

Changing income shocks or changed insurance - what determines consumption inequality?

Johannes Ludwig
Ruhr Graduate School in Economics & Ruhr-Universität Bochum

Abstract

Contrary to the implications of economic theory, consumption inequality in the US does not react to the increases in income inequality during the last three decades. This paper investigates if a change in the type of income inequality - from permanent to transitory - or a change in the ability to insure income shocks is responsible for this. A new method to impute household consumption into the PSID is used to create panel data on income and consumption for the period 1980-2010. The minimum distance investigation of covariance relationships shows that both explanations work together: the share of transitory shocks increases over time, but the capability to insure permanent and transitory shocks to income also improves. Together, these phenomena can explain the lack of an increase in consumption inequality.

JEL classification: D12, D31, E21

Johannes Ludwig
Chair of Macroeconomics
Building GC 2/149
Ruhr-Universität Bochum
D-44780 Bochum, Germany

Phone: +49 234 32 26912
Fax: +49 234 32 14453
E-Mail: johannes.ludwig@rgs-econ.de

1 Introduction

Much research has been devoted to the problem of increasing income inequality in the US since the beginning of the 1980s. Far less authors have analyzed the development of consumption inequality over the same time frame. In case it is examined, consumption inequality has mostly been found to increase surprisingly little (see e.g. Heathcote et al. 2010, for a different view see Aguiar & Bils 2013, Attanasio et al. 2013). From a theoretical perspective this lack of a response to increasing income inequality is stunning: When the variance of income increases in a permanent income-setting, households should adjust their consumption expenditure and consumption inequality should increase proportionally.

Two explanations for this phenomenon have been suggested: Krueger & Perri (2006) assume that over time households have increased their ability to insure consumption expenditures against shocks to income with the help of credit. Blundell et al. (2008), on the contrary, argue that it is the type of the income shocks that has changed - from less insurable permanent shocks to better insurable transitory shocks. The aim of this paper is to test these two explanations against each other and investigate by how much changing insurance possibilities and changing income shocks can explain the transition of income inequality to consumption inequality during the last three decades.

The development of consumption inequality over time is of its own interest, since consumption is generally considered to be a better indicator for household welfare compared to income (Meyer & Sullivan 2013b). Thus, to better understand how unequal welfare is distributed across the economy a closer look at consumption inequality is necessary. Most researchers use only income as an indicator, but to be able to determine the welfare consequences of increasing income inequality it is necessary to know how changes in income inequality translate into consumption. Therefore, this paper analyzes how well households are insured against an increase in inequality due to permanent and due to transitory income shocks. The analysis is also important for policy advice. Knowing how permanent and transitory shocks to income have affected consumption over the last three decades, allows, e.g., to analyze whether the expansion of credit to poorer households has helped to alleviate the consequences of increasing income inequality or not. Finally, estimating consumption insurance parameters is also relevant for computational macroeconomists, since Kaplan & Violante (2010) have argued that the parameters estimated by Blundell et al. should become a benchmark for macroeconomic incomplete-markets models.

Evidence on consumption inequality is scarce mainly due to data restrictions. Up to

now, the literature has heavily relied on consumption data from the Consumer Expenditure Survey (CE). However, the quality of CE data has recently been under discussion. This paper circumvents CE-related problems by exploiting information on consumption from the Panel Study of Income Dynamics (PSID) collected between 1999 and 2011. This information is used to impute consumption into the PSID for all sample years following the example of Attanasio & Pistaferri (2014). The resulting dataset is a household panel that includes information on income and consumption which is needed to study the effects of transitory and permanent income shocks on consumption.

The empirical framework of Blundell et al. (2008) is used to identify the size of transitory and permanent income shocks and the consumption insurance parameters. This framework is parsimonious and flexible - it allows that the variances of the shocks as well as the insurance parameters change over time. I use their techniques to check if the size of transitory and permanent income inequality has changed and whether the degree of insurance against permanent and transitory shocks has remained constant over the last thirty years.

