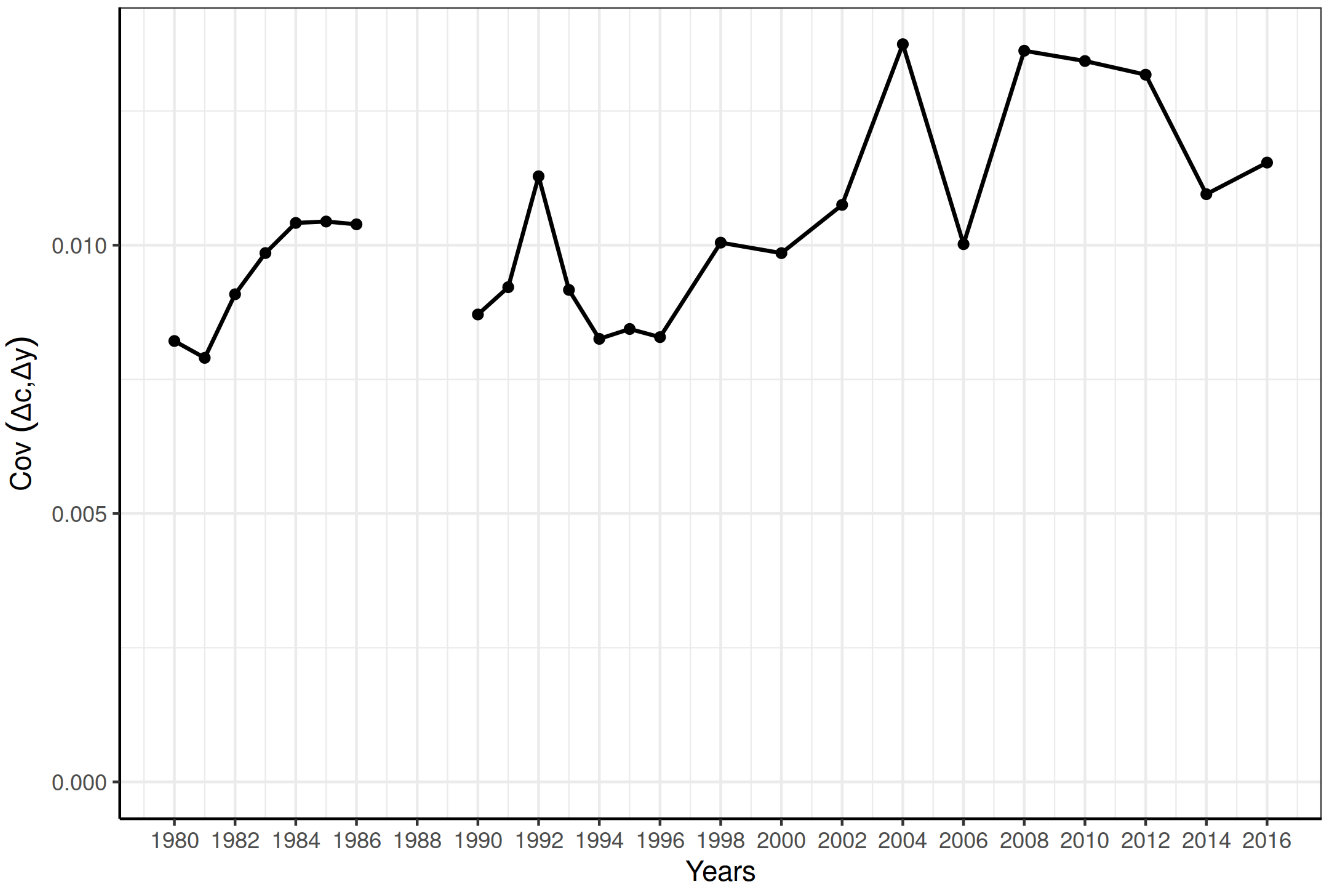

Households' ability to smooth consumption appears to have slightly deteriorated since the 1980s. However, I know from my previous results that not all households have the same risk-sharing coefficients. In figure (1.5), I plot the risk-sharing coefficients of participating and nonparticipating households. I can see that the coefficients for nonparticipating households always lie above the participating households after 1998. Prior to 1998, the differences are not significant. Concentrating on the 1998-2016 period, I clearly see that nonparticipating households have seen their ability to risk-share deteriorate. Participating households has a relatively flat trend. Interestingly, both participating and nonparticipating households have a large increase in the coefficient around the financial crisis. However, this spike does not occur at the same time. Participating households see their spike in 2008, then the coefficients subsequently decline. Non- participating households' coefficient increase in 2008 and 2010 again. This would indicate that nonparticipating households were more affected by the ensuing recession than participating households. The slight increase in risk-sharing in figure (1.3) can be partially explained by non participating households.

Figure 1.5: Time varying risk-sharing - Participating and nonparticipating households

I then plot the risk-sharing coefficients for entrepreneurs and employees. Indeed, I have seen that households receiving less income in the form of wages tend to share more risk. In figure (1.6), the differences between the two groups are striking. Indeed, employees' coefficients always lie above the coefficients for entrepreneurs. In 2006, I see the β for entrepreneur is negative, which violates theoretical bounds[21]. However, I do see that entrepreneurs' ability to smooth consumption has generally increased (i.e lower coefficient) since the 1980s. Employees' consumption smoothing ability is increasing. Interestingly, I also see that entrepreneurs experienced a large shock to risk-sharing during the financial crisis (not observed for employees).

[21]Note that a dip in 2006 is observed throughout the graphs. This is caused by a smaller variance of idiosyncratic income. I fail to provide any explanation as to why this happens.

Figure 1.6: Time varying risk-sharing - Entrepreneurs and Employees

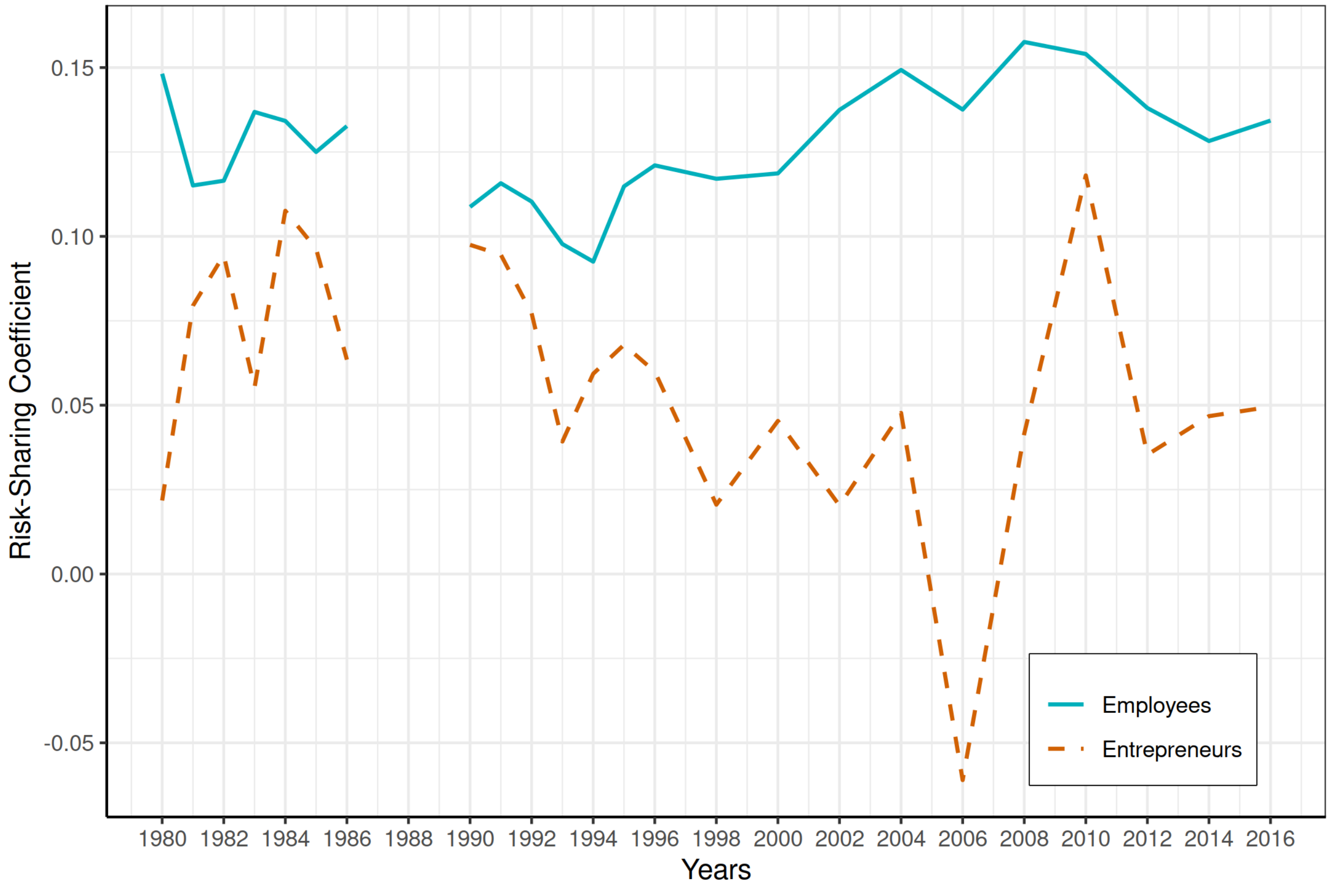

Finally, I plot the risk-sharing coefficients for households above and below the median level of wealth. Figure (1.7) is relatively similar to figure (1.5). Wealthy households' risk-sharing coefficients are relatively stable. These households experience a very moderate shock around the financial crisis compared to households below median levels of wealth. Furthermore, households that are below median levels of wealth experienced a sharp decrease in risk-sharing since the 1980s.

I have not detected the improvement in risk-sharing suggested by the increase in financial innovation. Indeed, the risk-sharing coefficients seem to remain constant at best, increasing at worst, in direct contradiction of the traditional view. However, the transmission mechanism studied so far is simple. In the following section, I investigate a more complete and robust framework of consumption insurance.

Figure 1.7: Time varying risk-sharing - Above and below median levels of wealth

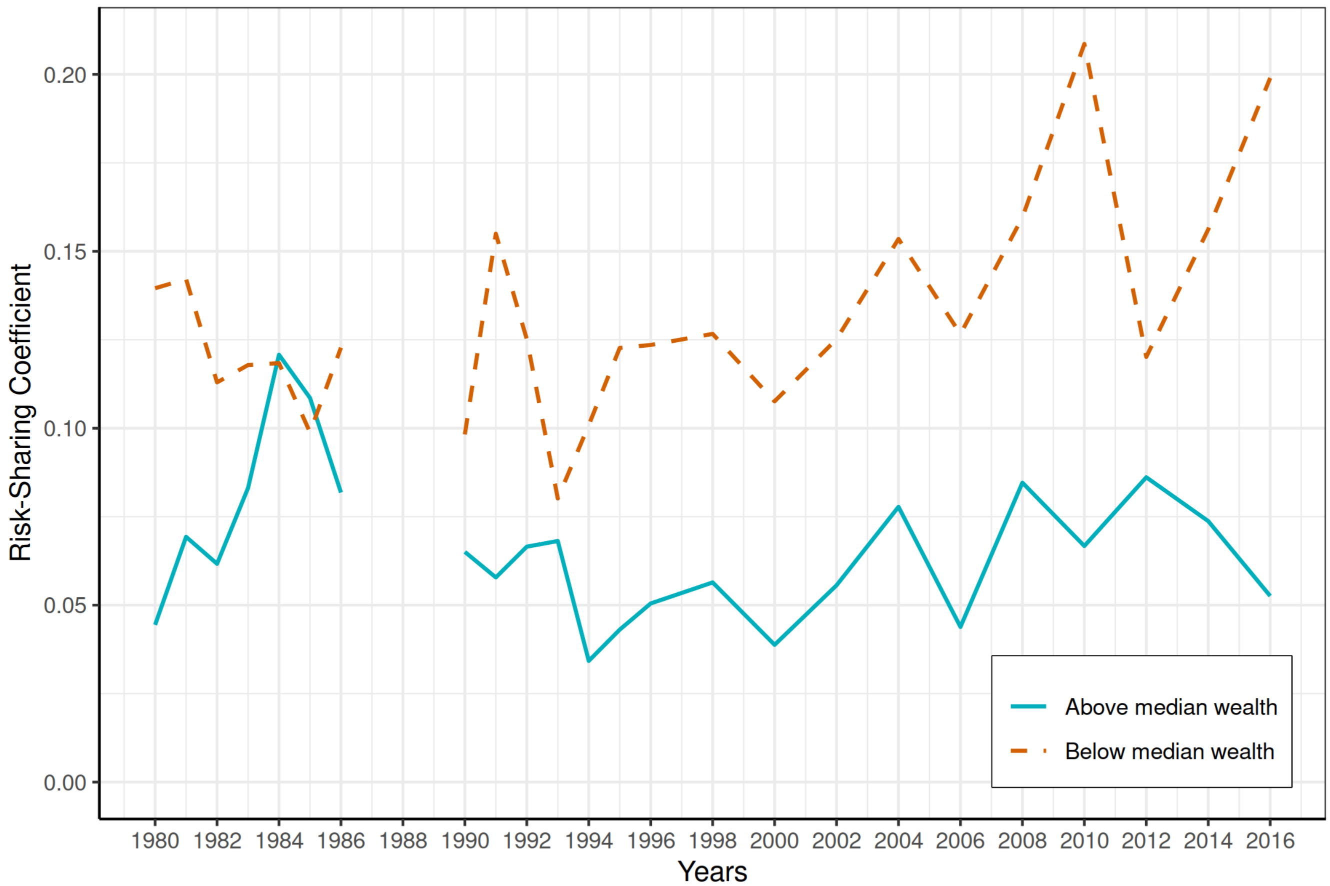

1.4.3 Joint Distribution of Consumption, Permanent, and Transitory Income

Cochrane (1991) uses PSID proxies to model shocks. He considers involuntary job loss or number of days looking for jobs. The length of those variables can be construed as proxies for temporary or permanent shocks. This is indeed one of the empirical strategies used in the literature (Misra and Surico, 2014; Kan, Peng, and Wang, 2017; Fagereng, Holm, and Natvik, 2019). Empirical proxies are often transitory in nature, thus the transitory consumption response has been largely investigated. Alternatively, one can study the joint distribution of permanent and transitory income shocks by imposing a certain structure on the income and consumption response. This second method has the added benefit one not requiring proxies.

Blundell, Pistaferri, and Preston (2008) present a partial insurance insurance model that allows me to differentiate between permanent and transitory income shocks and the degree of transmission of those shocks. The flexibility of the framework lets me estimate time-varying shocks and insurance parameters. This framework is used by Casado (2011), Ludwig (2016), or Santaeulalia-Llopis and Zheng (2018). In this model, the log unexplained annual income (y_t) allows for a permanent component P and a transitory component ν:

$$y_{i,t} = P_{i,t} + \nu_{i,t} \tag{1.6}$$

Following Meghir and Pistaferri (2004), I define the permanent component as a martingale process with serially uncorrelated innovations $\zeta_{i,t}$:

$$P_{i,t} = P_{i,t-1} + \zeta_{i,t} \tag{1.7}$$

Modelling the permanent component in this fashion ensures that the innovations do not disappear. In addition, I model the transitory component to ensure that these innovations do disappear. I represent ν as an MA(q) process with serially uncorrelated innovations. I do not assume the order q of the process, rather I determine it empirically.

$$\nu_{i,t} = \sum_{j=0}^{q} \theta_j \varepsilon_{i,t-j} \tag{1.8}$$

The shocks to the permanent and transitory components are both independent and identically distributed across time and households with mean 0 and their respective variance $\sigma^2_{\zeta_t}$ and $\sigma^2_{\nu_t}$: $\zeta_t \sim i.i.d(0, \sigma^2_{\zeta_t})$ and $\zeta_t \sim i.i.d(0, \sigma^2_{\nu t})$. I can therefore write the growth in unexplained income as:

$$\Delta y_{i,t} = \zeta_{i,t} + \Delta\nu_{i,t} \tag{1.9}$$

I use equation (1.9) to determine the order (q) of the transitory process. I assume that the permanent innovations $\zeta_{i,t}$ and transitory innovations $\varepsilon_{i,t}$ are uncorrelated. Thus, for the true order (q), the autocovariance between $\Delta y_{i,t}$ and $\Delta y_{i,t+s}$ with $s \leq q+1$ are significantly different from zero and the autocovariance with $s > q+1$ are equal to zero. The order q is determined by estimating the autocovariance of $\Delta y_{i,t}$ with various values of s and testing their significance. Table (1.5) reports the results for $s = 0, 1, 2, 3$ over the entire sample period.

The variance of income goes from 0.0718 in 1980 to 0.0873 in 2016 with a highest point of 0.1098 in 2008. The variance is consistently and unsurprisingly significant at the 1% level. The first-order autocovariance is, as expected, negative, and consistently significant at the 1% level. Note that after 1996, I take the autocovariance as $\Delta y_t * \Delta y_{t+2}$[22]. Although an

[22]Similar approximation are used for $s = 2$, and $s = 3$.

Table 1.5: Autocovariance of Idiosyncratic Income Growth

This table reports the autocovariances of unexplained income growth $Cov(\Delta y_{i,t}, \Delta y_{i,t+s})$ *for values of* $s = 0, 1, 2, 3$. *I use all households in the calculation. I treat the post 1996 waves as if consecutive. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

Year	$Var(\Delta y_t)$	$Cov(\Delta y_t, \Delta y_{t+1})$	$Cov(\Delta y_t, \Delta y_{t+2})$	$Cov(\Delta y_t, \Delta y_{t+3})$
1980	0.0718***	-0.0211***	-0.0021	-0.0029
1981	0.0756***	-0.0210***	-0.0066***	0.0032
1982	0.0705***	-0.0199***	-0.0046**	0.0029
1983	0.0803***	-0.0237***	-0.0030	-0.0033
1984	0.0812***	-0.0275***	0.0004	-0.0015
1985	0.0872***	-0.0245***	-0.0020	-0.0004
1986	0.0830***	-0.0223***	-0.0036	0.0001
1987	0.0760***	-0.0281***	0.0009	-0.0014
1988	0.0922***	-0.0291***	-0.0044**	0.0013
1989	0.0794***	-0.0236***	-0.0042**	0.0045*
1990	0.0823***	-0.0233***	-0.0032	-0.0008
1991	0.0785***	-0.0293***	0.0023	0.0009
1992	0.1003***	-0.0314***	-0.0058**	-0.0051**
1993	0.0959***	-0.0340***	-0.0007	0.0037
1994	0.0974***	-0.0262***	-0.0023	0.0043
1995	0.0890***	-0.0266***	-0.0076***	0.0041
1996	0.0936***	-0.0304***	-0.0014	-0.0002
1998	0.1083***	-0.0291***	-0.0027	0.0025
2000	0.1042***	-0.0322***	-0.0028	-0.0025
2002	0.1036***	-0.0344***	-0.0012	-0.0049*
2004	0.1098***	-0.0316***	-0.0012	-0.0028
2006	0.0979***	-0.0263***	0.0000	-0.0017
2008	0.1012***	-0.0315***	-0.0015	-0.0038
2010	0.1023***	-0.0276***	0.0024	-0.0013
2012	0.1008***	-0.0288***	-0.0022	
2014	0.0952***	-0.0276***		
2016	0.0873***			

approximation of the true autocovariance, I cannot get around this hurdle. I do address the biennial structure in the estimation process. The second and third order autocovariances are small and not consistently significant. In fact they are significant for a few years only. This result supports a MA(1) process for the transitory component. Although I adopt the MA(1) process in the estimation process, the rest of this section will present a special case of the process with MA(0) as in Heathcote, Perri, and Violante (2010) where the transitory shocks ν and the transitory innovations ε are not differentiated[23].

In order to determine how the consumption process responds to the income process, Blundell, Pistaferri, and Preston (2008) introduce two parameters ϕ and ψ that measures how much the consumption process reacts to permanent and transitory shocks, respectively. These two parameters are often referred to as partial insurance parameters. I describe the unexplained consumption growth process as:

$$\Delta c_{i,t} = \phi_{i,t}\zeta_{i,t} + \psi_{i,t}\varepsilon_{i,t} + \xi_{i,t} \tag{1.10}$$

ξ measures changes to the consumption process that are unrelated to changes in the income process, such as heterogeneous shifts in preferences. It is distributed as $\xi \sim i.i.d(0, \sigma_\xi^2)$. The four main parameters of interest are ϕ, ψ, ζ, and ε. The first two account for the degree of insurance. A lower loading factor means a higher degree of insurance. Indeed, in perfectly complete markets, both ψ and ϕ would be equal to 0. In perfectly incomplete markets, both parameters would be equal to 1. The consumption process described in equation 1.8 allows me to test whether the rapid growth in financial innovation has improved households' ability to insure against permanent or transitory shocks.

While I do observe $\Delta y_{i,t}$ in the PSID, I do not $\Delta c_{i,t}$. Indeed, I do not have a true consumption series, rather an imputed one. Following Blundell, Pistaferri, and Preston (2008), I add a measurement error parameter $u_{i,t}$ to account for potential biases introduced in the imputation. I can re-write equation (1.10) as:

$$\Delta c^*_{i,t} = \phi_{i,t}\zeta_{i,t} + \psi_{i,t}\varepsilon_{i,t} + \xi_{i,t} + u^c_{i,t} - u^c_{i,t-1} \tag{1.11}$$

While the theory on financial innovation affirms that it will increase risk-sharing, it does not differentiate between its permanent and transitory components. Furthermore, Krueger and Perri (2006) hypothesize that an increase in the variance of shocks should provide households more incentives to insure themselves. Blundell, Pistaferri, and Preston

[23] The MA notation can get cumbersome, I present this simpler process in the interest of space.

(2008) does find increased variance for the 1980s. The decomposition of the response to income shocks can perhaps explain why I observed the trends in overall risk-sharing in figures (1.3) and (1.4). Indeed, it is possible that the exogenous financial market development hypothesized by Krueger and Perri (2006) improved households' ability to insure against one type of shock while their ability to insure against another deteriorated, thus resulting in no observable trend in a simple transmission process.

1.4.4 Minimum Distance Estimation

To estimate the parameters outlined in the previous section, I follow the literature standard: the Minimum Distance Estimator used by Blundell, Pistaferri, and Preston (2008), Casado (2011), Ludwig (2016), Santaeulalia-Llopis and Zheng (2018), and Kubota (2020). The parameters are derived from the variances, covariances, and autocovariances of the unexplained income and consumption growth. I then estimate these parameters by minimizing the distance between the empirical moments and the moments predicted by the parametric consumption-income model described below. I start by describing the implications of equation (1.9). I write the variance and autocovariance of idiosyncratic income as:

$$Var(\Delta y_{i,t}) = Var(\zeta_{i,t}) + Var(\varepsilon_{i,t}) + Var(\varepsilon_{i,t-1}) \tag{1.12}$$

$$Cov(\Delta y_{i,t}, \Delta y_{i,t+1}) = -Var(\varepsilon_{i,t}) \tag{1.13}$$

Equation (1.13) explains why I expected the first autocovariance to be negative. As noted in Blundell, Pistaferri, and Preston (2008), equations (1.12) and (1.13) are sufficient to identify the variance of transitory and permanent income shocks. Adding consumption moments refines the estimation. I then describe the implications of equation (1.11). I can write the variance and autocovariance of idiosyncratic consumption as:

$$Var(\Delta c_{i,t}) = \phi_t^2 Var(\zeta_{i,t}) + \psi_t^2 Var(\varepsilon_{i,t}) + Var(\xi_i) + Var(u^c_{i,t}) + Var(u^c_{i,t-1}) \tag{1.14}$$

$$Cov(\Delta c_{i,t}, \Delta c_{i,t+1}) = -Var(u^c_{i,t}) \tag{1.15}$$

$\Delta c_{i,t}$ is linked to $\Delta c_{i,t+s}$ only through $u^c_{i,t}$. In Table 1.6, I calculate the variance and autocovariance of consumption for $s = 1, 2, 3$. The variance of consumption increases from 0.0348 in 1980 to 0.0438 in 2016 (growth that is similar to that of idiosyncratic income). The observation of note in Table 1.6 is the covariance between Δy_t and Δy_{t+1}. Firstly, it is negative as predicted by equation (1.15). Secondly, it is consistently significant, thus justifying the introduction of $u^c_{i,t}$ in the estimation. However, it is relatively small and consistent throughout the sample period. It does not invalidate the imputation procedure. Furthermore, the higher order autocovariances are small and not significant.

I can use equation (1.15) to identify the variance of the measurement error. However, I still cannot identify the partial insurance parameters and the preferences shock using consumption moments alone. I write the covariances between income and consumption as:

$$Cov(\Delta y_{i,t}, \Delta c_{i,t}) = \phi_t Var(\zeta_{i,t}) + \psi_t Var(\varepsilon_{i,t}) \tag{1.16}$$

$$Cov(\Delta y_{i,t+1}, \Delta c_{i,t}) = -\psi_t Var(\varepsilon_{i,t}) \tag{1.17}$$

From equation (1.17), I identify the transitory insurance parameter which is then used in equation (1.16) to identify the permanent insurance parameter. Although all parameters can be time-varying, note that the parameter describing shifts in preferences ξ will be restricted to be the same for all periods. This is done for two reasons. First, there are no theoretical justifications as to why it should vary. Second, it reduces the number of estimated parameters.

The first differenced moments and the lagged values are all defined for annual data. However, post-1996 PSID data are only available biennially. Ludwig (2016) shows that I can rely on second seasonal differences and adjust the moments to fit the data structure[24]. Given the yearly variances cannot be estimated, I assume that they are equal for the measured year and the gap year. The variances then become an average of two years. I can rewrite the moments in equations (1.12), (1.14), and (1.16) as:

$$Var(\Delta y_{i,t}) = 2Var(\zeta_{i,t}) + Var(\varepsilon_{i,t}) + Var(\varepsilon_{i,t-2}) \tag{1.18}$$

$$Var(\Delta c_{i,t}) = \phi_t^2 2Var(\zeta_{i,t}) + \psi_t^2 2Var(\varepsilon_{i,t}) + 2Var(\xi_i) + Var(u^c_{i,t}) + Var(u^c_{i,t-2}) \tag{1.19}$$

$$Cov(\Delta y_{i,t}, \Delta c_{i,t}) = \phi_t 2Var(\zeta_{i,t}) + \psi_t Var(\varepsilon_{i,t}) \tag{1.20}$$

In Table 1.7, I calculate the autocovariances of unexplained consumption and income growth. The covariance at $s = 0$ is the series plotted in figure (1.4). The main interest of this table is the $Cov(\Delta c_t, \Delta y_{t+2})$. A significant covariance between today's consumption and future income shocks would suggest superior information from households. Kaufmann and Pistaferri (2009) argue that the identification strategy requires a clear separation between unforeseen shocks and anticipated events. They use a micro-level household dataset which

[24]Similar adjustments are used by Santaeulalia-Llopis and Zheng (2018) to account for gaps in their data.

Table 1.6: Autocovariance of Idiosyncratic Consumption Growth

This table reports the autocovariances of unexplained consumption growth $Cov(\Delta c_{i,t}, \Delta c_{i,t+s})$ *for values of* $s = 0, 1, 2, 3$. *I use all households in the calculation. I treat the post 1996 waves as if consecutive. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

Year	$Var(\Delta c_t)$	$Cov(\Delta c_t, \Delta c_{t+1})$	$Cov(\Delta c_t, \Delta c_{t+2})$	$Cov(\Delta c_t, \Delta c_{t+3})$
1980	0.0348***	-0.0133***	-0.0009	0.0006
1981	0.0342***	-0.0122***	-0.0001	0.0012
1982	0.0329***	-0.0096***	-0.0009	-0.0005
1983	0.0317***	-0.0121***	0.0006	0.0001
1984	0.0357***	-0.0124***	-0.0014	
1985	0.0331***	-0.0103***		
1986	0.0330***			
1987				
1988				
1989				
1990	0.0357***	-0.0135***	-0.0010	0.0014
1991	0.0374***	-0.0130***	-0.0007	-0.0001
1992	0.0337***	-0.0104***	0.0005	0.0000
1993	0.0333***	-0.0116***	-0.0005	-0.0003
1994	0.0311***	-0.0116***	0.0004	-0.0011
1995	0.0345***	-0.0128***	0.0010	-0.0014*
1996	0.0348***	-0.0133***	-0.0004	0.0006
1998	0.0417***	-0.0110***	-0.0009	0.0003
2000	0.0375***	-0.0128***	-0.0010	-0.0000
2002	0.0443***	-0.0137***	-0.0000	-0.0016
2004	0.0483***	-0.0174***	-0.0007	0.0005
2006	0.0462***	-0.0165***	0.0007	0.0002
2008	0.0495***	-0.0186***	-0.0017	-0.0009
2010	0.0499***	-0.0157***	-0.0010	0.0004
2012	0.0464***	-0.0177***	-0.0011	
2014	0.0484***	-0.0150***		
2016	0.0438***			

contains expectations of future income and employment to disentangle insurance and information. They find that households do have some superior information which may overstates the degree of insurance. The PSID does not contain households' income expectations thus rendering me unable to differentiate the two effects. However, I find that the covariance between $\Delta c_{i,t}$ and $\Delta y_{i,s}$ is not significant for $s = 2$ and $s = 3$. This result implies that households are not likely to possess superior information and adjust consumption ahead of shocks. I can be reasonably certain that the parameters I measure are indeed insurance parameters.

To estimate the parameters, I first collect the residual consumption and income growth for each household. The sample period in first difference is 1980 to 2016. I therefore have 27 possible income growth data points and 24 possible consumption growth data points for each households[25]. For each household, I stack the consumption values (Δc_i) and the income moments (Δy_i) in the vector x_i. I create two vectors for each households d_i^c and d_i^y containing 1 for non-missing years, 0 otherwise, for consumption growth and income growth respectively. These two vectors are stacked into d_i which has the same dimensions as x_i.

$$x_i = \begin{pmatrix} \Delta c_i \\ \Delta y_i \end{pmatrix} \quad \text{and} \quad d_i = \begin{pmatrix} d_i^c \\ d_i^y \end{pmatrix}$$

I calculate the variance-covariance matrix of unexplained consumption and income growth M:

$$M = \left(\sum_i x_i x_i' \right) \oslash \left(\sum_i d_i d_i' \right)$$

where $\oslash$ represents an element-wise division. M is a 51 by 51 matrix (24 observations from consumption, 27 from income). M contains all estimates of the variance, covariances, and autocovariances of unexplained consumption and income growth. This formulation allows me to get around the unbalanced nature of the PSID data in a simple manner and account for years where I am unable to impute consumption data or households dropping out of the sample. I define m as the half-vectorized matrix M. m contains all unique second moments of the data. I define the objective function as $f(\Lambda)$ implied by the consumption income process defined above. Λ contains the parameters that I am identifying (insurance parameters, variance of shocks, measurement errors, and preferences shift). I estimate the parameters by simply minimizing the distance between:

[25] 27 years results from the biennial structure of the PSID, and 24 years results from missing consumption in 1987,1988, and the inability to first difference 1989.

Table 1.7: Covariance of Idiosyncratic Consumption Growth and Idiosyncratic Income Growth

This table reports the covariances of unexplained consumption growth and income growth $Cov(\Delta c_{i,t}, \Delta y_{i,t+s})$ *for values of* $s = 0, 1, 2, 3$. *I use all households in the calculation. I treat the post 1996 waves as if consecutive. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

Year	$Cov(\Delta c_t, \Delta y_t)$	$Cov(\Delta c_t, \Delta y_{t+1})$	$Cov(\Delta c_t, \Delta y_{t+2})$	$Cov(\Delta c_t, \Delta y_{t+3})$
1980	0.0082***	0.0007	-0.0019	-0.0002
1981	0.0079***	-0.0003	-0.0014	-0.0006
1982	0.0091***	-0.0031**	-0.0005	0.0009
1983	0.0099***	-0.0019	0.0011	-0.0005
1984	0.0104***	-0.0049***	0.0013	-0.0013
1985	0.0104***	-0.0014	-0.0002	0.0021
1986	0.0104***	-0.0023*	-0.0030**	0.0035**
1987				
1988				
1989				
1990	0.0087***	-0.0003	-0.0004	0.0004
1991	0.0092***	-0.0042***	0.0014	0.0007
1992	0.0113***	-0.0002	-0.0014	-0.0004
1993	0.0092***	-0.0018	0.0005	-0.0012
1994	0.0083***	-0.0024**	0.0015	0.0007
1995	0.0084***	-0.0006	0.0002	0.0007
1996	0.0083***	-0.0044***	0.0005	0.0003
1998	0.0100***	-0.0004	0.0002	0.0022
2000	0.0099***	-0.0047***	0.0023	0.0020
2002	0.0108***	-0.0042***	0.0015	-0.0019
2004	0.0137***	-0.0016	-0.0015	-0.0022
2006	0.0100***	-0.0001	-0.0006	-0.0010
2008	0.0136***	-0.0053***	0.0034**	0.0017
2010	0.0134***	-0.0034**	-0.0024	0.0008
2012	0.0132***	-0.0027	-0.0001	
2014	0.0110***	-0.0033*		
2016	0.0115***			

$$\min_{\Lambda}(m - f(\Lambda))' A(m - f(\Lambda)), \tag{1.21}$$

For inference purposes, I rely on the variance-covariance matrix of m defined as:

$$V = \left(\sum_{i=1}^{N} ((m_i - m)(m_i - m)') \otimes vech(d_i d_i') \right) \oslash \left(\sum_{i=1}^{N} vech(d_i d_i') \right),$$

and standard errors following Chamberlain (1984):

$$\widehat{Var(\hat{\theta})} = (G'AG)^{-1} G'AVAG(G'AG)^{-1}, \tag{1.22}$$

with $G = \dfrac{\partial f(\theta)}{\partial \theta}|_{\theta=\hat{\theta}}$ is the Jacobian matrix evaluated at the estimated parameters $\hat{\theta}$. Note that the matrix A in equations (1.21) and (1.22) is an identity matrix in Equally Weighted Minimum Distance (EWMD), V^{-1} for Optimal Minimum Distance (OMD), or $diag(V^{-1})$ for Diagonally Weighted Minimum Distance (DWMD).

As certain parameters depend on autocovariances, I must make certain assumptions to ensure stability. I will assume that the variance of transitory shocks of the last measurable period is equal to the variance of transitory shocks for the last available year. This allows me to identify the variance of permanent shocks in the last available year. I make the same assumption for the variance of the measurement error. I also assume that the lag variance of the measurement error in the first year is equal to the variance of the measurement error in the first available year. I must assume that the last available year and the first available year are equal to their respective lags for the missing intervals[26]. Blundell, Pistaferri, and Preston (2008) replace missing moments with the whole sample average. Considering the panel has a longer time dimension and I am interested in the time series properties of the estimates, replacing missing moments with the most recent non-missing values is more appropriate.

1.4.5 Results of the MDE

I first restrict the partial insurance parameters to be the same for the whole sample period. I estimate ϕ and ψ for the same groups used in table 1.2. The results are presented in table 1.8. Panel A reports the diagonally-weighted estimates using imputed data while panel B reports the diagonally-weighted estimates using PSID measured data. When estimating with the PSID measured data, I maintain the measurement errors term $u_{i,t}$ in the moments.

[26] For the sample period, I have $Var(\varepsilon_{2016}) = Var(\varepsilon_{2014})$, $Var(u_{1979}) = Var(u_{1980})$, $Var(u_{1986}) = Var(u_{1985})$, $Var(u_{1989}) = Var(u_{1990})$, and $Var(u_{2016}) = Var(u_{2014})$.

Although the term is added to account for measurement error arising from imputation, I feel it is prudent to leave it in. Indeed, the PSID data remains subject to potential measurement errors. Before going into the results of table 1.8, I can mention results not reported[27]. The measurement error parameters are precisely measured indicating that they absorb a fair amount of the cross-section of consumption. The parameter measuring preference ξ is consistently significant and varies between 0.03 and 0.08 across specifications. The serial correlation of the transitory shock (as I estimate a MA(1) process) is also significant and varies between 0.09 and 0.10 between specifications.

As pointed out by Commault (2020), estimating subgroups presents several advantages. In the baseline estimation, I assume that all households have the same income and consumption process and the same parameters. When estimating subgroups separately, I can relax the assumption of homogeneous parameters. Indeed, this separate estimation yields more precise estimates of θ and ξ. Furthermore, the comparison between groups is made easier as differences observed cannot be caused by different income processes.

My estimates are consistent with the literature. I find that across specifications, the permanent insurance parameters is always larger than the transitory insurance parameter. This result suggests that transitory shocks are easier to insure than permanent ones. For all households, the permanent insurance parameter ϕ is 0.2586 and significantly from 0 at the 1% level. This implies that a 10% shock to permanent income will cause a 2.586% shift in permanent consumption. The transitory insurance parameter ψ is 0.0677 and is significantly different from 0 at the 1% level, indicating that that a 10% shock to transitory income will cause a 0.677% shift in transitory consumption. This is consistent with the result of (1.2). If we think of β in equation (1.3) as a weighted average of ϕ and ψ with the respective variances as weights, then the variance of transitory shocks should account for approximately 75% of the variance of total shock[28].

$$\hat{\beta} = \frac{cov(\Delta c, \Delta y)}{\sigma_{\Delta y}} = 0.1123 \approx \frac{\phi_t Var(\zeta_{i,t}) + \psi_t Var(\varepsilon_{i,t})}{\sigma_{\Delta y}}$$

Unlike Blundell, Pistaferri, and Preston (2008) who fail to reject perfect transitory insurance, the partial insurance parameter is significantly different from 0 at the 1% level for all specifications except for entrepreneurs. Recall from tables (1.2) and (1.3) that entrepreneurs have the lowest risk-sharing coefficients. Furthermore from table (1.2), I show that different groups have significant differences in risk-sharing coefficients. From table (1.8),

[27]Full results are available upon request.

[28]This is of course an approximation as the estimation of the partial parameters depend of the value of ξ, and θ absent of the estimation of equation (1.3).

Table 1.8: Minimum Distance Partial Insurance Estimates

This table reports the results from the minimum distance estimation presented in section 1.4.4.The parameters are estimated using equation (1.20) and the standard errors using equation (1.21). Only the variance of shocks and the variance of the measurement errors are time-varying. I use the diagonal of the inverse of the variance-covariance matrix of the empirical moments as the weighting matrix. Panel A uses the imputed consumption series while Panel B uses the PSID measured data. I still include the term $u_{i,t}$ *when estimating Panel B. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	All households	Employees	Entrepreneurs	Market participants	No participation	Stock owners	Non stock owners	No transfer receipts	Transfer recipients	High Wealth	Low Wealth
Panel A. Diagonally weighted estimates - Imputed Data											
ϕ	0.2586***	0.2865***	0.1974***	0.1643***	0.3847***	0.1475***	0.3263***	0.2080***	0.3573***	0.1625***	0.3291***
Perm. ins.	(0.0129)	(0.0152)	(0.0321)	(0.0212)	(0.0291)	(0.0222)	(0.0217)	(0.0176)	(0.0401)	(0.0173)	(0.0238)
ψ	0.0677***	0.0683***	0.0142	0.0642***	0.0423***	0.0796***	0.0552***	0.0683***	0.0767***	0.0382***	0.0602***
Trans. ins.	(0.0107)	(0.0136)	(0.0196)	(0.0182)	(0.0154)	(0.0232)	(0.0142)	(0.0175)	(0.0172)	(0.0136)	(0.0170)
Panel B. Diagonally weighted estimates - Real Data											
ϕ	0.2730***	0.3352***	0.1905**	0.1716***	0.3598***	0.1703***	0.2878***	0.2509***	0.3219***	0.1325***	0.3081***
Perm. ins.	(0.0252)	(0.0323)	(0.0789)	(0.0385)	(0.0535)	(0.0438)	(0.0345)	(0.0321)	(0.0629)	(0.0304)	(0.0448)
ψ	0.1481***	0.1197***	0.0291	0.1031**	0.1986***	0.1283*	0.1667***	0.0542	0.2142***	0.1087***	0.1915***
Trans. ins.	(0.0276)	(0.0330)	(0.0468)	(0.0451)	(0.0413)	(0.0696)	(0.0357)	(0.0511)	(0.0421)	(0.0380)	(0.0474)

the differences appear to be caused by differences in permanent insurance. Indeed, the levels of transitory insurance appears to be similar across groups (employees/entrepreneurs being the exception). A similar relationship is found for direct participation where the permanent insurance parameter of stockowners is 0.1475 vs 0.3263 for non stockonwers.

Transfer recipients' transitory insurance parameter is 0.0683 versus 0.0767 for households who do not receive any transfer income, suggesting the difference in risk-sharing is coming from differences in permanent insurance. However, the permanent insurance parameters are different across the two groups (0.2080 for households not receiving transfers vs 0.3573 for households who do). This would suggest that formal government transfers are effective at providing consumption insurance for transitory shocks; less so for permanent shocks. This is somewhat not surprising. Indeed a transitory shock such as a short period of unemployment would be covered by unemployment insurance. However, unemployment insurance would run out and not cover long-term unemployment, which I could characterize as a permanent shock to income.

High wealth households have a lower permanent insurance parameter (0.1625 vs 0.3291) than low wealth households. The difference in transitory insurance is also more pronounced for high and low wealth households with wealthy households being less sensitive to transitory shocks. These patterns are similar for entrepreneurs/employees and participants/nonparticipating households. This casts some doubt on the precise nature and channel of the observed differences. Indeed, entrepreneurs are more likely to have higher wealth and/or participate. Similarly, high wealth individuals are more likely to participate in the stock market. Unfortunately, data constraints make it impossible to disentangle the precise effects[29].