The imputation procedure of Attanasio & Pistaferri (2014) yields a consumption variable that shows only a slight increase in consumption inequality compared to the drastic increase in income inequality during the same time. The minimum distance estimation of empirical covariance relationships shows that the variance of transitory shocks to household income is increasing significantly whereas the variance of permanent shocks is rather flat over the last three decades. Moreover, the parameters indicating insurance against income shocks also change substantially over time. Both, the insurance against permanent and transitory shocks improve. Permanent shocks, though, are found to be less well insured throughout the sample period, but its insurance parameter improves to a larger extent. Thus, the type of shock that is better insured becomes more important over time and any shock to household income translates less and less into consumption over time. Combining these two phenomena can explain why consumption inequality does not react much to income inequality.

I also analyze whether consumption insurance parameters vary across demographic subgroups. The results indicate that differences in insurance parameters between single and couple households and educated and less educated households are surprisingly small. Differences in the size and composition of income shocks seem to be more important.

The paper is organized as follows: Section 2 gives an overview over the related literature. The sources of data and the imputation procedure are explained in Section 3 while Section 4 introduces the methodological framework. The results are presented

in Section 5 whereas Section 6 discusses them in the light of the related literature. Finally, Section 7 concludes.

2 Related Literature

A first pioneering empirical contribution on the effect of permanent and transitory income changes on consumption has been the paper by Hall & Mishkin (1982). The authors use PSID data on income and food consumption to show that consumption reacts more strongly to permanent income changes compared to transitory ones. However, due to the data limitations of that time their study is limited to the reaction of food consumption which is a necessity and is expected to show less of a reaction compared to overall (nondurable) household consumption. Moreover, Hall & Mishkin are not able to show any trends over time.

Krueger & Perri (2006) were the first to highlight the different dynamics of income and consumption inequality in the US since the 1980s. To explain this phenomenon they develop a model where financial markets react endogenously to increasing income inequality by extending credit supply to households that are hit by large shocks. Thus, when income inequality increases over time, credit supply reacts and households do not need to decrease spending on consumption. In this way households are able to insure consumption against larger and larger income shocks. Krueger & Perri use repeated cross-section data on income and consumption from the CE survey to estimate inequality trends in the period 1980-2003 and then try to match the trend in consumption inequality via simulation. They simulate several models with different specifications for the financial sector and show that the development of consumption inequality is determined by the amount of insurance that the financial sector offers.

Blundell et al. (2008), on the contrary, emphasize that the type of income shock is crucial for the reaction of consumption. They criticize Krueger & Perri for not differentiating between transitory and permanent shocks to income which have different effects on consumption. Blundell et al. first generate a new panel dataset on income and consumption by imputing information on consumption from CE data into the PSID and then estimate empirically how consumption reacts to the two kinds of income shocks. Their results show that household consumption does not react at all to transitory income shocks, but that it is only partially insured against permanent shocks. Contrary to the assumptions of Krueger & Perri they do not find the degree of insurance to change over time. Blundell et al. show that during their period under review it is a change in the composition of aggregate income inequality - from permanent to transitory - that drives

the response of consumption inequality.

It has to be noted that Blundell et al. (2008) only examine the period 1980-1992. It is intuitive that changes in the degree of consumption insurance cannot be detected if the observation period is sufficiently short. The first contribution of this paper is to demonstrate that the results of Blundell et al. break down when the observation period is prolonged. By examining PSID data up to 2010 it becomes possible to show that both the insurance against permanent and against transitory shocks improve significantly at the aggregate level. Second, in comparison to Blundell et al. who only consider continuously married couples I use a much broader sample. I also include households headed by singles - a steadily increasing group in the economy - and analyze whether there are differences between demographic groups. Finally, I use a new imputed measure for consumption proposed by Attanasio & Pistaferri (2014) which allows to check if the results of Blundell et al. are robust to other measures of household consumption.

An important discussion in the literature on income and consumption inequality is how to distinguish consumption insurance from superior information. If an individual knows that her income will decrease at some point in the future, she is able to accumulate savings before so that at the point of time the income shock hits, it has no effect on the consumption level. Usually, the econometrician will not have information on foreseen shocks and will classify the lack of a response to the shock as a sign for well-functioning consumption insurance mechanisms. This point has been made very clear by Primiceri & van Rens (2009) who argue that most of the permanent income shocks are foreseen since consumption inequality does not react appropriately to increasing permanent income shocks. However, they also cannot identify information and insurance separately as data on income expectations are needed for that. Kaufmann & Pistaferri (2009) use an Italian dataset that contains subjective income expectations to distinguish foreseen and unforeseen shocks from each other. They can show that a huge part of transitory and about a third of permanent income shocks are indeed foreseen. In the following, I will not be able to differentiate between insurance and information due to data limitations. Thus, it is important to keep in mind that everything that is identified as consumption insurance could also be a sign of superior information. I will get back to this point in Section 6, for a detailed overview of the insurance vs. information-debate see Meghir & Pistaferri (2011).

3 The Data

One major problem encountered when trying to study empirically the response of consumption to permanent and transitory income changes is the lack of panel data on income *and* consumption. Although containing some panel elements, the Consumer Expenditure Survey that contains information on household income and consumption is primarily a repeated cross-section data set. The Panel Study of Income Dynamics - a panel data set renowned for high quality information on household income - only includes continuous information on food expenditure, but not on other consumption items for most of the time. The majority of studies concerned with consumption inequality therefore uses information from CE data (e.g. Fisher et al. 2013, Meyer & Sullivan 2013a). However, the quality of CE data has recently been under discussion. It has been shown that CE data and tables from the National Income and Product Accounts (NIPA) on aggregate consumption diverge over time. The CE Survey seems to suffer from an increasingly worsened coverage of nondurable consumption. Therefore, Aguiar & Bils (2013) as well as Attanasio et al. (2013) argue that the use of raw data from the CE yields biased results on consumption inequality. Meyer & Sullivan (2013b), in contrast, state that the largest expenditure categories in the CE are measured well and that most of the differences to NIPA data stem from definitional differences.

Attanasio & Pistaferri (2014) have recently shown how to circumvent CE-related problems by exploiting information from the PSID on consumption collected between 1999 and 2011. They use this information to impute consumption “backwards” for all sample years. In the remainder of this paper, I will follow their example and generate a panel dataset on income and consumption for the years 1978-2010.

3.1 The Panel Study of Income Dynamics

The Michigan Panel Study of Income Dynamics (PSID) is the world’s longest running household panel dataset. Since its start in 1968 it has collected very detailed information on household income as well as on socioeconomic characteristics, household composition and health status. Interviews have been conducted yearly until 1996 and biennially since 1997. For most of the time information on consumption expenditure has been confined to few variables on food expenditure (food at home, restaurant meals and value of food stamps). Since 1999, though, the PSID has started to collect more detailed information on nondurable consumption, namely expenditure on household utilities, gasoline, car maintenance, transportation, health, education and child care. These variables sum up to around 70% of NIPA nondurable consumption. In 2005 the PSID has added further

categories like expenditure on clothing, vacation and entertainment. Blundell et al. (2012) have shown that the new PSID consumption variables align well with aggregate NIPA data and do not seem to suffer from shrinking coverage.

My sample includes all years between 1978 and 2011. Households with a head who is not in working age (i.e. younger than 25 or older than 65 years) were dropped from the dataset. The central income variable used in the analysis is household net income since it is the type of income that determines household consumption and savings. To create a variable for household net income information on federal and state tax payments are necessary which are not available in the PSID in many cases. They therefore have to be simulated with the help of NBER TaxSim (Appendix A1 contains further information on sample selection, tax simulation and the construction of the central variables). All income variables are first deflated using Consumer Price Index variables from the U.S. Bureau of Labor Statistics and then equivalence-weighted according to the modified OECD scale. The final sample comprises 12,029 households and 130,692 observations.