In panel B of table (1.8), I only consider the sample period 2000 to 2016 and estimate the partial insurance model using the actual consumption variables measured in the PSID. The permanent insurance parameters are reasonably similar across groups (0.2586 vs 0.2730 for all households). However, I do obverse large differences in the transitory insurance parameter, more specifically an increase. I fail to find a significant transitory parameter for nontransfer recipients in panel B, though the point estimate is similar to panel A. These results are, for the most part, consistent with the results of table (1.2) where the risk-sharing coefficient went from 0.1123 to 0.1508. This could indicate two things. First, the autocovariance between income and consumption is smaller for imputed consumption than for actual consumption, thus overstating the degree of transitory insurance. Second, transitory

[29]Ideally, I could condition market participation or entrepreneurship on wealth level. This conditioning results in very small groups rendering my estimates unstable.

insurance has deteriorated in the second half of the sample period considered. The second explanation seems more likely. Indeed, I have observed a slight increase in the risk-sharing coefficient. I formally test change in insurance parameters in table (1.9).

Blundell, Pistaferri, and Preston (2008) did not find a significant difference of insurance in their sample period. However, the sample period in question (1980-1992) is relatively short and not particularly characterized by changes in macroeconomic conditions. My sample period is longer and is characterized by profound changes notably to financial markets. I therefore test whether I can observe changes in households' ability to smooth out consumption shocks. The empirical exercise is similar to Santaeulalia-Llopis and Zheng (2018). They test whether Chinese households' ability to insure permanent and transitory risk has changed between the 1989-1997 and 1998-2009 periods. The periods are chosen as they are characteristic of different economic environments marked by reform. The determination of a "pre" and "post" period is not as clear in my sample period. I do not have exogenous shocks that would help make a determination so I pick 1980-1996 as my preperiod and 1998-2016 as my postperiod. Although arbitrary, I have several reasons for choosing this breakpoint. Firstly, it marks the first year where the wealth module is consistently administered and I am not relying on an imperfect identification of market participants. If nothing else, it will make participation and wealth estimation more robust. Secondly, it also marks the beginning of the PSID biennial structure, thus making it a natural breakpoint. Finally, and perhaps more importantly, these two periods appear to be characterized by different regimes of permanent and transitory income shocks. I argue that if differences are to be found in the degree of insurance, it is most likely to be around the chosen breakpoints.

In figure (1.8), I plot the estimates of $Var(\zeta)$ and $Var(\varepsilon)$. I see that the variance of transitory shocks is increasing from 1980 to 1996 when it then stabilizes for the 1998-2016 period. The variance of permanent shocks increases through the first period, then sharply falls and stabilizes over the second period. Note that I estimate these parameters while restricting the insurance parameters to be different in the pre and post period[30]. The variance of shocks are similar up to 1993 when they diverge.

I therefore estimate 2 transitory and 2 permanent insurance parameters. I report the result in table 1.9 along with the p-value for a t-test of difference in parameters. I find that households' ability to insure against permanent shocks seems to have increased. ϕ_{Pre97} is equal to 0.2856 vs 0.2371 for the post period. This result is significant at the 5% level. The point estimate of ψ_{Post97} is larger than the preperiod coefficient, though the difference

[30]Similar patterns can be seen when estimating static insurance parameters, pre and post parameters, or simply estimating the variances alone using only income moments

Figure 1.8: Variance of permanent and transitory income shocks - All households

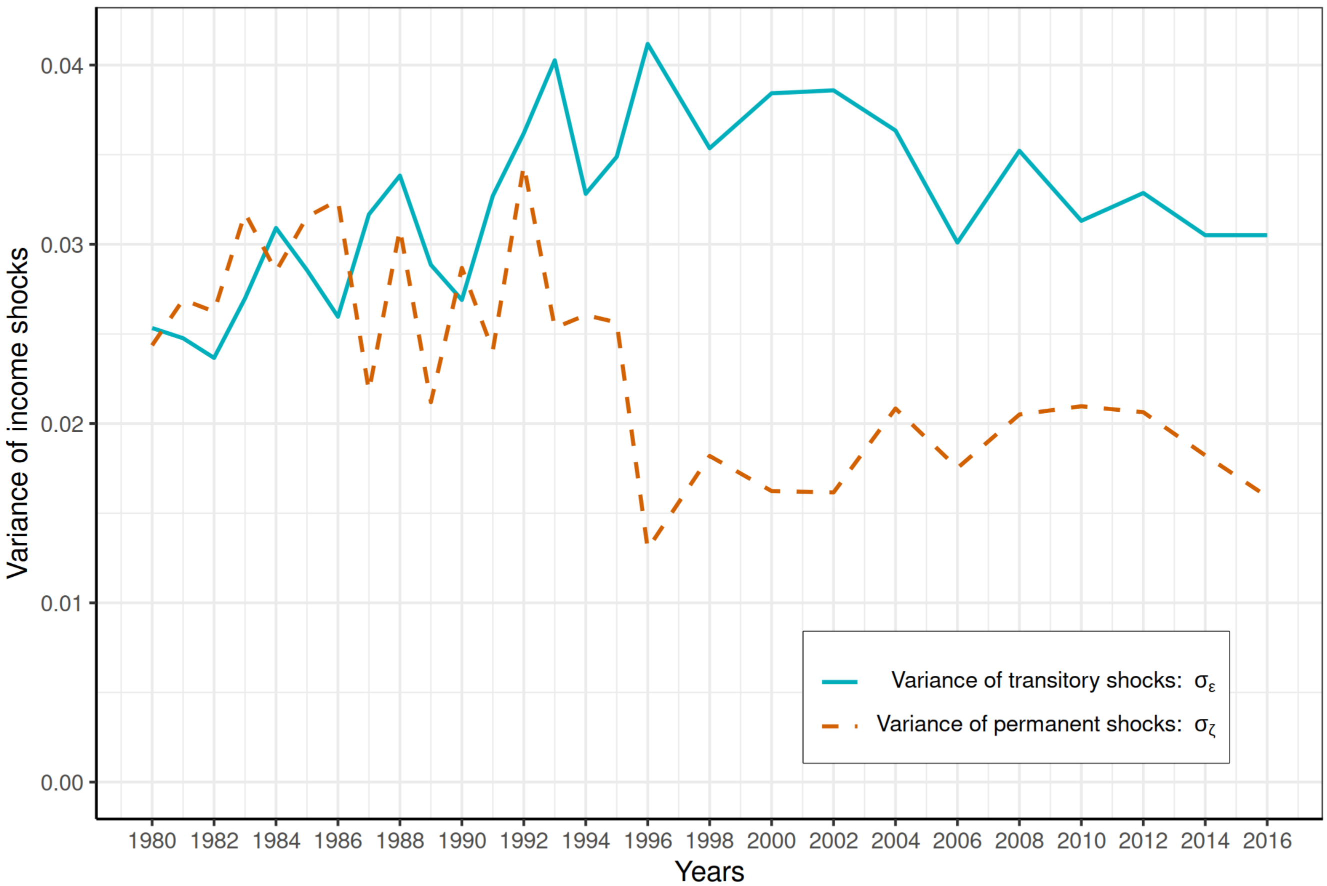

is statistically insignificant. The post period coefficients are different from the coefficients estimated in table (1.8) panel. It is to be noted that table (1.8) is estimated only with 2000 to 2016 data. Thus, the time-invariant parameters (ξ and θ) are different. The impact of past transitory shocks is likely to cause the observed differences. Although different point estimates, the general patterns observed in panel B of table (1.8) and table (1.9) are similar.

Looking at subgroups, I find that market participants experienced a large increase in their permanent insurance; 0.2340 vs 0.1312 with a p-value of 0.0151. nonparticipating households pre-1997 ϕ parameter is 0.3856 and the post-1997 parameter is 0.3844. The difference is not statistically significant. Both, however, have experienced a decrease in their ability to smooth transitory shocks. The transitory parameters are statistically insignificant prior to 1997 and close to 0 but highly significant afterwards and equal to 0.0880 and 0.0729 for participating and nonparticipating households respectively. In figure (1.9), I plot the the variances of income shocks for participating and nonparticipating households. The y-axis for panels A and B have the same scale to make easy comparisons. The variance of

Table 1.9: Minimum Distance Time Varying Partial Insurance Estimates

*This table reports the results from the minimum distance estimation presented in section 1.4.4.The parameters are estimated using equation (1.20) and the standard errors using equation (1.21). The variance of shocks and the variance of the measurement errors vary with each wave. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	All households	Employees	Entrepreneurs	Market participants	No participation	Stock owners	Non stock owners	No transfer receipts	Transfer recipients	High Wealth	Low Wealth
ϕ_{pre97}	0.2856***	0.3069***	0.1892***	0.2340***	0.3856***	0.2127***	0.3675***	0.2837***	0.3948***	0.2079***	0.3237***
	(0.0174)	(0.0207)	(0.0370)	(0.0377)	(0.0327)	(0.0374)	(0.0294)	(0.0277)	(0.0600)	(0.0235)	(0.0275)
ϕ_{post97}	0.2371***	0.2722***	0.1234***	0.1312***	0.3844***	0.1166***	0.2992***	0.1606***	0.3440***	0.1119***	0.3181***
	(0.0166)	(0.0193)	(0.0383)	(0.0226)	(0.0409)	(0.0256)	(0.0256)	(0.0209)	(0.0476)	(0.0203)	(0.0321)
p-value test of equal ϕ	0.0298	0.1819	0.2011	0.0151	0.9789	0.0312	0.0485	0.0002	0.4659	0.0011	0.8862
ψ_{pre97}	0.0528***	0.0500***	0.0316	0.0238	0.0197	0.0527*	0.0182	0.0541**	0.0665***	0.0174	0.0232
	(0.0151)	(0.0189)	(0.0324)	(0.0292)	(0.0208)	(0.0316)	(0.0203)	(0.0225)	(0.0242)	(0.0190)	(0.0242)
ψ_{post97}	0.0794***	0.0858***	0.0149	0.0880***	0.0729***	0.0956***	0.0884***	0.0696**	0.0847***	0.0574***	0.1074***
	(0.0149)	(0.0194)	(0.0276)	(0.0238)	(0.0222)	(0.0343)	(0.0186)	(0.0274)	(0.0243)	(0.0191)	(0.0237)
p-value test of equal ψ	0.2036	0.1844	0.6975	0.0882	0.0783	0.3616	0.0095	0.6650	0.5966	0.1328	0.0127

permanent shocks for both subgroups averages around 0.02. Furthermore, both groups have increasing transitory shocks variances. However, nonparticipating households experience larger transitory shocks. Once again, these results are in direct contradiction of Guvenen (2007). Not only are market participants sharing risk more completely (through better permanent insurance against permanent shocks), the gap is getting wider. Furthermore, market participants are not exposed to larger income shocks.

Figure 1.9: Variance of permanent and transitory income shocks for participating and non-participating households

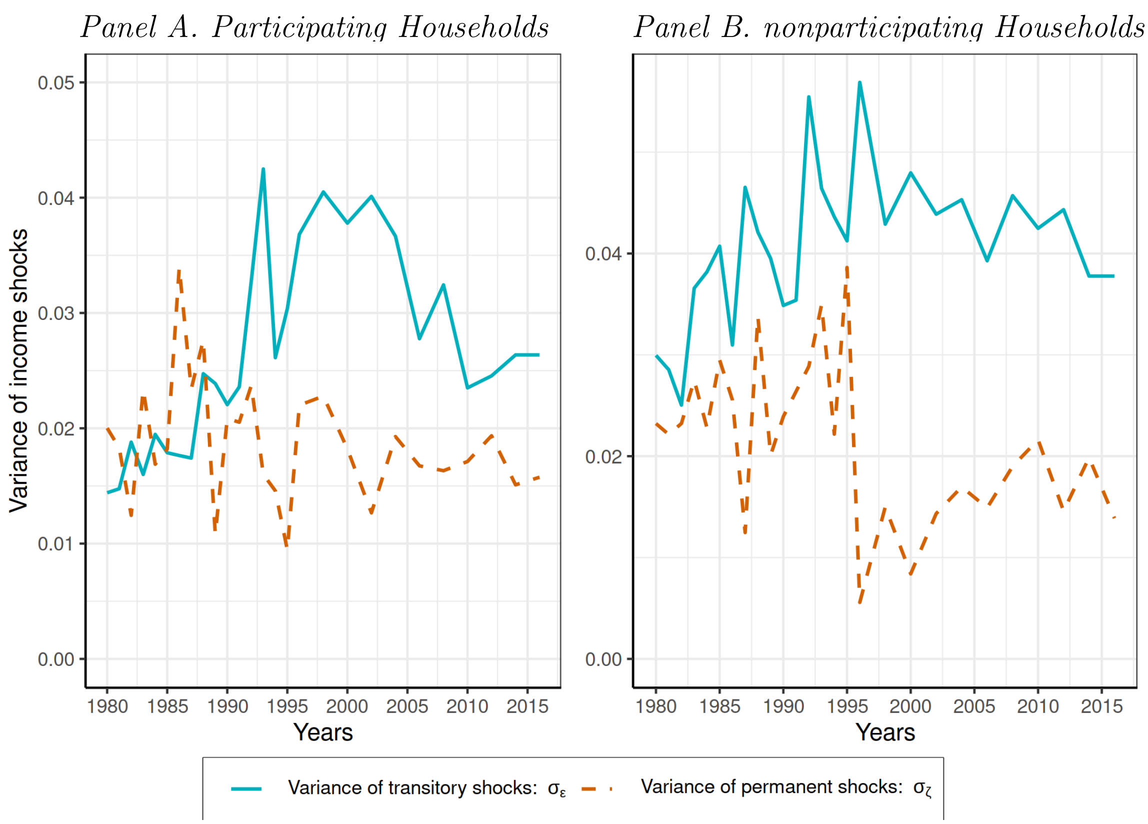

High wealth households experience a significant increase in permanent insurance and no significant change in transitory insurance. The opposite is true for low-wealth households that experience no change in permanent insurance and a significant decrease in transitory insurance. I fail to detect any change in the insurance environment of transfer recipients. Both permanent and transitory tests have high p-values. Non transfer recipients have a significantly lower permanent loading but no transitory insurance changes. Entrepreneurs and employees are two subgroups that stand out in that neither experience any changes to their ability to smooth shocks. Indeed, p-values for either group are extremely large. Entrepreneurs are also the only group for which I fail to reject the null of full transitory insurance.

In figure (1.10), I plot the variances of permanent and transitory shocks for entrepreneurs and employees. I see that the variances of shocks for employees are smaller and smoother. Relative to entrepreneurs, employees' variance of transitory shocks is barely increasing. However, entrepreneurs have a variance of permanent shocks twice as large as employees in the early years of my sample and experience a dramatic decrease throughout the sample. I can see that employees' shocks are similar to that of the whole sample. This is not surprising considering entrepreneurs make up a small proportion of my sample.

Figure 1.10: Variance of permanent and transitory income shocks for Entrepreneurs and Employees

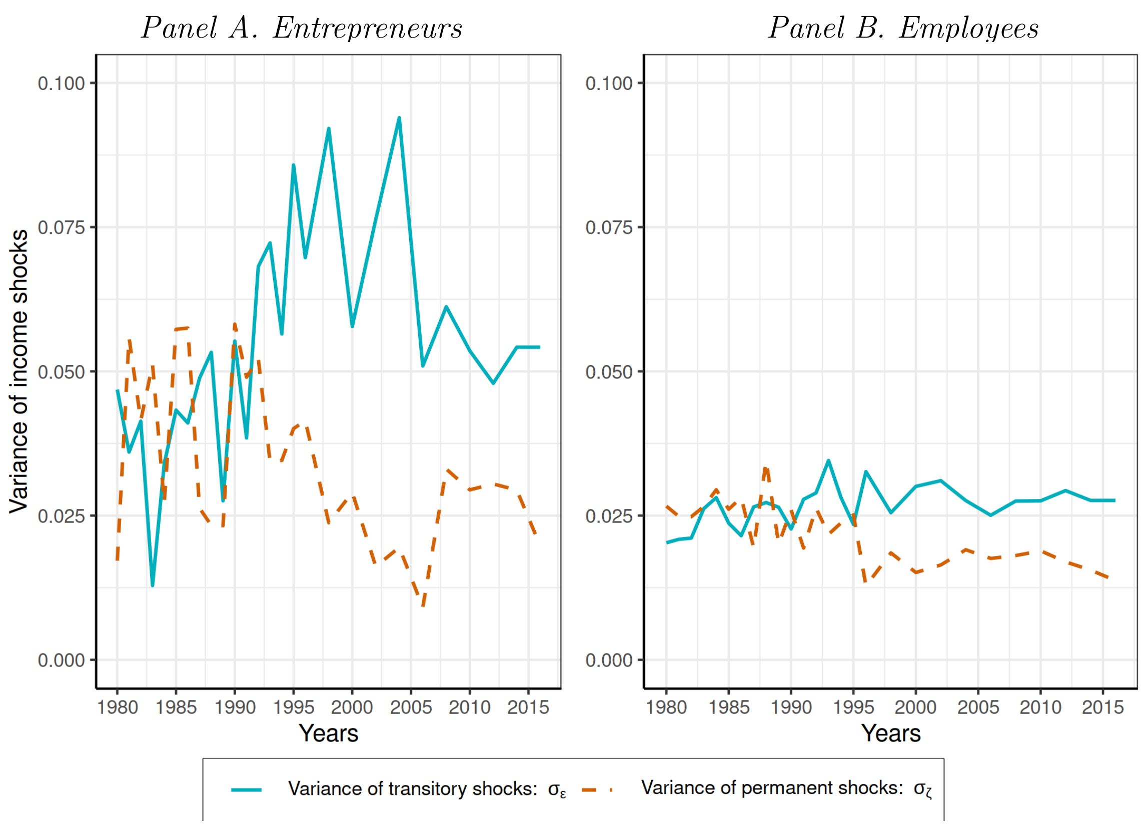

To get a clearer picture of the trend in permanent and transitory insurance parameters, I let both parameters vary through time and plot the resulting parameters. To ease computation, I group parameters with 2 to 3 years groups. While I failed to find a significant difference in transitory insurance, figure (1.11) tells a different story. $\phi_{1980-1983}$ is not significantly different from 0 while the others are. I do see an increasing trend in the transitory insurance parameters. However, the increase is gradual and thus not picked up by the test around the 1997 breakpoint. There is also a large increase for $\phi_{2010-2016}$. This is consistent with the financial crisis. This spike also explains why some groups have significant differences in transitory insurance in table (1.9). From previous results in figure (1.7), I know that low-wealth households were more affected by the financial crisis explaining the deterioration of transitory insurance. From figure (1.5), I know that both participating and nonparticipating households were similarly affected by the financial crisis with different timing. The test in table (1.9) would be unable to detect such differences.

Figure 1.11: Partial Insurance Parameters - All households

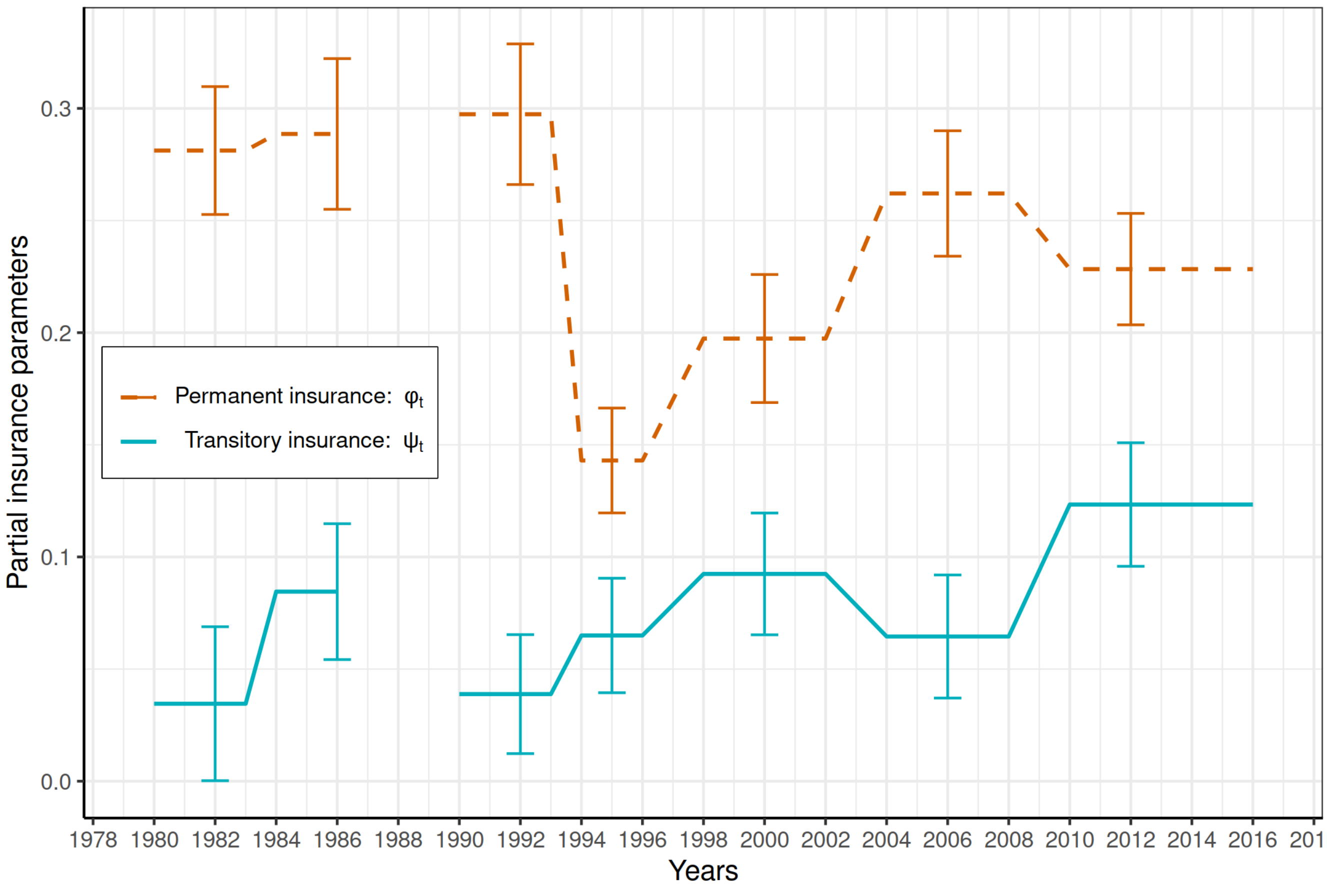

The trend in permanent insurance is however similar to table (1.9) results The parameters are on average 0.28 in the preperiod and 0.23 in the postperiod. There is a large drop in permanent insurance in the 1994-1996 year group. This significant shift is likely caused by a small spike in the variance of income. It is unlikely to have affected my results. Indeed, if anything, it is pushing the pre-1997 average permanent insurance down thus reducing the difference between the two period. Overall, I find a small improvement in permanent insurance and a more significant decline in transitory insurance. The combination of these two trends are consistent with the trend observed in figure (1.3).

As pointed out before, subgroups are not exclusive and the precise insurance mechanisms are difficult to highlight. To provide alternative explanations for my previous results, I consider different measures of income. Specifically, I repeat the analysis of table (1.9) with disposable income net of financial income, disposable income net of business income, and disposable income net of transfer income. I do not have information on the specific contribution of income categories on a household's tax liability, so I must assume that taxes paid on subcategories of income are a linear function of total income. In other words I make the following adjustments:

$$Y_{adjusted} = (Y_{Total} - Y_{Subcategory}) - Taxes * \left(\frac{Y_{Total} - Y_{Subcategory}}{Y_{Total}}\right) \tag{1.23}$$

I then re-estimate $\Delta y_{i,t}$ using each new income variable[31]. Financial income is the sum of income from dividends, interest, trust funds, and royalties. Removing this stream of income allows me to see if the observed differences across participating and nonparticipating households are indeed caused by financial income as a smoothing component. I recognize the limitation of this proxy. It does include revenues streams that are available to non-market participants. However, dividend income, as a separate item, is not consistently measured. Furthermore, the proportion of firms paying dividends has declined (Fama and French, 2001). Thus the proportion income stemming from dividends must also be declining. However imperfect, this measure is the best PSID has to offer for my test. Business income is the sum of labor and asset income derived from unincorporated business or farm for all household members. Removing business income from total family income allows me to test whether the lower marginal propensity to consume out of permanent or transitory shocks of market participants or high wealth individuals is caused by them being entrepreneurs. Transfer income[32] is also removed to investigate how formal government transfers contribute to risk-sharing.

The effect of removing these components of income is not necessarily straightforward. In the case of transfers, the variance of permanent and transitory shocks would increase and the insurance parameters would decrease (i.e., appear better). This is because I am artificially introducing shocks (removing buffers) but consumption still reflects the smoothing benefits (Blundell, Pistaferri, and Preston, 2008). Government transfers are designed to be received when households experience shocks to their income. Business and financial income are not. On one hand, if financial and business income are countercyclical to wage risk,

[31] I require households to have at least \$100 in $Y_{adjusted}$

[32] Note that I include social security income to the PSID transfer income variable (when not included) to have a consistent measure throughout the sample

then removing them would increase the variance of permanent and transitory shocks. The insurance parameters would appear lower. On the other hand, if these streams of income are sources of income risk, then the variance shocks would be lower than the baseline levels of table (1.9). This lower variance would result in higher (i.e worst insurance parameters). This is easy to see in my sample transmission framework where $\beta = cov(\Delta c, \Delta y)/\sigma_{\Delta y}$. A lower $\sigma_{\Delta y}$ results in a higher β. In the partial insurance framework, I can easily show that the parameters ϕ and ψ are:

$$\psi = \frac{E[\Delta c_t \Delta y_t]}{E[\Delta y_t \Delta y_{t+1}]}$$
$$\phi = \frac{E[\Delta c_t(\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})]}{E[\Delta y_t(\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})]}$$

where, $E[\Delta y_t \Delta y_{t+1}]$ and $E[\Delta y_t(\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})]$ are the respective variance of transitory (σ^2_ε) and permanent shocks (σ^2_ζ). Thus, lower variances would also result in larger insurance parameters. The results are presented in table (1.10). I first consider nonfinancial income in panel A. In the first column, for all households, I find that the permanent insurance parameters are higher than their baseline counterparts. The pre-1997 coefficient is .3094 (vs 0.2856) and the post-1997 coefficient is 0.2511 (vs 0.2371). Both are statistically significant at the 1% level and their difference is statistically significant at the 5% level. The higher insurance parameters would imply lower variance of shocks. Financial may increase households' total income risk. The transitory insurance parameters for both periods are also higher (0.0612 vs 0.0528 and 0.0833 vs 0.0794). While the parameters are statistically significant, their difference is not. It must be noted that the differences in point estimates between tables (1.9) and (1.10) are indeed small when estimating with the whole sample. That is caused by the fact that financial income, on average, makes a small proportion of households' income. Thus, its removal was unlikely to cause a big change in the estimation.

The main group of interest in panel A is participating households. I find that when excluding financial income, ϕ_{pre97} goes from 0.2340 to 0.3199. This would indicate that financial income is a significant source of income risk for market participants. However, ϕ_{post97} remains unchanged (0.1307 vs .1312). Furthermore, the transitory parameters are also unchanged from table (1.9). This result indicates that financial income is a source of permanent shocks for participating households and does not contribute to transitory income risk. Furthermore, this particular source of risk does not have the same effect in the second half of the sample. I have mentioned above the imperfect nature of this proxy. Consider nonparticipating households. Their permanent parameters have also increased, albeit to

Table 1.10: Partial Insurance Estimates Net of Financial, Business, or Transfer Income

*This table reports the results from the minimum distance estimation presented in section 1.4.4.The parameters are estimated using equation (1.20) and the standard errors using equation (1.21). The variance of shocks and the variance of the measurement errors vary with each wave. Financial income is measured as the sum of income from dividends, interest, trust funds, and royalties. Business income is measured as the sum of business and labor income for non-incorporated business and farms. Transfer income is the sum of all transfer categories plus social security income. Refer to the PSID codebook for more detail. All household members' income are taken into account. I assume that each subcategory of income contributes equally to the overall tax liability of the household. I require households to have more than $100 in the income category to be included in the estimation. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	All households	Employees	Entrepreneurs	Market participants	No participation	Stock owners	Non stock owners	No transfer receipts	Transfer recipients	High Wealth	Low Wealth
Panel A. Disposable nonfinancial Income											
ϕ_{pre97}	0.3094***	0.3122***	0.2306***	0.3199***	0.4194***	0.2930***	0.3991***	0.3365***	0.3101***	0.2392***	0.3237***
	(0.0188)	(0.0212)	(0.0455)	(0.0453)	(0.0351)	(0.0445)	(0.0318)	(0.0301)	(0.0512)	(0.0244)	(0.0281)
ϕ_{post97}	0.2511***	0.2863***	0.1348***	0.1307***	0.4114***	0.1363***	0.3202***	0.1724***	0.3565***	0.1183***	0.3246***
	(0.0175)	(0.0204)	(0.0363)	(0.0227)	(0.0443)	(0.0276)	(0.0271)	(0.0215)	(0.0494)	(0.0210)	(0.0325)
p-value test of equal ϕ	0.0141	0.3404	0.0909	0.0001	0.8700	0.0021	0.0315	0.0000	0.4806	0.0001	0.9824
ψ_{pre97}	0.0612***	0.0651***	0.0186	0.0296	0.0226	0.0540*	0.0226	0.0394*	0.0840***	0.0205	0.0262
	(0.0152)	(0.0194)	(0.0311)	(0.0276)	(0.0208)	(0.0299)	(0.0203)	(0.0228)	(0.0261)	(0.0198)	(0.0242)
ψ_{post97}	0.0833***	0.0878***	0.0140	0.0870***	0.0749***	0.0806**	0.0845***	0.0648**	0.0854***	0.0618***	0.1077***
	(0.0147)	(0.0192)	(0.0277)	(0.0232)	(0.0214)	(0.0324)	(0.0183)	(0.0271)	(0.0241)	(0.0191)	(0.0238)
p-value test of equal ψ	0.2928	0.4050	0.9137	0.1166	0.0772	0.5530	0.0207	0.4777	0.9696	0.1332	0.0162
Panel B. Disposable Non-Business Income											
ϕ_{pre97}	0.2006***			0.1176***	0.3361***	0.0887***	0.3092***	0.1525***	0.3259***	0.1028***	0.2992***
	(0.0152)			(0.0263)	(0.0317)	(0.0256)	(0.0281)	(0.0196)	(0.0513)	(0.0164)	(0.0296)
ϕ_{post97}	0.2285***			0.1066***	0.3729***	0.1155***	0.3020***	0.1371***	0.3792***	0.0740***	0.3174***
	(0.0169)			(0.0199)	(0.0414)	(0.0231)	(0.0274)	(0.0186)	(0.0515)	(0.0178)	(0.0329)
p-value test of equal ϕ	0.1912			0.7270	0.4203	0.4163	0.8375	0.5574	0.4227	0.2124	0.6537
ψ_{pre97}	0.0576***			0.0310	0.0123	0.0754***	0.0168	0.0568**	0.0627***	0.0281	0.0334
	(0.0150)			(0.0272)	(0.0198)	(0.0287)	(0.0195)	(0.0228)	(0.0228)	(0.0191)	(0.0227)
ψ_{post97}	0.0695***			0.0481**	0.0807***	0.0267	0.0838***	0.0256	0.0623***	0.0340*	0.1114***
	(0.0143)			(0.0229)	(0.0221)	(0.0310)	(0.0177)	(0.0264)	(0.0237)	(0.0190)	(0.0236)
p-value test of equal ψ	0.5600			0.6161	0.0211	0.2247	0.0101	0.3762	0.9903	0.8206	0.0167

a lesser degree than it did for participating households. This indicates that this source of income is not strictly coming from financial markets, although it seems to affect participating households to a stronger degree.

Financial income being a significant source of risk is a puzzling result. Consider high wealth households. The pattern found is similar to that of market participants with a smaller change. Indeed, while the permanent insurance parameter prior to 1997 does increase relative to table (1.9), its change is small (0.2392 vs 0.2079). This is consistent with financial innovation improving households' ability to share risk. Process innovation, mainly the creation of ETFs, allows for cheaper diversification. As Guvenen (2007) points out, under-diversification is a significant source of risk. However, wealthy households would have the means to diversify more completely by having access to reserved financial products. The introduction of ETFs allows for households of modest financial means to fully diversify and reduce their exposure to market risk.

I also find that that financial income has a moderate to no impact on employees, entrepreneurs, low wealth households, transfer recipients and nontransfer recipients. These results indicate that entrepreneurs higher consumption smoothing abilities are not caused by financial income. The lack of impact of financial income on the transfer groups is simply caused by the fact that nonfinancial income would be small even in the nontransfer group.

I then consider business income in panel B of table (1.10). Note that employees and entrepreneurs are not estimated in this panel. Employees, by construction, receive no business income, so the results are identical to the baseline estimation of table (1.9). Entrepreneurs are not estimated simply because I cannot. As opposed to market participants whose portion of income derived from financial assets is relatively small, entrepreneurs derived most of their income from their business. Thus removing this income stream results in a large portion of the sample dropping out. The resulting sample (households who receive a small portion of income from business) is too small to estimate. Nonetheless, the patterns emerging from panel B are very interesting.

In the whole sample, I find that the pre-1997 ϕ parameter is lower than the baseline point estimate (0.2006 vs .2856). The post-1997 parameter is relatively similar. The difference is no longer significant implying no change in households' ability to insure against permanent shocks. Removing business income from the estimation has no impact on transitory insurance. These results indicate that the improvement documented in table (1.9) is caused (in part) by business income. Subgroups are not all equally affected by removing business income. The parameters for permanent insurance for participating and high wealth households are the most impacted in Panel B. The post 1997 parameters are lower, although less so than the preperiod parameters. This would indicates that a large portion of consumption insurance observed in these groups is in fact caused by households in these groups receiving business income. Households that do not participate or do receive transfer have small changes, suggesting that business income is not a significant source of consumption insurance. Similarly to financial income, I find the results on transitory insurance to be similar to their baseline level suggesting that business income is not used to smooth transitory variation in income shocks anymore than other sources of income.

In panel C, households that do and do not receive transfers are excluded from the estimation. Non-recipient, by construction, are unaffected. The recipient sample is too small to estimate. In the first column of panel C, I find point estimates for permanent insurance to be smaller than the baseline parameters for both periods. Indeed, ϕ_{pre97} is 0.2087 vs 0.2856 and ϕ_{post97} is 0.1475 vs 0.2371. Furthermore, while the difference in parameters was significant at the 5% in the baseline estimation, it is now significant at the 0.1%. Formal government transfer income is a significant source of permanent insurance and is now more effective than in the past. These patterns are extended to employees, nonparticipating, and low-wealth households. Market participants and entrepreneurs also have lower permanent insurance parameters, although to a lesser degree. This indicates that while these groups do receive government transfers, transfers represent a lower proportion of their total income. High wealth households also have lower permanent insurance parameters, especially ϕ_{post97}. This result is surprising. This could indicate that wealthy households do in fact receive enough transfers to increase their ability to smooth out shocks. Alternatively, it could indicate that households above median levels of wealth still are not particularly wealthy and require government assistance[33]. I also do find the same level of differences for transitory parameters. This is also surprising as I would think that government transfers are an excellent source of transitory insurance. One potential explanation can be found in Commault (2020). She argues that constraining the idiosyncratic consumption process to be a random walk creates a downward bias on transitory insurance thus overstating the marginal propensity of households to consume out of transitory shocks. The question remains in the data: what type of income source is generating transitory insurance?

[33] Although I would like to test high/low wealth with more restrictive separation, I am constrained by the data.

Table 1.11: Partial Insurance Estimates Net of Financial, Business, or Transfer Income (Continued)

	All households	Employees	Entrepreneurs	Market participants	No participation	Stock owners	Non stock owners	No transfer receipts	Transfer recipients	High Wealth	Low Wealth
Panel C. Disposable nontransfer Income											
ϕ_{pre97}	0.2087***	0.2184***	0.1792***	0.1931***	0.2035***	0.1940***	0.2155***			0.1799***	0.2159***
	(0.0142)	(0.0158)	(0.0304)	(0.0290)	(0.0219)	(0.0310)	(0.0202)			(0.0202)	(0.0208)
ϕ_{post97}	0.1475***	0.1532***	0.0886***	0.0833***	0.2108***	0.0715***	0.1609***			0.0645***	0.2002***
	(0.0110)	(0.0121)	(0.0283)	(0.0133)	(0.0236)	(0.0181)	(0.0156)			(0.0117)	(0.0202)
p-value test of equal ϕ	0.0004	0.0007	0.0252	0.0004	0.8088	0.0005	0.0237			0.0000	0.5572
ψ_{pre97}	0.0678***	0.0644***	0.0269	0.0000	0.0738***	0.0132	0.0631***			0.0259	0.0398*
	(0.0141)	(0.0174)	(0.0308)	(0.0289)	(0.0188)	(0.0293)	(0.0189)			(0.0165)	(0.0240)
ψ_{post97}	0.0953***	0.1050***	0.0214	0.0767***	0.1019***	0.0779***	0.1258***			0.0510***	0.1219***
	(0.0157)	(0.0204)	(0.0273)	(0.0232)	(0.0234)	(0.0293)	(0.0205)			(0.0196)	(0.0269)
p-value test of equal ψ	0.1893	0.1275	0.8985	0.0405	0.3428	0.1275	0.0210			0.3292	0.0224
Panel D. Disposable nonfinancial, nontransfer, and non-business Income											
ϕ_{pre97}	0.1545***	0.2105***		0.0945***	0.1877***	0.0777***	0.1982***	0.1799***		0.0720***	0.1958***
	(0.0121)	(0.0154)		(0.0231)	(0.0194)	(0.0241)	(0.0190)	(0.0206)		(0.0132)	(0.0206)
ϕ_{post97}	0.1259***	0.1513***		0.0695***	0.1725***	0.0615***	0.1417***	0.1331***		0.0441***	0.2074***
	(0.0100)	(0.0121)		(0.0131)	(0.0203)	(0.0155)	(0.0143)	(0.0182)		(0.0109)	(0.0216)
p-value test of equal ϕ	0.0588	0.0016		0.3411	0.5678	0.5642	0.0123	0.0804		0.1019	0.6795
ψ_{pre97}	0.0700***	0.0585***		0.0270	0.0631***	0.0527**	0.0522***	0.0280		0.0712***	0.0532**
	(0.0146)	(0.0175)		(0.0286)	(0.0194)	(0.0257)	(0.0189)	(0.0239)		(0.0183)	(0.0243)
ψ_{post97}	0.1001***	0.1089***		0.0543**	0.1252***	0.0383	0.1238***	0.0439		0.0522***	0.1203***
	(0.0164)	(0.0200)		(0.0233)	(0.0268)	(0.0298)	(0.0212)	(0.0272)		(0.0202)	(0.0271)
p-value test of equal ψ	0.1636	0.0554		0.4616	0.0555	0.7124	0.0093	0.6602		0.4850	0.0642

In panel D of table (1.11), I repeat the analysis by removing the effect of financial, business, and transfer income. Entrepreneurs are transfer recipients are excluded from this analysis. In the whole sample, I find that the permanent insurance parameters are lower than their baseline levels (0.1545 vs 0.2856 for the pre-1997 coefficients and 0.1259 vs 0.2371 for the post-1997 coefficients). The difference is significant at the 10% suggesting a moderate improvement in households' ability to insurance against permanent shocks.

The transitory parameters are slightly higher than their baseline levels (0.0700 vs 0.0528 for the pre-1997 period and 0.1001 vs 0.0794 for the post-1997 period). While both parameters are significant at the 1% level, their difference is statistically insignificant. These results indicate that the alternative streams of income are both a source of permanent consumption insurance and a source of transitory risk. The observed for employees is similar. I find that the permanent insurance parameters are lower and the transitory insurance parameters are higher than their respective baseline levels. The test of differences of transitory insurance reveals that the marginal propensity to consume out of transitory shocks has increased (at the 5% level). In the previous 3 tests, I fail to detect a significant difference. This could indicate that separately, each income stream did not generate significant changes to the variance of transitory risk. Taken together, they do add significant transitory risk in the post-1997 period. In the case of market participants, I find that neither permanent nor transitory insurance has changed from the pre-1997 to the post-1997 period. while I found significant permanent insurance improvements and significant transitory insurance deterioration. The lack of difference may indicate that the income streams play opposite roles.

To get a clearer picture of the contribution of each income category to consumption insurance, I repeat the analysis and allow insurance parameters to vary across groups of years. I plot the results alongside the baseline results (figure 1.11) in figure (1.12). The top panel plots the permanent insurance parameters ϕ_t; the bottom panel plots the transitory insurance parameters ψ_t. nonfinancial income is above the baseline result. The two are close, suggesting that financial income is neither used for consumption smoothing or a significant source of risk. This is not surprising. Indeed, financial income is small (or zero) for most households in the sample. Non-business income is below the baseline results until 1996. In the second half of the sample, non-business income is similar to the total income and nonfinancial income baseline levels. nontransfer income is consistently below the baseline level, although relatively close for the 1994-2002 period. The wage income line is approximately the sum of all deviations from the baseline level. It shows that, on average, nonwage income generates significant consumption insurance (mainly in the form of government transfers). Furthermore, I can see

Figure 1.12: Partial Insurance Parameters for Different Income Streams - All households

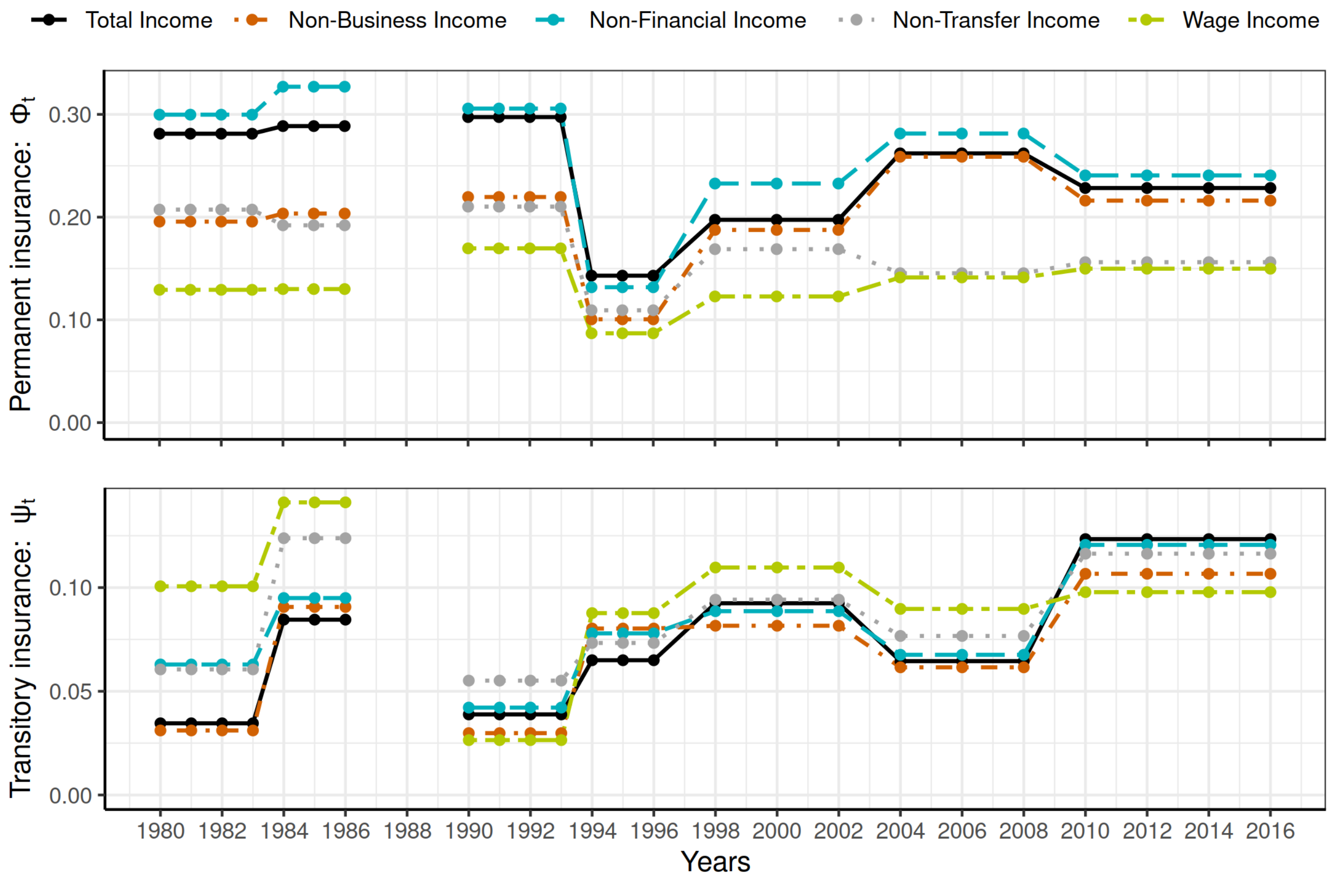

that the decrease in permanent insurance is caused by an increase in risk caused by financial and business income. In the bottom panel, I see that the trends in nonbusiness, nonfinancial, nontransfer, and wage income are relatively similar. There is no noticeable pattern in terms of movement or magnitude. The deterioration of transitory insurance observed in figure (1.11) is robust to measurement of the income variable.

In figure (1.13), I plot the same series for participating households only. I see that separate sources of income are an important smoothing mechanism for market participants. Indeed, permanent insurance parameters of wage income are consistently lower than baseline levels. In the first year group (1980-1984), ϕ_t is 0 indicating full permanent insurance. As the permanent insurance baseline trends downward, the permanent trend of nonfinancial, business, and transfer income trends upward. As similar trend is observed for non-business income. Non-business income appears to be a significant determinant of the marginal propensity to consume out of permanent shocks. Its importance seems to be decreasing throughout the sample. Nonfinancial income tracks the baseline level almost perfectly. It could be easy to assume that financial income represents a proportion of households' income too small to have an effect. However, I do observe stronger differences in transitory insurance relative to nonfinancial income. However, it appears that the major source of transitory insurance is business income.

Figure 1.13: Partial Insurance Parameters for Different Income Streams - Market Participants

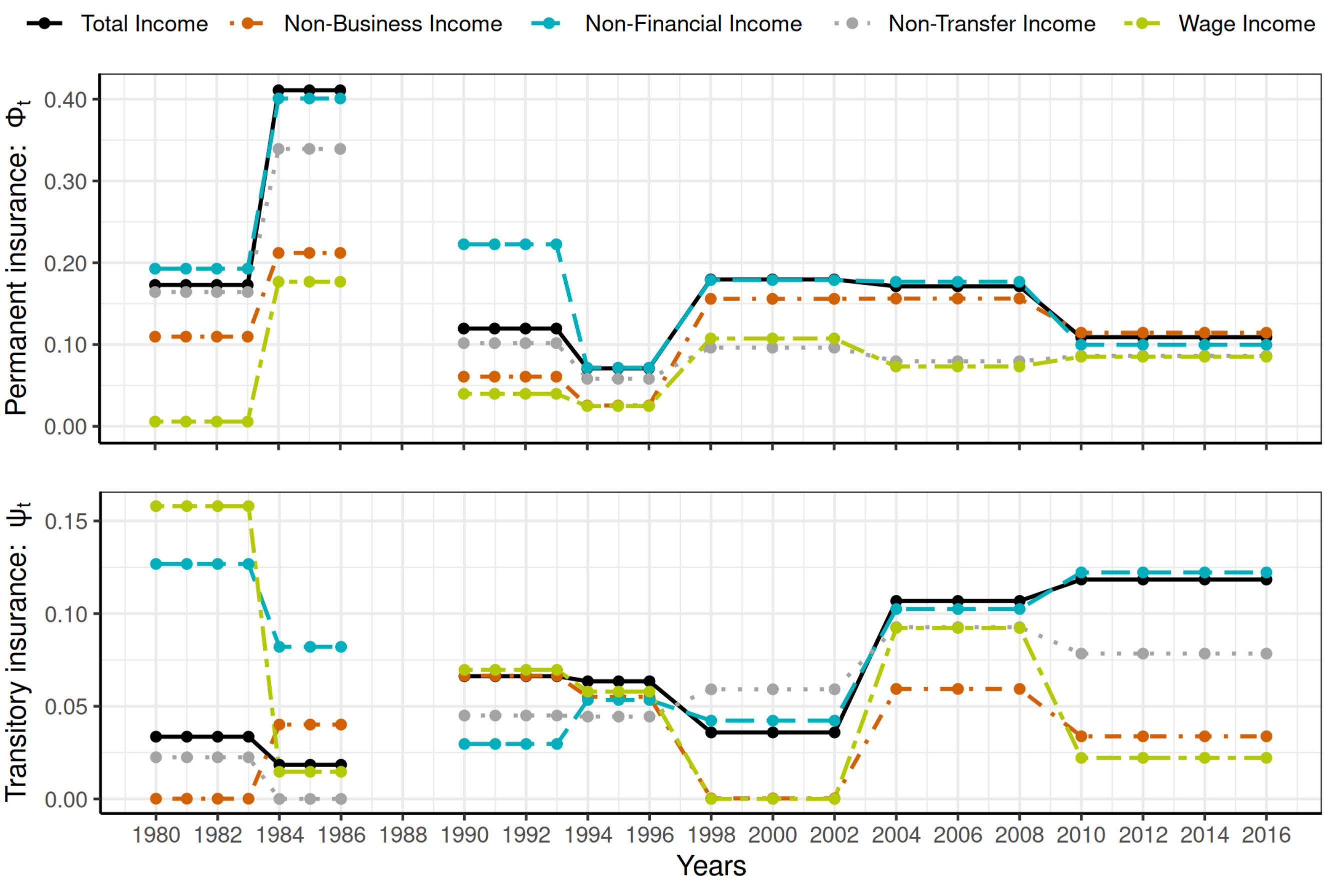

As a final robustness check, I reestimate idiosyncratic income and consumption shocks by restricting the set of demographics variables detailed in section (1.4.1). I consider 4 sets of variables to exclude and their year-dummy interaction terms. First, I exclude the employment status of the reference person. As pointed out by Commault (2020), adding the employment status could (at least in part) capture the effect of unexpected job loss (or hiring). By adding the variable, I may be implicitly reducing the variance of shocks. Furthermore, the annual nature of the data makes it impossible to distinguish between being unemployed or losing one's job which would illicit different consumption responses. I then exclude the dummy variables for the presence of extra income earners and whether the spouse earns income. Blundell, Pistaferri, and Saporta-Eksten (2016) shows that female labor supply is an important determinant of consumption insurance. However, it is possible that the inclusion of the two dummy variables partially captures shocks to household's income stemming for income earners other than the reference person. The third exclusion is the region of residence. A change in residence might reflect a shift in preferences. However, it could also reflect changes in job status, tax situation, and standards of living, all of which could be construed as permanent shocks. Finally, I also exclude the households' marital status. Blundell, Pistaferri, and Preston (2008) required all households to be continuously married. I do not. Thus, controlling for whether the reference person is single, married, or divorced might capture changes that would induce permanent shocks to income and consumption. The results are presented in table (1.12).

In panel A, I find that removing the employment status of the reference person has a limited impact on the permanent insurance parameters. In the whole sample, the parameters net of employment dummies are 0.2782 and 0.2119 vs their baseline levels of 0.2856 and 0.2371. The one exception is transfer recipients. The pre-1997 parameter is 0.3083 in table (1.12) vs 0.3948 in table (1.9). This may indicate that the inclusion of this variable captures unexpected shocks to employment and understates risk-sharing in table (1.9). However, these concerns cannot be alleviated given the data constraints. Removing the employment dummy has a significant impact on transitory insurance. Across the board, ψ pre and post 1997 are consistently insignificant (with the exception of market participants and transfer recipients) and in most cases actually 0 (that is, the point estimate is too small for me to precisely measure). At face value, this could indicate that transitory insurance is complete. However, this would contradict the literature that uses natural experiments to estimate the marginal propensity to consume transitory income shocks (Fagereng, Holm, and Natvik, 2019). Their estimate is not only significant but also higher than ours. The notion that transitory income shocks are fully insured is therefore unlikely. It is, however likely that the employment dummy variable does capture some shocks to income and these shocks are of a transitory nature.

The general trends observed in table (1.9) are consistent with results in table (1.12). Indeed, entrepreneurs appear to share permanent and transitory income risk more than employees; market participants smooth shocks more completely as well. The difference in parameters are however not significant. The small changes in the point estimate are most likely caused by measurement of shocks instead of the dummy variable on extra income earners wrongly capturing shocks. Similar patterns are found in panel C, and D. The result do show that the estimation is robust to small variation in the set of controls used in the estimation of the deterministic component of income and consumption. In figure (1.14), I plot

Table 1.12: Partial Insurance Estimates - Variation in the set of demographics variables

*This table reports the results from the minimum distance estimation presented in section 1.4.4.The parameters are estimated using equation (1.20) and the standard errors using equation (1.21). The variance of shocks and the variance of the measurement errors vary with each wave. I estimate $\Delta c_{i,t}$ and $\Delta y_{i,t}$ by excluding some of variables described in section (1.4.2) and their year-dummy interaction. In panel A, I remove the employment status of the reference person. In panel B, I remove dummy variables for extra income earners and whether the spouse earns income. In panel C, I remove the region of residence dummy variables. In panel D, I remove the marital status indicator. I use the diagonally weighted minimum distance estimator. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	All households	Employees	Entrepreneurs	Market participants	No participation	Stock owners	Non stock owners	No transfer receipts	Transfer recipients	High Wealth	Low Wealth
Panel A. Net of Employment Shocks											
ϕ_{pre97}	0.2782***	0.2856***	0.1936***	0.2337***	0.2896***	0.2204***	0.2940***	0.2909***	0.3083***	0.1961***	0.2654***
	(0.0171)	(0.0198)	(0.0397)	(0.0364)	(0.0279)	(0.0373)	(0.0275)	(0.0279)	(0.0416)	(0.0220)	(0.0258)
ϕ_{post97}	0.2119***	0.2388***	0.1137***	0.1276***	0.3289***	0.1275***	0.2747***	0.1529***	0.2428***	0.1123***	0.3125***
	(0.0149)	(0.0171)	(0.0338)	(0.0207)	(0.0345)	(0.0270)	(0.0229)	(0.0205)	(0.0355)	(0.0182)	(0.0305)
p-value test of equal ϕ	0.0019	0.0576	0.1116	0.0083	0.3417	0.0342	0.5591	0.0000	0.1925	0.0021	0.2060
ψ_{pre97}	0.0000	0.0000	0.0449	0.0097	0.0000	0.0343	0.0000	0.0342	0.0000	0.0000	0.0000
	(0.0153)	(0.0200)	(0.0293)	(0.0293)	(0.0235)	(0.0316)	(0.0224)	(0.0231)	(0.0285)	(0.0192)	(0.0278)
ψ_{post97}	0.0153	0.0069	0.0189	0.0519**	0.0014	0.0559	0.0011	0.0554*	0.0426	0.0329	0.0000
	(0.0158)	(0.0219)	(0.0291)	(0.0258)	(0.0248)	(0.0344)	(0.0204)	(0.0289)	(0.0296)	(0.0203)	(0.0267)
p-value test of equal ψ	0.4862	0.8149	0.5256	0.2848	0.9658	0.6422	0.9710	0.5693	0.3034	0.2306	1.0000
Panel B. Net of Extra Income Earners											
ϕ_{pre97}	0.2590***	0.2596***	0.1932***	0.1773***	0.3313***	0.1283***	0.3345***	0.2286***	0.2731***	0.2005***	0.2977***
	(0.0170)	(0.0189)	(0.0402)	(0.0350)	(0.0292)	(0.0334)	(0.0275)	(0.0240)	(0.0407)	(0.0237)	(0.0271)
ϕ_{post97}	0.2493***	0.2807***	0.1364***	0.1337***	0.4058***	0.1141***	0.3188***	0.1618***	0.3397***	0.1098***	0.3384***
	(0.0168)	(0.0195)	(0.0448)	(0.0220)	(0.0407)	(0.0255)	(0.0260)	(0.0195)	(0.0473)	(0.0211)	(0.0329)
p-value test of equal ϕ	0.6604	0.3997	0.3375	0.2758	0.0935	0.7348	0.6462	0.0250	0.2555	0.0028	0.2924
ψ_{pre97}	0.0489***	0.0507**	0.0230	0.0415	0.0274	0.0988***	0.0166	0.0485**	0.0654**	0.0228	0.0196
	(0.0156)	(0.0202)	(0.0307)	(0.0302)	(0.0221)	(0.0335)	(0.0211)	(0.0244)	(0.0275)	(0.0186)	(0.0245)
ψ_{post97}	0.0814***	0.0871***	0.0210	0.0737***	0.0790***	0.0925***	0.0928***	0.0704***	0.0948***	0.0598***	0.1028***
	(0.0144)	(0.0189)	(0.0264)	(0.0237)	(0.0220)	(0.0320)	(0.0181)	(0.0271)	(0.0232)	(0.0182)	(0.0240)
p-value test of equal ψ	0.1216	0.1825	0.9615	0.3961	0.0949	0.8900	0.0053	0.5499	0.4150	0.1489	0.0150

Table 1.13: Partial Insurance Estimates - Variation in the Demographics Set (Continued)

	All households	Employees	Entrepreneurs	Market participants	No participation	Stock owners	Non stock owners	No transfer receipts	Transfer recipients	High Wealth	Low Wealth
Panel C. Net of Geographic Shocks											
ϕ_{pre97}	0.2877***	0.2943***	0.1813***	0.2564***	0.3908***	0.2427***	0.3868***	0.2693***	0.3769***	0.2097***	0.3088***
	(0.0176)	(0.0203)	(0.0366)	(0.0385)	(0.0325)	(0.0393)	(0.0303)	(0.0264)	(0.0555)	(0.0237)	(0.0268)
ϕ_{post97}	0.2440***	0.2702***	0.1406***	0.1298***	0.3873***	0.1120***	0.3165***	0.1555***	0.3357***	0.1120***	0.3072***
	(0.0170)	(0.0195)	(0.0419)	(0.0228)	(0.0421)	(0.0258)	(0.0271)	(0.0206)	(0.0442)	(0.0204)	(0.0304)
p-value test of equal ϕ	0.0534	0.3534	0.4443	0.0033	0.9412	0.0045	0.0526	0.0004	0.5290	0.0011	0.9669
ψ_{pre97}	0.0528***	0.0555***	0.0240	0.0108	0.0208	0.0314	0.0184	0.0539**	0.0682***	0.0163	0.0263
	(0.0150)	(0.0194)	(0.0323)	(0.0284)	(0.0208)	(0.0306)	(0.0200)	(0.0235)	(0.0253)	(0.0186)	(0.0247)
ψ_{post97}	0.0758***	0.0870***	0.0088	0.0857***	0.0836***	0.0948***	0.0895***	0.0715**	0.0786***	0.0564***	0.1176***
	(0.0148)	(0.0195)	(0.0274)	(0.0236)	(0.0219)	(0.0336)	(0.0182)	(0.0285)	(0.0239)	(0.0189)	(0.0246)
p-value test of equal ψ	0.2702	0.2514	0.7178	0.0423	0.0364	0.1578	0.0074	0.6374	0.7685	0.1251	0.0087
Panel D. Net of Marital Status											
ϕ_{pre97}	0.3149***	0.3484***	0.2134***	0.2796***	0.3914***	0.2671***	0.3971***	0.3005***	0.3574***	0.2075***	0.3312***
	(0.0180)	(0.0219)	(0.0381)	(0.0408)	(0.0330)	(0.0419)	(0.0300)	(0.0265)	(0.0515)	(0.0237)	(0.0278)
ϕ_{post97}	0.2630***	0.2951***	0.1293***	0.1241***	0.4105***	0.1148***	0.3310***	0.1805***	0.3615***	0.1198***	0.3123***
	(0.0174)	(0.0202)	(0.0365)	(0.0213)	(0.0417)	(0.0282)	(0.0274)	(0.0216)	(0.0489)	(0.0215)	(0.0302)
p-value test of equal ϕ	0.0248	0.0512	0.1013	0.0005	0.6805	0.0021	0.0650	0.0002	0.9500	0.0040	0.6187
ψ_{pre97}	0.0847***	0.0831***	0.0471	0.0350	0.0566***	0.0596**	0.0467**	0.0837***	0.1009***	0.0281	0.0437*
	(0.0150)	(0.0184)	(0.0315)	(0.0285)	(0.0212)	(0.0294)	(0.0208)	(0.0233)	(0.0263)	(0.0183)	(0.0251)
ψ_{post97}	0.1019***	0.1176***	0.0019	0.1038***	0.0979***	0.1057***	0.1140***	0.0947***	0.1067***	0.0616***	0.1449***
	(0.0145)	(0.0187)	(0.0280)	(0.0246)	(0.0221)	(0.0322)	(0.0183)	(0.0281)	(0.0235)	(0.0187)	(0.0245)
p-value test of equal ψ	0.4058	0.1912	0.2863	0.0730	0.1771	0.2948	0.0141	0.7641	0.8718	0.1947	0.0040

Figure 1.14: Partial Insurance Parameters Net of Demographic Characteristics

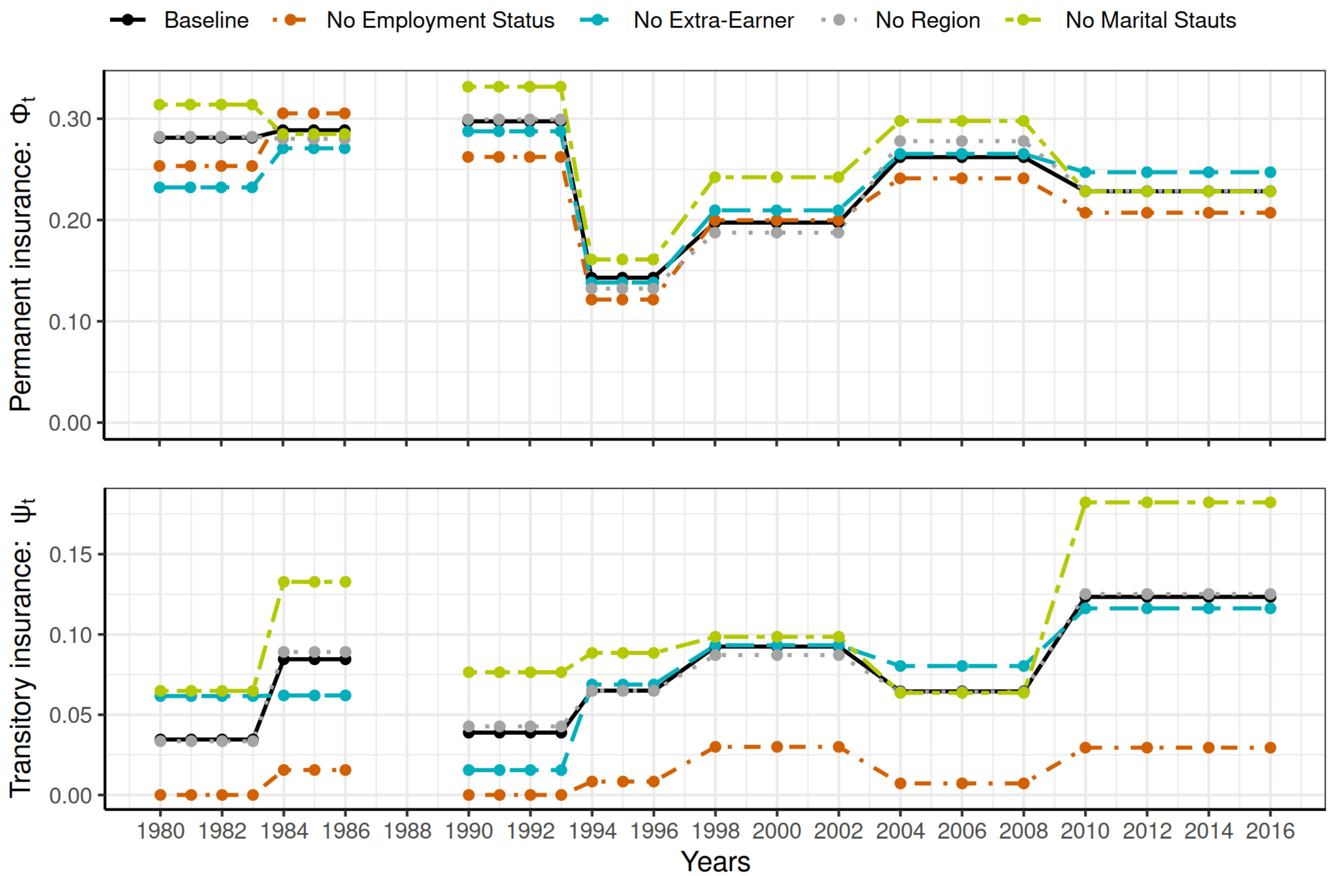

the time-varying insurance parameters for all households. These parameters are estimated with shocks net of demographics variables. In the top panel, I see that the exclusion of the various demographic components has a limited impact of permanent insurance parameters. Furthermore, except for employment status, the other variables seem to have a limited impact on ψ_t as well. Overall, the baseline results appear to be robust.

1.5 Conclusion

Financial innovation should improve risk-sharing. Indeed, the incentives for financial innovation rely on insuring more states of nature. However, Simsek (2013a) and Simsek (2013b), or Buss and Uppal (2017) argues that in the event of disagreement, speculation could overtake the risk-sharing benefits thus reducing welfare. While the literature agrees that markets are incomplete and households cannot fully insure their idiosyncratic risk, it has yet to document how risk-sharing evolves through time. This gap in the literature is all the more important considering the recent financial product innovation (Huang, O'Hara, and Zhong, 2021; Iachan, Nenov, and Simsek, 2021) or process innovation (Turley, 2012).

Using PSID data and an imputation methodology from Attanasio and Pistaferri (2014), I create a panel of idiosyncratic consumption and income. I also use the identification technique by Guvenen (2007) to identify financial market participants. By extension, I am also able to identify households in the bottom and top of the wealth distribution. I also identify entrepreneurs (as opposed to employees), and transfer recipients as further sources of household heterogeneity.

I use two separate testing framework. First, I use a simple transmission model where risk-sharing is the extent to which idiosyncratic consumption is affected by idiosyncratic income. The results are in line with the literature. I find that a fair portion of income shocks are insured and I reject full insurance. However, trends in risk-sharing show that households' ability to smooth shocks has decreased or at the very least remained constant. I also find that participating households smooth shocks more completely than nonparticipating households in direct contradiction of results by Guvenen (2007). However, given data constraints, I am unable to differentiate between households in a precise way. Thus, the observed differences between participating and nonparticipating households may be caused by other household characteristics: mainly wealth and entrepreneurial status.

I then turn to a more complete income and consumption process. I use the process and minimum distance estimator in Blundell, Pistaferri, and Preston (2008) to decompose risk-sharing into its permanent and transitory component. I find that the variance of transitory shocks has increased while the variance of permanent shocks has decreased. Consistent with the literature, I find that transitory shocks are transmitted to a lesser degree than permanent shocks.

I test whether the transmission parameters have shifted over time. I find that permanent insurance has increased or at the very least remained constant. Market participants, wealth households, and entrepreneurs are not only better insured than their counterparts (non participating, low wealth, and employees respectively) but have seen their consumption insurance increase. Transitory insurance has in some cases deteriorated though for most subgroups has remained constant. I test whether the estimate are robust to different income specifications. While trends are relatively similar to baseline results across groups, I do find that business and transfer income are stronger sources of both permanent and transitory insurance than financial income. The baseline results are also robust to changes in the estimation of shocks.

Overall, I find the impact of financial innovation on risk-sharing to be empirically underwhelming. Indeed, the development of financial markets over the sample period is undeniable. Furthermore, the traditional view is clear on the supposed effect of financial innovation on risk-sharing. The traditional view may thus need to be revisited. While I do not find downward trends in risk-sharing, I do find significant heterogeneity in households' ability to smooth out shocks. These sources of heterogeneity are likely to have welfare and portfolio choice consequences, two questions I leave for further research.

References

Abowd, John M and David Card (1989). "On the covariance structure of earnings and hours changes". In: *Econometrica: Journal of the Econometric Society*, pp. 411–445.

Aguiar, Mark and Mark Bils (2015). "Has consumption inequality mirrored income inequality?" In: *American Economic Review* 105.9, pp. 2725–56.

Allen, Franklin and Douglas Gale (1988). "Optimal security design". In: *The Review of Financial Studies* 1.3, pp. 229–263.

Allen, Franklin, Douglas Gale, et al. (1994). *Financial innovation and risk sharing*. MIT press.

Attanasio, Orazio and Luigi Pistaferri (2014). "Consumption inequality over the last half century: some evidence using the new PSID consumption measure". In: *American Economic Review* 104.5, pp. 122–26.

Bai, Jennie, Thomas Philippon, and Alexi Savov (2016). "Have financial markets become more informative?" In: *Journal of Financial Economics* 122.3, pp. 625–654.

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). "Consumption inequality and partial insurance". In: *American Economic Review* 98.5, pp. 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten (2016). "Consumption inequality and family labor supply". In: *American Economic Review* 106.2, pp. 387–435.

Blundell, Richard and Ian Preston (1998). "Consumption inequality and income uncertainty". In: *The Quarterly Journal of Economics* 113.2, pp. 603–640.

Buss, Adrian and Raman Uppal (2017). "Financial Innovation and Asset Prices". In:

Calvet, Laurent, Martin Gonzalez-Eiras, and Paolo Sodini (2004). "Financial Innovation, Market Participation, and Asset Prices". In: *Journal of Financial and Quantitative Analysis*, pp. 431–459.

Casado, Jose Maria (2011). "From income to consumption: measuring households partial insurance". In: *Empirical Economics* 40.2, pp. 471–495.

Chamberlain, Gary (1984). *Panel data', in (Z. Griliches and M. Intriligator, eds.) Handbook of Econometrics.*

Cochrane, John H (1991). "A simple test of consumption insurance". In: *Journal of political economy* 99.5, pp. 957–976.

Cochrane, John H (2009). *Asset pricing: Revised edition*. Princeton university press.

Commault, Jeanne (2020). "Does Consumption Respond to Transitory Shocks?" In: *American Economic Journal: Macroeconomics*.

Constantinides, George M and Darrell Duffie (1996). "Asset pricing with heterogeneous consumers". In: *Journal of Political economy* 104.2, pp. 219–240.

Constantinides, George M and Anisha Ghosh (2017). "Asset pricing with countercyclical household consumption risk". In: *The Journal of Finance* 72.1, pp. 415–460.

Cutler, David M et al. (1991). "Macroeconomic performance and the disadvantaged". In: *Brookings papers on economic activity* 1991.2, pp. 1–74.

Dynan, Karen E, Douglas W Elmendorf, and Daniel E Sichel (2006). "Can financial innovation help to explain the reduced volatility of economic activity?" In: *Journal of monetary Economics* 53.1, pp. 123–150.

Elul, Ronel (1997). "Financial innovation, precautionary saving and the risk-free rate". In: *Journal of Mathematical Economics* 27.1, pp. 113–131.

Fagereng, Andreas, Martin Blomhoff Holm, and Gisle James James Natvik (2019). "MPC heterogeneity and household balance sheets". In: *Available at SSRN 3399027*.

Fama, Eugene F and Kenneth R French (2001). "Disappearing dividends: changing firm characteristics or lower propensity to pay?" In: *Journal of Financial economics* 60.1, pp. 3–43.

Feenberg, Daniel and Elisabeth Coutts (1993). "An introduction to the TAXSIM model". In: *Journal of Policy Analysis and management* 12.1, pp. 189–194.

Fisher, Jonathan D and David S Johnson (2020). "Inequality and Mobility over the Past Half Century using Income, Consumption and Wealth". In: *NBER Chapters*.

Flavin, Marjorie and Takashi Yamashita (2002). "Owner-occupied housing and the composition of the household portfolio". In: *American Economic Review* 92.1, pp. 345–362.

Frame, W Scott and Lawrence J White (2004). "Empirical studies of financial innovation: lots of talk, little action?" In: *Journal of Economic Literature* 42.1, pp. 116–144.

Gervais, Martin and Paul Klein (2010). "Measuring consumption smoothing in CEX data". In: *Journal of Monetary Economics* 57.8, pp. 988–999.

Guvenen, Fatih (2007). "Do stockholders share risk more effectively than nonstockholders?" In: *The Review of Economics and Statistics* 89.2, pp. 275–288.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L Violante (2010). "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006". In: *Review of Economic dynamics* 13.1, pp. 15–51.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante (2014). "Consumption and labor supply with partial insurance: An analytical framework". In: *American Economic Review* 104.7, pp. 2075–2126.

Hryshko, Dmytro, Maria Jose Luengo-Prado, and Bent E Sørensen (2010). "House prices and risk sharing". In: *Journal of Monetary Economics* 57.8, pp. 975–987.

Huang, Shiyang, Maureen O'Hara, and Zhuo Zhong (2021). "Innovation and informed trading: Evidence from industry ETFs". In: *The Review of Financial Studies* 34.3, pp. 1280–1316.

Iachan, Felipe S, Plamen T Nenov, and Alp Simsek (2021). "The choice channel of financial innovation". In: *American Economic Journal: Macroeconomics* 13.2, pp. 333–72.

Johnson, David S and Jonathan D Fisher (2020). "Inequality and Mobility over the Past Half Century using Income, Consumption and Wealth". In: *Measuring and Understanding the Distribution and Intra/Inter-Generational Mobility of Income and Wealth*. University of Chicago Press.

Kan, Kamhon, Shin-Kun Peng, and Ping Wang (2017). "Understanding consumption behavior: Evidence from consumers' reaction to shopping vouchers". In: *American Economic Journal: Economic Policy* 9.1, pp. 137–53.

Kaplan, Greg and Giovanni L Violante (2010). "How much consumption insurance beyond self-insurance?" In: *American Economic Journal: Macroeconomics* 2.4, pp. 53–87.

Kaufmann, Katja and Luigi Pistaferri (2009). "Disentangling insurance and information in intertemporal consumption choices". In: *American Economic Review* 99.2, pp. 387–92.

Kimberlin, Sara, Jiyoun Kim, and Luke Shaefer (2014). "An updated method for calculating income and payroll taxes from PSID data using the NBER's TAXSIM, for PSID survey years 1999 through 2011". In: *Unpublished manuscript, University of Michigan. Accessed May* 6, p. 2016.

Krueger, Dirk and Fabrizio Perri (2006). "Does income inequality lead to consumption inequality? Evidence and theory". In: *The Review of Economic Studies* 73.1, pp. 163–193.

Kubota, Kohei (2020). "Partial insurance in Japan". In: *The Japanese Economic Review*, pp. 1–30.

Ludwig, Johannes (2016). "Income Shocks or Insurance-What Determines Consumption Inequality?" In: *Research in Labor Economics* 43, pp. 131–166.

Mace, Barbara J (1991). "Full insurance in the presence of aggregate uncertainty". In: *Journal of Political Economy* 99.5, pp. 928–956.

Mankiw, N Gregory (1986). "The equity premium and the concentration of aggregate shocks". In: *Journal of Financial Economics* 17.1, pp. 211–219.

Meghir, Costas and Luigi Pistaferri (2004). "Income variance dynamics and heterogeneity". In: *Econometrica* 72.1, pp. 1–32.

Mehra, Rajnish and Edward C Prescott (1985). "The equity premium: A puzzle". In: *Journal of monetary Economics* 15.2, pp. 145–161.

Michaud, Amanda and David Wiczer (2018). "Occupational hazards and social disability insurance". In: *Journal of Monetary Economics* 96, pp. 77–92.

Miller, Merton H (1986). "Financial innovation: The last twenty years and the next". In: *Journal of Financial and Quantitative Analysis* 21.4, pp. 459–471.

Misra, Kanishka and Paolo Surico (2014). "Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs". In: *American Economic Journal: Macroeconomics* 6.4, pp. 84–106.

Nelson, Julie A (1994). "On testing for full insurance using Consumer Expenditure Survey data". In: *Journal of Political Economy* 102.2, pp. 384–394.

Olovsson, Conny (2010). "Quantifying the risk-sharing welfare gains of social security". In: *Journal of Monetary Economics* 57.3, pp. 364–375.

Ortigueira, Salvador and Nawid Siassi (2013). "How important is intra-household risk sharing for savings and labor supply?" In: *Journal of Monetary Economics* 60.6, pp. 650–666.

Piketty, Thomas (2015). "About capital in the twenty-first century". In: *American Economic Review* 105.5, pp. 48–53.

Rubin, Donald B (2004). *Multiple imputation for nonresponse in surveys*. Vol. 81. John Wiley & Sons.

Saez, Emmanuel and Gabriel Zucman (2016). "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data". In: *The Quarterly Journal of Economics* 131.2, pp. 519–578.

Santaeulalia-Llopis, Raul and Yu Zheng (2018). "The price of growth: consumption insurance in China 1989–2009". In: *American Economic Journal: Macroeconomics* 10.4, pp. 1–35.

Simsek, Alp (2013a). "Financial innovation and portfolio risks". In: *American Economic Review* 103.3, pp. 398–401.

— (2013b). "Speculation and risk sharing with new financial assets". In: *The Quarterly journal of economics* 128.3, pp. 1365–1396.

Turley, Robert (2012). "Informative prices and the cost of capital markets". In: *Harvard University Working Paper*.

Weil, Philippe (1992). "Equilibrium asset prices with undiversifiable labor income risk". In: *Journal of Economic Dynamics and Control* 16.3-4, pp. 769–790.

Wolff, Edward N (2016). "Household Wealth Trends in the United States, 1962 to 2013: What Happened over the Great Recession?" In: *RSF: The Russell Sage Foundation Journal of the Social Sciences* 2.6, pp. 24–43.

Chapter 2

RISK-SHARING, HOUSEHOLD WELFARE AND SUBJECTIVE WELL-BEING

Abstract

I investigate the welfare and life satisfaction consequences of incomplete markets in a subset of US households. I use a set of parameters describing households' economic environments in terms of income growth, income risk, and transmission risk. I find that changes in risk-sharing have significant implications for household welfare. Cross-sectional differences in risk-sharing environment result in significantly different welfare criteria. I then use IV-regressions to separate the impact of permanent and transitory income and consumption shocks on life satisfaction. As suggested by consumption insurance theory, I find that transitory shocks have no effect on life satisfaction while permanent shocks do. This result suggest that risk-sharing environments have important consequences for households' well-being as well as a significant degree of insurance despite incomplete markets.

JEL classification: D12, D52, E21, I31, P36

Keywords: Idiosyncratic Risk, Welfare, Life Satisfaction

2.1 Introduction

What are the welfare effects of incomplete markets? What are the well-being effects of incomplete markets? A large body of economic literature is dedicated to quantifying the degree of market completeness by measuring the transmission of income shocks to consumption shocks (Cochrane, 1991; Krueger and Perri, 2006; Blundell, Pistaferri, and Preston, 2008; Kaplan and Violante, 2010; Gervais and Klein, 2010; Commault, 2020). Estimates differ, the consensus appears to show that although incomplete, there exists a significant degree of consumption insurance. Furthermore, using Blundell, Pistaferri, and Preston (2008) framework, Casado (2011), Kubota (2020), and Santaeulalia-Llopis and Zheng (2018) show that permanent shocks are less insured than transitory shocks.

Another strand of literature investigates the welfare implications of idiosyncratic income risk. Floden and Linde (2001) and Floden (2001) show that idiosyncratic income risk is a significant driver of household welfare. Heathcote, Storesletten, and Violante (2008) echo these results by showing that eliminating wage risk would improve household welfare to a larger extent than eliminating business cycle risk. However, Karahan and Ozkan (2013) find different welfare benefits for agents with different risk profiles. De Nardi, Fella, and Paz-Pardo (2020) augments the traditional transitory/permanent income process with nonlinear dynamics and show that welfare changes significantly depend on income process specifications.

In this paper, I investigate the welfare differences in US households from a consumption insurance perspective. First, I compare welfare estimates across time for groups of representative households. Second, I compare welfare estimates across groups, and groups and time. The first exercise is motivated by recent developments in the literature by Simsek (2013a) and Simsek (2013b). In these papers, he argues that financial innovation's impact on risk-sharing benefits depends on the degree of price agreement between market participants. His results go against the traditional view of financial innovation which suggests improving market completeness and providing consumption insurance is the main driver of financial innovation (Cochrane, 2009, p. 56). Despite the undeniable development of financial markets in the past 40 years (Iachan, Nenov, and Simsek, 2021; Huang, O'Hara, and Zhong, 2021), Duvernois (2021) does not find significant improvements in the degree of consumption insurance since 1980[1]. However, he does find that the variance of transitory shocks is significantly larger than the variance of permanent shocks post 1996. Furthermore, he finds that different groups of households share risk differently and experience different income shocks.

I use Duvernois (2021)'s minimum distance parameters and the welfare framework

[1]He uses the minimum distance estimator by Blundell, Pistaferri, and Preston (2008)

from Santaeulalia-Llopis and Zheng (2018). In this framework, I define welfare as the annual consumption equivalent change from a baseline estimation. Welfare is defined to be a function of income growth, income risk, and transmission risk. I successively replace parameters with counterfactual parameters which can be taken from the same group in a different period or in the same period from a different group of households. I use the Panel Survey of Income Dynamics (PSID) to estimate households' income growth rate, saving rate, and risk aversion.

I find that the changes in economic environment between 1980-1996 and 1996-2016 have significant welfare implications. The slowdown in income growth dominates the welfare calculations. However, households experience an overall lower risk and lower transmission which generates significant welfare gains. With high levels of risk aversion[2], the welfare benefits can be larger than the welfare costs of slower income growth. Furthermore, I find that the differences in economic environment between types of households can generate large welfare benefits. Indeed, I find that entrepreneurs, market participants, nontransfer recipients, and wealthy households[3] have significantly larger annual equivalent consumption than their counterparts.

This paper also makes a contribution to the literature on subjective well-being. Bayer and Juessen (2015) show that idiosyncratic income shocks have a significant impact on households' life satisfaction. I use the recently added life satisfaction question[4] to investigate whether income and consumption shocks have an impact on US household life satisfaction. I find that total and permanent income shocks have a statistically significant impact while transitory shocks do not. Furthermore, as suggested by Brown and Gathergood (2017), consumption shocks have a stronger effect on life satisfaction. Transitory consumption shocks do not consistent with the idea that transitory income shocks are fully insured. Risk-sharing does have significant implications for household welfare and well-being.

The rest of the chapter is organized as follows. Section 2.2 will connect this paper to the related literature. Section 2.3 will present the main welfare framework. Section 2.4 will present the data. Section 2.5 will present the welfare results while section 2.6 will present the subjective well-being results. Section 2.7 will conclude.

2.2 Literature Review

Whether idiosyncratic income risk are insured and persistent is debated in the literature (Guvenen, 2007b; Guvenen, 2009). Whether idiosyncratic income risk matter for household welfare is also debated in the literature. Floden (2001) shows that there are

[2]implied from PSID 1996 risk attitude questionnaire.