3.2 The Imputation Procedure

Attanasio & Pistaferri (2014) have recently shown how to use the new PSID variables on nondurable consumption to create a consistent panel data set that includes income and consumption. They first pool all new consumption expenditure variables that are continuously available for the period 1999-2011 to create a measure of “net consumption” $N_{i,t}$ for household i in year t . Attanasio & Pistaferri then regress this indicator on a set of variables that are available for all years in the PSID according to:

$$\ln N_{i,t} = Z'_{i,t}\beta + p'_t\gamma + g(F_{i,t};\theta) + u_{i,t}. \quad (1)$$

$Z_{i,t}$ denotes observable socio-economic characteristics, p_t is a set of different prices and $F_{i,t}$ stands for the consumption variables that are available for all years in the broad sample (i.e. food consumption). Note that all of the variables used here need to be available for all sample years in order to be able to impute consumption backwards¹. The imputed value for total household consumption $\widehat{C}_{i,t}$ is then equal to the sum of the consumption variables that are continuously available ($F_{i,t}$) and the predictions for net consumption ($\widehat{N}_{i,t}$):

$$\widehat{C}_{i,t} = F_{i,t} + \exp \left\{ Z'_{i,t}\widehat{\beta} + p'_t\widehat{\gamma} + g \left(F_{i,t}; \widehat{\theta} \right) \right\}. \quad (2)$$

¹The PSID did not collect information on food expenditure in 1988 and 1989. Thus, total consumption expenditure cannot be imputed in these years.

For further details on the regression results and the variables included see Appendix A2.

Attanasio & Pistaferri (2014) argue that this procedure has three important advantages compared to other approaches. First of all, there is no need to rely on information from the CE which is valuable in the light of the debates on data quality. Second, due to the availability of some consumption variables over the whole time span it is possible to control for preference heterogeneity to some extent. Finally, compared to other methods imputation quality can be assessed conveniently for the last decade by comparing the forecasts with PSID values.

4 Methodological Framework

4.1 The Income Process

The empirical framework of this paper is built around the seminal work of Blundell et al. (2008). The income process of the model is specified in a well-known permanent-transitory fashion - the log of real household net income $Y_{i,t}$ can be separated into a part that is explained by observable household characteristics $X_{i,t}$ and two unexplained components:

$$\log Y_{i,t} = X'_{i,t}\varphi_t + P_{i,t} + \nu_{i,t}. \quad (3)$$

The first of these components $P_{i,t}$ collects all shocks to income that exert a permanent influence on household income whereas the latter $\nu_{i,t}$ includes all shocks that only affect $Y_{i,t}$ temporarily. Therefore, $P_{i,t}$ is modeled as a random walk with serially uncorrelated innovations $\zeta_{i,t}$:

$$P_{i,t} = P_{i,t-1} + \zeta_{i,t}. \quad (4)$$

This assures that permanent innovations to household income $\zeta_{i,t}$ do not vanish over time. The transitory component, on the contrary, needs to be modeled in a way that the effect of the innovations on household income disappears after some time. Hence, $\nu_{i,t}$ is represented by an MA(q)-process where the innovations $\epsilon_{i,t}$ are also serially uncorrelated:

$$\nu_{i,t} = \epsilon_{i,t} + \sum_{j=1}^q \theta_j \epsilon_{i,t-j}. \quad (5)$$

The order q of the MA-process is a priori unknown and will be estimated empirically.

As the two residual components are of main interest, the remainder of this article is concerned with residual income $y_{i,t}$ and the growth in residual income $\Delta y_{i,t}$:

$$y_{i,t} = \log Y_{i,t} - X'_{i,t}\varphi_t = P_{i,t} + \nu_{i,t}, \quad (6)$$

$$\Delta y_{i,t} = \zeta_{i,t} + \Delta \nu_{i,t}. \quad (7)$$

The variable $\Delta y_{i,t}$ is also needed to identify the order q of the MA-process. By assumption, the innovations $\zeta_{i,t}$ and $\epsilon_{i,t}$ are mutually uncorrelated. Then, all covariances of $\Delta y_{i,t}$ and $\Delta y_{i,t+s}$ with $s > q + 1$ are equal to zero, whereas for $s \leq q + 1$ all covariances need to be significantly different from zero. The order q can, thus, be identified by estimating $Cov(\Delta y_{i,t}, \Delta y_{i,t+s})$ for $s = 1, \dots, T - t$ and checking at which value for s the covariances cease to be significantly different from zero. Table 1 depicts the covariances for $s = 0, 1, 2, 3$ for the years 1979 - 1996. After 1996 the frequency of PSID income data is biennial so that first differences of income cease to be available.

Table 1: Variance and Autocovariance of Income Growth

Year	$Var(\Delta y_{i,t})$	$Cov(\Delta y_{i,t}, \Delta y_{i,t+1})$	$Cov(\Delta y_{i,t}, \Delta y_{i,t+2})$	$Cov(\Delta y_{i,t}, \Delta y_{i,t+3})$
1978	-	-	-	-
1979	0.1120***	-0.0377***	-0.0007	-0.0032
1980	0.1079***	-0.0460***	0.0032	-0.0010
1981	0.1182***	-0.0475***	-0.0034	0.0013
1982	0.1203***	-0.0483***	-0.0002	0.0013
1983	0.1236***	-0.0457***	-0.0021	-0.0002
1984	0.1228***	-0.0512***	-0.0018	-0.0013
1985	0.1461***	-0.0488***	-0.0081**	0.0004
1986	0.1356***	-0.0507***	-0.0030	-0.0037
1987	0.1351***	-0.0517***	0.0060*	-0.0032
1988	0.1358***	-0.0567***	-0.0029	-0.0018
1989	0.1324***	-0.0450***	0.0007	0.0014
1990	0.1302***	-0.0494***	-0.0055	-0.0041
1991	0.1389***	-0.0541***	-0.0012	0.0044
1992	0.1734***	-0.0764***	-0.0029	-0.0015
1993	0.2017***	-0.0825***	0.0016	0.0002
1994	0.1930***	-0.0680***	-0.0058	-
1995	0.1828***	-0.0664***	-	-
1996	0.1880***	-	-	-

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The values in Table 1 show that - as expected - the variances of the first difference of residual income are always positive and significantly different from zero. The first

autocovariance of income change is always negative and significantly different from zero. The second and the third autocovariance are very close to zero, sometimes positive and sometimes negative. Statistically, they cannot be distinguished from being zero in nearly all cases. This pattern is in line with $q = 0$, i.e. with a transitory income component that is a simple idiosyncratic shock and does not show any moving average-behavior. Thus, PSID data in my sample can be described best by the following income process:

$$\log Y_{i,t} = X'_{i,t}\varphi_t + P_{i,t} + \epsilon_{i,t}. \quad (8)$$

This income process has often been used for PSID data (e.g. Heathcote et al. 2010). It is a special case of the more general income process of Blundell et al. (2008) since there is no difference between the transitory shock $\nu_{i,t}$ and the transitory innovation $\epsilon_{i,t}$.

4.2 The Response of Consumption to Income Shocks

The second central equation in the model of Blundell et al. (2008) is the equation that describes how transitory and permanent shocks are transmitted into (residual) household consumption. Economic theory predicts that the reaction of household consumption is driven by the state of the financial markets of the economy. In the case of perfect insurance markets it is assumed that households can buy insurance against any kind of shock and, thus, their consumption does not react at all to transitory or permanent income shocks. In a situation where the financial markets allow only self-insurance via a non-contingent bond, households accumulate savings that enable them to buffer temporary shocks to income, but permanent shocks will affect the level of consumption. Finally, when credit markets are imperfect and households face tight borrowing constraints, transitory and permanent shocks will both be fully transmitted into consumption.

Blundell et al. (2008) do not want to presuppose any kind of financial market, but prefer to allow the amount of insurance with respect to both kinds of income shocks to be determined freely. Hence, two parameters $\phi_{i,t}$ and $\psi_{i,t}$ are introduced that govern how much of the permanent and transitory shocks is transmitted into consumption. Therefore, the change in residual consumption is determined by

$$\Delta c_{i,t} = \phi_{i,t}\zeta_{i,t} + \psi_{i,t}\epsilon_{i,t} + \xi_{i,t} \quad (9)$$

where $c_{i,t}$ is residual real household consumption, $\phi_{i,t}$ is the parameter for insurance against a permanent income shock and $\psi_{i,t}$ is the counterpart with respect to a transi-

tory shock. Both parameters theoretically lie between zero and one: If a parameter is zero, the shock does not affect consumption at all. If a parameter is one, the shock fully translates into consumption. If the parameter lies between zero and one, Blundell et al. speak of "partial insurance". Finally, the term $\xi_{i,t}$ allows for unobserved heterogeneity and captures all innovations to residual consumption growth that are not related to income changes.