[3]relative to employees, non-participants, transfer recipients, and low wealth households

[4]added to the PSID in 2009

significant welfare effects due to the levels and changes in idiosyncratic uncertainty. They consider an economy where agents are impacted by different idiosyncratic wage shocks and whose ability to smooth shocks differ. Risk-sharing can be improved with public debt by way of increased liquidity. They find that increased risk-sharing improves welfare by achieving better consumption/leisure allocation. This result goes against Aiyagari and McGrattan (1998) who argue that welfare gains from more complete markets are negligible.

Karahan and Ozkan (2013) show that the persistence and size of shocks is not consistent over households' life cycle. The variance of permanent shocks is U-shaped with age. They find that assuming flat profiles understates welfare gains as younger households enjoy a higher degree of consumption insurance. However, De Nardi, Fella, and Paz-Pardo (2020) allows for richer, non-linear dynamics of income processes (suggested by Arellano, Blundell, and Bonhomme, 2017) and calibrate a model which implies earnings risk generating lower welfare costs than the traditionally used income processes.

Chetty and Looney (2006) further question the assumptions that as consumption is relatively smooth, changes in risk-sharing have little welfare consequences. They argue that consumption may be smooth because welfare consequences of fluctuating consumption are in fact quite high. Indeed, they show that marginal welfare gains of a dollar of insurance depends on the size of consumption shocks[5] as well as how much households value a smoother consumption path[6]. Floden and Linde (2001) find significant difference in idiosyncratic wage risk between the United States and Sweden. Furthermore, Duvernois (2021) finds significant differences in idiosyncratic income risk and transmission risk across American households.

Buera and Shin (2011) use a generalized Bewley model (as defined by Ljungqvist and Sargent, 2018) to show that welfare costs of missing consumption insurance and incomplete markets depend on the degree of shock persistence and the relative impact of the opposing effects. The literature on incomplete-market heterogeneous models tend to focus on the consumption side in which welfare costs increase in the persistence of shocks. Buera and Shin (2011) show that on the production side, welfare costs decrease in the persistence of shocks due to improved self-financing for entrepreneurs.

Heathcote, Storesletten, and Violante (2008) investigate the welfare impact of three characterization of incomplete markets. They show that eliminating idiosyncratic risk generates welfare gains twice as large as welfare gains from eliminating business cycle risk. De Santis (2007) focuses on consumption risk instead of income risk. However, they look at the welfare impact of aggregate risk. They find that removing aggregate risk, thus producing a

[5]How much smoother consumption now is.
[6]i.e the relative risk aversion coefficient

smoother consumption path, has significant welfare consequences. Heathcote, Storesletten, and Violante (2008) also point out that removing risk yields lower welfare improvement than insuring risk as the former also removes opportunities for labor productivity. As such, the investigation of welfare and risk-sharing is opportune.

The literature above tends to measure welfare as costs as a ratio of two utility measures (one without insurance and one with insurance). Another strand of the literature investigates whether macroeconomic conditions or microevents impact households' subjective well-being. Subjective well-being is being increasingly substituted for welfare in economic and policy research (Nikolova and Graham, 2020). Economists have augmented their studies with happiness data to study a wide array of issues. Kimball and Willis (2006) make a clear distinction between utility and happiness whereas utility is the extent to which individuals acquire what they desire as evidenced by choices. Happiness is how individuals feel at a given moment. Although different, Kimball and Willis (2006) argue that both are distinct yet valid empirical candidates of welfare measures.

See A. E. Clark (2018) or Frey and Stutzer (2002) for an excellent review of the literature on the economics of happiness. For the purpose of this study, let's focus on income risk and happiness. Deaton (2008) finds that income is related to well-being. However, Kahneman and Deaton (2010) show that income effects on well-being are smaller in comparison to other life events. Indeed, A. E. Clark and Oswald (1996), Ferrer-i-Carbonell (2005), and Caporale et al. (2009) find that happiness depends upon a comparison level of income rather that just income itself. Nonetheless, Frijters, Haisken-DeNew, and Shields (2004) show that East-Germans' life satisfaction levels converged to those of West-Germans along with the increase in real household income brought by the reunification[7].

A. E. Clark and Oswald (1994) find that unemployment has strong effects on British households' self-reported life satisfaction. The effect is stronger than the effect of divorce or marital separation. Furthermore, they find that the short-term unemployed (households that just lost their job) are more affected than long-term unemployed households. L. Winkelmann and R. Winkelmann (1995) use West-Germany panel data and similarly find that unemployment has strong negative effect on life satisfaction. Additionally, they find that income support program has a small effect on the unemployed happiness. This would suggest that income shocks are not the main cause of disutility. A. E. Clark (2003) show that the effect of unemployment on life satisfaction is heavily influenced by a reference group unemployment; that is heavily influenced by social norms. A. Clark, Georgellis, and Sanfey (2001) further

[7]Although the West and East Germany reunification may have had non-income related impacts on happiness such as higher personal liberties or increased savings as pointed out by Fuchs-Schundeln (2008).

show that individuals that experience a large happiness shock upon entering unemployment are less likely to remain unemployed for longer than one year.

Headey and Wooden (2004) show that wealth has a complementary effect to income in increasing households' life satisfaction. One explanation as to why can be found in Bayer and Juessen (2015). They investigate the impact of idiosyncratic income shocks on German households' happiness levels. They find that income shocks have a significant impact on happiness. Decomposing the income process into permanent and transitory components, they show that only permanent income shocks have a significant effect on happiness shocks.

Some have argued that the low impact of income on satisfaction is caused by a satiation around $75,000 Deaton, 2008. Stevenson and Wolfers (2013) and Lien, Hu, and Liu (2017) challenge the satiation hypothesis and find no evidence of a satiation point. Around this debate, Brown and Gathergood (2017) argue that consumption should matter more than income for life satisfaction. After all, utility is measured as a function of consumption. They find that consumption does indeed matter more for life satisfaction in US households. Furthermore, they find that conspicuous consumption[8] has a stronger impact alluding to similar social effects documented by A. E. Clark (2003). Consumption levels do not appear to have any level of satiation.

Brown and Gathergood (2017) and Bayer and Juessen (2015)'s result make interesting contributions to another area of the literature: consumption insurance. Indeed, if consumption idiosyncratic income shocks matters for life satisfaction, there must be some degree of transmission of income shocks to consumption. Although the literature does not agree on a precise quantity of transmission, there is clear evidence of incomplete markets. Blundell, Pistaferri, and Preston (2008) develop a partial insurance framework to measure the transmission of permanent and transitory income shocks to consumption. They find that permanent shocks are less insurable than transitory shocks. This model has become the workhorse of the consumption insurance literature and has been used to a variety of national household panels. Kubota (2020) applies the partial insurance framework to Japanese data and finds that transitory shocks are almost fully insured despite increasing over time. On the other hand, permanent shocks are insured by half and remain constant throughout the sample period. Casado (2011) uses the Spanish Household Budget Continuous Survey and find similar results. He also finds that home-owners, high-wealth households, and college educated households have higher degree of permanent insurance. Santaeulalia-Llopis and Zheng (2018) use a longitudinal panel dataset of Chinese households. They find that Chinese households experienced a decline in permanent insurance combined with increased

[8]which can be easily observed by other households

levels of income risk accompanying the rapid economic growth during the past 30 years. Furthermore, they show that the welfare effects of growth can be as large as the welfare effects of risk and insurance. Santaeulalia-Llopis and Zheng (2018) further use the partial insurance parameters to estimate a welfare framework similar to the one presented by Floden and Linde (2001). The following section presents that framework.

2.3 Welfare Framework

The main question I ask is: what are the welfare costs of changes in household income risk and risk-sharing abilities. Furthermore, what are the welfare costs of heterogeneity in risk and transmission? Indeed, Duvernois (2021) find significant changes in the income risk environment and in their ability to smooth out these shocks for different households. The welfare framework is adapted from Santaeulalia-Llopis and Zheng (2018); further details can be found in their paper and the online appendix.

Consider a representative household income and consumption. The income and consumption process of the representative agent can be described as the sum of a deterministic component and a stochastic component. Let us start by describing the income process. Although, the change in lifetime consumption is the main measure of welfare as is typical in the literature (Chetty and Looney, 2006; Floden, 2001; Heathcote, Storesletten, and Violante, 2008), most of the parameters are derived from the income process. Let $\bar{y_t}$ be the deterministic component of income, y_t be the stochastic component of income and Y_t be total household income[9]. Empirically, I use disposable income, as opposed to labor income, as the main measure of Y_t since it reflects what households can use for consumption (Catherine, Sodini, and Zhang, 2020). The deterministic component of income is not of particular interest in my framework. Instead, the focus is on the stochastic component y_t as it generates the income risk parameters. To match the consumption insurance literature (Blundell, Pistaferri, and Preston, 2008; Santaeulalia-Llopis and Zheng, 2018), I model the stochastic component as the sum of a permanent and transitory component, where the permanent component is a martingale process with serially uncorrelated innovations (ζ_t), and the transitory process is an MA(0) process with serially uncorrelated innovations (ε_t). Insofar as the deterministic component is predictable, income shocks are changes in the unpredictable component:

$$\Delta y_t = \zeta_t + \varepsilon_t \tag{2.1}$$

Similarly, the consumption process combines deterministic and stochastic compo-

[9] Although, I do not include the individual dimension of these components, it is understood that different agents have different processes.

nents. I assume the deterministic consumption growth rate γ_c to be determined by the interest rate (r), the discount rate (δ), and the agent's risk aversion (η):

$$\gamma_c = (\delta(1+r))^{1/\eta} \tag{2.2}$$

The initial level of consumption is a function the household's income at $t = 0$ and the saving rate s. The deterministic consumption path is solved assuming no uncertainty given the deterministic income growth rate γ_y. The risk component will reside in the stochastic component of income. I can therefore write:

$$\bar{c_t} = (1-s)y_0 \left[\frac{\sum_t^T (\gamma_y/(1+r))^t}{\sum_t^T (\gamma_c/(1+r))^t} \right] \tag{2.3}$$

The deterministic consumption growth rate γ_c is common to all households. The parameters generating difference are s, y_0, and η. To simplify the notation, I write $\bar{c_t}$ as $c_0(\gamma_y)$. Differences in welfare are thus generated by changes in the saving rate, initial income level, deterministic income growth rate, and risk aversion.

The stochastic component of consumption is modeled to reflect uncertainty. I adopt the process in (Blundell, Pistaferri, and Preston, 2008) and used by Duvernois (2021):

$$\Delta c_t = \phi_t \zeta_t + \psi_t \varepsilon_t + \xi_t \tag{2.4}$$

Note that I take the change in stochastic consumption as the innovations in the permanent and transitory components reflect the shocks. ϕ and ψ in equation (2.4) what Blundell, Pistaferri, and Preston (2008) refer to as partial insurance parameters. They reflect the share of the income shock variance that gets transmitted to consumption. ϕ quantifies the transmission of permanent shocks while ψ quantifies the transmission of transitory shocks. ξ measures consumption shocks that are unrelated to income shocks and can be understood as a shift in preference.

ζ and ε respectively represent permanent and transitory income shocks. In the welfare representation, risk is the main concept. I define permanent and transitory income risk as the cross-sectional variance of ζ and ψ written as σ_ζ and σ_ε. To identify these parameters, Duvernois (2021)[10] rely on the minimum distance estimator developed by Blundell, Pistaferri, and Preston (2008).

[10] More detail is provided in the data section of this paper

Having described the income and consumption process of the representative household, I now turn to the welfare framework adopted from Santaeulalia-Llopis and Zheng (2018). Consider a representative household with time-separable constant-relative-risk-aversion (CRRA) utility function. Let δ be the discount factor and η the risk-aversion coefficient. This representative agent's ex-ante welfare can be written as:

$$E\sum_{t=0}^{\infty}\delta^t u(C_t) = E\sum_{t=0}^{\infty}\delta^t \frac{C_t^{1-\eta}}{1-\eta} \tag{2.5}$$

Welfare, in this case, is quantified as the lifetime utility of consumption for the representative agent. The welfare can be characterized by the income growth rate γ_y, the total income risk $\sigma = (\sigma_\zeta, \sigma_\varepsilon)$, and the degree of transmission $\Phi = (\phi, \psi)$. While I suppress the time dimension of the parameters for space, the parameters in the welfare environment can be time varying. Furthermore, I suppress the cross-sectional dimension; the parameters can however be different for agents being representative of different groups of households. Indeed, as shown by Santaeulalia-Llopis and Zheng (2018), urban and rural households in China are characterized by different income, risk, and transmission environment. I can write the household utility with the following notation:

$$E\sum_{t=0}^{\infty}\delta^t u(C_t) \equiv E\sum_{t=0}^{\infty}\delta^t u(C_t; \gamma_y, \sigma, \Phi) \tag{2.6}$$

Following Santaeulalia-Llopis and Zheng (2018)[11] and the definition of the consumption path, the expected sum of discounted utility can be written as:

$$E\sum_{t=0}^{\infty}\delta^t u(C_t) = E\sum_{t=0}^{\infty}\delta^t u(\bar{c}_t \cdot c_t)$$

$$E\sum_{t=0}^{\infty}\delta^t u(C_t) = \frac{(\bar{c_0}\gamma_y)^{1-\eta}}{1-\eta}c_0^{1-\eta}\sum_{t=1}^{T}(\gamma_c^{1-\eta}\delta)^t exp\left(\frac{1}{2}(1+\eta)^2(\phi^2\sigma_\zeta^2 + \psi^2\sigma_\varepsilon^2 + \sigma_\xi^2)t\right) \tag{2.7}$$

Equations (2.6) and (2.7) are identical, the notation is simply more tractable and illustrative of the following exercise. The expected total welfare can be written to reflect the representative agent's current state of nature. The current state of nature i is characterized by an income growth rate (γ_y^i), income risk (σ^i), and insurance (Φ^i). There also exists another state of nature j characterized by different parameters: $(\gamma_y^j, \sigma^j, \Phi^j)$. The expected lifetime consumption for the representative agent in environment i will be different from the

[11]Full derivations are found in their internet appendix.

expected lifetime consumption for the representative agent in environment j. I can write the difference in total welfare as $1+\omega_T$ in the following equation:

$$E\sum_{t=1}^{T}\delta^t u((1+\omega_T)C_t;\gamma_y^i,\sigma^i,\Phi^i) = E\sum_{t=1}^{T}\delta^t u(C_t;\gamma_y^j,\sigma^j,\Phi^j) \tag{2.8}$$

$(1+\omega_T)$ in equation (2.8) measures the total percentage change in lifetime equivalent consumption from switching from environment i to environment j. In other words, I ask: what would have been the welfare equivalent of a representative household i if their environment was characterized by j. Note that the welfare change is calculated for T periods while the subscript T in ω represents the total change. The total change can be decomposed into a growth, risk, and insurance effect. Indeed, from equation (2.8), it is easy to see that each parameter in the i^{th} environment can be individually replaced by the j^{th} environment parameter. The growth effect, noted as $1+\omega_G$ is the percentage change in consumption equivalent for switching from $(\gamma_y^i,\sigma^i,\Phi^i)$ to $(\gamma_y^j,\sigma^i,\Phi^i)$. I can re-write equation (2.7) to reflect the change in environment and the growth effect:

$$(1+\omega_G)\frac{(\bar{c_0}^i)^{1-\eta}}{1-\eta}E\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t(c_{i,t})^{1-\eta} = \frac{(\bar{c_0}^j)^{1-\eta}}{1-\eta}E\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t(c_{i,t})^{1-\eta} \tag{2.9}$$

Note that $c_{i,t}$ is on both sides of equation (2.9). That is the stochastic component of the consumption path and only depends on the risk and transmission parameters. Thus, the growth effect simplifies to:

$$(1+\omega_G)^{1-\eta} = \left(\frac{\bar{c_0}^j}{\bar{c_0}^i}\right)^{1-\eta} \tag{2.10}$$

In equation (2.3), $\bar{c}$ is shown to depend on initial (at $t=0$) levels of income, the income growth rate, and the saving rate. This means that the growth effect from changing to environment j also reflects a change in the representative saving rating. Santaeulalia-Llopis and Zheng (2018) use an average saving rate between the beginning and end of their sample period. However, in the United States, the personal saving rate has been steadily declining. Figure (2.1) shows that personal saving rate[12] went from 11% in 1980 to under 7.5% in 2016. Furthermore, Dynan, Skinner, and Zeldes (2004) find differences in saving behavior across groups, mainly focusing on income and wealth groups. Similarly, Juster et al. (2006) find a decline in the saving rate caused by capital gains, thus affecting households based on their participation status. It is easy to see from equation (2.3) that a higher saving

[12]The Bureau of Economic Analysis defines the personal saving rate as the ratio of personal saving to personal disposable income.

rate will lower deterministic consumption path, ceteris paribus. The saving rate plays a role in a household's ability to insure so the insurance effect will also be affected by changes in the saving rate. However, the effect of the saving rate on the partial insurance parameters cannot be accounted for in this framework. Time-series and cross-sectional differences in savings play a role which I will account for in the growth effect.

Figure 2.1: Personal Saving Rate

The risk effect can be understood as the percentage change in lifetime consumption equivalent of a representative agent in environment i caused by them experiencing the risk environment j. In other words, the incremental change from $(\gamma_y^j, \sigma^j, \Phi^i)$ to $(\gamma_y^j, \sigma^j, \Phi^i)$. From equation (2.7) and (2.9), I write the risk effect $(1+\omega_R)$ as:

$$
\begin{aligned}
(1+\omega_R)\frac{(\bar{c_0}^j)^{1-\eta}}{1-\eta} c_0^{1-\eta}\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_i^2\sigma_{\zeta,j}^2+\psi_i^2\sigma_{\varepsilon,j}^2+\sigma_\xi^2)t\right) \\
= \frac{(\bar{c_0}^j)^{1-\eta}}{1-\eta} c_0^{1-\eta}\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_i^2\sigma_{\zeta,i}^2+\psi_i^2\sigma_{\varepsilon,i}^2+\sigma_\xi^2)t\right) \quad (2.11)
\end{aligned}
$$

Note that σ_ξ^2 does not have an i or j subscript. σ_ξ^2 is the variance of consumption shocks not caused by income shocks[13]. It is therefore not attributable to risk or insurance effects. In that regards, it is similar to risk aversion. It may nonetheless cause changes in equivalent consumption. Indeed, there may be cross-sectional differences across environments. In practice, risk-aversion and preference shifts are different across groups. While differences in risk preferences are likely to cause differences, the variances of preference shifts are too small to cause significant changes in lifetime consumption equivalents. Solving for the risk effect in equation (2.11):

$$(1+\omega_R)^{1-\eta} = \frac{\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_i^2\sigma_{\zeta,j}^2 + \psi_i^2\sigma_{\varepsilon,j}^2 + \sigma_\xi^2)t\right)}{\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_i^2\sigma_{\zeta,i}^2 + \psi_i^2\sigma_{\varepsilon,i}^2 + \sigma_\xi^2)t\right)} \tag{2.12}$$

The insurance effect is similar to the risk effect. It is the welfare change caused by switching from environment i to j and experiencing the transmission parameters of representative household j. It is written as:

$$(1+\omega_I)^{1-\eta} = \frac{\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_j^2\sigma_{\zeta,j}^2 + \psi_j^2\sigma_{\varepsilon,j}^2 + \sigma_\xi^2)t\right)}{\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_i^2\sigma_{\zeta,j}^2 + \psi_i^2\sigma_{\varepsilon,j}^2 + \sigma_\xi^2)t\right)} \tag{2.13}$$

The total welfare effect is described in simple notation in equation (2.8). Santaeulalia-Llopis and Zheng (2018) show that it want be written in the following form:

$$(1+\omega_T)^{1-\eta} = \left(\frac{\bar{c_0}^j}{\bar{c_0}^i}\right)^{1-\eta} \frac{\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_j^2\sigma_{\zeta,j}^2 + \psi_j^2\sigma_{\varepsilon,j}^2 + \sigma_\xi^2)t\right)}{\sum_{t=1}^{T}(\delta\gamma_c^{1-\eta})^t exp\left(\frac{1}{2}(1+\eta)^2(\phi_i^2\sigma_{\zeta,i}^2 + \psi_i^2\sigma_{\varepsilon,i}^2 + \sigma_\xi^2)t\right)} \tag{2.14}$$

which simplifies to:

$$(1+\omega_T)^{1-\eta} = (1+\omega_G)^{1-\eta}(1+\omega_R)^{1-\eta}(1+\omega_I)^{1-\eta} \tag{2.15}$$

The advantages of this decomposition are twofold. Firstly, it allows for a differentiation of the risk and insurance effect. Indeed, Krueger and Perri (2006) argue that an exogenous increase in income risk incentivizes households to increase their insurance ability. The net welfare changes thus depend on the size of the income shock variance and the share

[13] Note that it is not a measurement error, simply a preference shift. Measurement error is captured by $u_{i,t}$ in the minimum distance estimator. See Blundell, Pistaferri, and Preston (2008) for details.

that gets transmitted to consumption. Furthermore, Iachan, Nenov, and Simsek (2021) find that while participation has increased (indicating an increase in financial innovation which increases risk-sharing as predicted by theory[14]), the variance of households' portfolio returns has also increased (thus generating higher risks for households). The net welfare changes may be ambiguous in this case. This framework allows for a clear separation of the two effect as well as a total effect to quantify the net welfare gains or losses. Secondly, the framework is flexible and allows for multitudes of environment definitions. The framework allows for a decomposition of welfare effects for the same representative agent across time, or for representative agents of different groups in the same period.

2.4 Data

The data used come from the longitudinal survey, the Panel Survey of Income Dynamics (PSID). The PSID has been surveying households since 1968. The panel dimension of the PSID makes it unique in the US. Indeed, it tracks the original 5,000 households as well as households formed by their descendants. The PSID reports various household characteristics, food spending, and income. In 1999, the PSID was redesigned and now includes broader measures of consumption. This allows researchers to investigate more complex and robust relationships between households' consumption and income. Furthermore, a measure of consumption can be imputed backward using methodology developed by Attanasio and Pistaferri (2014) and Fisher and Johnson (2020).

The income risk and transmission parameters are gathered from Duvernois (2021). I provide a cursory description of their data and estimation process here; a detailed description can be found in their sections (3) and (4.3). The data in the calculations of certain parameters are also from Duvernois (2021) to precisely match the sample selection. I define a household by following the reference person through each wave. I drop the household-wave observation for households experiencing a major family composition change as defined by a change to the reference person or spouse. I assign these households a new unique identifier in the following wave, thus becoming a separate household. All households are required to be present for at least 4 consecutive waves. I restrict the reference person to be between 25 and 65 years old.

Households that received positive business income are considered as entrepreneurs. Business income is comprised of labor and asset income for non-incorporated business or farms. Transfer recipients are households that receive positive transfer income (including social security income; I also include other family members in the calculation of transfer income). Market participants are households that hold stocks directly or indirectly through

[14]Simsek (2013a) and Simsek (2013b)

IRAs. To identify stockholders in each wave, I follow Guvenen (2007a). High wealth household are household with above median total wealth[15].

The main measure of income is total family disposable income. Total family income is defined as the sum of wage, business, financial[16], and transfer income for all individual in the households[17]. I use NBER Taxsim and guidelines from Feenberg and Coutts (1993) and Kimberlin, Kim, and Shaefer (2014) to estimate household taxes for the 1992 wave onward and use the PSID variable for federal taxes for the prior waves. Household income, consumption, and wealth are deflated to 1980 dollars. Table 1 panel A presents summary statistics for the average disposable income in the sample.

I calculate the average income for all households in sample as well as each subgroup. Furthermore, I calculate the average income at three points in time, 1980, 1996, and 2016. 1980 and 2016 are the first and last waves in the sample. Duvernois (2021) tests whether insurance parameters are different around 1996. Although arbitrary, 1996 does appear to be a natural break point in the data and is represented by different risk and transmission environment. I then calculate the average yearly growth rate of disposable income for each subgroup, for each subsample, and over the entire sample period. The income growth rates are calculated as:

$$\gamma_y = exp\left(\frac{log(Y_t/Y_{t-1})}{T_t - T_{t-1}}\right) \tag{2.16}$$

Equation (2.16) is used to calculate the deterministic consumption path in equation (2.3). The average disposable income in the overall sample is \$26,900. The average annual growth rate over the entire sample period is 1.19%. However, the growth rate is concentrated on the first portion of the sample. Indeed, the average annual growth rate over 1980 to 1996 is 1.66%. The average income level was \$34,993 and grew to \$41,199 to 2016, implying a growth rate of only 0.82%. Income grew at a much slower pace in the second portion of the sample. This pattern is seen across almost all subgroups. Entrepreneurs have an income growth rate of 1.60% over the 1980-1996 sample period versus 1.17% for the 1996-2016 period. Non-entrepreneurs have a similar growth rate in the pre-1996 sample period (1.54%) but have a smaller growth rate in the post period relative to their counterparts. Not only do nonentrepreneurs have lower growth rates, their disposable income levels are also lower. Indeed, the average entrepreneur has disposable income of \$34,427 vs \$25,761 in 1980 and \$56,020 vs \$38,912 in 2016 for the average employed household. Households that do not receive transfers have unsurprisingly higher levels of disposable income. Over the entire

[15]These households are identified using a similar framework used to identify stockholders.
[16]Financial income is defined as the sum of income from dividends, interest, trust funds, and royalties
[17]This includes the reference person, spouse, and other family members.

Table 2.1: Parameters used in the welfare decomposition

This table reports the parameters used in the welfare decomposition presented in section (2.3). The parameters represent the average (per group and time period) which I take to be the parameters describing the representative household. Y is the representative household disposable income level in 1982 dollars. Disposable income is the sum of labor, business, transfer, social security, and financial income minus federal taxes. γ^Y *is the growth rate of income as described in equation (1.16). The saving rate for all households is retrieved from the Bureau of Economic Analysis. The saving rate per group is defined as the group's median saving rate calculated as the change in wealth over the average income during the same period.* η *is implied by the households' responses to the 1996 risk attitudes question. See section (2.4) and Hryshko, Luengo-Prado, and Sørensen (2011) for discussion. Participating households have at least 1$ directly or indirectly invested in the stock market, entrepreneurs receive at least 1$ in business income, transfer recipients receive at least 1$ in transfer income, wealthy households are above the median level of total wealth.*

	All	Participation ?		Entrepreneur?		Transfer Receipts?		Wealthy?	
	Households	Yes	No	Yes	No	Yes	No	Yes	No
Panel A. Household Parameters									
Y_{1980}	26900	35520	21955	34427	25761	22861	29164	29710	16315
Y_{1996}	34993	48018	24850	44381	32918	33234	36301	40919	22635
Y_{2016}	41199	55001	31919	56020	38912	34005	45090	47593	25556
$\gamma^Y_{1980-2016}$	1.19%	1.22%	1.04%	1.36%	1.15%	1.11%	1.22%	1.32%	1.25%
$\gamma^Y_{1980-1996}$	1.66%	1.90%	0.78%	1.60%	1.54%	2.37%	1.38%	2.02%	2.07%
$\gamma^Y_{1996-2016}$	0.82%	0.68%	1.26%	1.17%	0.84%	0.11%	1.09%	0.76%	0.61%
Saving $\text{Rate}_{1980-2016}$	7.39%	9.70%	1.96%	11.84%	3.94%	2.79%	5.69%	11.10%	-1.45%
Saving $\text{Rate}_{1980-1996}$	9.12%	15.66%	1.93%	15.64%	4.88%	3.88%	6.86%	12.48%	-0.72%
Saving $\text{Rate}_{1996-2016}$	5.82%	8.81%	1.98%	10.43%	3.70%	2.62%	5.31%	10.70%	-1.61%
Risk Aversion: η	11.71	10.47	12.74	10.95	11.88	11.76	11.68	12.26	11.47
Panel B. Partial Insurance Parameters									
$\phi_{1980-2016}$	0.285	0.190	0.404	0.208	0.297	0.337	0.249	0.196	0.334
$\psi_{1980-2016}$	0.070	0.064	0.053	0.015	0.079	0.084	0.055	0.035	0.066
$\phi_{1980-1996}$	0.320	0.318	0.408	0.221	0.309	0.308	0.330	0.245	0.331
$\phi_{1996-2016}$	0.251	0.125	0.408	0.137	0.288	0.347	0.177	0.114	0.324
$\psi_{1980-1996}$	0.057	0.009	0.030	0.028	0.074	0.086	0.040	0.017	0.034
$\psi_{1996-2016}$	0.082	0.102	0.081	0.013	0.083	0.086	0.067	0.058	0.106
Panel C. Average Permanent and Transitory Income Shocks									
$\sigma^2_{\zeta,1980-2016}$	0.022	0.017	0.020	0.030	0.021	0.021	0.018	0.018	0.024
$\sigma^2_{\varepsilon,1980-2016}$	0.032	0.025	0.041	0.054	0.026	0.053	0.021	0.028	0.035
$\sigma^2_{\zeta,1980-1996}$	0.025	0.016	0.024	0.036	0.024	0.025	0.019	0.020	0.029
$\sigma^2_{\zeta,1996-2016}$	0.017	0.018	0.014	0.025	0.016	0.017	0.016	0.015	0.018
$\sigma^2_{\varepsilon,1980-1996}$	0.029	0.021	0.038	0.046	0.024	0.051	0.019	0.024	0.032
$\sigma^2_{\varepsilon,1996-2016}$	0.035	0.032	0.046	0.062	0.029	0.056	0.022	0.034	0.039

sample period, the income growth rates are relatively similar (1.11% for transfer recipients vs 1.22%). nontransfer recipients have a larger growth rate of income in the pre-period than in the postperiod. Transfer recipients have an income growth rate in the postperiod virtually equal to 0. Indeed, transfer recipients' income grew 2.37% annually on average between 1980 and 1996 but grew an average 0.11% between 1996 and 2016. Households above and below the median level of wealth have similar rates of income overall and the subsamples. However (and unsurprisingly), wealthier households have much larger levels of disposable income. The average household above the median have a disposable income of $29,710 in 1980 while households below the median have a disposable income of $25,556 in 2016. Households that do not participate are the only households that do not fit the general pattern of table (2.1). Indeed, participating households have a higher income growth rate prior to 1996. nonparticipating households, however, enjoy a faster rate of income growth between 1996 and 2016 than they did between 1980 and 2016 (1.26% vs 0.78%). Santaeulalia-Llopis and Zheng (2018) empirical framework is similar in that their sample period roughly match ours. However, their sample of Chinese households experienced an acceleration of income growth rates in the second period thus generating significant growth effect. The effect is likely to be reversed in my sample considering the observed slow down.

Panel A also presents the personal saving rate. In this particular empirical exercise, I need the personal saving rate for all households as well as the saving rate for each subgroup. For the overall sample, I use the personal saving rate calculated by the Bureau of Economic Analysis[18]. This measure is more robust to data errors than the saving rate I could infer from PSID. I take the average saving rate over the entire sample period and over each subsample. Similarly to the income growth rate, the saving rate has been declining. The average rate is 7.39% over the entire sample period but is 9.12% in the 1980-1996 and 5.28% in the 1996-2016 subperiods. These patterns are consistent with Dynan, Skinner, and Zeldes (2004) and Juster et al. (2006) who also document declines in the saving rate in their sample period. Calculating the saving rate in PSID can be tricky. Santaeulalia-Llopis and Zheng (2018) use the same saving rate in all period and for all groups. While this would certainly be easier, it would discount the heterogeneity observed across households.

There are several issues with estimating the saving rate. As pointed out by Dynan, Skinner, and Zeldes (2004), there are several ways to define savings but each presents a unique set of advantages and drawbacks. Savings can be defined to include all forms of savings such as the realized and unrealized capital gains on financial assets, businesses, non-housing and housing real estate. Another definition would exclude capital gains from income

[18]See figure (2.1).

and look at the difference between income and consumption. This is referred to as active savings. However, capital gains are uncertain and should not be reflected in the deterministic consumption path. An empirical issue also arises from using active savings. A measure of active savings exists in the PSID for the 1989 and 1994 waves. However, the required variables are not consistently available. Thus the measure of saving rate would be reflective of the situation between 1984 and 1994 and not reflect the decrease in savings observed in other time-series of saving rate. A simple definition of saving would be the difference between income and consumption. However, this definition requires an unbiased measure of income. Biases introduced by transitory income may be reflected in the saving rate across groups. A good proxy for permanent income is difficult to find and use in the above definition. To mitigate these problems, I calculate the personal saving rate as the change in wealth divided the average income multiplied by n, n being the number of years between the two wealth points. The average income is a proxy of permanent income (Catherine, Sodini, and Zhang, 2020). Indeed, prior to 1999, the wealth module is administered in the 1984, 1989, and 1994 waves. Thus I measure the change in wealth between 5 years intervals; the average income will reflect permanent income. After 1999, wealth data is available at each wave, but the survey is biennial making $n = 2$. I calculate the median saving rate per group to mitigate the effect of large changes in wealth and potential biases in income.

The subgroup saving rates are in line with the saving rate measure from the BEA. Indeed, for most groups, I find the saving rate to be lower in the second subsample. Market participants have a saving rate of 15.66% prior to 1996 and 8.81% after 1996. Entrepreneurs and wealth households have similar patterns. Non participants and non-entrepreneurs have much lower saving rates relative to their group counterparts and do not experience a large change between the two time samples. Households below the median level of wealth have a negative saving rate indicating a negative change in wealth.

Finally, the last parameter to calculate is η, the risk aversion coefficient. Santaeulalia-Llopis and Zheng (2018) use $\eta = 2$ and $\eta = 4$ in their welfare decomposition. However, these values are somewhat arbitrary and may not reflect the true risk aversion of the average representative household. Furthermore, applying a single risk aversion to all types of households may not be truly representative. Households with lower risk aversion may select themselves into groups characterized by higher income risk profiles. Estimating risk aversion at the household level is extremely challenging mainly due to lack of data. However, in 1996, the PSID included a series of questions designed to reveal households' preferences and attitude towards risk. If the reference person had been employed in 1995, households were asked a series of 6 questions presenting a hypothetical gamble. The gambles offered different income prospects designed to measure the reference person's willingness to accept risk. All hypotheticals have the same structure. The status quo (i.e., refusing the gamble) is a job guaranteeing lifetime income equal to the household's current total income. The alternative (ie. taking the gamble) is a job doubling the current income level with a 50% probability or cutting income by a $1-\lambda$ fraction.

$1-\lambda$ is equal to 33.33% in the first gamble. On one hand, if the reference person refuses, they are asked another hypothetical with $1-\lambda = 20\%$. If the reference person accepts, no other hypotheticals are presented; if they refuse, they are asked whether they would accept the job with $1-\lambda = 10\%$. On the other hand, if the reference person accepts the initial gamble, they are presented with another gamble with a higher λ fraction. If the reference person accepts the initial gamble, they are asked whether they would take the job with $1-\lambda = 50\%$[19]. If they accept, they are presented with a final hypothetical with $1-\lambda = 75\%$. Accepting the gamble simply means the following is true for the reference person: $\frac{1}{2}U(2c) + \frac{1}{2}U(\lambda)c \geq U(C)$.

There are therefore 6 categories of households characterized by different levels of risk aversion. On the two extreme sides of the spectrum, households that accept all gambles are the least risk averse, and households that refuse all gambles are the most risk averse. Hryshko, Luengo-Prado, and Sørensen (2011) use CRRA utily and methodology by Barsky et al. (1997) to bracket each groups risk aversion parameters and calculate their respective conditional means. From least to most risk averse, the relative risk aversion coefficients are 0.18, 0.63, 1.46, 2.83, 5.44, and 33.9. These coefficients are assigned to households based on the 1996 responses. I then take the average by group, and assign each group the same risk aversion for each wave.

The average risk aversion (as implied by Hryshko, Luengo-Prado, and Sørensen, 2011) is 11.71. This coefficient is higher than the values used in Santaeulalia-Llopis and Zheng (2018). However, in the overall sample, a more accurate measure of risk aversion is less meaningful. Indeed, a higher η parameter simply amplifies the risk and insurance effects. Granted, a more accurate risk aversion will result in a more accurate percentage change in expected lifetime utility. Picking a reasonably low and high coefficients would serve to bracket the true change and would be just as meaningful. Furthermore, in the overall sample, I simply investigate the sub-period effects and the coefficient of risk aversion is static. Indeed, while I can calculate the average risk aversion coefficient in 1980 or 2016 by carrying backward and forward the 1996 coefficient, this would be problematic for two reasons. First, it would require a significant number of households to be present at either point. Second, it would

[19] If they refuse, no further questions are asked.

require the risk aversion to remain constant. Hryshko, Luengo-Prado, and Sørensen (2010) show that older households are more risk-averse than younger households indicating that a change over the life-cycle. Furthermore, households that are present both in 1980 and 1996 are likely to be the 1996 older households; similarly, households present in 1996 and 2016 are likely to be the 1996 younger households. Thus, by construction, the average risk aversion will be higher in the pre-period and lower in the postperiod but may not reflect changes in the true risk-aversion parameter. I, on the other hand, argue that the demographic composition and risk aversion level are unlikely to vary through time for homogeneous groups.

In table 2.1, I find that participating households are less risk averse than their counterparts (10.47 vs 12.74). Entrepreneurs are also less risk averse than employees (10.95 vs 11.88). Whether households receive transfer income or not does not seem to play a role on risk aversion. Finally, wealthy individuals are more risk averse than households below median levels of wealth. Although surprising at first glance, wealthy individuals tend to be older. Overall, there are cross-sectional differences in risk aversion that will potentially affect the welfare decomposition. Panel B presents the partial insurance parameters estimated by Duvernois (2021) for several measures of income. These parameters are estimated using a Generalized Method of Moments (GMM) and moments implied by the variances, covariances, and autocovariances of the income and consumption process detailed in equations (2.1) and (2.4). See Duvernois (2021) for a thorough discussion. ϕ quantifies the degree of transmission of permanent shocks to consumption. ψ quantifies the degree of transitory insurance. 0 implies perfect risk-sharing or no transmission of shocks to households' consumption. Everything else being equal, a lower partial insurance parameter implies higher welfare.

In panel C, I report the average variance of permanent and transitory shocks. The averages are calculated over 3 different periods: the whole sample period (1980-2016), pre 1996 (1980-1996), and post 1996 (1996-2016). The variances for the 1980-2016 sample are estimated with static transmission parameters, while the pre and post variances are estimated with time-varying transmission parameters. In the all household sample, transitory shocks (σ_{ε}) are larger than permanent shocks (σ_{ζ}). However, when looking at subperiods, permanent shocks are larger in the pre-sample and smaller in the post-sample. I observe a similar pattern for all subgroups expect participating households. Indeed, participating households' transitory shocks are consistently larger than permanent shocks. Interestingly, panel C shows that entrepreneurs experience the largest shocks (both permanent and transitory).

2.5 Results

2.5.1 Welfare comparison across time

I present in the following section the results from the welfare decomposition. Empirically, I start with a baseline environment characterized by income growth, income risk, and income shock transmission parameters. I then successively replace each parameter with their counterpart environment parameters. Several empirical details are worth mentioning. Firstly, the income growth rate, saving rate, and risk aversion are not estimated but measured from the data. The sample selection process insures that these parameters are as unbiased as possible. The risk and transmission parameters are, on the other hand, estimated and may suffer from biases. These biases can be addressed using a bootstrapped estimation. By assumption, the shocks are *i.i.d.* I can therefore draw a random vector of parameters from the multivariate normal distribution characterized by the vector of estimated parameters and the variance-covariance matrix of the estimates. I draw a thousand parameters and estimate the welfare changes. The 95% confidence intervals are reported in each table. As the growth effect does not depend on any risk-sharing parameters, it remains unaffected by the bootstrap procedure. Secondly, the welfare path depends on average estimates of time-varying processes. Indeed, I use the average income growth rate over the sample period. Similarly, the transmission parameters quantify the amount of shocks being transmitted to consumption (on average). I do, however, have an estimates of the variance of shocks (both permanent and transitory) for each year. I resolve the time-series of income risk by taking the average. Finally, I take $\delta = 0.98$, and interest rate $= 2\%$.

Table 2.2 presents the welfare decomposition for all households as well as subgroups. In panel A, the baseline scenario is characterized by parameters from the 1980-1996 subsample. I then successively replace the baseline parameters with their 1996-2016 counterparts. The total effect measures the percentage change in consumption equivalent of a representative households had they experienced the 1996-2016 environment. I consider two levels of risk aversion (in line with Santaeulalia-Llopis and Zheng, 2018) of 2 and 4. Recall from table 2.1, the income growth parameters were lower in the second subsample. Thus, the growth effect is negative. The representative agent would have a consumption equivalent that is 6.28% lower had they received the 1996-2016 average growth rate of income. The negative effect of the deterministic component is unsurprising and not that interesting in this research framework. The stochastic component of income is the main point of research. The size of the growth effect is however relevant. Santaeulalia-Llopis and Zheng (2018) find that the rapid growth of the Chinese economy was accompanied by a deterioration of the income risk and transmission environment. Furthermore, the risk and insurance effect can compare the growth effect in magnitude. Duvernois (2021) documents a lower income risk environment and lower transmission (higher permanent insurance) in the post period. The PSID representative household had moderate welfare gains 0.33% and 1.07% with risk aversion coefficients of 2 and 4 respectively. These positive changes in annual consumption variations are further increased by a modest insurance effect accounting for 0.22% and 0.72%. The insurance and risk effect, though positive, are not large enough to compare to the large welfare loss caused by the slow down in income growth. Indeed, with low levels of risk aversion ($\eta = 2$), the total effect is -5.77%; with $eta = 4$, the total effect is -4.59%. There is an unambiguous welfare loss from moving from the PSID overall representative economic environment in the 80s/90s to that of the 21^{st} century.