Since it is very likely that the imputed consumption variable is measured with error, equation (9) needs to be expanded to allow for measurement error. Blundell et al. therefore use the following variant of equation (9) for estimation:

$$\Delta c_{i,t} = \phi_{i,t} \zeta_{i,t} + \psi_{i,t} \epsilon_{i,t} + \xi_{i,t} + u_{i,t}^c - u_{i,t-1}^c. \quad (10)$$

The advantage of this framework is its flexibility - both the variances of transitory and permanent shocks as well as the insurance parameters are in general allowed to change over time. Due to this fact the framework can also be used to analyze whether the prediction of Krueger & Perri (2006) on the development of consumption inequality can be confirmed by the data. Their model assumed that financial markets react endogenously to increasing income inequality by extending credit supply. Thus, they assume that over time the level of insurance to income shocks increases. As they do not differentiate between transitory and permanent shocks to income, it is not possible to tell which one of the two population mean values ϕ_t and ψ_t needs to decrease, but it is clear that if the mechanism that Krueger & Perri (2006) describe is present, a decrease in either ϕ_t or ψ_t needs to be found.

On the other hand, Blundell et al. (2008) stress the fact that it is the kind of shock that determines the evolution of consumption inequality. They explain the lack of a response of consumption inequality to the rise in income inequality by an ever larger fraction of well-insured transitory shocks. If their result for the 1980s is a general finding than we should find (i) $\phi_t > \psi_t$, i.e. transitory shocks are better insured than permanent shocks, and (ii) ϕ_t and ψ_t do not vary over time.

These are clearly extreme characterizations of the transmission of income inequality to consumption inequality. Of course, both explanations can be present at the same time and work together. However, this investigation makes it possible to analyze how much of the lack of response of consumption inequality to income inequality is due to changed insurance and how much is due to changing shocks.

4.3 Minimum Distance Estimation of the Model Parameters

The model parameters are estimated by collecting the variances, covariances and autocovariances of $\Delta y_{i,t}$ and $\Delta c_{i,t}$ for all t and then minimizing the distance between the empirical covariances and the parameters predicted by the parametric income-consumption model. Using the first differences of variables in estimation has a long tradition especially in labor economics that traces back to Abowd & Card (1989). Thus, it first needs to be analyzed what the two central equations (8) and (10) imply for the empirical variances and covariances. Specifying the income process with $q = 0$ has the consequence that only the first autocovariance of income growth is different from zero:

$$\begin{aligned} Var(\Delta y_{i,t}) &= Var(\zeta_{i,t}) + Var(\epsilon_{i,t}) + Var(\epsilon_{i,t-1}), \\ Cov(\Delta y_{i,t}, \Delta y_{i,t+1}) &= -Var(\epsilon_{i,t}), \\ Cov(\Delta y_{i,t}, \Delta y_{i,t+s}) &= 0 \quad \forall s > 1. \end{aligned}$$

The first autocovariance of income change identifies the variance of the transitory shocks and together with $Var(\Delta y_{i,t})$ also the variance of permanent income shocks can be identified. Hence, variances and autocovariances of the change in residual income suffice in general to identify $Var(\epsilon_{i,t})$ and $Var(\zeta_{i,t})$. However, adding also the consumption moments helps to improve the estimation.

The autocovariances of the change in residual consumption have a similar structure. Due to the consumption-martingale property changes in consumption today are related to changes in the past or the future only via measurement error. The autocovariances of (10) can be summarized by:

$$\begin{aligned} Var(\Delta c_{i,t}) &= \phi_t^2 Var(\zeta_{i,t}) + \psi_t^2 Var(\epsilon_{i,t}) + Var(\xi_{i,t}) + Var(u_{i,t}^c) + Var(u_{i,t-1}^c), \\ Cov(\Delta c_{i,t}, \Delta c_{i,t+1}) &= -Var(u_{i,t}^c), \\ Cov(\Delta c_{i,t}, \Delta c_{i,t+s}) &= 0 \quad \forall s > 1. \end{aligned}$$

The autocovariance with respect to $t + 1$ identifies the variance of measurement error in consumption $Var(u_{i,t}^c)$, but the variance of unobserved heterogeneity in consumption growth $Var(\xi_{i,t})$ can only be identified if the insurance parameters ϕ_t and ψ_t are known.

To identify ϕ_t and ψ_t the covariances between income and consumption changes are of crucial importance. The covariances between $\Delta c_{i,t}$ and $\Delta y_{i,s}$ are as follows:

$$\begin{aligned}
Cov(\Delta c_{i,t}, \Delta y_{i,t}) &= \phi_t Var(\zeta_{i,t}) + \psi_t Var(\epsilon_{i,t}), \\
Cov(\Delta c_{i,t}, \Delta y_{i,t+1}) &= -\psi_t Var(\epsilon_{i,t}), \\
Cov(\Delta c_{i,t}, \Delta y_{i,t+s}) &= 0 \quad \forall s > 1 \text{ and } s < 0.
\end{aligned}$$

It can be seen that the three sets of second moments $Cov(\Delta y_{i,t}, \Delta y_{i,t+s})$, $Cov(\Delta c_{i,t}, \Delta c_{i,t+s})$ and $Cov(\Delta c_{i,t}, \Delta y_{i,t+s})$ for all t and s suffice to identify the key model parameters ϕ_t , ψ_t , $Var(\zeta_{i,t})$, $Var(\epsilon_{i,t})$, $Var(u_{i,t}^c)$ and $Var(\xi_{i,t})$. As in principle all parameters of the model are allowed to vary over time², this method enables us to analyze how much consumption inequality is driven by changes in the insurance parameters and how much by changes in the income shock variances.

Note that the covariance relationships are set up for annual data. As PSID data has a biennial frequency after 1997, the procedure has to be slightly adjusted. The first differences $\Delta c_{i,t}$ and $\Delta y_{i,t}$ cease to be available and it has to be relied on second seasonal differences $\Delta_2 c_{i,t} = c_{i,t} - c_{i,t-2}$ and $\Delta_2 y_{i,t} = y_{i,t} - y_{i,t-2}$. However, this change does not cause major problems to the model. Appendix A3 contains more details on the estimation process and shows how to ensure identification of the parameters with biennial data.


5 Results

Figure 1 shows the development of income and consumption inequality over time where inequality is measured as the standard deviation of log-variables. In line with much of the literature income and consumption inequality are both increasing over time, but inequality of imputed consumption increases far less than the inequality of household net income. This is not only true for overall income and consumption inequality represented by solid lines, but also for their residual measures, i.e. the share of inequality that cannot be explained by observable characteristics, which are depicted by the dotted lines.

To generate income and consumption residuals an OLS-regression is run that controls for a set of observable characteristics including the number of persons in the household, employment status and education of head (and spouse if present), state of residence, ethnicity, sex and age of the household head and presence of other income earners in the household. The effect of most right-hand side variables is allowed to vary

²Since there is no economic intuition why $Var(\xi_{i,t})$ should vary over time, this parameter is treated as stationary to speed up computation.

Figure 1: Inequality of household income and consumption, overall and residual



with the respective year. R^2 for the income regression is 0.570, for the consumption regression it is 0.746 showing that a larger amount of household consumption can be explained by observable characteristics.

It can also be seen in Figure 1 that the increase in overall income inequality is nearly fully driven by an increase in residual income inequality. Overall income inequality increases by 0.232 between 1978 and 2010 whereas residual income inequality increases by 0.220. Residual consumption inequality even increases faster than overall inequality. Overall consumption inequality rises by 0.040 while residual inequality rises by 0.060 over the same time span. That is why the analysis in the following is concentrated on the residual measures of income and consumption that are the driving forces of the increase in overall income and consumption inequality in the last three decades.