There is a similar trend across subgroups. Households that do not receive business income have a negative growth effect of 5.26%. They experience a positive risk effect in line with the average sample (0.31% vs 1.03% for the two risk aversion levels). However, the insurance effect is 0.06% and 0.20% suggesting that employees experienced a slight improvement in their risk environment but virtually no changes in their insurance environment. Indeed, while the permanent pass-through coefficient is slightly lower, the transitory component is slightly higher, thus counteracting each other. Entrepreneurs also have a negative growth effect, albeit lower than their employed counterparts. With a risk aversion coefficient of 4, the combined risk and insurance effects reduce the negative growth effect to -1.60%; a welfare loss with an upper bound of -0.29%.

Households that receive transfer income have the worse total effect of all subgroups. Indeed, they experienced a deterioration in the insurance environment (-0.18% and -0.59%) in the post period and experienced only a modest improvement in the risk environment (0.30% and 0.98%). Not only do these two effects cancel out, but they also pale in comparison to the growth effect of -16.02%. Their counterpart fare much better. Households that do not receive any transfers have a total effect of -0.25% for $\eta = 4$. The insurance effect (1.50%) does much in offsetting the negative growth effect of 2.22%. High and low wealth households have a similar negative growth effect, but the welfare effects of their respective risk and insurance environments differ. High wealth households have a small but positive risk and insurance effect. Low wealth households have a negative insurance effect but have a significant welfare improvement from their lower risk-environment in the post 1996 sample.

Market participants have a large negative growth effect. The risk effect is insignificantly different from 0. With a risk aversion coefficient of 4, the insurance effect is 1.56%, reducing the negative total welfare effect to -7.74%. nonparticipating households are the only group that experience a positive growth effect in their alternative environment. A positive growth effect of 3.76% compounded by 0.72% and 2.38% (with $\eta = 2$ and $\eta = 4$ respectively).

Table 2.2: Welfare Effects of Growth, Risk, and Insurance: Comparison Across Time

This table shows the effects on household welfare in percentage annual consumption equivalent from a baseline level in which I replace income growth parameters, income risk parameters, and transmission parameters from one period with their alternative period counterpart. The two sub-periods are 1980-1996 and 1996-2016. I use CRRA utility and risk aversion coefficient of 2 and 4. The income growth parameters are estimated using PSID data. The growth, income, insurance effects are estimated using equations (2.10), (2.12), (2.13), and (2.15). In panel A, 1980-1996 is the baseline; in panel B, 1996-2016 is the baseline. 95% confidence intervals are reported below and are based on 1000 bootstrap replicas.

	All Households		Employees		Entrepreneurs		Mkt. participants		No participation		No transfer		Transfer rcpt.		High wealth		Low wealth	
	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$	$\eta=2$	$\eta=4$
Panel A. 1980-1996																		
Growth Effect $\{\gamma^{post},\sigma^{pre},\Phi^{pre}\}$	-6.28%	-6.28%	-5.26%	-5.26%	-3.28%	-3.28%	-9.00%	-9.00%	3.76%	3.76%	-2.22%	-2.22%	-16.02%	-16.02%	-9.30%	-9.30%	-10.69%	-10.69%
Risk Effect $\{\gamma^{post},\sigma^{post},\Phi^{pre}\}$	0.33% [0.22 0.42]	1.07% [0.73 1.38]	0.31% [0.20 0.41]	1.03% [0.67 1.33]	0.21% [-0.05 0.43]	0.67% [-0.15 1.39]	-0.05% [-0.19 0.15]	-0.17% [-0.61 0.49]	0.72% [0.41 0.97]	2.38% [1.35 3.21]	0.15% [0.04 0.27]	0.50% [0.11 0.87]	0.30% [0.01 0.52]	0.98% [0.02 1.71]	0.11% [0.03 0.18]	0.35% [0.10 0.57]	0.52% [0.30 0.71]	1.74% [0.98 2.36]
Insurance Effect $\{\gamma^{post},\sigma^{post},\Phi^{post}\}$	0.22% [0.04 0.39]	0.72% [0.13 1.27]	0.06% [-0.12 0.23]	0.20% [-0.40 0.76]	0.33% [-0.15 0.69]	1.05% [-0.50 2.22]	0.48% [0.05 0.82]	1.56% [0.13 2.66]	-0.11% [-0.56 0.29]	-0.35% [-1.82 0.95]	0.46% [0.21 0.69]	1.50% [0.65 2.23]	-0.18% [-0.73 0.25]	-0.59% [-2.41 0.82]	0.25% [0.09 0.40]	0.80% [0.29 1.26]	-0.13% [-0.48 0.19]	-0.42% [-1.58 0.61]
Total Effect	-5.77% [-5.96 -5.59]	-4.59% [-5.23 -4.00]	-4.90% [-5.11 -4.72]	-4.09% [-4.79 -3.48]	-2.76% [-3.25 -2.36]	-1.60% [-3.24 -0.29]	-8.61% [-8.92 -8.35]	-7.74% [-8.74 -6.91]	4.39% [3.92 4.83]	5.86% [4.26 7.29]	-1.61% [-1.84 -1.38]	-0.25% [-1.02 0.51]	-15.92% [-16.41 -15.53]	-15.69% [-17.32 -14.40]	-8.97% [-9.13 -8.83]	-8.26% [-8.76 -7.80]	-10.33% [-10.70 -10.02]	-9.51% [-10.75 -8.46]
Panel B. 1996-2016																		
Growth Effect $\{\gamma^{post},\sigma^{pre},\Phi^{pre}\}$	7.55%	7.55%	6.25%	6.25%	3.81%	3.81%	11.17%	11.17%	-4.05%	-4.05%	2.55%	2.55%	21.64%	21.64%	11.58%	11.58%	13.52%	13.52%
Risk Effect $\{\gamma^{post},\sigma^{post},\Phi^{pre}\}$	-0.21% [-0.30 -0.09]	-0.69% [-0.98 -0.31]	-0.30% [-0.42 -0.14]	-0.98% [-1.37 -0.47]	-0.09% [-0.19 0.15]	-0.29% [-0.62 0.48]	0.06% [0.01 0.11]	0.20% [0.02 0.37]	-0.77% [-1.18 -0.14]	-2.55% [-3.90 -0.46]	-0.04% [-0.08 0.02]	-0.14% [-0.27 0.07]	-0.42% [-0.77 0.16]	-1.41% [-2.56 0.53]	-0.01% [-0.05 0.03]	-0.03% [-0.15 0.11]	-0.52% [-0.79 -0.13]	-1.75% [-2.64 -0.44]
Insurance Effect $\{\gamma^{post},\sigma^{post},\Phi^{post}\}$	-0.40% [-0.67 -0.13]	-1.31% [-2.20 -0.43]	-0.12% [-0.43 0.19]	-0.39% [-1.41 0.61]	-0.50% [-1.10 0.10]	-1.62% [-3.58 0.30]	-0.54% [-0.86 -0.17]	-1.73% [-2.75 -0.56]	0.10% [-0.82 0.89]	0.33% [-2.83 2.95]	-0.64% [-0.92 -0.36]	-2.07% [-2.96 -1.18]	0.29% [-0.67 1.13]	0.99% [-2.31 3.80]	-0.39% [-0.57 -0.20]	-1.24% [-1.81 -0.62]	0.09% [-0.57 0.67]	0.29% [-1.94 2.25]
Total Effect	6.90% [6.67 7.14]	5.42% [4.68 6.20]	5.81% [5.58 6.07]	4.80% [4.04 5.65]	3.20% [2.72 3.78]	1.83% [0.29 3.70]	10.64% [10.29 11.05]	9.46% [8.35 10.78]	-4.70% [-5.14 -4.22]	-6.19% [-7.64 -4.62]	1.84% [1.58 2.11]	0.28% [-0.58 1.14]	21.48% [20.84 22.26]	21.11% [18.98 23.69]	11.14% [10.94 11.35]	10.17% [9.54 10.83]	13.03% [12.58 13.54]	11.87% [10.38 13.57]

The insurance effect is slightly negative, but the total welfare effect is significantly positive.

Panel B presents the results with the baseline environment being the 1996-2016 period. The signs are reversed and the magnitudes are different, but the results are similar and can be interpreted similarly. The magnitudes are different for two reasons. First, the length of each period is different, so the sums in equations (2.3), (2.12), and (2.13) are different. Furthermore, I use the saving rate of the postperiod in panel B, thus resulting in a different growth effect. The results are simply presented in this fashion for completeness. In table 3, I consider the same welfare decomposition with the risk aversion coefficients implied by the 1996 questionnaire.

Table 2.3 considers the pre-1996 environment as the baseline while panel B considers the post-1996 environment as the baseline. The growth effect across groups remains unchanged and the risk and insurance effects have been amplified. One could simply assume that the much larger risk aversion coefficients considered will be sufficient to generate positive total welfare changes. Indeed, the smallest η considered is 10.47 for market participants. However, the negative growth effects still outweigh the combined risk and insurance effects in some cases. In the PSID representative sample, with a risk aversion coefficient of 11.71, the percentage change in annual consumption equivalent accounted by the lower risk environment of the post-1996 environment is 6.32%. Furthermore, the lower (i.e better) transmission parameters in the post sample generate welfare improvements of 4.21%. The total welfare effect is 3.83% (with a 95% confidence interval of 0.19 and 7.48). With the large negative growth effect, it takes a large risk aversion coefficient to generate a net welfare gain.

With a lower negative negative growth effects, the lower risk environment and better insurance of the post sample generates large total welfare changes of 5.63% for entrepreneurs. Interestingly, with a large enough risk aversion coefficient, the percentage change in annual equivalent consumption generated by a lower income risk environment is larger than the growth effect for entrepreneurs (in absolute value). This is similar for the insurance effect, implying that the total effect would be positive (although closer to 0) even if one of the two effects remained constant between the two environments.

The higher risk aversion coefficient also intensifies the differences between risk and insurance effect. The risk effect was larger than the insurance effect for employees in table 2.2. However, the pattern is more noticeable in table 2.3. The risk effect is equal to 6.20% vs 1.16% for the insurance effect, more than 5 times larger. Similarly, the insurance effect generates large welfare gains for participating households (7.54% with $\eta = 10.47$). However, market participants also experience more risk in the post-1996 period, decreasing the welfare

Table 2.3: Welfare Effects of Growth, Risk, and Insurance: Comparison Across Time with PSID Estimated Risk Aversion

This table shows the effects on household welfare in percentage annual consumption equivalent from a baseline level in which I replace income growth parameters, income risk parameters, and transmission parameters from one period with their alternative period counterpart. The two sub-periods are 1980-1996 and 1996-2016. The constant relative risk aversion coefficient are implied from 1996 risk attitude questions and methodology by Hryshko, Luengo-Prado, and Sørensen (2011). The CRRA coefficient is measure for each household present in 1996 and average over groups. The income growth parameters are estimated using PSID data. The growth, income, insurance effects are estimated using equations (2.10), (2.12), (2.13), and (2.15). In panel A, 1980-1996 is the baseline; in panel B, 1996-2016 is the baseline. 95% confidence intervals are reported below and are based on 1000 bootstrap replicas.

	All Households $\eta = 11.71$	Employees $\eta = 11.88$	Entrepreneurs $\eta = 10.95$	Mkt. participants $\eta = 10.47$	No participation $\eta = 12.74$	No transfer $\eta = 11.68$	Transfer rcpt. $\eta = 11.76$	High wealth $\eta = 12.26$	Low wealth $\eta = 11.47$
Panel A. 1980-1996									
Growth Effect $\{\gamma^{post}, \sigma^{pre}, \Phi^{pre}\}$	-6.28%	-5.26%	-3.28%	-9.00%	3.76%	-2.22%	-16.02%	-9.30%	-10.69%
Risk Effect $\{\gamma^{post}, \sigma^{post}, \Phi^{pre}\}$	6.32%	6.20%	3.52%	-0.79%	16.46%	2.87%	5.83%	2.08%	10.21%
	[4.22 8.19]	[3.98 8.07]	[-0.89 7.29]	[-2.91 2.30]	[8.65 22.46]	[0.63 4.99]	[-0.17 10.18]	[0.57 3.44]	[5.52 13.90]
Insurance Effect $\{\gamma^{post}, \sigma^{post}, \Phi^{post}\}$	4.21%	1.16%	5.49%	7.54%	-2.23%	8.75%	-3.36%	4.77%	-2.34%
	[0.64 7.38]	[-2.44 4.42]	[-3.23 11.59]	[0.02 12.92]	[-11.95 5.69]	[3.48 13.11]	[-14.13 4.33]	[1.59 7.54]	[-8.91 3.20]
Total Effect	3.83%	1.78%	5.63%	-2.91%	18.14%	9.39%	-14.11%	-2.99%	-3.88%
	[-0.19 7.48]	[-2.63 5.54]	[-3.62 12.65]	[-8.05 1.13]	[6.09 28.13]	[4.56 14.03]	[-24.22 -6.72]	[-6.23 -0.13]	[-11.53 2.25]
Panel B. 1996-2016									
Growth Effect $\{\gamma^{pre}, \sigma^{post}, \Phi^{post}\}$	7.55%	6.25%	3.81%	11.17%	-4.05%	2.55%	21.64%	11.58%	13.52%
Risk Effect $\{\gamma^{pre}, \sigma^{pre}, \Phi^{post}\}$	-3.93%	-5.69%	-1.48%	0.97%	-15.53%	-0.79%	-7.99%	-0.19%	-9.54%
	[-5.56 -1.80]	[-7.90 -2.83]	[-3.21 2.40]	[0.08 1.75]	[-23.47 -4.41]	[-1.51 0.38]	[-14.50 2.08]	[-0.90 0.64]	[-14.29 -2.93]
Insurance Effect $\{\gamma^{pre}, \sigma^{pre}, \Phi^{pre}\}$	-7.39%	-2.28%	-8.24%	-8.06%	2.17%	-11.33%	5.95%	-7.24%	1.68%
	[-12.40 -2.69]	[-8.44 3.40]	[-18.23 0.89]	[-12.73 -2.79]	[-20.62 18.40]	[-16.12 -6.83]	[-16.04 22.12]	[-10.55 -3.80]	[-11.89 12.45]
Total Effect	-4.30%	-2.08%	-6.15%	3.19%	-17.20%	-9.79%	18.57%	3.31%	4.42%
	[-8.28 -0.25]	[-6.37 2.53]	[-13.84 2.47]	[-1.97 9.12]	[-25.93 -8.56]	[-14.37 -5.45]	[5.80 32.65]	[-0.29 7.11]	[-3.70 13.12]

gains. Furthermore, the positive effect of higher consumption insurance is not sufficient to offset the negative effect of higher income risks and lower income growth. There is therefore a significant welfare loss from moving from the participating income/consumption environment in the 80s/90s to that of the 21^{st} century. As the only group with positive welfare changes in table 2.2, nonparticipating households have the highest positive effect in table 2.3. This is further exacerbated by the fact that nonparticipating households have the highest risk aversion coefficient.

Transfer recipients' welfare decomposition remains largely unaffected by the change in risk aversion coefficient. Indeed, the risk and insurance effect have the opposite effect. The lower risk environment is generating positive welfare changes while the lower insurance environment is generating welfare losses. Although the risk effect is slightly stronger in magnitude, it is not sufficient to make any kind of dent in the large negative growth effect. The difference between ω_T with $\eta = 4$ and ω_T with $\eta = 11.76$ is 15.69% vs -14.11%. High and low wealth households still have a negative welfare total effect despite the larger risk and insurance effects.

In an unreported test, I solve for the risk aversion coefficient that makes the total welfare effect equal to 0. For all household, η would need to equal 8.75 to force the total welfare effect to equal 0. For employees and entrepreneurs, the value is 9.75 and 6.20, respectively. For market participants, the risk aversion coefficient would need to be as high as 13.50. nonparticipating households would need to be risk-loving and have a negative η to have a net welfare change of 0. nontransfer recipients would be in line with the estimation in table 2.2 with $\eta = 4.20$. High and low-wealth households are similar with $\eta = 16$. This exercise yields unrealistic results for transfer recipients. Indeed, with a large negative growth effect, the combined effect of risk and insurance need to be substantial. However, the effects counteract each other, thus requiring a very large risk aversion coefficient of 77.50 to generate a risk effect sufficiently large to compensate for the insurance and growth effect.

2.5.2 Welfare comparison across households

Until this point, the analysis focuses on how the change of the environment through time affects the welfare of households. The flexibility of the welfare framework allows to consider differing environments across households. Indeed, they are cross-sectional differences in households' risk and transmission environment. I keep the same format for all tables reporting the welfare implications of cross-sectional differences in risk-sharing environments. Similarly to the time-varying welfare implications, I start with a base group. In this case, the base group is not common to all columns. The base group is the opposite of the column header. The parameters are then successively replaced by the parameters of the column header. Thus, the base group has the following environment: $\{\gamma^{No}, \sigma^{No}, \Phi^{No}\}$. Let us take the first group as an example: Participation. The base group is nonparticipating households, hence the No superscript describing the economic environment. As I successively replace the parameters, the superscript change to Yes, indicating the parameters used are those of participating households. I use the exogenous parameters[20] from the base group unless stated otherwise. In the previous section, the saving rate emphasized the growth effect but did not have a changing impact on the overall estimation. When comparing across households, the situation is different. From equation (2.3), a higher saving rate mechanically reduces the deterministic component of consumption. The question becomes: which saving rate should be used in equation (2.10). Santaeulalia-Llopis and Zheng (2018) use an average saving rate through time and space focusing solely on the income growth rate and discount heterogeneity in savings. There is, in my sample, large differences in saving rate and insurance across subgroups. Insofar as savings are used to smooth income shocks (Mazzocco, 2004), there is a trade-off between a lower saving rate (resulting in a higher growth effect) and higher insurance (resulting in a higher insurance effect). To reflect these assumptions, I use each groups' respective saving rate in calculating their deterministic consumption path, thus a higher income growth may not necessarily result in a positive growth effect if the saving is large[21].

In table 2.4, I consider the entire sample period with one transmission parameters[22] and annualized variances of shocks over the 36 years period. Similarly, the income growth rates are taken over the 36 years period. The table reports the welfare decomposition for three levels of risk aversion, 2, 4, and the PSID implied value. From table 2.1, entrepreneurs have a higher annual income growth rate than employed households (1.36% vs 1.15%). However, entrepreneurs' saving rate is much higher than that of employed households (11.84% vs 3.94%). Despite the higher rate of growth, the growth effect is negative indicating that the annual consumption equivalent of entrepreneurs is 4.90% lower than that of employees. Furthermore, entrepreneurs experience slightly higher permanent shocks (0.030 vs 0.021) and much higher transitory shocks (0.054 vs 0.026). Thus, the risk effect is significantly negative even at low levels of risk aversion. Employees also have the highest risk aversion coefficients across subgroups. Although small with $\eta = 2$ (-0.83%), the risk effect is -16.90% with $\eta = 12.74$. Despite higher income risk[23], entrepreneurs have lower permanent insurance and

[20] Preference shocks ξ and risk aversion coefficient η

[21] I also repeat the analysis by keeping the savings constant between base and counterfactual groups to highlight the income growth effect alone as in Santaeulalia-Llopis and Zheng (2018).

[22] one for permanent and transitory shocks

[23] Krueger and Perri (2006) argue that higher income risk increase the incentives to ensure so according

Table 2.4: Welfare Effects of Growth, Risk, and Insurance: Comparison Across Groups

This table shows the effects on household welfare in percentage annual consumption equivalent from a baseline level in which I replace income growth parameters, income risk parameters, and transmission parameters from one group with their counterpart group parameters. The base groups are non-entrepreneurs, non-participants, transfer receipts, and low wealth households. Groups are not mutually exclusive. The decomposition is over the entire 1980-2016 period. I use CRRA utility and risk aversion coefficient of 2, 4, PSID implied. The income growth parameters are estimated using PSID data. The growth, income, insurance effects are estimated using equations (2.10), (2.12), (2.13), and (2.15). In panel A, I use base group risk-aversion parameters. In panel B, I use counterfactual group parameters and assume similar savings rate across groups. 95% confidence intervals are reported below and are based on 1000 bootstrap replicas.

	Entrepreneurs?			Participation?			No Receipts?			High Wealth?		
	$\eta = 2$	$\eta = 4$	$\eta = 12.74$	$\eta = 2$	$\eta = 4$	$\eta = 11.88$	$\eta = 2$	$\eta = 4$	$\eta = 11.68$	$\eta = 2$	$\eta = 4$	$\eta = 11.47$
Panel A. 1980-2016 - Base Group Parameters												
Growth Effect $\{\gamma^{Yes}, \sigma^{No}, \Phi^{No}\}$	-4.90%	-4.90%	-4.90%	-5.06%	-5.06%	-5.06%	-1.17%	-1.17%	-1.17%	-11.32%	-11.32%	-11.32%
Risk Effect $\{\gamma^{Yes}, \sigma^{Yes}, \Phi^{No}\}$	-0.83%	-3.01%	-16.90%	0.41%	1.54%	10.52%	0.51%	1.90%	11.83%	0.61%	2.31%	13.99%
	[-1.31 -0.28]	[-4.78 -1.03]	[-26.51 -7.22]	[0.06 0.88]	[0.19 3.28]	[0.59 22.39]	[0.20 0.97]	[0.75 3.61]	[4.12 22.55]	[0.38 0.87]	[1.44 3.28]	[8.30 20.06]
Insurance Effect $\{\gamma^{Yes}, \sigma^{Yes}, \Phi^{Yes}\}$	1.41%	5.26%	36.68%	1.82%	6.70%	54.55%	0.85%	3.12%	20.35%	1.16%	4.31%	27.84%
	[0.49 2.33]	[1.69 8.69]	[5.98 61.77]	[0.96 2.54]	[3.48 9.34]	[20.04 78.98]	[-0.04 1.56]	[-0.16 5.75]	[-5.07 37.68]	[0.60 1.61]	[2.19 6.01]	[11.83 39.55]
Total Effect	-4.35%	-2.90%	8.02%	-2.94%	2.85%	62.17%	0.18%	3.85%	33.03%	-9.74%	-5.37%	29.23%
	[-4.84 -3.77]	[-4.71 -0.78]	[-4.88 21.76]	[-3.51 -2.36]	[0.66 5.06]	[37.42 83.95]	[-0.48 0.84]	[1.40 6.32]	[12.62 51.12]	[-10.20 -9.30]	[-7.08 -3.67]	[14.91 42.15]
Panel B. 1980-2016 - Counterfactual Group Parameters												
Growth Effect $\{\gamma^{Yes}, \sigma^{No}, \Phi^{No}\}$	3.62%	3.62%	3.62%	3.08%	3.08%	3.08%	1.87%	1.87%	1.87%	1.20%	1.20%	1.20%
Risk Effect $\{\gamma^{Yes}, \sigma^{Yes}, \Phi^{No}\}$	-0.83%	-2.97%	-15.34%	0.41%	1.51%	8.11%	0.51%	1.85%	11.64%	0.61%	2.17%	14.95%
	[-1.31 -0.28]	[-4.72 -0.98]	[-24.16 -6.40]	[0.06 0.88]	[0.18 3.23]	[0.60 17.28]	[0.20 0.96]	[0.73 3.52]	[4.01 22.18]	[0.38 0.86]	[1.33 3.09]	[8.73 21.45]
Insurance Effect $\{\gamma^{Yes}, \sigma^{Yes}, \Phi^{Yes}\}$	1.41%	5.19%	32.43%	1.81%	6.58%	40.00%	0.84%	3.04%	19.96%	1.15%	4.04%	29.67%
	[0.49 2.33]	[1.59 8.61]	[5.92 54.63]	[0.95 2.54]	[3.24 9.28]	[15.66 57.66]	[-0.05 1.55]	[-0.38 5.61]	[-5.65 36.97]	[0.59 1.60]	[1.96 5.65]	[11.91 42.26]
Total Effect	4.22%	5.77%	16.18%	5.38%	11.52%	56.01%	3.25%	6.91%	36.43%	2.98%	7.57%	50.83%
	[3.69 4.85]	[3.83 8.05]	[3.77 29.51]	[4.75 6.01]	[9.06 13.92]	[37.46 72.71]	[2.57 3.94]	[4.29 9.46]	[15.35 54.81]	[2.46 3.49]	[5.62 9.40]	[32.18 66.79]

insignificantly different from 0 transitory insurance (implying perfect insurance). This better insurance environment results in welfare gains of 1.41% to 36.68%. Recall from equation (2.15) that the total welfare effect is the product of the growth, risk, and insurance effects. Ignoring the growth effect, I can define the net effect as:

$$(1+\omega_{Net}) = (1+\omega_G)(1+\omega_R)(1+\omega_I) \quad with \quad \omega_G = 0 \tag{2.17}$$

By setting $\omega_G = 0$, I assume that the two groups have the same income growth rate and the same saving rate. Obviously unrealistic, it allows me to isolate the combined effect of risk and insurance. This is particularly interesting in the entrepreneurs vs. employees case as the two effects differ in direction and magnitude. The net effect is 0.57%, 8.31%, and 13.58% for risk aversions levels of 2, 4, and 11.88, respectively. Differences in risk and transmission environment do cause significant differences in household welfare. At all levels of risk aversion, the insurance effect overtakes the risk effect and generates welfare gains. However, the total effect is significantly negative with $\eta = 2$ or $\eta = 4$. The total effect is only positive with the high coefficient of risk aversion implied by the PSID questionnaire. Assuming that consumption smoothing is the only motive for households' saving, these results would imply that entrepreneurs save too much for the level of insurance they are getting. Setting entrepreneurs' saving rate to 7.75% results in both groups having the same consumption equivalent. Of course, this does not take into account the unrealized benefits of higher savings (i.e in deferred consumption or bequest).

Turning to participating households, the growth effect is negative (-5.06%) because of a high saving rate (despite a higher income growth rate). Participating households have a higher annual income growth rate (1.22% vs 1.04%), and a higher saving rate (9.70% vs 1.96%). Market participants have lower transitory income risk and similar permanent income risk. At low levels of risk aversion, the risk effect is 0.41% with the 95% confidence bound at 0.06% and 0.88% indicating the differences are negligible. However with increasing levels of risk aversion, the risk effects gets larger (1.54% with $\eta = 4$ and 10.52% with $\eta = 11.88$). The insurance effect is significant at low levels of risk aversion. Indeed, while both groups have similar levels of transitory shock transmission, participating households have a much lower permanent insurance coefficient. With $\eta = 2$, the insurance effect is 1.82% and 54.55% with $\eta = 11.88$. With low levels of risk aversion, the risk and insurance effects are not extremely relevant to households' welfare compared to the growth effect. The total effect is negative with $\eta = 2$ at -2.94% but is positive with $\eta = 4$ and $\eta = 11.88$ at 2.85% and 62.17% respectively. With moderate levels of risk aversion, the risk and insurance effect

to them, the correct grammatical term would be in spite of outweigh the growth effect. With the growth effect carrying unrealized welfare gains of participating households through higher saving rate, let us concentrate on the net effect of risk and insurance. The annual consumption equivalent of participating households is 2.24%, 8.34%, and 70.81% (with $\eta = 2; 4; 11.88$ respectively) higher than that of nonparticipating households over the 1980-2016 sample period[24].

The double negative in the third group implies that transfer recipients are the base group. Households that do and do not receive transfers have similar annual income growth rates over the entire sample period. The higher saving rate of nonrecipient causes a slightly negative growth effect. However, the effect is small enough to get balanced out by the risk and insurance effect. Indeed, the total effect is 0.18%, 3.85%, and 33.03%. Although the total effect is not significantly different from 0 with $\eta = 2$, higher levels of risk aversion generate significant welfare gains.

Comparing high and low wealth households, I find the net effect to be positive reflecting welfare gains from the lower risk and better insured environment. Wealthy households have an annual consumption equivalent that is 1.78%, 6.72%, and 45.72% higher than that of low wealth households that is directly attributable to their stochastic consumption path. However, relative to households below the median level of wealth, high wealth households have a deterministic consumption equivalent 11.32% lower. Although both groups have similar income growth rate, the difference in saving rate is large. This figure would imply that high wealth households have lower welfare. This results serves to highlight the limitations of the welfare framework. To circumvent this issue, I consider in panel B the welfare decomposition with exogenous parameters from the counterfactual group. That is the growth effect is solely a function of the growth rate[25]. It is noteworthy to mention that I substitute counterfactual preference shocks and risk aversion. As such, the risk and insurance effects are changed[26] for all levels of risk aversion.

The risk and insurance effects with $\eta = 2, 4$ remains relatively unchanged. Households that do or do not receive transfer income have relatively similar levels of risk aversion. The risk and insurance effects remain stable from panel A to panel B. However, when only considering the difference in income growth rate, nontransfer recipients' annual consumption equivalent is 1.87% higher than that of transfer recipients. The counterfactual groups all have higher income growth rates than their respective base groups resulting in positive growth effects. The total effects are smaller for entrepreneurs and participating households.

[24]A saving rate of 7.5% would equalize both groups' consumption equivalent.

[25]The saving rate is that of the counterfactual group. However, since it affects both groups similarly, the results remain unchanged should I use the saving rate of the base group.

[26]Though marginally as the cross-sectional standard deviation in preference shocks is small.

Indeed, these counterfactual groups are less risk averse than the base groups. Despite the higher growth effect, the risk and insurance effects are lower[27]. Overall, table 6 shows that entrepreneurs, market participants, wealthy households, and households that do not receive transfers have higher consumption equivalent relative to employed, nonparticipating, low wealth households, and those that do receive transfers as a result of the better combined risk and insurance effects. The growth effect is ambiguous. On one hand, the counterfactual groups have higher income growth, but they also have higher saving rate, thus resulting in a negative growth effect. Households can trade off consumption for higher insurance or later consumption. As such, the welfare measure does take into account future consumption resulting from higher saving past my sample period.

I then decompose welfare across groups and time by considering the economic environment of 1980-1996 and 1996-2016 separately. Although I don't expect the results to be different from table 2.4 (groups with higher consumption equivalent will still have higher consumption equivalents), the decomposition across time and groups may reveal that the economic environment did not change uniformly across households. In panels A and B of table 2.5, I present the welfare decomposition with 1980-1996 and 1996-2016 subperiods as the respective economic environment. These panels consider both income growth and saving rate in the growth effect.

The income effect is negative for entrepreneurs, not as result of a difference in growth rate but rather as a result of different saving rate. Indeed, the saving rate of employed households in the 80s and 90s is 4.88% compared to 15.64% for entrepreneurs. On the other hand, the better insurance environment is generating welfare gains for entrepreneurs. Entrepreneurs do experience high income risk than employed households, both permanent and transitory. Thus, I see a significantly negative risk effect. I find a similar pattern in panel B in the 1996-2016 sample with different magnitudes. To compare whether the change in environment affected both groups equally, I calculate the combined effect of risk and insurance. In the first subsample, the total effects for all three levels of risk aversion are 0.27%, 0.86%, and 5.1%; and 0.47%, 1.52%, and 9.34% in the second subsample. In either period, with low levels of risk aversion, entrepreneurs do not have a higher consumption equivalent. With high levels of risk aversion, entrepreneurs have annual consumption equivalent that are 5.1% and 9.34% higher than employees in the first and second subsample respectively. In panels C and D, I assume both groups that have the same saving rate to highlight the income growth effect. Both groups have a similar average annual growth rate in the first subsample (resulting in a growth effect of 0.47%) and entrepreneurs have a higher income

[27] Only when considering the PSID implied risk aversion.

Table 2.5: Welfare Effects of Growth, Risk, and Insurance: Comparison Across Groups - Pre and Post 1996

This table shows the effects on household welfare in percentage annual consumption equivalent from a baseline level in which I replace parameters from one group with their counterpart group parameters. The base groups are non-entrepreneurs, non-participants, transfer receipts, and low wealth households. In panels A and B, only the 1980-1996 period is considered; in panels C and D, only the 1996-2016 sub-period is considered. 95% confidence intervals are reported below and are based on 1000 bootstrap replicas.

	Entrepreneurs?			Participation?			No Receipts?			High Wealth?		
	$\eta=2$	$\eta=4$	$\eta=12.74$	$\eta=2$	$\eta=4$	$\eta=11.88$	$\eta=2$	$\eta=4$	$\eta=11.68$	$\eta=2$	$\eta=4$	$\eta=11.47$
Panel A. 1980-1996 - Base Group Parameters												
Growth Effect $\{\gamma^{Yes},\sigma^{No},\Phi^{No}\}$	-10.90%	-10.90%	-10.90%	-6.21%	-6.21%	-6.21%	-10.32%	-10.32%	-10.32%	-13.45%	-13.45%	-13.45%
Risk Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{No}\}$	-0.50% [-0.87 -0.03]	-1.64% [-2.84 -0.11]	-9.31% [-16.00 -1.28]	0.54% [0.28 0.83]	1.79% [0.90 2.77]	12.18% [5.72 18.97]	0.33% [0.14 0.58]	1.09% [0.46 1.91]	6.53% [2.58 11.44]	0.42% [0.25 0.60]	1.41% [0.82 2.01]	8.20% [4.61 11.74]
Insurance Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{Yes}\}$	0.77% [-0.00 1.42]	2.54% [-0.07 4.72]	15.89% [-2.57 29.68]	0.45% [-0.16 0.97]	1.46% [-0.52 3.18]	9.80% [-4.59 21.27]	-0.06% [-0.65 0.39]	-0.20% [-2.13 1.27]	-1.18% [-12.75 6.99]	0.41% [-0.00 0.76]	1.36% [-0.01 2.50]	7.88% [-0.48 14.47]
Total Effect	-10.67% [-11.02 -10.19]	-10.14% [-11.31 -8.59]	-6.35% [-13.69 2.76]	-5.29% [-5.73 -4.79]	-3.14% [-4.61 -1.51]	15.52% [3.76 27.52]	-10.08% [-10.55 -9.66]	-9.53% [-11.11 -8.14]	-5.59% [-15.55 2.49]	-12.72% [-13.05 -12.39]	-11.03% [-12.14 -9.92]	1.03% [-6.22 7.90]
Panel B. 1996-2016 - Base Group Parameters												
Growth Effect $\{\gamma^{Yes},\sigma^{No},\Phi^{No}\}$	-4.30%	-4.30%	-4.30%	-11.50%	-11.50%	-11.50%	5.72%	5.72%	5.72%	-10.98%	-10.98%	-10.98%
Risk Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{No}\}$	-0.46% [-0.76 -0.14]	-1.50% [-2.49 -0.45]	-8.59% [-14.19 -2.93]	-0.26% [-0.55 0.23]	-0.86% [-1.81 0.76]	-5.47% [-11.57 4.14]	0.20% [0.01 0.51]	0.66% [0.03 1.71]	3.93% [0.05 10.15]	0.13% [-0.01 0.30]	0.42% [-0.03 0.99]	2.41% [-0.24 5.68]
Insurance Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{Yes}\}$	0.93% [0.47 1.37]	3.07% [1.54 4.57]	19.62% [8.73 29.59]	1.16% [0.48 1.69]	3.86% [1.51 5.63]	28.02% [7.95 41.59]	0.66% [0.13 1.05]	2.19% [0.37 3.46]	13.51% [1.13 21.55]	0.77% [0.46 1.03]	2.56% [1.52 3.44]	15.40% [8.62 20.86]
Total Effect	-3.86% [-4.04 -3.58]	-2.85% [-3.43 -1.94]	4.64% [0.91 10.28]	-10.71% [-11.00 -10.43]	-8.87% [-9.84 -7.95]	7.10% [-0.42 13.97]	6.63% [6.24 7.00]	8.74% [7.40 9.97]	24.71% [15.48 32.76]	-10.18% [-10.41 -9.97]	-8.32% [-9.08 -7.59]	5.21% [0.13 9.83]
Panel C. 1980-1996 - Counterfactual Group Parameters												
Growth Effect $\{\gamma^{Yes},\sigma^{No},\Phi^{No}\}$	0.47%	0.47%	0.47%	9.06%	9.06%	9.06%	-7.45%	-7.45%	-7.45%	-0.39%	-0.39%	-0.39%
Risk Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{No}\}$	-0.50% [-0.87 -0.03]	-1.64% [-2.84 -0.11]	-8.35% [-14.36 -1.04]	0.54% [0.28 0.83]	1.79% [0.90 2.77]	9.12% [4.43 14.16]	0.33% [0.14 0.58]	1.09% [0.46 1.91]	6.47% [2.56 11.33]	0.42% [0.25 0.60]	1.41% [0.82 2.01]	9.00% [5.04 12.90]
Insurance Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{Yes}\}$	0.77% [-0.00 1.42]	2.54% [-0.07 4.72]	14.04% [-2.05 26.21]	0.45% [-0.16 0.97]	1.46% [-0.52 3.18]	7.29% [-3.02 15.83]	-0.06% [-0.65 0.39]	-0.20% [-2.13 1.27]	-1.17% [-12.63 6.94]	0.41% [-0.00 0.76]	1.36% [-0.01 2.50]	8.67% [-0.63 15.94]
Total Effect	0.73% [0.33 1.26]	1.32% [0.01 3.07]	5.01% [-2.26 14.16]	10.13% [9.62 10.70]	12.63% [10.92 14.52]	27.67% [17.96 37.81]	-7.20% [-7.69 -6.77]	-6.63% [-8.26 -5.20]	-2.62% [-12.79 5.65]	0.44% [0.06 0.83]	2.39% [1.11 3.67]	18.00% [8.66 26.75]
Panel D. 1996-2016 - Counterfactual Group Parameters												
Growth Effect $\{\gamma^{Yes},\sigma^{No},\Phi^{No}\}$	2.89%	2.89%	2.89%	-4.87%	-4.87%	-4.87%	8.72%	8.72%	8.72%	1.29%	1.29%	1.29%
Risk Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{No}\}$	-0.46% [-0.76 -0.14]	-1.50% [-2.49 -0.45]	-7.69% [-12.73 -2.57]	-0.26% [-0.55 0.23]	-0.86% [-1.81 0.76]	-4.19% [-8.86 3.35]	0.20% [0.01 0.51]	0.66% [0.03 1.71]	3.89% [0.06 10.06]	0.13% [-0.01 0.30]	0.42% [-0.03 0.99]	2.64% [-0.27 6.22]
Insurance Effect $\{\gamma^{Yes},\sigma^{Yes},\Phi^{Yes}\}$	0.93% [0.47 1.37]	3.07% [1.54 4.57]	17.28% [7.79 26.04]	1.16% [0.48 1.69]	3.86% [1.51 5.63]	20.51% [6.27 30.32]	0.66% [0.13 1.05]	2.19% [0.37 3.46]	13.38% [1.12 21.34]	0.77% [0.46 1.03]	2.56% [1.52 3.44]	16.99% [9.42 23.06]
Total Effect	3.36% [3.17 3.66]	4.45% [3.83 5.43]	11.38% [7.88 16.73]	-4.02% [-4.33 -3.72]	-2.05% [-3.09 -1.05]	9.84% [4.16 15.17]	9.66% [9.25 10.04]	11.83% [10.46 13.10]	28.07% [18.68 36.26]	2.20% [1.94 2.44]	4.32% [3.45 5.15]	21.64% [15.21 27.48]

growth rate in the second subsample (resulting in a growth effect of 2.89%). Overall, it appears the welfare benefits of entrepreneurship are concentrated in the second subsample.

Similarly to entrepreneurs, the growth effect is negative for participating households in the first 2 panels of table 2.5 despite a higher income growth rate. In panel C, the effect is positive as market participants have a higher growth rate than nonparticipating households. The effect is negative in panel D as market participants have now a lower rate of income growth. The slow down in income growth is affecting participating households more severely, generating welfare losses of 4.87% compared to the welfare gains of 9.06% in the first subsample. Prior to 1996, participating households have a favorable risk and insurance environment resulting in combined welfare gains of 0.99%, 3.28%, and 23.17% relative to nonparticipating households. However, post 1996, the income risk of participating households increases generating significant welfare losses. The insurance environment does ameliorates generating net gains of 0.9%, 2.97%, and 21.07%. Market participants compensate the welfare losses of higher risk with better insurance. Although the net effect of risk and insurance is stable across the welfare advantage of participating households has changed in nature.

2.6 Subjective Well-Being

2.6.1 Data Description

In the previous section, I show that the changes in economic environments, as it relates to the size and transmission of idiosyncratic shocks, generates significant differences in consumption equivalents. In turn, the differences in consumption equivalent are, within the framework, interpreted as differences in household utility. In this section, I investigate the impact of households' idiosyncratic shocks on subjective well-being (SWB). Starting in 2009, households (reference people) are asked to categorize their life satisfaction on a scale from 1 to 5, 1 being "completely satisfied", 5 being "not at all satisfied"[28]. This measure is used by Brown and Gathergood (2017) and a similar measure is used by Bayer and Juessen (2015) in the German annual socio-economic panel (SOEP).

The objective is to quantify the impact of income and consumption idiosyncratic shocks on households' life satisfaction. Shocks to consumption are constructed following Duvernois (2021). Consider the following process for life satisfaction:

$$h_{i,t} = h^*_{i,t} + f(z_{i,t}) \tag{2.18}$$

[28]The question asked is :"Please think about your life as a whole. How satisfied are you with it? Are you completely satisfied, very satisfied, somewhat satisfied, not very satisfied, or not at all satisfied?"

$z_{i,t}$ is the set of observable characteristics used in the income and consumption shock regressions. h^* can be further written as:

$$h^*_{i,t} = h^{**}(\psi_{i,t}, \phi_{i,t}) + \mu_i + \zeta_{i,t} \tag{2.19}$$

μ is a household fixed-effect and ζ is the independently and identically distributed influences on survey responses. These two terms are removed by first-differencing h^* into Δh^*. I follow Bayer and Juessen (2015) to construct h^*. I first estimate equation (2.18) using an order-probit regression of the following form:

$$h_{i,t} = j \quad if \quad v^{**}_{i,t} \in (\bar{c}_j, \bar{c}_{j+1}] \tag{2.20}$$

Bayer and Juessen (2015) infer an interval $V_{i,t} = (\tilde{\bar{c}}_{h_{i,t}} - \widetilde{f(z_{i,t})}, \tilde{\bar{c}}_{h_{i,t}+1} - \widetilde{f(z_{i,t})}]$ where h^* exists. Assuming normality, the conditional expected value $\bar{v^*_{i,t}}$ is:

$$\bar{h^*_{i,t}} = \frac{\int_{h \in H_{i,t}} h\phi(h)}{\Phi(H_{i,t})}, \tag{2.21}$$

where ϕ and $\Phi(H)$ are, respectively, the density and the probability of H. By assuming a standard normal distribution, h^* and Δh^{**} can be used in OLS without bias to the estimator. In table 2.6, I present the summary statistics for the sample used. I use the PSID 2007 waves onward[29]. The last available wave is 2019 and is used for instruments' creation. Disposable income is the sum of labor, business, financial income for all households' members net of federal taxes. Total consumption is the sum of all expenditure categories (including the categories added after 2005) including imputed rent expenditure[30]. I decompose consumption in two categories: conspicuous and inconspicuous consumption following Brown and Gathergood (2017). Conspicuous consumption is visible by individuals outside of households (food away from home, clothing, holidays,...). Both income and consumption are deflated to 1982 dollars.

From table 2.6, 94% of the sample is employed. A household is considered employed if the reference person (or spouse if present) report more than 13 weeks of works. The average life satisfaction is 2.11, suggesting households in the PSID are generally satisfied with their lives. The average respondent is 45 years old, male (88.2% of respondent are male), white (88.2% of respondent are white), and college-educated (70% of respondent are college-educated). The average household size is 2.849 with 0.886 child.

[29] Although life satisfaction is measured starting 2009, 2007 is retained for conditioning purposes.

[30] Rent expenditure is imputed as 6% of house value for homeowners (See Attanasio and Pistaferri (2014) and Flavin and Yamashita (2002)).

Table 2.6: Summary Statistics - Life Satisfaction Sample

This tables presents summary statistics for the PSID sample. Data are from the 2009 to 2019 waves. Disposable income is the sum of labor, business, financial income for all household members minus federal taxes (estimated using TAXSIM). Total consumption is the sum of all expenditures categories provided by the PSID plus rent imputation (as Flavin and Yamashita (2002) for homeowners.). Conspicuous and inconspicuous consumption are calculated following Brown and Gathergood (2017). Employment is a dummy variable that takes the value of 1 of the reference person reports more than 13 weeks of work. K6-Index refers to the Kessler index modified to be bounded between 0 and 10. Gender is a dummy variable that takes the value of 1 for Male reference person. Race is a dummy variable that takes the value of 1 if the reference person is white. College education is a dummy variable that takes the value of 1 if the reference person has completed some college or higher.

	Obs.	Mean	St. Dev.	25^{th} Pctl.	Median	75^{th} Pctl.
Disposable Income	11,221	33,060.730	21,343.290	20,670.670	30,821.470	41,801.450
Total Consumption	11,221	21,820.470	14,641.280	13,056.460	19,189.980	26,619.160
Conspicuous Consumption	11,221	5,626.563	5,894.363	2,445.782	4,063.959	6,689.771
Inconspicuous Consumption	11,221	16,193.910	10,467.250	9,873.916	14,498.170	19,941.630
Employment	11,221	0.933	0.250	1	1	1
Life Satisfaction	11,221	3.889	0.770	3	4	4
K6-Index	11,221	2.335	1.748	1	2	3
Age of Reference Person	11,221	45.479	10.261	36	45	55
Household Size	11,221	2.849	1.439	2	3	4
Number of Children	11,221	0.886	1.184	0	0	2
Gender of Reference Person	11,221	0.842	0.365	1	1	1
Race of Reference Person	11,221	0.882	0.323	1	1	1
College Education	11,221	0.692	0.462	0	1	1

In figure (2.2), I plot life satisfaction against income, expenditure[31], or wealth. Although it shows that neither of the measures alone explain life satisfaction, it reveals interesting patterns. Higher levels of consumption, income, and wealth are associated with higher levels of satisfaction. A fuller bar indicates there are more households in the category of life satisfaction. In all three metric, higher levels of satisfactions have fuller bars at higher levels (of income, consumption, or wealth). However, there are very satisfied low-income households. This indicates the need to look at shocks to get meaningful determinants of life satisfaction.

Figure (2.2) plots all households in the sample. I average life satisfaction and consumption for all households and all years at the state level and plot the relationship in figure (2.3). The expected relationship is higher consumption means higher satisfaction.

[31]total household expenditure is divided by the OECD adult equivalent scale used by Attanasio and Pistaferri (2014).

Figure 2.2: Life satisfaction as a function of income, consumption, and wealth

The map's color should move from the bottom left quadrant (mostly green colors) to the top right quadrant (purple colors). Redish and blueish colors should be less prevalent. Green, dark green, and brownish colors are the most prevalent colors in the map, suggesting an average upward relationship between consumption and life satisfaction. Again, these are levels are not shocks. Thus any causation may not be inferred.

2.6.2 Results

I first investigate the impact of shocks on life satisfaction. I regress income shocks on shocks to household satisfaction. β is the coefficient on the independent variable Δy, the sum of a permanent and transitory component as described in equation (2.1). Shocks to happiness h^{**} can also be understood to be the sum of permanent and transitory shocks as the notation in equation (2.19) suggests. Transitory shocks are more easily smoothed (Blundell, Pistaferri, and Preston, 2008; Duvernois, 2021) and should have less of an impact on life satisfaction than permanent shocks. These shocks are not directly observable. However, Blundell, Pistaferri, and Preston (2008) propose an instrumental variable regression

Figure 2.3: Life satisfaction and consumption - state level averages

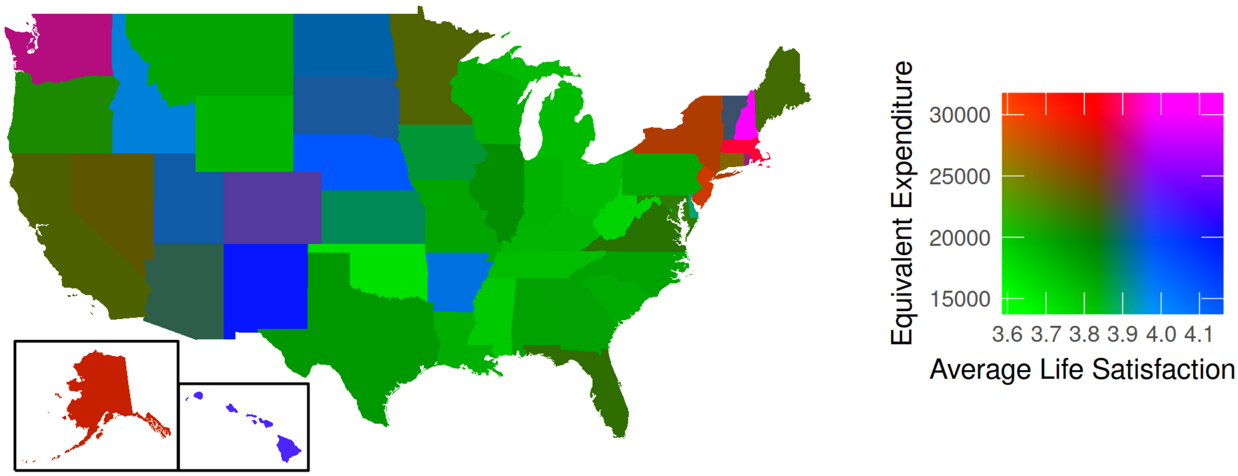

interpretation of their minimum distance estimator. Transitory and permanent shocks can be identified with the following moment restrictions:

$$E[(\Delta h^*_{i,t} - \Delta y^*_{i,t})\Delta y^*_{t+1}] = 0 \tag{2.22}$$

$$E[(\Delta h^*_{i,t} - \Delta y^*_{i,t})(y^*_{i,t+1} + y^*_{i,t} + y^*_{i,t-1})] = 0 \tag{2.23}$$

Instruments for transitory shocks are therefore the income shocks at $t+1$; instruments for permanent shocks are the sum of the residuals at time $t, t+1$, and $t-1$. The estimation results are presented in table 2.7. There is a significant degree of transmission of income shocks to household life satisfaction. The positive coefficient of 0.142 is in line with results by Bayer and Juessen (2015) using German longitudinal survey.

In columns 3 and 4, the results show that permanent shocks to income have an impact on household satisfaction while transitory shocks do not. Indeed, the coefficient on the projection of transitory shocks is not statistically significant. Furthermore, not only is the coefficient on permanent shocks significantly different, it is 3 times larger than the

Table 2.7: Life Satisfaction and Income Shocks

*This tables reports the effect of income shocks (permanent and transitory) on life satisfaction shocks. IV regressions are estimated using equation (2.23) and equation (2.22) for permanent and transitory shocks. Life satisfaction, income, and employment have been regressed on the same set of control variables. Income and employment shocks are first-differenced residuals from OLS regressions. Life satisfaction shocks are estimated using an ordered-probit regression. Standard errors are reported in parentheses below the coefficients. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	OLS	OLS	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Income Shock:						
All: β	0.142***	0.130***				
	(0.035)	(0.036)				
Permanent: ψ			0.527**	0.589**		
			(0.213)	(0.256)		
Transitory: ϕ					0.163	0.338
					(0.104)	(0.301)
Employment Shock: Δe		0.070		-0.596		-0.709
		(0.058)		(1.233)		(1.168)
Constant: α	-0.003	-0.003	-0.006	-0.007	-0.003	-0.005
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Observations	8,916	8,916	8,916	8,916	8,916	8,916

average shock. In both IV regressions, weak-instruments F-tests are rejected at the 1% level suggesting the instruments are appropriate. Furthermore, I fail to reject Wu-Hausman test at reasonable significance level, suggesting IV regressions are consistent with OLS.

Bayer and Juessen (2015) argue that as permanent and transitory shocks are latent variable, the OLS regressions used are inherently biases caused by omitted variables. More specifically, the literature points to the effect of employment on satisfaction (Grün, Hauser, and Rhein, 2010; L. Winkelmann and R. Winkelmann, 1998). I therefore augment each specification with a shock to employment variable Δe. Δe is constructed as the first-differenced residual of an employment dummy variable on household characteristics. The employment variable is used in OLS specification. In IV regression, the effect is instrumented by the lagged employment shocks. The results of table 2.7, columns 2, 4, and 6 show that the estimation is robust to the omitted variable bias.

As suggested by the permanent income hypothesis, transitory income shocks do not have an impact on latent happiness. However, this is the case because transitory shocks are better insured than permanent shocks in that the transmission to consumption shock is lower. What of consumption shocks? Brown and Gathergood (2017) shows that while income does have an impact on life satisfaction, the impact of consumption is stronger. Furthermore, they do find evidence of satiation in consumption (unlike income). In table 2.8, I repeat a similar analysis as table 2.7, in that I regress consumption shocks on happiness shocks. In unreported results, I regress total, conspicuous, and inconspicuous consumption shocks on happiness shocks. The coefficients are respectively 0.19, 0.08, and 0.16 and are all significant at the 1% level. The point estimates are very similar to the point estimate in table 2.7. This is surprising considering the significant amount of consumption insurance. However, consider equation (2.4) describing consumption shocks. $\zeta_{i,t}$ reflect consumption shocks unrelated to income shocks. Despite limited transmission of income shocks, the coefficients on $\Delta c^*_{i,t}$ would incorporate preference shocks. In the previous section, the welfare difference in preferences are incorporated in the different effects as the preference parameter is estimated with the minimum distance estimator. It is however not observable in the data. To net these effects, I first regress income shocks on consumption shocks and collect the fitted values. This projection is the proportion of income shocks that translate to consumption shocks and is then used as independent variable against satisfaction shocks.

I use total income shocks as instrument in the first two columns of table 2.8. In panel A, I use total consumption shocks; in panel B, I use conspicuous consumption; in panel C, I use inconspicuous consumption shocks. The consumption shocks have a stronger effect on life satisfaction than income shocks, even controlling for employment effects[32]. Indeed, the point estimates are 0.571, 0.459. and 0.581 for total, conspicuous, and inconspicuous shocks respectively. Brown and Gathergood (2017) suggest that conspicuous consumption affects life satisfaction more than non-visible consumption. This does not seem to be the case here as the point estimates are comparable. If anything, the coefficient is lower for visible consumption suggesting that these shocks affect households to a lesser degree.

To capture permanent and transitory shocks to consumption, I use the same of instruments as in table 2.7. I project permanent or transitory income shocks onto consumption and use the fitted values from these IV regressions in columns (3)-(6)[33]. Permanent consumption

[32]In include employment shocks in column 2

[33]Note that odd-numbered columns include employment shocks in the instrument set.

Table 2.8: Life Satisfaction and Consumption Shocks

*This tables reports the effect of consumption shocks (permanent and transitory) on life satisfaction shocks. Income shocks are first regressed on consumption shocks. Fitted values are then used in a regression with life satisfaction shocks as dependant variables. Only the second stage is reported in the table. Panel A reports the results with total consumption shocks as the main independent variable. Panel B reports the results with conspicuous consumption shocks as the main independent variable. Panel C reports the results with inconspicuous consumption shocks as the main independent variable. Standard errors are reported in parentheses below the coefficients. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	IV	IV	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Total Consumption						
All: β	0.571***	0.520***				
	(0.138)	(0.145)				
Permanent: ψ			0.151**	0.092**		
			(0.037)	(0.033)		
Transitory: ϕ					0.997***	0.001
					(0.242)	(0.045)
Constant: α	−0.003	−0.003	−0.003	−0.002	−0.004	−0.002
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Panel B. Conspicuous Consumption						
All: β	0.459***	0.417***				
	(0.110)	(0.115)				
Permanent: ψ			0.128**	0.034		
			(0.031)	(0.022)		
Transitory: ϕ					0.724***	-0.005
					(0.173)	(0.024)
Constant: α	−0.004	−0.004	−0.003	−0.002	−0.005	−0.002
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Panel C. Inconspicuous Consumption						
All: β	0.581***	0.530***				
	(0.141)	(0.147)				
Permanent: ψ			0.158***	0.113**		
			(0.038)	(0.036)		
Transitory: ϕ					1.100***	0.006
					(0.267)	(0.055)
Constant: α	−0.003	−0.003	−0.003	−0.003	−0.004	−0.002
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Employment Shock Instrument	No	Yes	No	Yes	No	Yes
Observations	8,916	8,916	8,916	8,916	8,916	8,916

shocks are consistently significant and robust to the inclusion of employment shocks (with the exception of panel B). While transitory consumption shocks have a strong significant impact on happiness, the statistical effect disappears when controlling for employment shocks. This suggests that consumption is sensitive to employment shocks. Furthermore, it suggests that transitory shocks to consumption not caused by changes in employment status do not have an impact on households' life satisfaction. In panel B, I find that permanent shocks to conspicuous consumption do not affect life satisfaction. This result is surprising in the light of Brown and Gathergood (2017)'s result. As social comparison of consumption seems to be a significant driver of life satisfaction, permanent shocks should affect households' happiness which is not the case here. However, as visible consumption represents a smaller part of a household's budget as shown in table 2.6, shocks are likely to be small, thus not registering a significant impact on happiness.

As an alternative definition of subjective well being, I consider the K6 Index of Non-specific Psychological Distress developed by Kessler et al. (2002). The K6 is not a measure of subjective well being but rather measures negative emotions. It has been included in the PSID since 2001[34]. PSID respondent are asked to answer 6 with "All of the time" (4 points), "Most of the time" (3 points), "Some of the time" (2 points), "A little of the time" (1 point), or "None of the time" (0 points). The points are totaled for each household. The K6 is therefore measured between 0 and 24 with a higher score indicating severe psychological distress. The K6 index is an imperfect measure of subjective well-being as it only includes negative emotions. Furthermore, as the question asks respondent to consider the past 30 days, it makes it complicated to establish an empirical link with income shocks. However, Kahneman and Deaton (2010) show that income correlates similarly to positive and negative emotions. Clingingsmith (2016) finds that higher income reduces the experience of negative emotions. Furthermore, he finds that the income effect on serious mental illness is more pronounced in low-income households.

In table 2.6, I report summary statistics for the K6 Index. Note that although the K6 Index is measured from 0 to 24, I group responses into 11 tranches from 0 to 11[35]. This is done to simplify pseudo-residuals estimations but limited information is lost as few households have K6 scores above 13. It is also to be noted that a low life satisfaction score indicates high well-being and low K6-score indicates high emotional well-being. Both scales are inverted to reflect the direction of shocks (i.e a positive shock to income and consumption suggests higher than predicted income or consumption). The correlation between the two measures is 0.39 (consistent with Kahneman and Deaton, 2010 and Clingingsmith, 2016).

[34] To match with the life satisfaction customer, I do not include waves prior to 2009.

[35] (0-1):0;(2-3):1;(4-5):2;(6-7):3;(8-9):4;(10-11):5;(12-13):6;(14-15):7;(16-17):8;(18-20):9;(21-24):10

Although similar, it is not unsurprising to have different information encapsulated in these measures. As suggested by Kimball and Willis (2006), any good happiness measure should reflect both positive and negative emotions.

In table 2.9, I repeat the same empirical framework as in table 2.7 and table 2.8. In panel A, I consider the impact of income shocks on negative emotion shocks. Average income shocks have a significant impact on negative emotions. This result is robust to the inclusion of employment shocks. In panel B, the overall shocks on consumption also have a significant impact on negative emotions. The coefficients (0.112 and 0.446) are similar to the coefficients found in table 2.7 and table 2.8 (0.130 and 0.520).

The results on transitory and permanent income shocks are also similar. Transitory shocks are found to have no statistically significant impact on negative emotions, while permanent shocks have a significant impact in column (4). It is to be noted that the significance of ψ is revealed with the inclusion of employment shocks casting doubt on the effect. Similarly, in panel B, both permanent and transitory consumption shocks are significant with the significance fading with the inclusion of employment shocks instruments.

It is important to note the differences between life satisfaction and K6-index. The life satisfaction question asks households to gauge their overall satisfaction. While it may be biased towards how a household feels during the interview, it does allow the household to consider shocks beyond a certain time frame. The K6 Index frames the emotions around the past 30 days at the time of the interview. As shocks are measured annually, the timing of these shocks may not necessarily coincide with negative emotions, thereby reducing the effect. Furthermore, the effect of ξ in equation (2.19) is likely to be stronger in the K6-index. That is, day, month of the interview, current weather, interviewees' mood are more likely to impact the measurement of the K6-index than in the life satisfaction question. Nonetheless, income and consumption shocks have a significant impact on life satisfaction and negative emotions. Furthermore, permanent shocks have a stronger impact compared to transitory shocks. These results complement the welfare framework in showing the importance of consumption insurance.

2.7 Conclusion

The lack of complete markets has significant welfare implications. Floden and Linde (2001), Buera and Shin (2011), or Chetty and Looney (2006) use different frameworks and show that even-though consumption is relatively smooth, idiosyncratic wage risk can have significant welfare costs. Heathcote, Storesletten, and Violante (2008) show that insuring wage risk yield welfare improvements that are higher compared to the elimination of wage

Table 2.9: K6 Index, Income and Consumption Shocks

*This tables reports the effect of income and consumption shocks (permanent and transitory) on negative emotion shocks shocks. Negative emotions shocks are estimated from the K6 Index using an ordered-probit regression. Panel A reports the results with income shocks as the main independent variable. Panel B reports the results with consumption shocks as the main independent variable. Standard errors are reported in parentheses below the coefficients. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Income Shocks						
All: β	0.141***	0.112***				
	(0.029)	(0.030)				
Permanent: ψ			0.281	0.472*		
			(0.174)	(0.272)		
Transitory: ϕ					0.125	0.458
					(0.085)	(0.326)
Employment Shock: Δe		0.165***		−1.421		−1.425
		(0.049)		(1.386)		(1.357)
Constant: α	−0.003	−0.003	−0.003	−0.002	−0.004	−0.002
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Panel B. Consumption Shocks						
All: β	0.560***	0.446***				
	(0.110)	(0.115)				
Permanent: ψ			0.155***	0.034		
			(0.031)	(0.022)		
Transitory: ϕ					0.862***	-0.060*
					(0.174)	(0.032)
Constant: α	−0.002	−0.002	−0.001	−0.001	−0.003	−0.001
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Employment Shock Instrument	No	Yes	No	Yes	No	Yes
Observations	11,221	11,221	11,221	11,221	11,221	11,221

risk. In this paper, I investigate the welfare consequences of changes in the risk-sharing environment of US households. Indeed, Duvernois (2021) finds that households' ability to insure against permanent shocks slightly increased since the 1980s while transitory insurance has deteriorated. Households experience stronger transitory shocks since the turn of the 21^{st} century while the variance of permanent shocks has fallen over the same period.

I use the minimum distance parameters from Duvernois (2021) and the welfare framework from Santaeulalia-Llopis and Zheng (2018). I use PSID data from 1980 to 2016 to calculate income growth rate and saving rate. I also use the 1996 questionnaire on attitude towards risk and methodology by Hryshko, Luengo-Prado, and Sørensen (2011) to estimate households' risk aversion level. I average the risk aversion measure per group and assume it remains constant over my sample period. The welfare framework characterizes household utility as a function of a deterministic income growth rate, stochastic income risk parameters (variance of permanent and transitory shocks), and stochastic transmission parameters (partial insurance parameters as defined by Blundell, Pistaferri, and Preston, 2008). The welfare measure represents the percentage change in annual consumption equivalent caused by a household theoretically stitching from one environment (baseline environment) to another environment (counterfactual environment) characterized by different parameters.

I first estimate the framework over time and split the sample period into two blocks: 1980-1996 and 1996-2016. The average annual income growth rate is lower in the second subsample which has significant welfare costs across all households. However, the lower risk and better risk-sharing environment of the 21^{st} century generates welfare gains. However, the risk and insurance effects outweigh the growth effect only with high levels of risk aversion. Transfer recipients and low-wealth households do have negative welfare gains as their risk-sharing abilities deteriorate between the two sample periods. Comparing across groups, the results are unequivocal. Entrepreneurs, market participants, high wealth, and nontransfer recipients have significantly higher annual consumption equivalents than their respective counterparts as a result of their economic environments. This is true even for entrepreneurs who experience larger risk effects. The counterfactual groups have negative growth effects but also much higher saving rate suggesting unrealized insurance effects, which would further increase counterfactual groups' welfare gains.

This welfare framework is without doubt weakened by the Lucas critique as pointed out by Santaeulalia-Llopis and Zheng (2018). To strengthen my argument that the transmission of income shocks is a significant determinant of household welfare, I make use of the recently added life satisfaction question. Although Kimball and Willis (2006) argue that utility and happiness are not perfect proxy of each other, Bayer and Juessen (2015) show that income shocks have an impact on life satisfaction. I use the same empirical strategy as Bayer and Juessen (2015) with IV regression to investigate the impact of income shocks on life satisfaction shocks. I find that permanent shocks have a statistically significant impact on life satisfaction while transitory shocks do not. These results are robust to employment shocks. Furthermore, to the extent that consumption is a stronger predictor of life satisfaction (Brown and Gathergood, 2017), I replace income shocks by consumption in the IV regressions. Since this study aims to measure the impact of risk-sharing on household welfare, I need to identify the proportion of consumption shocks caused by income shocks. I project income shocks on consumption shocks using the same IV regressions and use the fitted values as permanent and transitory consumption shocks. Similarly to income shocks, transitory shocks do not have a significant impact on life satisfaction while permanent shocks do. Furthermore, the effects of consumption shocks on life satisfaction are larger than the effects of income.

This paper provides evidence that the size and transmission of income shocks are significant determinants of household welfare and life satisfaction. The relatively short time span in which life satisfaction data is available in the PSID prevents me from investigating how shocks affect different groups of households. This question is left for further research to be investigated with richer datasets.

References

Aiyagari, S Rao and Ellen R McGrattan (1998). “The optimum quantity of debt”. In: *Journal of Monetary Economics* 42.3, pp. 447–469.

Arellano, Manuel, Richard Blundell, and Stephane Bonhomme (2017). “Earnings and consumption dynamics: a nonlinear panel data framework”. In: *Econometrica* 85.3, pp. 693–734.

Attanasio, Orazio and Luigi Pistaferri (2014). “Consumption inequality over the last half century: some evidence using the new PSID consumption measure”. In: *American Economic Review* 104.5, pp. 122–26.

Barsky, Robert B et al. (1997). “Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study”. In: *The Quarterly Journal of Economics* 112.2, pp. 537–579.

Bayer, Christian and Falko Juessen (2015). “Happiness and the persistence of income shocks”. In: *American Economic Journal: Macroeconomics* 7.4, pp. 160–87.

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). “Consumption inequality and partial insurance”. In: *American Economic Review* 98.5, pp. 1887–1921.

Brown, Gordon DA and John Gathergood (2017). “Consumption and life satisfaction: A micro panel data study”. In: *Available at SSRN 2962837*.

Buera, Francisco J and Yongseok Shin (2011). “Self-insurance vs. self-financing: A welfare analysis of the persistence of shocks”. In: *Journal of Economic Theory* 146.3, pp. 845–862.

Caporale, Guglielmo Maria et al. (2009). “Income and happiness across Europe: Do reference values matter?” In: *Journal of Economic Psychology* 30.1, pp. 42–51.

Casado, Jose Maria (2011). “From income to consumption: measuring households partial insurance”. In: *Empirical Economics* 40.2, pp. 471–495.

Catherine, Sylvain, Paolo Sodini, and Yapei Zhang (2020). “Countercyclical Income Risk and Portfolio Choices: Evidence from Sweden”. In: *Available at SSRN*.

Chetty, Raj and Adam Looney (2006). “Consumption smoothing and the welfare consequences of social insurance in developing economies”. In: *Journal of public economics* 90.12, pp. 2351–2356.

Clark, Andrew, Yannis Georgellis, and Peter Sanfey (2001). "Scarring: The psychological impact of past unemployment". In: *Economica* 68.270, pp. 221–241.

Clark, Andrew E (2003). "Unemployment as a social norm: Psychological evidence from panel data". In: *Journal of labor economics* 21.2, pp. 323–351.

— (2018). "Four decades of the economics of happiness: Where next?" In: *Review of Income and Wealth* 64.2, pp. 245–269.

Clark, Andrew E and Andrew J Oswald (1994). "Unhappiness and unemployment". In: *The Economic Journal* 104.424, pp. 648–659.

— (1996). "Satisfaction and comparison income". In: *Journal of public economics* 61.3, pp. 359–381.

Clingingsmith, David (2016). "Negative emotions, income, and welfare: Causal estimates from the PSID". In: *Journal of Economic Behavior & Organization* 130, pp. 1–19.

Cochrane, John H (1991). "A simple test of consumption insurance". In: *Journal of political economy* 99.5, pp. 957–976.

— (2009). *Asset pricing: Revised edition*. Princeton university press.

Commault, Jeanne (2020). "Does Consumption Respond to Transitory Shocks?" In: *American Economic Journal: Macroeconomics*.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo (2020). "Nonlinear household earnings dynamics, self-insurance, and welfare". In: *Journal of the European Economic Association* 18.2, pp. 890–926.

De Santis, Massimiliano (2007). "Individual consumption risk and the welfare cost of business cycles". In: *American Economic Review* 97.4, pp. 1488–1506.

Deaton, Angus (2008). "Income, health, and well-being around the world: Evidence from the Gallup World Poll". In: *Journal of Economic perspectives* 22.2, pp. 53–72.

Duvernois, Nicolas (2021). "Financial Innovation and Risk-Sharing". In: *University of New Orleans Working Paper*.

Dynan, Karen E, Jonathan Skinner, and Stephen P Zeldes (2004). "Do the rich save more?" In: *Journal of political economy* 112.2, pp. 397–444.

Feenberg, Daniel and Elisabeth Coutts (1993). "An introduction to the TAXSIM model". In: *Journal of Policy Analysis and management* 12.1, pp. 189–194.

Ferrer-i-Carbonell, Ada (2005). "Income and well-being: an empirical analysis of the comparison income effect". In: *Journal of public economics* 89.5-6, pp. 997–1019.

Fisher, Jonathan D and David S Johnson (2020). "Inequality and Mobility over the Past Half Century using Income, Consumption and Wealth". In: *NBER Chapters*.

Flavin, Marjorie and Takashi Yamashita (2002). "Owner-occupied housing and the composition of the household portfolio". In: *American Economic Review* 92.1, pp. 345–362.

Floden, Martin (2001). "The effectiveness of government debt and transfers as insurance". In: *Journal of Monetary Economics* 48.1, pp. 81–108.

Floden, Martin and Jesper Linde (2001). "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?" In: *Review of Economic dynamics* 4.2, pp. 406–437.

Frey, Bruno S and Alois Stutzer (2002). "What can economists learn from happiness research?" In: *Journal of Economic literature* 40.2, pp. 402–435.

Frijters, Paul, John P Haisken-DeNew, and Michael A Shields (2004). "Money does matter! Evidence from increasing real income and life satisfaction in East Germany following reunification". In: *American Economic Review* 94.3, pp. 730–740.

Fuchs-Schundeln, Nicola (2008). "The response of household saving to the large shock of German reunification". In: *American Economic Review* 98.5, pp. 1798–1828.

Gervais, Martin and Paul Klein (2010). "Measuring consumption smoothing in CEX data". In: *Journal of Monetary Economics* 57.8, pp. 988–999.

Grün, Carola, Wolfgang Hauser, and Thomas Rhein (2010). "Is any job better than no job? Life satisfaction and re-employment". In: *Journal of Labor Research* 31.3, pp. 285–306.

Guvenen, Fatih (2007a). "Do stockholders share risk more effectively than nonstockholders?" In: *The Review of Economics and Statistics* 89.2, pp. 275–288.

— (2007b). "Learning your earning: Are labor income shocks really very persistent?" In: *American Economic Review* 97.3, pp. 687–712.

— (2009). "An empirical investigation of labor income processes". In: *Review of Economic dynamics* 12.1, pp. 58–79.

Headey, Bruce and Mark Wooden (2004). "The effects of wealth and income on subjective well-being and ill-being". In: *Economic record* 80, S24–S33.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante (2008). “Insurance and opportunities: A welfare analysis of labor market risk”. In: *Journal of Monetary Economics* 55.3, pp. 501–525.

Hryshko, Dmytro, Maria Jose Luengo-Prado, and Bent E Sørensen (2010). “House prices and risk sharing”. In: *Journal of Monetary Economics* 57.8, pp. 975–987.

— (2011). “Childhood determinants of risk aversion: The long shadow of compulsory education”. In: *Quantitative Economics* 2.1, pp. 37–72.

Huang, Shiyang, Maureen O’Hara, and Zhuo Zhong (2021). “Innovation and informed trading: Evidence from industry ETFs”. In: *The Review of Financial Studies* 34.3, pp. 1280–1316.

Iachan, Felipe S, Plamen T Nenov, and Alp Simsek (2021). “The choice channel of financial innovation”. In: *American Economic Journal: Macroeconomics* 13.2, pp. 333–72.

Juster, F Thomas et al. (2006). “The decline in household saving and the wealth effect”. In: *Review of Economics and statistics* 88.1, pp. 20–27.

Kahneman, Daniel and Angus Deaton (2010). “High income improves evaluation of life but not emotional well-being”. In: *Proceedings of the national academy of sciences* 107.38, pp. 16489–16493.

Kaplan, Greg and Giovanni L Violante (2010). “How much consumption insurance beyond self-insurance?” In: *American Economic Journal: Macroeconomics* 2.4, pp. 53–87.

Karahan, Fatih and Serdar Ozkan (2013). “On the persistence of income shocks over the life cycle: Evidence, theory, and implications”. In: *Review of Economic Dynamics* 16.3, pp. 452–476.

Kessler, R C et al. (2002). “Short screening scales to monitor population prevalences and trends in non-specific psychological distress”. In: *Psychological medicine* 32.6, p. 959.

Kimball, Miles and Robert Willis (2006). “Utility and happiness”. In: *University of Michigan* 30.

Kimberlin, Sara, Jiyoun Kim, and Luke Shaefer (2014). “An updated method for calculating income and payroll taxes from PSID data using the NBER’s TAXSIM, for PSID survey years 1999 through 2011”. In: *Unpublished manuscript, University of Michigan. Accessed May* 6, p. 2016.

Krueger, Dirk and Fabrizio Perri (2006). “Does income inequality lead to consumption inequality? Evidence and theory”. In: *The Review of Economic Studies* 73.1, pp. 163–193.

Kubota, Kohei (2020). “Partial insurance in Japan”. In: *The Japanese Economic Review*, pp. 1–30.

Lien, Donald, Yue Hu, and Long Liu (2017). “Subjective Well-Being and Income: A Re-Examination of Satiation Using the Regression Kink Model With an Unknown Threshold”. In: *Journal of Applied Econometrics* 32.2, pp. 463–469.

Ljungqvist, Lars and Thomas J Sargent (2018). *Recursive macroeconomic theory*. MIT press.

Mazzocco, Maurizio (2004). “Saving, risk sharing, and preferences for risk”. In: *American Economic Review* 94.4, pp. 1169–1182.

Nikolova, Milena and Carol Graham (2020). *The Economics of Happiness*. Tech. rep. GLO Discussion Paper.

Santaeulalia-Llopis, Raul and Yu Zheng (2018). “The price of growth: consumption insurance in China 1989–2009”. In: *American Economic Journal: Macroeconomics* 10.4, pp. 1–35.

Simsek, Alp (2013a). “Financial innovation and portfolio risks”. In: *American Economic Review* 103.3, pp. 398–401.

— (2013b). “Speculation and risk sharing with new financial assets”. In: *The Quarterly journal of economics* 128.3, pp. 1365–1396.

Stevenson, Betsey and Justin Wolfers (2013). “Subjective well-being and income: Is there any evidence of satiation?” In: *American Economic Review* 103.3, pp. 598–604.

Winkelmann, Liliana and Rainer Winkelmann (1995). “Happiness and unemployment: a panel data analysis for Germany”. In: *Applied Economics Quarterly* 41.4, pp. 293–307.

— (1998). “Why are the unemployed so unhappy? Evidence from panel data”. In: *Economica* 65.257, pp. 1–15.

Chapter 3

INCOME AND CONSUMPTION SHOCKS, AND PORTFOLIO CHOICE

Abstract

What are the portfolio choice implications of idiosyncratic income shocks and their transmission? Household consumption is well insured against income shocks. As idiosyncratic income shocks are not consumed, I investigate whether these shocks have implications for portfolio choice. Using household level income, consumption, and wealth data, I find that positive income shocks are associated with an increase in participation, risky allocation, and a decrease in unsecured debts. Using a two-stage instrumental variable approach, I find that permanent shocks are invested in housing while transitory shocks are used to reduce households' short-term liabilities. While these results are robust to definition change, the economic impact is obscured by the latent variable approach. As such, I can only measure the direction of the relationship.

JEL classification: D12, D14, G11, G51

Keywords: Idiosyncratic Risk, Labor Income, Portfolio Choice, Debt

3.1 Introduction

What are the portfolio choice implications of idiosyncratic income shocks and their transmission? A large body of household finance literature is dedicated to understanding the low participation puzzle. Indeed, financial markets offer investors a change to diversify their sources of income, yet participation rates remain low. In the Panel Survey of Income Dynamics (PSID), the average direct participation is 25% while indirect is 50%[1]. Drivers of portfolio choice are a relevant question considering the development of financial markets (Frame and White, 2004), increased exchange-traded assets (Iachan, Nenov, and Simsek, 2021), and reduced transaction cost (Turley, 2012). Nonetheless, participation and portfolio allocation appears suboptimal for most households.

The prevailing view argues that participation in financial market participation and allocation to risky assets depends on households' exposure to uninsurable risk (Cocco, Gomes, and Maenhout, 2005; Benzoni, Collin-Dufresne, and Goldstein, 2007; Lynch and Tan, 2011). Theoretical models (such as Viceira, 2001) argue that investors rationally invest in assets that are uncorrelated with their labor income risk. Households try to hedge their uninsurable wage risk by making portfolio decisions. There is empirical support in the literature (Guiso, Jappelli, and Terlizzese, 1996; Betermier et al., 2012; Bonaparte, Korniotis, and Kumar, 2014) that households reduce their risky share in the presence of background risk. Furthermore, the cyclicality of background risk is critical. Palia, Qi, and Y. Wu (2014) find that participation increases as the correlation between risk and market returns decreases. Catherine, Sodini, and Zhang (2020) show that higher moments of risk (skewness of household idiosyncratic income risk) dominate portfolio decisions.

Fagereng, Guiso, and Pistaferri (2018) argue that proxies of background risk are incomplete and incorporate large heterogeneity components. They isolate a better proxy of wage risk by matching households with the firm that employs them and find statistically and economically significant support for the hedging hypothesis. Their results are in contradiction of Massa and Simonov (2006) who find that households do not hedge. They instead invest in stocks with low information acquisition costs, familiar stocks that tend to be linked to their wage risk.

The econometric concerns raised by Fagereng, Guiso, and Pistaferri (2018) are indeed non-trivial. Their results are economically more robust than previous results. More importantly, there is an important bundle of papers looking at determinants of portfolio composition (Barasinska, Schäfer, and Stephan, 2012; Dimmock and Kouwenberg, 2010; Guiso,

[1] Direct participation entails owning shares directly as opposed to indirectly through Individual Retirements Accounts (I.R.A).

Sapienza, and Zingales, 2008; Renneboog and Spaenjers, 2012). Risk, loss-aversion, trust, financial literacy, and education are strong determinant of portfolio choice.

While this paper is related to household risk, I concentrate on households' income shocks and their transmission to income shocks. In that regards, this paper is related to Addoum, Delikouras, and Korniotis (2019) who show that income-consumption sensitivities are important determinants of portfolio choice. However, I ask a simpler question. There is significant evidence that households are well insured against income shocks; that is consumption is not highly sensitive to income shocks. Blundell, Pistaferri, and Preston (2008) and Commault (2020) document low marginal propensity to consume out of shocks both transitory and permanent. However, there is a gap in the literature in that I am the first to ask how do households' portfolio adjust to the realization of shocks, not their risk. This paper is therefore more related to Basten, Fagereng, and Telle (2016). They show that households' deplete their savings in direct response to a negative employment shocks. Surprisingly, they find that households increase savings three to four years prior to the labor shock indicating households have significant information on future shocks. Primiceri and Van Rens (2009) argue that households do in fact possess superior information about shocks that appears unpredictable to the econometrician. Basten, Fagereng, and Telle (2016) also find that households reduce risky asset participation and shift composition to safer assets following the labor shocks[2]. Fagereng, Holm, and Natvik (2019) use lottery wins to investigate the impact of positive income shocks. They find a significant marginal propensity to consume out of the shock but limited portfolio changes. They find that while savings do increase slightly, they are consumed within the year. Similar results are found for stock and bond investments, or changes in household debts.

I use data from the Panel Survey of Income Dynamics which contains rich data on a representative sample of US households. I concentrate on the dynamics of income, consumption, and wealth. Following the redesign of the PSID, detailed consumption and wealth variables are available at every wave[3]. I use labor income to avoid income shocks being correlated with portfolio performance. I calculate income and consumption shocks following Blundell, Pistaferri, and Preston (2008). I use a latent variable approach to deal with the two main problems in using survey based wealth data. I first calculate several ratios indicative of household portfolio choices: stock to financial wealth, stock to total wealth, stock to total wealth and home equity, secured debt to total debt or total wealth, unsecured debt to total debt or total wealth, total debt to total wealth. This ratios are typically used in the literature to describe portfolio allocation (Palia, Qi, and Y. Wu, 2014; Brown, Garino, and Taylor,

[2]Their sample period is not long enough to determine whether the effect is permanent or transitory.
[3]Waves are available biennially.

2013). However, these ratios are likely to contain measurement errors (Neri and Rannalli, 2011; Vermeulen, 2016). To reduce the impact of measurement error, I assign households to an 11-point scale from 0 to 10 depending on their respective percentages. This reduces mismeasurement problems but not truncation problems. I use the methodology proposed by Bayer and Juessen (2015) to estimate pseudo residuals from an ordered probit model. The change in residuals can be understood as the change in latent utility from portfolio allocation. This methodology has two advantages. Firstly, it removes household observed and unobserved heterogeneity in portfolio decision. By removing household characteristics and fixed effects, I can identify the impact of income or consumption shocks on portfolio choices. Secondly, I can use simple linear regression models and do not have to use truncated models to account for low participation.

The main independent variable becomes idiosyncratic labor income shocks and the left hand side variables become latent portfolio choices. Another benefit of such a framework is that it allows from more complex income and transmission processes. Indeed, income shocks, defined as changes in residuals, are a weighted average of permanent and transitory shocks. While latent and unobserved variables themselves, they play an important role. Indeed, Blundell, Pistaferri, and Preston (2008), Kaplan and Violante (2010), or Bayer and Juessen (2015) show that transitory shocks have a much lower (if any) impact on consumption or happiness than permanent shocks do. I therefore use a set of instruments describe by Blundell, Pistaferri, and Preston (2008) and formally used by Bayer and Juessen (2015) to measure the impact of different shocks on household portfolio choices.

A drawback of this framework is that it obscures the economic impact. While less important in happiness studies, quantifying the actual shift in portfolio composition would have been interesting. I can nonetheless look at the direction of the impact of shocks. I find that positive income shocks to household labor increase their latent portfolio participation. However, only permanent shocks that are transmitted to consumption have a statistically significant impact on portfolio choice. This is caused by the lack of permanent insurance. Households consume a significant portion of these shocks. On the other hand, transitory income shocks are significantly impactful for household participation; only for stocks and do not impact I.R.As allocation. These results hold true for different measures of risky allocations as well. Furthermore, the results are robust to the definition of permanent shocks.

I also investigate the impact of shocks on household debt. I find that households experiencing positive labor income shocks reduce their share of unsecured debt. However, the effect only holds true for transitory shocks. Permanent shocks induce an increase in secured debt. Secured debt is comprised of mostly mortgage debt. This indicate that households use permanent shocks to lever up and increase their housing investment. This result is supported by Cocco (2005) who show that housing and stock investments are substitutes.

The rest of the paper is organized as follows. Section 3.2 will connect this work to the related literature. Section 3.3 will describe the data I use. Section 3.4 will describe the methodology used to estimate shocks. Section 3.5 will present results and section 3.6 will conclude.

3.2 Literature Review

The literature on household portfolio choice tends to focus on uninsurable risk. While household risk is generate via a myriad of sources, the most prominent source across a majority of households is wage risk or labor income risk. The empirical consensus is that households' avoid unnecessary risks and will reduce their ownership of risky assets in the presence of uninsurable wage risk. This is often referred to as the income hedging motive. Cocco, Gomes, and Maenhout (2005) develop a model to show that in the presence of labor risk, agents will increase their investments in risky assets, assuming labor income is uncorrelated with equity returns. Benzoni, Collin-Dufresne, and Goldstein (2007) look at portfolio choice in a market when labor income and dividends are cointegrated. Their model fit empirical facts that older households hold more of their wealth in stocks contrary to conventional wisdom. Lynch and Tan (2011) argue that countercyclical variation in labor income growth plays a significant role in portfolio dynamics. This result is supported by Bonaparte, Korniotis, and Kumar (2014) who find that individuals with low income-return correlation have a higher propensity to participate in financial markets and higher allocation toward risky assets in samples of Dutch and American households.

Guiso, Jappelli, and Terlizzese (1996) find that Italian households will hold a lower proportion of risky asset when experiencing higher background risk. Furthermore, when faced with idiosyncratic risk, households become concerned with borrowing constraints and allocate wealth to more liquid assets. Guiso, Jappelli, and Terlizzese (1996) rely on subjective expectations on risk which is appealing to the econometrician as it encompasses only information available to households[4]. However, as all survey measurements, the expectations are prone to errors. This is further compounded by measurement errors in wealth (Neri and Rannalli, 2011; Vermeulen, 2016). To mitigate measurement errors, researchers can use administrative records. Betermier et al. (2012) use Swedish household records and find that households do adjust their portfolio depending on human capital risk. Their study is robust to heterogeneity mismeasurement. Indeed, Fagereng, Guiso, and Pistaferri (2018) point out

[4] Hochguertel (2003) uses similar expectations in Dutch panel and Kézdi and Willis (2009) in an American panel and finds similar results.

that unobserved heterogeneity may contribute to both background risk and portfolio allocation. Betermier et al. (2012) circumvent this issue by looking at households that switch industries creating an exogenous shock to wage risk.

It is important to point out that the hedging motive is not the only driving force of portfolio choice. Massa and Simonov (2006) argue that portfolio choice models often only consider three components: uninsurable risk, a risky asset, and a risk-free asset. These considerations are done to match survey data which seldom provides specific holdings. They use Swedish data which allows them to match income and wealth data to specific holdings. Their findings are contrary to hedging. Indeed, they find that households invest in stocks which are linked to their labor income. They coin the term familiarity to explain this phenomenon; households invest in stocks in which information is cheaper to acquire. Another mechanism at play is the transmission of these income shocks to consumption shocks. Addoum, Delikouras, and Korniotis (2019) find that households' whose consumption is more sensitive to income shocks have stronger incentives to hedge and thus invest more in risky assets.

The studies cited above all share a common econometric issue in that the proxy for background risk is not necessarily robust. Risk is often defined as residuals (or change in residuals) from regressions removing household characteristics from income. Primiceri and Van Rens (2009) argue that these unpredictable changes are, in fact, known to households and only unknown to the observer. Fagereng, Guiso, and Pistaferri (2018) use administrative data from Norway and match households' income and wealth to the companies that employ them. This rich data allow them to isolate heterogeneous wage risk that depends on workers earnings' variation. They find that the effect of wage risk on portfolio choice is significant. Indeed, the effect is much larger than estimates from previous studies ignoring econometric issues.

The small economic impact of covariance risk on portfolio choice along with papers showing that cyclical skewness in labor income risk fits equity premium data (Constantinides and Ghosh, 2017; Schmidt, 2016; Guvenen et al., 2015) has motivated researchers to look at higher moments' impact on portfolio choice. Catherine, Sodini, and Zhang (2020) find that households facing skewed income risk in low market returns environments participate less and invest lower amounts in stocks. Shen (2018) shows that cyclical skewness in idiosyncratic income risk is transmitted to both consumption growth and portfolio allocation.

In addition to the impact of wage risk on portfolio choice, another strand of literature focuses on the heterogeneous determinants mentioned above. Barasinska, Schäfer, and Stephan (2012) use a reliable measure of risk preferences from the German Socioeconomic Panel and find that more risk averse individuals hold fewer assets and concentrate their holdings in risk-free assets. Furthermore, they find that borrowing constraints and precautionary motives are significant predictors of portfolio choice, echoing results by Guiso, Jappelli, and Terlizzese (1996). However, they conclude that the low participation puzzle[5] cannot be explained alone by risk attitudes and argue that financial literacy or participation costs may play an important role. Similarly, Dimmock and Kouwenberg (2010) use data from the The Netherlands and estimate households' loss aversion. They find that loss averse households have a lower probability of participation. When they do participate, they are more likely to hold mutual funds rather than individual stock, suggesting heterogeneity in households affects both choice and allocation. Guiso, Sapienza, and Zingales (2008) show that trust in the stock market is also a strong determinant of participation. Trust in stock markets is also a central theme of Renneboog and Spaenjers (2012) who show that religious attitudes are related to household finance behaviors.

Risk and loss preferences are difficult concepts to measure. There are other easier to measure variables affecting portfolio choice. Branikas, Hong, and Xu (2020) show that household residence location is linked to portfolio choice and the demand for stocks. They argue that their results are driven by economic prospects expectation of the target location. However latent, a household's choice of residence may reflect a certain demand for stocks. Rosen and S. Wu (2004) find that households with poorer health are less likely to hold stocks. However, Love and Smith (2010) argue that a proper econometric treatment renders any effect insignificant. Wiemann and Lumsdaine (2020) find that health care uncertainty alone is enough to drive portfolio choice. In additional to health and health care risk, Bogan and Fertig (2013) show that poor mental health reduces risky holdings and increases risk-free holdings. Spaenjers and Spira (2015) use a subjective life horizon question in the 1995 Survey of Consumer Finances and find that households with longer life horizons invest more in risky assets. Love (2010) shows that family composition matters for portfolio choice, mainly marital status and the presence of children. Alessie, Hochguertel, and Soest (2004) further the argument on heterogeneity in household and portfolio decisions by showing it correlates to ownership in stocks or mutual funds. They find that unobserved heterogeneity influences how much to invest but also which assets to invest in.

On the other side of asset allocation, there is debt allocation. As borrowing constraints appear to be driving portfolio choice in the presence of uninsurable wage risk, a small portion of the literature focuses on household debts. As risk attitudes are linked to asset allocation, so is debt allocation. Brown, Garino, and Taylor (2013) find that highly risk averse households tend to accumulate less debt. Jiang and Lim (2018) show that trusting households hold

[5] Another econometric issues which forces the research to truncate a significant portion of the sample. Note that this low participation puzzle is common to many countries and datasets.

more debt but are also less likely to default. Agarwal, Chomsisengphet, and Liu (2011) find that many households characteristics affecting financial market participation (marital status, children, education, etc.) also affects debt levels. Almenberg et al. (2020) find that household characteristics can be proxies for attitudes towards debt. Furthermore, they argue that debt attitudes have social and cultural dimensions.

However, to the best of my knowledge, I am the first to investigate a direct link between household income shocks and debt allocation. The literature focuses on choice in the presence of risk. I chose to focus on households' response to shocks. That is, I ask what is the marginal allocation (in either assets or debt) of income shocks? In other words, what is the marginal propensity to invest out of labor income shocks. The following section will describe the data and empirical strategy.

3.3 Data

I use data from the Panel Survey of Income Dynamics (PSID). The PSID is a survey of American households covering consumption, income, and wealth[6]. Furthermore, the survey is longitudinal and follows the same households wave after wave. The panel dimension allows the use of more robust econometric estimators. While the PSID started in 1968 with a representative sample of 5,000 households, I use data from the 1999-2019 waves. The PSID underwent redesign in 1999 in two impact ways for this paper. Firstly, the wealth module is administered at every wave as opposed to every 5 years prior to 1999. This allows me to track households' financial market participation status, financial holdings, wealth, and debts without resorting to imperfect identification techniques as in Guvenen (2007). Secondly, the PSID expanded the expenditure categories to include items beyond food consumption. These items now cover to 70% of the categories covered by the CEX[7] (Blundell, Pistaferri, and Saporta-Eksten, 2016). This allows me to estimate consumption shocks without resorting to imputation techniques as in Attanasio and Pistaferri (2014), Blundell, Pistaferri, and Preston (2008), or Fisher and Johnson (2020). One drawback of using the redesigned PSID is the frequency of the data. Surveys are now administered every other year, which may obscure certain transmission mechanisms between waves.

Although wealth and expenditure data starts in 1999, I will use first differenced data and thus require previous wave data. My sample period is thus 1997-2019, which covers 11 distinct surveys. The main measure of income is household disposable labor income. Labor

[6]The PSID actually contains a wide range of topics but the main interest of this papers focuses on these three dimensions of the survey

[7]The Consumer Expenditure Survey is the survey of reference when it comes to household consumption patterns. However, it lacks a panel dimension.

income is defined as the sum of labor income of reference person, spouse (if present), and other family members[8] (if present). I estimate households' federal taxes using NBER Taxsim and guidelines by Feenberg and Coutts (1993) and Kimberlin, Kim, and Shaefer (2014). This tax liability is estimated over all sources of income. I therefore calculate the ratio of labor income to taxable income and assume that ratio is the proportion of taxes paid out of labor income.

Consumption is defined as the sum of all 1999 expenditure categories. Further categories where added in 2005 but I only consider items present in all waves for consistency. The items are home insurance, electricity, heating, water, miscellaneous utilities, car insurance, car repairs, gasoline, parking, bus fares, taxi fares, other transportation, school tuition, other school expenses, child care, health insurance, out-for pocket health, rent, food at home, food away from home, and food delivered. For home owners, I follow Flavin and Yamashita (2002) and Attanasio and Pistaferri (2014) and impute a rent expenditure equal to 6% of the reported house value.

In the wealth modules of 1984, 1989, and 1994, financial assets are grouped into one category containing direct ownership of stocks, mutual funds, or through I.R.As. After 1999, these two categories are separated, allowing me to identify households that participate directly or indirectly in financial markets. These separation is important as Bonaparte, Korniotis, and Kumar (2020) point out; there is limited trading activity in I.R.A accounts. I define stock owners as households with a positive value of reported stock ownership; I define I.R.A owners as households with a positive value of reported I.R.A ownership. I also identify households that own positive values in either accounts or both. Indeed, Alessie, Hochguertel, and Soest (2004) identify state dependence between the ownership of the two types of assets suggesting switching costs. I then define several ratios of households risky asset share. I scale the value of stocks with financial wealth, total wealth without home equity, and total wealth with home equity following Palia, Qi, and Y. Wu (2014). I define financial wealth as the sum of stocks, I.R.As, and cash, savings account, and treasury bills. This ratio is the risky share of financial assets. Scaling financial wealth by total wealth considers further investment vehicles available to households such as non-public business or non-residential real estate investments. I consider total wealth without home equity as imputed rent expenditure is part of consumption. The inclusion of house value explicitly in wealth or implicitly in consumption may induce measurement errors. I also consider these three ratios with the sum of stocks and I.R.As as the numerator. I.R.As are not riskless

[8] Labor income of other family members is not consistently measured across waves. For consistency, I use OFUMs taxable income. Although taxable income may contain financial asset income which may bias the estimation. However, taxable OFUM income represents a small portion of household income.

assets and the choice to invest and how much to invest may be depend on the households' labor income risk profile.

The wealth module does not only consider households' assets but also debts. I define two types of debt following Brown, Garino, and Taylor (2013): secured and unsecured debts. Unsecured debt is the sum of debt categories such as for credit card charges, student loans, medical bills, legal bills, and loans from relatives. I define secured debt as the sum of debt categories such as mortgages[9], debt on business or farms, and debt on other real estate holdings. Household total debt is the sum of secured and unsecured debt. I calculate 5 debt ratios, secured debt to total debt, unsecured debt to total debt, secured debt to total wealth, unsecured debt to total wealth, and total debt to total wealth. Note that these ratios are truncated at 1; that is, households whose debts are worth more than their assets have a maximum debt ratio of 100%. Although incorrect, this is an issue for a very small number of households.

All income, consumption, and wealth data are deflated to 2010 dollars. Although data quality is not necessarily an issue in the PSID, it is nonetheless survey-based and not administrative records. I therefore apply several filters to ensure data quality. I exclude households with 0 food consumption[10]. I exclude households with missing observations on education, region of residence, age of reference person, and employment status. I exclude households with a retired reference person. Furthermore, if the reference person retires at any point during the sample period, previous observations for these households are also dropped. I require the reference person to be between 25 and 65 years old. I only consider households from the Survey of Research Center sample and exclude households from the Survey of Economic Opportunities and immigrant refresher samples. I remove households whose income is less than $100, grows by more than 500% or falls by more than 80%. Finally, I require all households to have at least 4 continuous observations to be included.

The final sample is composed of 3,035 unique households and 21,771 household-year observations. Table 3.1 reports summary statistics for the final sample. The average and median disposable labor income are $98,595 and $76,919 respectively. While this seems higher than nationally representative statistics, the sample selection criteria are constructed in a way that exclude low-income households. Indeed, by considering labor income, households whose main income source is government transfers are excluded from the sample. Household consumption is on average $38,541. 22.1% of the sample own stocks, 35.3% own I.R.As, 14.5% own both, and 42.9% own either. The average financial portfolio is composed of

[9]Reference persons are asked about their first and second mortgages.

[10]This filter also excludes households with 0 total consumption

Table 3.1: Summary Statistics

This presents summary statistics for the final sample used. This sample covers the 1997-2019 PSID surveys. 1997 and 1999 are only kept for conditioning, first-differencing, and instrumentation. It is composed of 21,771 household year observations and 3,035 unique households. Labor income is the sum of labor income for all household members (including reference person, spouse, and other family members). Federal taxes, estimating with NBER Taxsim, are deducted from labor income to calculate disposable labor income. The proportion of labor to total income is used to adjusted households' tax liabilities. Households must have \$100 of income and income growth between -80% and 500% to be included. Total household consumption is the sum of all expenditure categories in the 1999 survey. Household are required to have positive food and total consumption. Household with non-zero stock or I.R.A ownership have a value of 1. Financial wealth is defined as the sum of stock, I.R.As, checking and savings accounts, CDs, and treasury bills. Secured debt is the sum of mortgages, business or farm debt, and other real estate debt. Unsecured debt is the sum of credit card debt, legal and medical bills, student loans, and loans from relatives.

	Mean	St. Dev.	$25^{th}Pctl.$	Median	$75^{th}Pctl.$
Disposable Labor Income	93,595	105,419	48,276	76,919	113,402
Total Household Consumption	38,541	23,770	23,642	33,474	46,819
Stock Ownership	0.221	0.415	0	0	0
I.R.As Ownership	0.353	0.478	0	0	1
Stock and I.R.As Ownership	0.145	0.352	0	0	0
Stock or I.R.As Ownership	0.429	0.495	0	0	1
Stk./Financial Wealth	0.095	0.229	0	0	0
Stk&IRA./Financial Wealth	0.302	0.392	0	0	0.7
Stk./Wealth w/out home equity	0.066	0.184	0	0	0
Stk&IRA./Wealth w/out home equity	0.207	0.327	0	0	0.4
Stk./Wealth with home equity	0.043	0.133	0	0	0
Stk&IRA./Wealth with home equity	0.136	0.240	0	0	0.2
Secured/Total Debt	0.591	0.450	0	0.9	1.0
Unsecured/Total Debt	0.268	0.396	0	0.04	0.3

10% of stocks and 30.2% of financial risky assets. When considering a broader definition of wealth, the financial risky share of household portfolio is 13.6%. The average household holds 60% of its debt in mortgages.

3.4 Shocks to Income, Consumption, and Portfolio

Estimating idiosyncratic shocks to households' income and consumption is straight forward. I regress log income (or consumption) on a set of household characteristics, collect the residuals, and first difference them. For example, income shocks can be written as:

$$\Delta y_{i,t} = log(y_{i,t+2})^* - log(y_{i,t})^* \quad , \tag{3.1}$$

where $\log(y_{i,t})^*$ is:

$$log(y_{i,t})^* = log(y_{i,t}) - f\widehat{(t, Z_{i,t})} \tag{3.2}$$

Note that equations (3.1) and (3.2) are used to estimate consumption shocks noted as $\Delta c_{i,t}$. Equation (3.2) represents the difference between the actual and fitted values from a pooled OLS regression with log income or log consumption as the dependant variable. $f(t, Z_{i,t})$ is a vector of household characteristics containing family size, number of children, and dummy variables for education[11], region of residence, year of birth, year, employment status, marital status, race of reference person[12], presence of income earning spouse, presence of income earning other family members, gender of the reference person, and whether the household supports children outside of the household. These variables are allowed to vary with time. Similar regressions are used in Blundell, Pistaferri, and Preston (2008), or Gervais and Klein (2010).

As seen in table 3.1, the wealth variables are skewed towards 0. Indeed, a majority of households do not hold any financial assets or debts. This fact creates estimation issues. This is usually resolved in the literature by employing a 2-stage Heckman (Heckman, 1979) to account for the probability of participation. However, this not does solve all estimation issues in this case. Indeed, participation and portfolio allocation is extremely likely to be affected by a series of unobservable household characteristics. The most likely are risk aversion and future expectations (Bucciol, Miniaci, and Pastorello, 2017; Barasinska, Schäfer, and Stephan, 2012; Kézdi and Willis, 2009). Furthermore, while I can observe the decision to participate in financial markets or not, or the allocation of risky assets, I cannot observe the

[11] Less than high school, high school diploma, some college or more

[12] white, black, or other.

net utility of said decision. I am therefore interested in the latent variable and the impact of shocks on the latent variable. This is a well-known problem in the subjective well-being literature. I adopt the methodology by Bayer and Juessen (2015) to estimate the effect of shocks on latent participation and portfolio allocation. Consider a household's share of risky assets $s_{i,t}$ generated by an ordered Probit model such as:

$$s_{i,t} = s^*_{i,t} + f(t, z_{i,t}) \tag{3.3}$$

$$s^*_{i,t} = s^{**}(\Psi_{i,t}) + \mu_i^s + \xi_{i,t} \tag{3.4}$$

I assume that latent portfolio allocation is additively separable in household characteristics ($z_{i,t}$), household fixed effects (μ_i^s), and an i.i.d error term capturing survey measurement error ($\xi_{i,t}$). I can remove the effect of household observable characteristics in a first-stage regression. Latent portfolio allocation is the residual from this regression. I can further remove the household fixed effects by first-differencing $s^*_{i,t}$. I can then estimate the marginal effect of shocks ($\Psi_{i,t}$ could be income or consumption shocks) on latent portfolio choice or portfolio allocation[13].

Bayer and Juessen (2015) use this type of framework on a happiness variable that is measured on a scale from 0 to 10. I assume that the portfolio allocation is generated by an ordered Probit but the actual measured variable is not. In fact, it is a continuous variable between 0 and 1. I thus bucket the ratios presented in table 3.1 on a 11-point scale between 0 and 1. Households with 0 values are given 0 point; households with a ratio between 0.01% and 9.99% are given 1 point; and so on. This re-measuring serves two purpose. First, I can then estimate an ordered Probit. Second, it reduces a well-known measurement error. As pointed by Neri and Rannalli (2011) and Vermeulen (2016), measuring wealth in surveys is a tricky endeavor. Heaton and Lucas (2000) show that portfolio composition and the impact of demographics is sensitive to wealth measurements and definitions. Vermeulen (2016) further emphasizes measurement issues for financial wealth which tends to be under-reported. By segmenting portfolio allocation in large increments, the effect of mismeasuring wealth is likely to be attenuated. I recognize that this could increase measurement error in $\xi_{i,t}$ if households get assigned a non-continuous value that is above or below their true allocation. However, this is only likely for households near breakpoints and thus not likely to be a significant issue.

[13]Note that I use portfolio choice for participation, and portfolio allocation for the amount invested in risk assets (or debts).

I therefore follow Bayer and Juessen (2015) in estimating latent allocation from observed portfolio allocation. I assume that allocation $s_{i,t}$ is determined by:

$$s_{i,t} = j \quad \text{if} \quad s^{**}_{i,t} \in (\bar{c_j}, \bar{c_{j+1}}] \tag{3.5}$$

I estimate the effects of household characteristics and cutoff values $\bar{c_j}$ with a standard ordered Probit estimator. Note that $f(t, z_{i,t})$ is the same set of variables and interactions used in equation (3.2) when estimation income and consumption shocks. Thus, the same effect of risk aversion impacting portfolio choice and household self-selecting in industry with smoother labor income profiles can be removed. Given the estimated $\widehat{f(t, z_{i,t})}$, I infer an interval $S_{i,t} = (\tilde{\bar{c}}_{sit} - \widehat{f_{t,z_{i,t}}}, \tilde{\bar{c}}_{sit+1} - \widehat{f_{t,z_{i,t}}}]$, where the true latent allocation $s^*_{i,t}$ is. I then assume normality for $v^*_{i,t}$ are calculate the conditional expected value of residual latent portfolio allocation $(\overline{s^*_{i,t}})$:

$$\overline{s^*_{i,t}} = \frac{\int_{s \in S_{i,t}} s\phi(s)}{\Phi(S_{i,t})}, \tag{3.6}$$

where ϕ is the density and $\Phi(S)$ the probability of S for a standard normal distribution. I use $\overline{s^*_{i,t}}$ as a measure of latent allocation $s^*_{i,t}$. Bayer and Juessen (2015) argue that this method is an extension and generalization of Van Praag (2004) probit-OLS procedure. Furthermore, working with first differenced-residuals on both sides allows me to use simple linear regression to estimate the effects of shocks. This procedure is applied to all financial wealth and debt ratios of table 3.1. Each of the continuous variables is sorted into tranches of 10%. It is also applied to participation despite the ordinal variable being either 0 or 1. This results in only one cutoff point being equal to the constant term from a simple probit model.

3.5 Results

The main goal of this analysis is to investigate the impact of idiosyncratic income and consumption shocks on portfolio decisions both choice and allocation. As pointed out in the previous section, the ordered probit removal of characteristics allows for simple linear regression of the following form:

$$\Delta\overline{s^*_{i,t}} = \alpha + \beta\Delta y_{i,t} + \varepsilon \tag{3.7}$$

$\Delta\overline{s^*_{i,t}}$ can be calculated using any of the wealth and debt ratios. Furthermore, $\Delta y_{i,t}$ can be replaced by $\Delta c_{i,t}$ or any other shock proxies. I first consider the impact of average income and consumption shocks on latent portfolio choice. That is, how the latent ownership status of households responds to idiosyncratic income or consumption shocks. The results are presented in table 3.2. I consider the impact of income shocks in panel A and the impact of consumption shocks in panel B. In panel C, I first regress income shocks and consumption shocks and use the fitted values as the measure of shocks. Indeed, there is significant evidence in the literature[14]. The fitted values are therefore the share of income shocks that get transmitted to income shocks.

The dependant variables are the latent participation variables. I consider four choices made by households: invest directly in stocks or mutual funds (1), invest indirectly through investments vehicles such as I.R.As (2), invest in both (3), and invest in either (4). In panel A, all coefficients are significantly positive. These results suggest that income shocks have a positive impact on households' participation decisions. The point estimate for specification (1) is larger than the point estimate for specification (2) (0.083 vs 0.052). While both positive and statistically significant at the 1% level, this would suggest that income shocks have a stronger impact on the decision to directly invest in the stock markets. The coefficients for specifications (3) and (4) are similar to the ones for (1) and (2), suggesting that the two investment vehicles are not substitute. The decision to invest in either is separate.

In panel B, the coefficients are also statistically positive at the 1% level. However, the point estimates are now larger for investing in I.R.As in specifications (2) and (4). Indeed, the coefficient for stocks is 0.038 vs 0.086 for I.R.As. Unexplained consumption Δc can be written as the sum of transmitted income shocks and changes in consumption independent of income shocks such as preference shocks, or innovation to higher moments of the income process (Blundell, Pistaferri, and Preston, 2008)[15]:

$$\Delta c_{i,t} = \beta \Delta y_{i,t} + \Xi_{i,t} \tag{3.8}$$

In panel C, I remove the effect of Ξ in equation (3.8) by estimating the fitted values of Δy on Δc. A small proportion of income shocks (11%) get transmitted to consumption shocks. The results are therefore similar to the patterns observed in panel A. Income shocks that are transmitted to consumption shocks impact the decision to participated in financial markets directly more so than the decision to indirectly participate. These results, taken with the results of panel B, suggest that is partly driven by consumption shocks independent of income shocks. Δc may be capturing preference shifts not removed in the first stage

[14] See Mace, 1991; Cochrane, 1991; Blundell, Pistaferri, and Preston, 2008; Gervais and Klein, 2010

[15] Blundell, Pistaferri, and Preston (2008) argue that Ξ may also contain measurement error in consumption. They add a specific term $u_{i,t}$ to catch the measurement error. As I use measured expenditure as opposed to imputed consumption, measurement error is less likely to be an issue.

Table 3.2: Latent Portfolio Choice and Average Idiosyncratic Shocks

*This table reports the OLS estimation of equation (3.7). The dependant variables are latent participation estimated with the framework presented in section (3.3) and equations (3.3) to (3.6). Stocks refers to latent direct participation. I.R.As refer to latent indirect participation. Both refers to households who have non-zero funds in both stocks and I.R.A accounts. Either refers to households who have non-zero funds in both stocks and I.R.A accounts. The independent variables are Δy and Δc estimated using equations (3.1) and (3.2) for panels A and B respectively. In panel C, the effect of Ξ in equation (3.8) is removed. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

Ownership of :	Stocks (1)	I.R.A (2)	Both (3)	Either (4)
Panel A. Income Shocks				
Shocks : Δy	0.083^{***}	0.052^{***}	0.098^{***}	0.054^{***}
	(0.017)	(0.018)	(0.018)	(0.018)
Constant : α	0.016^{***}	-0.003	0.009	0.003
	(0.006)	(0.007)	(0.006)	(0.007)
Panel B. Consumption Shocks				
Shocks : Δc	0.038^{**}	0.086^{***}	0.061^{***}	0.077^{***}
	(0.019)	(0.021)	(0.020)	(0.020)
Constant : α	0.016^{***}	-0.003	0.010	0.003
	(0.006)	(0.007)	(0.006)	(0.007)
Panel C. Consumption Shocks Projection				
Shocks : $\widehat{\Delta c}$	0.775^{***}	0.490^{***}	0.916^{***}	0.510^{***}
	(0.155)	(0.172)	(0.169)	(0.171)
Constant : α	0.013^{**}	-0.004	0.007	0.002
	(0.006)	(0.007)	(0.007)	(0.007)

estimation. However, these preference shifts are unlikely impact risk aversion. Indeed, I.R.As are likely to be safer investment vehicles than individual investment accounts. I.R.As are managed by professionals and are diversified. On the other hand, individuals investors are seldom fully diversified (Ivković, Sialm, and Weisbenner, 2008).

The income shocks described by equation (3.1) can be understood as the sum of two components: permanent and transitory. Modeling income as such is common is this literature (Kaplan and Violante, 2010):

$$y_{i,t} = P_{i,t} + v_{i,t} \quad , \tag{3.9}$$

where $P_{i,t}$, the permanent component, is a martingale process with serially uncorrelated innovations $\zeta_{i,t}$. The transitory component is a MA(0) process with serially uncorrelated innovations $\varepsilon_{i,t}$. The innovations in permanent and transitory components are independently and identically distributed. I can therefore write the growth in unexplained income growth as:

$$\Delta y_{i,t} = \zeta_{i,t} + \varepsilon_{i,t} \tag{3.10}$$

The transmission of shocks to latent portfolio variables can be written as:

$$\Delta \overline{s^*_{i,t}} = \phi \zeta_{i,t} + \psi \varepsilon_{i,t} \tag{3.11}$$

Unfortunately, just as latent portfolio choice is unobserved, neither permanent nor transitory shocks are observed. Fortunately, Blundell, Pistaferri, and Preston (2008) offer an instrumental variable approach to the Minimum Distance Estimator (also used in Bayer and Juessen, 2015). I follow their example and use $y^*_{i,t+1} - y^*_{i,t-2}$ as an instrument for permanent shocks, and $\Delta y^*_{i,t+1}$ as an instrument for transitory shocks. See Blundell, Pistaferri, and Preston (2008) or Kaplan and Violante (2010) for a detailed explanation of the proposed instruments. Permanent and transitory are transmitted to consumption (or $\Delta \overline{s^*_{i,t}}$) with parameters ϕ and ψ respectively with $\phi \neq \psi$. I employ a two-stage process to measure the impact of "permanent" and "transitory" consumption shocks. I first use the IV framework with consumption shocks on the left hand-side and estimate the fitted values. I then use the fitted values as independent variables in a simple OLS regression with latent variables on the left-hand side. I repeat the analysis of table 3.2 with permanent shocks in table 3.3.

In panel A, I use the IV regression with latent participation as the dependent variable. The coefficients on permanent income shocks are not significant for the first 3 specification

Table 3.3: Latent Portfolio Choice and Permanent Idiosyncratic Shocks

This table presents the impact of permanent shocks on latent portfolio choice. In panel A, I use an IV regression with $y^*_{i,t+1} - y^*_{i,t-2}$ *as an instrument for permanent shocks and latent choice as the dependent variable. Only the first stage is reported. In panel B, I first use the IV regression with consumption shocks as the dependent variable and estimate the fitted values. These fitted values, representing permanent consumption shocks, are then used in a simple linear regression with latent choice as the dependant variable. Stocks refers to latent direct participation. I.R.As refer to latent indirect participation. Both refers to households who have non-zero funds in both stocks and I.R.A accounts. Either refers to households who have non-zero funds in both stocks and I.R.A accounts. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Stocks (1)	I.R.A (2)	Both (3)	Either (4)
Panel A. Income Shocks				
Shocks : Δy	0.137	0.147	0.059	0.239***
	(0.085)	(0.092)	(0.093)	(0.092)
Constant : α	0.009	−0.006	0.001	−0.003
	(0.007)	(0.008)	(0.008)	(0.008)
Panel B. Consumption Shocks Projection				
Shocks : $\widehat{\Delta c}$	0.352***	0.239**	0.345***	0.275***
	(0.086)	(0.093)	(0.094)	(0.093)
Constant : α	0.009	−0.005	−0.000	−0.001
	(0.007)	(0.008)	(0.008)	(0.008)

(Stocks, I.R.As. both) but is significantly positive for the latent participation in either. These results might suggest that permanent income shocks have no impact on the decision to participate in financial markets which would be surprising. However, consider the results presented in panel B. The transmission of permanent income shocks to consumption is approximately 39% which is consistent with estimates in the literature. The projection of permanent income shocks on consumption on latent portfolio choice yields statistically significant and positive coefficients. Permanent income shocks have a significant impact on portfolio choice if they are transmitted to consumption shocks. As in table 3.1, the point estimate for stock participation appears larger than the point estimate for I.R.As. The result that permanent income shocks do not have an impact on participation but permanent consumption shocks do is puzzling. I offer two potential explanations. The first one relates to Campbell and Cochrane (1999) habit formation hypothesis. Positive income shocks that do not get transmitted to consumption shocks do not affect the level of habit. As such the agent's risk aversion remains unchanged and so does their portfolio. On the other hand, income shocks that do get transmitted to consumption shocks result in the agent being above their habit level, reduces risk aversion, and increases the participation. The second alternative is more mechanical. Households are not perfectly insured against permanent shocks and thus consume a significant portion of income shocks. A smaller portion of these shocks remain for households to allocate to assets and debts. The projection is not the proportion of shocks allocated to investments but rather the consumed portion. From an estimation perspective, it is simply $(1 - \psi) * \Delta y$ as opposed to $\psi \Delta y$ and are therefore perfectly correlated. In panel B, I therefore measure the impact of shocks that remains after they have been transmitted to consumption.

I repeat the same analysis with transitory shocks in table 3.4. Unlike the coefficients on permanent shocks, some of the coefficients on transitory income are statistically significant. More specifically, the coefficients on stock participation and participation in both. The coefficients on I.R.As and either (specifications 2 and 4) are not, suggesting that the significance of the coefficient is the result of stock participation. The positive coefficients indicate that participation increases following positive transitory shocks. In contrast to table 3.3, this is surprising. Indeed, one would expect transitory shocks to have less of an impact on participation. The significant impact may indicate that investors have short term horizons. However, this would imply that households have exact information on the nature of shocks their incur. Kaufmann and Pistaferri (2009) do find that households have some superior information.

Looking at panel B, I infer a high degree of transitory insurance. The share of

Table 3.4: Latent Portfolio Choice and Transitory Idiosyncratic Shocks

This table presents the impact of transitory shocks on latent portfolio choice. In panel A, I use an IV regression with $\Delta y^*_{i,t+1}$ *as an instrument for transitory shocks and latent choice as the dependent variable. Only the first stage is reported. In panel B, I first use the IV regression with consumption shocks as the dependent variable and estimate the fitted values. These fitted values, representing transitory consumption shocks, are then used in a simple linear regression with latent choice as the dependant variable. Stocks refers to latent direct participation. I.R.As refer to latent indirect participation. Both refers to households who have non-zero funds in both stocks and I.R.A accounts. Either refers to households who have non-zero funds in both stocks and I.R.A accounts. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Stocks (1)	I.R.A (2)	Both (3)	Either (4)
Panel A. Income Shocks				
Shocks : Δy	0.123^{**}	-0.044	0.120^{**}	-0.032
	(0.056)	(0.061)	(0.061)	(0.061)
Constant : α	0.015^{**}	0.000	0.002	0.008
	(0.007)	(0.007)	(0.007)	(0.007)
Panel B. Consumption Shocks Projection				
Shocks : $\widehat{\Delta c}$	2.636^{***}	2.036^{***}	2.923^{***}	2.151^{***}
	(0.545)	(0.596)	(0.595)	(0.596)
Constant : α	0.002	-0.011	-0.012	-0.004
	(0.007)	(0.008)	(0.008)	(0.008)

transitory variance transmitted to consumption is approximately 5%[16]. This would suggest that the marginal propensity to consume out of transitory shocks is low, thus leaving a larger share of shocks to invest in stocks. Alternatively, a negative shocks would decrease latent participation. In this case, a household would exit the market, divest assets, and use the proceeds to fund consumption, explaining the high degree of transitory consumption insurance. While transitory income shocks do not affect latent I.R.A participation, it appears that transitory consumption shocks does. This would imply that given a high enough degree of consumption insurance, households are more likely to participate in retirement accounts following positive transitory shocks to their income. This implies that habit formation and changes to risk aversion are not likely to drive these results. Indeed, as transitory shocks are well insured and have limited impact on consumption shocks, there is no reason to believe a positive shock would shift consumption above habit levels. However, the low marginal propensity to consume out of transitory shocks implies households with positive transitory shocks have more disposable income not to consume with but to invest with.

I have shown that average, permanent, and transitory shocks have an impact on households' portfolio choice. I then test whether these same shocks affect portfolio allocation. I test whether households allocate more to their risky assets relative to three definitions of wealth: financial wealth, wealth excluding, and including home equity. Scaling the value of stocks with financial wealth reflects the share of risk assets invested in financial assets. Scaling the value of stocks with total wealth reflects the decision to invest in risky financial assets relative to the investment set available to the household. Indeed, households may choose to invest in non-dwelling real estate or businesses with the lack of investment in stocks being the portfolio allocation choice. I consider two wealth ratios including and excluding housing wealth as Cocco (2005) show that housing investments serve as a substitute for stock investment. I first test the impact of average idiosyncratic income and consumption shocks in table 3.5. In panel A, I use idiosyncratic income shocks as the independent variable. I find the coefficients on wealth ratios using only stocks in the numerator to be statistically significant and positive. The coefficients with the value of investments in retirement accounts in the numerator are all insignificant. The significant coefficients indicate that households receiving positive idiosyncratic shocks increase their allocation into risky financial assets. This result is robust to the definition of wealth.

In panel B, I use Δc as the measure of shocks. Odd-numbered specifications are not significant while even-numbered are. Positive idiosyncratic consumption shocks being associated with an increase in stocks and I.R.A investment but not stocks investment would

[16]Furthermore, Blundell, Pistaferri, and Preston (2008) find that transitory shocks are fully insured.

Table 3.5: Latent Portfolio Allocation and Average Idiosyncratic Shocks

*This table reports the OLS estimation of equation (3.7). The dependant variables are latent allocation estimated with the framework presented in section (3.4) and equations (3.3) to (3.6). Financial wealth (in columns 1 and 2) is the sum of stocks, IRAs, checking and savings accounts, CDs, and treasury bills. Wealth (columns 3 and 4) excludes home equity while total wealth in (columns 5 and 6) includes home equity. The independent variables are Δy and Δc estimated using equations (3.1) and (3.2) for panels A and B respectively. In panel C, the effect of Ξ in equation (3.8) is removed. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Stocks / Fin. Wealth (1)	Stocks + IRA / Fin. Wealth (2)	Stocks / Wealth (3)	Stocks + IRA / Wealth (4)	Stocks / Tot. Wealth (5)	Stocks + IRA / Tot. Wealth (6)
Panel A. Income Shocks						
Shocks : Δy	0.058***	0.017	0.071***	0.030	0.062***	0.031
	(0.019)	(0.023)	(0.017)	(0.023)	(0.016)	(0.021)
Constant : α	0.011	0.003	0.013**	0.002	0.014**	−0.003
	(0.007)	(0.008)	(0.006)	(0.008)	(0.006)	(0.007)
Panel B. Consumption Shocks						
Shocks : Δc	0.015	0.068***	0.028	0.061**	0.009	0.029
	(0.021)	(0.026)	(0.019)	(0.026)	(0.018)	(0.023)
Constant : α	0.012*	0.003	0.014**	0.001	0.015**	−0.004
	(0.007)	(0.008)	(0.006)	(0.008)	(0.006)	(0.007)
Panel C. Consumption Shocks Projection						
Shocks : $\widehat{\Delta c}$	0.542***	0.158	0.665***	0.281	0.583***	0.294
	(0.179)	(0.214)	(0.162)	(0.217)	(0.154)	(0.193)
Constant : α	0.009	0.003	0.011*	0.001	0.013**	−0.004
	(0.007)	(0.008)	(0.006)	(0.008)	(0.006)	(0.007)

indicate the effect is driven by the allocation into I.R.As. This effect is likely to be driven by preference shifts. To test this assertion, I consider the transmission of income shocks to consumption shocks and use the fitted values against latent allocation in panel C. As expected, the even-numbered specification are not significantly from 0. These results are partly suggested by Bonaparte, Korniotis, and Kumar (2020). They argue that trading activity in retirement accounts is lower. It is thus unlikely to respond to shocks. However, changes in preference can be significant enough to shift the allocation into retirement accounts.

In tables 3.6 and 3.7, I consider the impact of permanent and transitory shocks on portfolio allocation. Table 3.6 shows the results of the IV regressions with permanent shocks instrumentation. It appears that permanent income shocks have no impact on the allocation of risky assets into stocks. Indeed, specifications (1), (3), and (5), which only consider the allocation of stocks, are statistically insignificant. However, the allocations ratios that take into account retirement accounts are statistically significant at the 5% level. It would appear that positive idiosyncratic labor income shocks get invested into I.R.As as opposed to individual stock accounts. This contrasts the results from table 3.5 which showed that average shocks had no impact on I.R.A allocation. I then consider the projection of permanent income shocks onto consumption shocks. In panel B, only the ratios considering stock allocation alone are statistically significant. This result suggests that the transmission of permanent shocks is an important factor of portfolio allocation.

Table 3.6: Latent Portfolio Allocation and Permanent Idiosyncratic Shocks

*This table presents the impact of permanent shocks on latent portfolio allocation. In panel A, I use an IV regression with $y^*_{i,t+1} - y^*_{i,t-2}$ as an instrument for permanent shocks and latent choice as the dependent variable. Only the first stage is reported. In panel B, I first use the IV regression with consumption shocks as the dependent variable and estimate the fitted values. These fitted values, representing permanent consumption shocks, are then used in a simple linear regression with latent choice as the dependant variable. Financial wealth (in columns 1 and 2) is the sum of stocks, IRAs, checking and savings accounts, CDs, and treasury bills. Wealth (columns 3 and 4) excludes home equity while total wealth in (columns 5 and 6) includes home equity. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Stocks / Fin. Wealth (1)	Stocks + IRA / Fin. Wealth (2)	Stocks / Wealth (3)	Stocks + IRA / Wealth (4)	Stocks / Tot. Wealth (5)	Stocks + IRA / Tot. Wealth (6)
Panel A. Income Shocks						
Shocks : Δy	0.132	0.230**	0.130	0.229*	0.121	0.214**
	(0.099)	(0.116)	(0.089)	(0.118)	(0.084)	(0.105)
Constant : α	0.007	−0.000	0.006	−0.008	0.008	−0.010
	(0.008)	(0.010)	(0.008)	(0.010)	(0.007)	(0.009)
Panel B. Consumption Shocks Projection						
Shocks : $\widehat{\Delta c}$	0.273***	0.099	0.332***	0.157	0.291***	0.165
	(0.100)	(0.117)	(0.090)	(0.119)	(0.085)	(0.106)
Constant : α	0.008	0.002	0.006	−0.006	0.008	−0.008
	(0.008)	(0.010)	(0.008)	(0.010)	(0.007)	(0.009)

Transitory shocks, however only seem to affect latent stock allocation. Indeed, the specifications that include I.R.As are not statistically significant. The coefficients on income shocks and the projection on consumption shocks are significant only for stock allocation regardless of the wealth definition. I find that average and transitory shocks do not have an impact on I.R.A allocation while permanent shocks do. This lack of significance would indicate that the majority component of labor income shock is transitory in nature. It nonetheless impacts portfolio allocation. It follows that it would only impact stocks. Indeed, I.R.As have a longer investment horizon. Permanent shocks may be associated with changes in the allocation into retirement accounts as it foretells future preference shifts.

Table 3.7: Latent Portfolio Allocation and Transitory Idiosyncratic Shocks

*This table presents the impact of transitory shocks on latent portfolio allocation. In panel A, I use an IV regression with $\Delta y^*_{i,t+1}$ as an instrument for transitory shocks and latent choice as the dependent variable. Only the first stage is reported. In panel B, I first use the IV regression with consumption shocks as the dependent variable and estimate the fitted values. These fitted values, representing transitory consumption shocks, are then used in a simple linear regression with latent choice as the dependant variable. Financial wealth (in columns 1 and 2) is the sum of stocks, IRAs, checking and savings accounts, CDs, and treasury bills. Wealth (columns 3 and 4) excludes home equity while total wealth in (columns 5 and 6) includes home equity. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Stocks / Fin. Wealth (1)	Stocks + IRA / Fin. Wealth (2)	Stocks / Wealth (3)	Stocks + IRA / Wealth (4)	Stocks / Tot. Wealth (5)	Stocks + IRA / Tot. Wealth (6)
Panel A. Income Shocks						
Shocks : Δy	0.116*	−0.033	0.121**	−0.042	0.095*	−0.060
	(0.065)	(0.076)	(0.059)	(0.077)	(0.056)	(0.069)
Constant : α	0.010	0.008	0.009	0.003	0.012*	−0.001
	(0.008)	(0.009)	(0.007)	(0.009)	(0.007)	(0.008)
Panel B. Consumption Shocks Projection						
Shocks : $\widehat{\Delta c}$	1.884***	0.783	2.268***	1.115	1.997***	1.354**
	(0.636)	(0.741)	(0.575)	(0.754)	(0.542)	(0.676)
Constant : α	0.002	0.004	−0.001	−0.003	0.002	−0.009
	(0.008)	(0.010)	(0.008)	(0.010)	(0.007)	(0.009)

These results (along with latent portfolio choice) suggest that households do participate more and increase their allocation into receiving positive idiosyncratic shocks. The transmission of these shocks is important as the proportion not consumed is allocated to risky assets. However, the investment decision and the allocation into assets is only one side of households' balance sheets. As Cocco (2005) points out, mortgage debt can be used by households as leverage to finance investments in stocks. As households' receive positive (negative) labor shocks, they may reduce (increase) their debt liability. However, not all debt have the same horizon period. As such, I consider two types of debts. First, I consider unsecured debt which is the sum of credit card charges, student loans, medical bills, legal bills, and loans from relatives. These debt items are directly related to consumption categories (education expenditure, health expenditures, etc.). As the variables are changes in residuals, this link is of no econometric consequence. The interest of this particular variable is the direction of the sign. Indeed, previous variables where expected to have positive signs. However, the nature of the variable construction means the portfolio choice and portfolio allocation variables carried similar information. Unsecured debt can be understood to be borrowing constraint variable. As such, households have the incentive to reduce their borrowing constraints as they experience positive shocks. I thus expect the impact of positive shocks to have a negative impact on latent debt allocation.

The second debt category is secured by assets and is the sum of all household mortgages as well as business, farm, and other real estate debts. The impact of shocks on secured debt is ambiguous. On one hand, positive shocks could have households pay down their mortgage with the additional income. On another hand, positive shocks could have households to increase their debts as they finance nonfinancial investments[17]. Furthermore, the implications of the transitory/permanent income process on latent debt allocation is also clearer. Indeed, I expect permanent shocks to have a limited (if any) impact on unsecured debt and transitory shocks to have a limited (if any) impact on secured debt. This can be expected for two reasons. First, transitory shocks may have the same horizon as unsecured debts. Consider a one time payment such as the 2020 Covid-19 relief payments[18]. Coibion, Gorodnichenko, and Weber (2020) and Baker et al. (2020) show that households consumed some of the payments but also used the rest to reduce their short-term debt overhang. While I do not rely on natural experiments, the IV framework I use is more generalizable to a variety of transitory shocks (bonuses, increased overtime, or negative shocks such as short-term unemployment, reduced hours, etc.).

I consider five debt ratios to reflect different balance sheets. I first scale each debt category by the household's total debt. This would reflect the households' allocation of debt regardless of debt levels. This might obscure some of the dynamics of transitory shocks on unsecured debts. Indeed, a household with high level of wealth might be less concerned by their credit card debts as low-wealth households. As such, I scale the debt variables with wealth[19]. Finally, I consider total debt to wealth. As the effect of shocks on secured and unsecured debts might have opposite effects, this ratio may reflect the relative strength of each. In table 3.8, I present the result of average idiosyncratic shocks.

In panel A of 3.8, I use the general idiosyncratic labor income shocks. The results are as expected. I find that positive idiosyncratic labor income shocks are used to reduce

[17] Implying housing investments as substitute to stock investments (Cocco (2005)).

[18] Although there were multiple payments, each payment considered as a separate event would be transitory.

[19] Note that debts are added back into the wealth definition.

Table 3.8: Latent Debt Allocation and Average Idiosyncratic Shocks

This table reports the OLS estimation of equation (3.7). The dependant variables are latent debt allocation estimated with the framework presented in section (3.4) and equations (3.3) to (3.6). Unsecured debt is the sum of credit card charges, student loans, medical bills, legal bills, and loans from relatives. Secured debt is the sum of mortgages, business, farm, and other real estate debt. Total debt is the sum of unsecured and secured debt. Total wealth is the total value of assets (not removing debt still left to pay). The independent variables are Δy *and* Δc *estimated using equations (3.1) and (3.2) for panels A and B respectively. In panel C, the effect of* Ξ *in equation (3.8) is removed. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Secured/ Total D. (1)	Unsecured/ Total D. (2)	Secured/ Total W. (3)	Unsecured/ Total W. (4)	Total D./ Total W. (5)
Panel A. Income Shocks					
Shocks : Δy	0.149^{***}	-0.107^{***}	0.075^{***}	-0.111^{***}	-0.063^{***}
	(0.016)	(0.019)	(0.023)	(0.019)	(0.021)
Constant : α	0.003	0.000	0.008	-0.004	-0.002
	(0.006)	(0.007)	(0.008)	(0.007)	(0.008)
Panel B. Consumption Shocks					
Shocks : Δc	0.256^{***}	0.024	0.156^{***}	0.013	0.110^{***}
	(0.018)	(0.022)	(0.026)	(0.021)	(0.023)
Constant : α	0.003	-0.000	0.007	-0.004	-0.003
	(0.006)	(0.007)	(0.008)	(0.007)	(0.008)
Panel C. Consumption Shocks Projection					
Shocks : $\widehat{\Delta c}$	1.394^{***}	-1.001^{***}	0.702^{***}	-1.043^{***}	-0.594^{***}
	(0.152)	(0.181)	(0.216)	(0.177)	(0.197)
Constant : α	-0.001	0.003	0.005	-0.001	-0.000
	(0.006)	(0.007)	(0.008)	(0.007)	(0.008)

household unsecured debts as suggested by Baker et al. (2020). Positive shocks, however, increase the latent allocation of secured debts. This result is consistent with Cocco (2005). I find that households use positive income shocks to lever up. The coefficient (0.149) is also almost 3 times larger than the one on stock allocation in table 3.5. This is also consistent with the fact that a large portion of households' portfolios is concentrated on household wealth. Looking at the ratio of total debt to total wealth, I find a negative coefficient consistent in magnitude with the third specification. This result suggests that average labor shocks allow household to reduce their debt liabilities. This further suggests that the effect on unsecured debt is larger. Comparing the second and fourth specifications of panel A, I find similar point estimates (-0.107 vs -0.111). This suggests that the effect of positive labor income shocks on unsecured debt is not affected by a household overall level of wealth. However, when comparing specifications (1) and (3), the wealth scale coefficient is nearly half of its counterpart (0.149 vs 0.075). This suggests that the effect on secured debt is indeed dependant on household.

In panel B, I use consumption shocks as the independent variable. Consumption shocks are significant determinants of latent secured debt allocation. The fifth specification is also significantly negative. This result suggests that unexplained consumption is not financed with unsecured debt. However, unexplained consumption growth is associated with an increase in secured debt liability. This is likely the result of preference shifts. Indeed, consumption shocks might be reflecting a change in family situation not captured by the first stage regression. The result in panel C support this claim. Indeed, I find that consumption shocks caused by labor income shocks are indeed consistent with the results of panel A. I then consider the impact of permanent and transitory shocks on latent debt allocation in tables 3.9 and 3.10.

The coefficients on permanent labor income shocks are as predicted. Indeed, I find that the impact on secured debt ratios are positive and highly significant. This indicate that households use permanent income shocks to invest in housing investment. Furthermore, I find that permanent labor income shocks have no impact on unsecured debt allocation. The effect on total debt is positive but not significantly different from 0. The point estimates in the first and third specifications are comparable (0.384 vs 0.350), suggesting wealth may not play an important role in households' debt allocation response to permanent income shocks.

In panel B, I find that not only income shocks matters, but their transmission to consumption shocks also matters. Indeed, while the statistically significant coefficients remain statistically significant, the coefficients in the unsecured specifications are now negative and statistically significant. These results suggest that permanent and positive consumption

Table 3.9: Latent Debt Allocation and Permanent Idiosyncratic Shocks

This table presents the impact of permanent shocks on latent debt allocation. In panel A, I use an IV regression with $y^*_{i,t+1} - y^*_{i,t-2}$ *as an instrument for permanent shocks and latent choice as the dependent variable. Only the first stage is reported. In panel B, I first use the IV regression with consumption shocks as the dependent variable and estimate the fitted values. These fitted values, representing permanent consumption shocks, are then used in a simple linear regression with latent choice as the dependant variable. The dependant variables are latent debt allocation estimated with the framework presented in section (3.4) and equations (3.3) to (3.6). Unsecured debt is the sum of credit card charges, student loans, medical bills, legal bills, and loans from relatives. Secured debt is the sum of mortgages, business, farm, and other real estate debt. Total debt is the sum of unsecured and secured debt. Total wealth is the total value of assets (not removing debt still left to pay). Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Secured/ Total D. (1)	Unsecured/ Total D. (2)	Secured/ Total W. (3)	Unsecured/ Total W. (4)	Total D./ Total W. (5)
Panel A. Income Shocks					
Shocks : Δy	0.384^{***}	-0.147	0.350^{***}	-0.138	0.113
	(0.081)	(0.094)	(0.114)	(0.092)	(0.102)
Constant : α	-0.003	-0.001	-0.000	-0.004	-0.012
	(0.007)	(0.008)	(0.010)	(0.008)	(0.009)
Panel B. Consumption Shocks Projection					
Shocks : $\widehat{\Delta c}$	0.514^{***}	-0.363^{***}	0.243^{**}	-0.411^{***}	-0.258^{**}
	(0.082)	(0.095)	(0.115)	(0.093)	(0.103)
Constant : α	-0.000	-0.001	0.003	-0.004	-0.010
	(0.007)	(0.008)	(0.010)	(0.008)	(0.009)

shocks are associated with an increase in secured debt and a decrease in unsecured debt. Furthermore, the effect on total debt is now statistically significant and negative indication that permanent consumption shocks are associated with a decrease in households' total liabilities.

The coefficients on transitory income are also as predicted. The coefficients in specifications (2) and (4) considering unsecured debt are statistically significant from 0 and negative. A household receiving positive transitory labor income shocks will reduce the unsecured debt liability. The coefficient in the first specification of table 3.10 is significantly positive. This result is unexpected. This would suggest that households use transitory income shocks to invest. However, the scaled by wealth coefficient is insignificant suggesting

Table 3.10: Latent Debt Allocation and Transitory Idiosyncratic Shocks

*This table presents the impact of transitory shocks on latent debt allocation. In panel A, I use an IV regression with $\Delta y^*_{i,t+1}$ as an instrument for transitory shocks and latent choice as the dependent variable. Only the first stage is reported. In panel B, I first use the IV regression with consumption shocks as the dependent variable and estimate the fitted values. These fitted values, representing transitory consumption shocks, are then used in a simple linear regression with latent choice as the dependant variable. The dependant variables are latent debt allocation estimated with the framework presented in section (3.4) and equations (3.3) to (3.6). Unsecured debt is the sum of credit card charges, student loans, medical bills, legal bills, and loans from relatives. Secured debt is the sum of mortgages, business, farm, and other real estate debt. Total debt is the sum of unsecured and secured debt. Total wealth is the total value of assets (not removing debt still left to pay). Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Secured/ Total D. (1)	Unsecured/ Total D. (2)	Secured/ Total W. (3)	Unsecured/ Total W. (4)	Total D./ Total W. (5)
Panel A. Income Shocks					
Shocks : Δy	0.108^{**}	-0.106^{*}	-0.069	-0.158^{**}	-0.144^{**}
	(0.054)	(0.064)	(0.078)	(0.063)	(0.070)
Constant : α	0.008	-0.008	0.016^{*}	-0.009	-0.005
	(0.006)	(0.008)	(0.009)	(0.007)	(0.008)
Panel B. Consumption Shocks Projection					
Shocks : $\widehat{\Delta c}$	4.106^{***}	-3.052^{***}	2.197^{***}	-3.347^{***}	-1.961^{***}
	(0.530)	(0.625)	(0.758)	(0.613)	(0.682)
Constant : α	-0.013^{*}	0.007	0.004	0.007	0.004
	(0.007)	(0.008)	(0.010)	(0.008)	(0.009)

that there is no impact on secured debt when taking into account wealth. Panel B of table 3.10 also show that the transmission of the shocks is an important factor of portfolio choice.

Overall, this results indicate that idiosyncratic labor income shocks have a significant impact on households' decision to participate in financial markets, how much to allocate to their assets, and how to manage their liabilities. Furthermore, the transmission of income shocks to consumption shocks play an integral role in these decisions. However, these results depend on the validity of the instruments chosen to represent permanent and transitory shocks. Although statistical tests show the instruments are valid, I use Kopczuk, Saez, and Song (2010), Bonaparte, Korniotis, and Kumar (2014), and Catherine, Sodini, and Zhang (2020) definition of permanent income. I calculate the permanent component of log labor disposable income as the average over three continuous data points:

$$\hat{y}_{P,it} = \frac{y_{i,t-1} + y_{i,t} + y_{i,t+1}}{3} \tag{3.12}$$

In the literature, this definition is often applied to yearly data and is thus calculated over a three-year window. The PSID data is biennial[20]. Permanent income is thus understood to be calculated over a 6-year window. This is undeniably a longer horizon. However, the instruments I use to proxy for permanent and transitory shocks are also calculated over continuous wave points as opposed to yearly data points. I use this permanent income proxy and re-estimate equations (3.1) and (3.2) to get permanent income shocks. I then repeat the analysis of table 3.3, 3.6, and 3.9 with OLS instead of IV regressions. The results are presented in table 3.11.

The results of table 3.11 are mostly consistent with previous results. The coefficients are statistically significant with expected signs. Overall, this table shows that the impact of permanent income shocks is robust to the definition of permanent shock. Panel A's results are the only one with major differences. Indeed, in table 3.3, I find that permanent income shocks do not have a statistically significant impact on latent choice. In panel of table 3.11, I do find that the coefficients are statistically significant. The fourth specification (which considers ownership in either investment vehicle) is significant in both table and the point estimates are similar.

In panel B, I use latent debt allocation as the dependant variables. As expected permanent, shocks do not impact unsecured debt allocation. Permanent shocks are however positively associated with an increase in secured debt. The coefficient on total debt is also positively significant, indicating the impact of permanent shocks on the secured portion of household debt is significant enough to show up despite the effect being diluted in overall debt. In table 3.6, only income shocks that are transmitted to consumption are found to be significantly different from 0 except for ratios including allocation into retirement accounts. In panel C of table 3.11, all latent allocation variables are found to be positively impacted by permanent shocks. In unreported results, I test the transmission of the alternatively defined permanent shocks to consumption shocks and found similar results as previously found. The transmission coefficient is found to be similar.

In a final robustness check, I sort households based on the reference person's industry. I do so for several reasons. The first one is simply because I do not control for industry in my estimations of residuals. Secondly, while I could simply reevaluate the change in residuals

[20] After 1997 which is the sample period I use.

Table 3.11: Alternative Permanent Shocks and Latent Choice, Asset, and Debt Allocation

*This table repeats the analysis of tables 3.3, 3.6, and 3.9 with the permanent income shocks proxy described by equations (3.12), (3.1), and (3.2). Panel A uses the dependent variables of table 3.3; panel B uses the dependent variables of table 3.9; and panel C uses the dependent variables of table 3.6. See the tables for variable definitions. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Participation						
	Stocks	I.R.A	Both	Either		
Shocks : Δy	0.219***	0.197***	0.197***	0.261***		
	(0.038)	(0.041)	(0.041)	(0.041)		
Constant : α	0.009	−0.007	−0.001	−0.003		
	(0.007)	(0.008)	(0.008)	(0.008)		
Panel B. Debt Shares						
	S./T.D.	U./T.D.	S./T.W.	U./T.W.	T.D./T.W.	
Shocks : Δy	0.421***	−0.011	0.316***	−0.052	0.111**	
	(0.036)	(0.042)	(0.050)	(0.041)	(0.045)	
Constant : α	−0.001	−0.004	0.001	−0.007	−0.011	
	(0.007)	(0.008)	(0.010)	(0.008)	(0.009)	
Panel C. Risky Asset Shares						
	S./F.W.	S&RA/F.W.	S./W.	S&RA/W.	S./W.+HE.	S&RA/W.+HE.
Shocks : Δy	0.197***	0.231***	0.214***	0.198***	0.179***	0.175***
	(0.044)	(0.051)	(0.040)	(0.052)	(0.037)	(0.046)
Constant : α	0.007	−0.001	0.006	−0.008	0.008	−0.009
	(0.008)	(0.010)	(0.008)	(0.010)	(0.007)	(0.009)

with industry dummies, separating households may reveal endogenous relationships. Indeed, industry selection may itself be a latent choice variable as opposed to a random selection. Consider a household with low risk aversion. They would, ceteris paribus, have higher risky asset shares. Furthermore, they may also choose a riskier career path or an industry characterized by riskier labor income paths. Betermier et al. (2012) and Fagereng, Guiso, and Pistaferri (2018) argue that these endogeneity concerns are likely to create a downward bias in estimations. By separating out industries, I can consider the specific labor-industry shocks and their respective effects on portfolio choice and allocation.

The PSID industry classification is based on the 4-digit 2010 Census Detailed Industry Codes. It is a 19 industry classification. As my sample size is limited, I regroup industries into 6 broader industries to ensure each has a sufficient number of households and observations. Industries are classified as follows: Agriculture & Fishing, Mining & Construc-

Table 3.12: Average shocks and Latent Choice, Asset, and Debt Allocation by Industries

	Agriculture (1)	Mining (2)	Manufacturing (3)	Trade (4)	Services 1 (5)	Services 2 (6)
Panel A. Stock Participation						
Shocks : Δy	0.030	0.041	0.081*	0.136***	0.066	0.097***
	(0.050)	(0.045)	(0.041)	(0.041)	(0.053)	(0.029)
Shocks : $\widehat{\Delta c}$	0.365	0.502	0.848*	1.169***	0.455	0.880***
	(0.607)	(0.543)	(0.437)	(0.352)	(0.360)	(0.259)
Panel B. Risky share						
Shocks : Δy	−0.030	0.021	0.098**	0.130***	−0.012	0.071**
	(0.058)	(0.052)	(0.049)	(0.047)	(0.062)	(0.032)
Shocks : $\widehat{\Delta c}$	−0.366	0.257	1.030**	1.119***	−0.080	0.645**
	(0.702)	(0.638)	(0.519)	(0.405)	(0.423)	(0.294)
Panel C. Secured debt share						
Shocks : Δy	0.035	0.184***	0.231***	0.128***	0.110**	0.154***
	(0.056)	(0.047)	(0.042)	(0.043)	(0.048)	(0.026)
Shocks : $\widehat{\Delta c}$	0.422	2.238***	2.433***	1.099***	0.751**	1.394***
	(0.675)	(0.574)	(0.446)	(0.366)	(0.332)	(0.234)
Panel D. Unsecured debt share						
Shocks : Δy	0.005	−0.072	−0.184***	−0.201***	−0.123**	−0.072**
	(0.061)	(0.056)	(0.048)	(0.052)	(0.057)	(0.032)
Shocks : $\widehat{\Delta c}$	0.066	−0.875	−1.933***	−1.728***	−0.843**	−0.649**
	(0.742)	(0.678)	(0.506)	(0.444)	(0.390)	(0.291)
Observations	1,944	1,818	3,117	2,543	1,681	6,266

tion, Manufacturing, Trade[21] & Food Services, Services 1[22], and Services 2[23]. In the interest of space, I consider the effect of income and consumption shocks on four variables: stock participation, risky financial asset allocation, secured debt share, and unsecured debt share. Table 3.12 present the results of average shocks.

In panel A, I consider the impact of income (and consumption shocks) on latent stock participation. The coefficients are positive across industries although not consistently significant. Agriculture, Mining and Services 1 are insignificant while coefficients on Manufacturing, Trade, and Services 2 are significant. While it could be that some industries are

[21]Wholesale and retail

[22]Finance, Real Estate, Management Support, and Warehousing

[23]Educational services, health care, arts & entertainment, and other services including public administration.

not impacted by these shocks, a keen observer will see that insignificant industries have the lowest number of observations. It is entirely possible that the small number of households renders the power of the test too low to make any determination. From the second regression of panel A (separated by the dashed line), I can infer the degree of transmission of income shocks to consumption shocks. The risk-sharing coefficients are consistent across industries and thus the results are not driven by heterogeneous insurance capabilities. The results of panel B are similar to those of the first panel and are consistent in direction and magnitude with previous results, thus alleviating some endogeneity concerns.

In panels C and D, I look at debt rations. The results statistically significant for all industries except for Agriculture in panel C. The coefficients on secured debt are positive and statistically significant at the 5% level. Manufacturing has the largest coefficient while Services 1 has the lowest (0.231 vs 0.110). Services 1 has, on the other hand, the largest transmission coefficient (14%). It is however complicated to argue for any kind of economically different impact. In panel D, the coefficients are negative and significant for four out of six industries. Overall, this table reinforces earlier results. I then use IV regressions to measure the impact of permanent and transitory shocks. Results are presented in table 3.13

As the instruments require a lead data point, the last observation of each household is drop, thus reducing further each sample group. The low power of table 3.12 is compounded in table 3.13. Some results are nonetheless significant and of expected signs. Furthermore, the transmission to consumption shocks yields the most significant results. Interestingly, panels C and D which consider debt ratios have the most significant results. This would indicate that income shocks have a stronger effect on households' debt consideration than on asset allocation, specifically financial assets. This is not entirely surprising considering the low participation puzzle.

3.6 Conclusion

The current economic literature has shown that insurance markets are neither complete nor incomplete. As such idiosyncratic income risk has important implications for consumption smoothing, household welfare, and portfolio choice. Not only do idiosyncratic shocks matter, their persistence is critical. As a consequence, the transmission mechanisms of income shocks on portfolio decisions and allocations depends on the nature of the income process.

This paper assesses the link between household income and consumption shocks and portfolio allocation. I use longitudinal survey data from the PSID to ask whether households' portfolio choices are impacted by income shocks. Households are, on average, well insured

Table 3.13: Permanent and Transitory shocks and Latent Choice, Asset, and Debt Allocation by Industries

*This table reports the IV estimation of shocks on latent variables with $y^*_{i,t+1} - y^*_{i,t-2}$ and $\Delta y^*_{i,t+1}$ as permanent and transitory instruments respectively. The dependant variables are estimated with the framework presented in section (3.4) and equations (3.3) to (3.6). Panel A uses latent stock participation as the dependent variable. Panel B uses the ratio of stocks to financial assets as the dependent variable. Panel C uses the ratio of secured debt to total debt. Panel D uses the ratio of unsecured debt to total debt. The industry classification is based on 4-digit 2010 Census Detail Industry Codes. Standard errors are in parentheses. *, **, and *** denotes statistical significance at the 10%, 5%, and 1% respectively.*

	Agriculture		Mining		Manufacturing		Trade		Services 1		Services 2	
	Permanent	Transitory	Permanent	Transitory	Permanent	Transitory	Permanent	Transitory	Permanent	Transitory	Permanent	Transitory
Panel A. Stock Participation												
Shock : Δy	0.210	-0.340^{**}	0.228	0.313^{**}	0.298	0.154	0.012	0.248^{*}	0.147	0.010	0.107	0.139
	(0.250)	(0.168)	(0.350)	(0.130)	(0.284)	(0.133)	(0.204)	(0.140)	(0.248)	(0.181)	(0.125)	(0.106)
Shock : $\widehat{\Delta c}$	-0.161	1.142	0.197	2.301	0.296^{**}	2.086^{**}	0.790^{***}	4.136^{***}	-0.013	0.943	0.548^{***}	2.852^{***}
	(0.480)	(9.911)	(0.156)	(1.612)	(0.135)	(0.828)	(0.302)	(1.434)	(0.216)	(2.931)	(0.166)	(0.805)
Panel B. Risk Share												
Shock : Δy	0.349	-0.463^{**}	0.112	0.390^{**}	0.274	0.234	0.070	0.262	0.133	-0.121	0.072	0.117
	(0.296)	(0.196)	(0.416)	(0.156)	(0.333)	(0.158)	(0.235)	(0.161)	(0.297)	(0.215)	(0.143)	(0.121)
Shock : $\widehat{\Delta c}$	-0.750	-9.445	0.125	1.353	0.482^{***}	2.688^{***}	0.756^{**}	3.803^{**}	-0.170	-2.072	0.393^{**}	1.993^{**}
	(0.564)	(11.556)	(0.187)	(1.923)	(0.158)	(0.981)	(0.348)	(1.650)	(0.258)	(3.488)	(0.190)	(0.924)
Panel C. Secured debt share												
Shock : Δy	-0.051	-0.067	0.452	0.171	0.611^{**}	0.197	0.789^{***}	0.168	0.371^{*}	-0.022	0.273^{**}	0.108
	(0.266)	(0.182)	(0.357)	(0.131)	(0.287)	(0.134)	(0.219)	(0.145)	(0.218)	(0.164)	(0.109)	(0.095)
Shock : $\widehat{\Delta c}$	-0.060	3.710	0.529^{***}	5.843^{***}	0.554^{***}	3.747^{***}	0.554^{*}	4.435^{***}	0.175	2.903	0.625^{***}	3.828^{***}
	(0.514)	(10.847)	(0.159)	(1.635)	(0.134)	(0.831)	(0.309)	(1.492)	(0.188)	(2.650)	(0.144)	(0.725)
Panel D. Unsecured debt share												
Shock : Δy	0.141	-0.006	-0.193	0.210	-0.214	-0.291^{*}	-0.188	-0.277	-0.530^{**}	-0.096	-0.060	-0.126
	(0.292)	(0.199)	(0.404)	(0.155)	(0.311)	(0.150)	(0.246)	(0.173)	(0.257)	(0.194)	(0.135)	(0.118)
Shock : $\widehat{\Delta c}$	0.301	3.097	-0.308^{*}	-3.045	-0.383^{***}	-3.179^{***}	-1.009^{***}	-7.389^{***}	-0.277	-3.514	-0.259	-1.568^{*}
	(0.564)	(11.900)	(0.181)	(1.919)	(0.148)	(0.933)	(0.364)	(1.775)	(0.220)	(3.144)	(0.179)	(0.897)
Observations	1,386	1,603	1,296	1,474	2,221	2,532	1,813	2,096	1,181	1,369	4,471	5,149

against income shocks. This limited ability indicates that households consume a non-zero amount of these shocks. What of the non-consumed share of idiosyncratic labor income shocks?

I circumvent low participation and low allocation issues in the data by estimating latent wealth variables. I tranche allocation ratios and assume portfolio choices are generated by an ordered probit. From this process, I can estimate pseudo-residuals and shocks to latent utility derived from these decisions. While economic impacts are obscured, I find the following results.

Positive income shocks are associated with positive shocks to latent participation variables. This would suggest that households receiving positive labor income shocks tend to participate in financial markets more. However, this effect is only noticeable for direct participation as latent indirect participation is not impacted by income shocks. Furthermore, transitory income shocks are associated with higher participation, only permanent consumption shocks are associated with significant changes in latent participation. This result suggests that transmission of shocks (i.e, insurance) is a significant driver of portfolio choice. Not only are households increasing their participation, they also increase their allocation towards risky assets. Negative shocks would reduce participation and reduce risk allocation; results observed in the literature.

Households invests some positive downfall towards risky assets. On the liability side, households use transitory income shocks to pay down unsecured debt. Positive shocks are associated with a decrease in latent unsecured debt allocation. However, permanent income shocks are associated with an increase in secured debt allocation. This result imply that households receiving permanent idiosyncratic labor income shocks will increase their housing investment, suggesting that stock and housing investments are substitutes. Overall, I find that as households' consumption is well insured, income shocks that are not consumed, are allocated towards a reduction of debt and investment in risky assets.

References

Addoum, Jawad M, Stefanos Delikouras, and George M Korniotis (2019). "Consumption-income sensitivity and portfolio choice". In: *The Review of Asset Pricing Studies* 9.1, pp. 91–136.

Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu (2011). "Consumer bankruptcy and default: The role of individual social capital". In: *Journal of Economic Psychology* 32.4, pp. 632–650.

Alessie, Rob, Stefan Hochguertel, and Arthur van Soest (2004). "Ownership of stocks and mutual funds: A panel data analysis". In: *Review of Economics and Statistics* 86.3, pp. 783–796.

Almenberg, Johan et al. (2020). "Attitudes toward debt and debt behavior". In: *The Scandinavian Journal of Economics*.

Attanasio, Orazio and Luigi Pistaferri (2014). "Consumption inequality over the last half century: some evidence using the new PSID consumption measure". In: *American Economic Review* 104.5, pp. 122–26.

Baker, Scott R et al. (2020). *Income, liquidity, and the consumption response to the 2020 economic stimulus payments*. Tech. rep. National Bureau of Economic Research.

Barasinska, Nataliya, Dorothea Schäfer, and Andreas Stephan (2012). "Individual risk attitudes and the composition of financial portfolios: Evidence from German household portfolios". In: *The Quarterly Review of Economics and Finance* 52.1, pp. 1–14.

Basten, Christoph, Andreas Fagereng, and Kjetil Telle (2016). "Saving and portfolio allocation before and after job loss". In: *Journal of Money, Credit and Banking* 48.2-3, pp. 293–324.

Bayer, Christian and Falko Juessen (2015). "Happiness and the persistence of income shocks". In: *American Economic Journal: Macroeconomics* 7.4, pp. 160–87.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S Goldstein (2007). "Portfolio choice over the life-cycle when the stock and labor markets are cointegrated". In: *The Journal of Finance* 62.5, pp. 2123–2167.

Betermier, Sebastien et al. (2012). "Hedging labor income risk". In: *Journal of Financial Economics* 105.3, pp. 622–639.

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). “Consumption inequality and partial insurance”. In: *American Economic Review* 98.5, pp. 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten (2016). “Consumption inequality and family labor supply”. In: *American Economic Review* 106.2, pp. 387–435.

Bogan, Vicki L and Angela R Fertig (2013). “Portfolio choice and mental health”. In: *Review of Finance* 17.3, pp. 955–992.

Bonaparte, Yosef, George M Korniotis, and Alok Kumar (2014). “Income hedging and portfolio decisions”. In: *Journal of Financial Economics* 113.2, pp. 300–324.

— (2020). “Income Risk, Ownership Dynamics, and Portfolio Decisions”. In:

Branikas, Ioannis, Harrison Hong, and Jiangmin Xu (2020). “Location choice, portfolio choice”. In: *Journal of Financial Economics* 138.1, pp. 74–94.

Brown, Sarah, Gaia Garino, and Karl Taylor (2013). “Household debt and attitudes toward risk”. In: *Review of Income and Wealth* 59.2, pp. 283–304.

Bucciol, Alessandro, Raffaele Miniaci, and Sergio Pastorello (2017). “Return expectations and risk aversion heterogeneity in household portfolios”. In: *Journal of Empirical Finance* 40, pp. 201–219.

Campbell, John Y and John H Cochrane (1999). “By force of habit: A consumption-based explanation of aggregate stock market behavior”. In: *Journal of political Economy* 107.2, pp. 205–251.

Catherine, Sylvain, Paolo Sodini, and Yapei Zhang (2020). “Countercyclical Income Risk and Portfolio Choices: Evidence from Sweden”. In: *Available at SSRN*.

Cocco, Joao F (2005). “Portfolio choice in the presence of housing”. In: *The Review of Financial Studies* 18.2, pp. 535–567.

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout (2005). “Consumption and portfolio choice over the life cycle”. In: *The Review of Financial Studies* 18.2, pp. 491–533.

Cochrane, John H (1991). “A simple test of consumption insurance”. In: *Journal of political economy* 99.5, pp. 957–976.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (2020). *How Did US Consumers Use Their Stimulus Payments?* Tech. rep. National Bureau of Economic Research.

Commault, Jeanne (2020). "Does Consumption Respond to Transitory Shocks?" In: *American Economic Journal: Macroeconomics.*

Constantinides, George M and Anisha Ghosh (2017). "Asset pricing with countercyclical household consumption risk". In: *The Journal of Finance* 72.1, pp. 415–460.

Dimmock, Stephen G and Roy Kouwenberg (2010). "Loss-aversion and household portfolio choice". In: *Journal of Empirical Finance* 17.3, pp. 441–459.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri (2018). "Portfolio choices, firm shocks, and uninsurable wage risk". In: *The Review of Economic Studies* 85.1, pp. 437–474.

Fagereng, Andreas, Martin Blomhoff Holm, and Gisle James James Natvik (2019). "MPC heterogeneity and household balance sheets". In: *Available at SSRN 3399027.*

Feenberg, Daniel and Elisabeth Coutts (1993). "An introduction to the TAXSIM model". In: *Journal of Policy Analysis and management* 12.1, pp. 189–194.

Fisher, Jonathan D and David S Johnson (2020). "Inequality and Mobility over the Past Half Century using Income, Consumption and Wealth". In: *NBER Chapters.*

Flavin, Marjorie and Takashi Yamashita (2002). "Owner-occupied housing and the composition of the household portfolio". In: *American Economic Review* 92.1, pp. 345–362.

Frame, W Scott and Lawrence J White (2004). "Empirical studies of financial innovation: lots of talk, little action?" In: *Journal of Economic Literature* 42.1, pp. 116–144.

Gervais, Martin and Paul Klein (2010). "Measuring consumption smoothing in CEX data". In: *Journal of Monetary Economics* 57.8, pp. 988–999.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (1996). "Income risk, borrowing constraints, and portfolio choice". In: *The American Economic Review*, pp. 158–172.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008). "Trusting the stock market". In: *the Journal of Finance* 63.6, pp. 2557–2600.

Guvenen, Fatih (2007). "Do stockholders share risk more effectively than nonstockholders?" In: *The Review of Economics and Statistics* 89.2, pp. 275–288.

Guvenen, Fatih et al. (2015). *What do data on millions of US workers reveal about life-cycle earnings risk?* Tech. rep. National Bureau of Economic Research.

Heaton, John and Deborah Lucas (2000). "Portfolio choice in the presence of background risk". In: *The Economic Journal* 110.460, pp. 1–26.

Heckman, James J (1979). "Sample selection bias as a specification error". In: *Econometrica: Journal of the econometric society*, pp. 153–161.

Hochguertel, Stefan (2003). "Precautionary motives and portfolio decisions". In: *Journal of Applied Econometrics* 18.1, pp. 61–77.

Iachan, Felipe S, Plamen T Nenov, and Alp Simsek (2021). "The choice channel of financial innovation". In: *American Economic Journal: Macroeconomics* 13.2, pp. 333–72.

Ivković, Zoran, Clemens Sialm, and Scott Weisbenner (2008). "Portfolio concentration and the performance of individual investors". In: *Journal of Financial and Quantitative Analysis*, pp. 613–655.

Jiang, Danling and Sonya S Lim (2018). "Trust and household debt". In: *Review of Finance* 22.2, pp. 783–812.

Kaplan, Greg and Giovanni L Violante (2010). "How much consumption insurance beyond self-insurance?" In: *American Economic Journal: Macroeconomics* 2.4, pp. 53–87.

Kaufmann, Katja and Luigi Pistaferri (2009). "Disentangling insurance and information in intertemporal consumption choices". In: *American Economic Review* 99.2, pp. 387–92.

Kézdi, Gábor and Robert J Willis (2009). "Stock market expectations and portfolio choice of American households". In:

Kimberlin, Sara, Jiyoun Kim, and Luke Shaefer (2014). "An updated method for calculating income and payroll taxes from PSID data using the NBER's TAXSIM, for PSID survey years 1999 through 2011". In: *Unpublished manuscript, University of Michigan. Accessed May* 6, p. 2016.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song (2010). "Earnings inequality and mobility in the United States: evidence from social security data since 1937". In: *The Quarterly Journal of Economics* 125.1, pp. 91–128.

Love, David A (2010). "The effects of marital status and children on savings and portfolio choice". In: *The Review of Financial Studies* 23.1, pp. 385–432.

Love, David A and Paul A Smith (2010). "Does health affect portfolio choice?" In: *Health economics* 19.12, pp. 1441–1460.

Lynch, Anthony W and Sinan Tan (2011). "Labor income dynamics at business-cycle frequencies: Implications for portfolio choice". In: *Journal of Financial Economics* 101.2, pp. 333–359.

Mace, Barbara J (1991). "Full insurance in the presence of aggregate uncertainty". In: *Journal of Political Economy* 99.5, pp. 928–956.

Massa, Massimo and Andrei Simonov (2006). "Hedging, familiarity and portfolio choice". In: *The Review of Financial Studies* 19.2, pp. 633–685.

Neri, Andrea and M Giovanna Rannalli (2011). "To misreport or not to report? The case of the Italian Survey on Household Income and Wealth". In: *Statistics in Transition new series* 12.2, pp. 281–300.

Palia, Darius, Yaxuan Qi, and Yangru Wu (2014). "Heterogeneous background risks and portfolio choice: Evidence from micro-level data". In: *Journal of Money, Credit and Banking* 46.8, pp. 1687–1720.

Primiceri, Giorgio E and Thijs Van Rens (2009). "Heterogeneous life-cycle profiles, income risk and consumption inequality". In: *Journal of monetary Economics* 56.1, pp. 20–39.

Renneboog, Luc and Christophe Spaenjers (2012). "Religion, economic attitudes, and household finance". In: *Oxford economic papers* 64.1, pp. 103–127.

Rosen, Harvey S and Stephen Wu (2004). "Portfolio choice and health status". In: *Journal of Financial Economics* 72.3, pp. 457–484.

Schmidt, Lawrence (2016). "Climbing and falling off the ladder: Asset pricing implications of labor market event risk". In: *Available at SSRN 2471342*.

Shen, Jialu (2018). "Countercyclical risks and portfolio choice over the life cycle: Evidence and theory". In: *9th Miami Behavioral Finance Conference*.

Spaenjers, Christophe and Sven Michael Spira (2015). "Subjective life horizon and portfolio choice". In: *Journal of Economic Behavior & Organization* 116, pp. 94–106.

Turley, Robert (2012). "Informative prices and the cost of capital markets". In: *Harvard University Working Paper*.

Van Praag, Bernard (2004). *The connexion between old and new approaches to financial satisfaction*. Tech. rep. CESifo Working Paper.

Vermeulen, Philip (2016). "Estimating the top tail of the wealth distribution". In: *American Economic Review* 106.5, pp. 646–50.

Viceira, Luis M (2001). “Optimal portfolio choice for long-horizon investors with nontradable labor income”. In: *The Journal of Finance* 56.2, pp. 433–470.

Wiemann, Thomas T and Robin L Lumsdaine (2020). “The Effect of Health Care Policy Uncertainty on Households’ Consumption and Portfolio Choice”. In: *Available at SSRN 3418411*.

VITA

Nicolas Duvernois was born in Agen, France, in November 1991. He received his Master of Science in Management from the ESCE Business School in May 2015. He received his Master of Business Administration from Nicholls State University in December of 2014. He joined the University of New Orleans Financial Economics Ph.D. in August 2016. He was awarded his Master of Science in Financial Economics from the University of New Orleans in May 2018 and his PhD in Financial Economics in August 2021.