5.1 Minimum Distance Results

Figure 2 depicts the development of $Var(\zeta_{i,t})$ and $Var(\epsilon_{i,t})$ at the aggregate level over time as estimated by the minimum distance procedure (estimation results and standard errors can also be found in Tables A3 and A4 in Appendix A3). It can be seen at first glance that the variance of transitory innovations which affect income only in the respective year is always larger in magnitude compared to the variance

of permanent innovations whose effect on household income is lasting. Moreover, the transitory variance is increasing steadily over time whereas it is harder to detect a general trend for the variance of permanent innovations. The latter increases during the first half of the 1980s but decreases again to the initial level by the end of the 1980s. In the 1990s the permanent variance increases substantially, but then drops below the starting level in 1998 and stays there until the end of the sample period. In total, the variance of permanent innovations is more or less flat or even slightly decreasing so that on average the difference in the levels of $Var(\zeta_{i,t})$ and $Var(\epsilon_{i,t})$ is increasing over time. Additionally, the variance of transitory innovations ends with a jump upward in 2010. This could be a consequence of the financial crisis which substantially increased (temporary) unemployment.




Figure 2 gives a first hint why consumption inequality does not increase as strongly as income inequality does. The composition of overall income inequality changes over time and the transitory part becomes more and more dominant. In case transitory shocks are easier to insure than permanent shocks, consumption will not be affected as much by an increase in the size of the former compared to a potential increase in permanent shocks by the same magnitude. However, this conclusion depends on the size of the insurance parameters. If these were constant over time, a steady increase in transitory inequality and a constant permanent inequality should lead to a steady

increase in consumption inequality. If the insurance parameters also changed over time, the development of consumption inequality could be different.

The development of the parameters of insurance against transitory and permanent shocks is shown in Figure 3 that also depicts 95% confidence bands for both parameters. To ease computation ϕ_t and ψ_t have not been estimated separately for every single year, but time groups of around 4 years have been used (1979-1982, 1983-1986, 1990-1993, 1994-1998, 1999-2002, 2003-2006, 2007-2010). These groups should suffice to determine general trends over time.

Figure 3: Development of insurance parameters over time with 95-percent confidence intervals




Figure 3 reveals some remarkable results: First, as expected above and also found by other studies, the insurance parameter ϕ_t always lies above ψ_t . This means that transitory shocks to income affect consumption to a lesser degree than permanent shocks to income in all years. Second, the insurance parameters are not constant over time at all. Both ϕ_t and ψ_t decrease steadily during the 1980s, the 1990s and most of the 2000s signaling an improvement in the insurance of income shocks. Estimated coefficients in the later periods lie outside the confidence bands of the initial period for both insurance parameters. Thus, the improvement in insurance is also statistically significant. The improvement is stronger in absolute terms for the insurance against permanent shocks whose parameter decreases from 0.325 to 0.186 between 1978 and

2010. The respective parameter for the insurance against transitory shocks declines from 0.143 to 0.064 between 1978 and 2006. Third, insurance against transitory income shocks jumps up drastically in the final period (2007-2010) that includes the financial crisis. This jump, though, indicates that the parameters also react to business cycle fluctuations.

Finally, although the degree of insurance against shocks is changing substantially over the course of the three decades, the level of ϕ_t and ψ_t between 1978 and 1993 is relatively stable. Both series are decreasing, but the differences in parameters are relatively small compared to the change over the whole sample period and they are not statistically significant. Hence, these results do not contradict the observations of Blundell et al. (2008) who find constant parameters in their study that covered the years between 1979 and 1992. Their result breaks down, though, when observations for the years 1994-2010 are added.

5.2 Subgroup Analysis


Krueger et al. (2010) and Meyer & Sullivan (2013b) have found that for households above the median consumption inequality follows income inequality closely while for households below the median there is nearly no reaction of consumption inequality to increasing income inequality. It is a priori unclear whether this divergence stems from differences in insurance or in the composition of shocks. It has been shown earlier that the composition of transitory and permanent income inequality differs significantly for demographic subgroups (Meghir & Pistaferri 2004, Ludwig 2014). However, it is intuitive to assume that demographic subgroups also differ in their ability to absorb shocks to income.

Households headed by couples usually have many different sources of income and when a shock hits e.g. the labor income of the head, this shock can be attenuated by adjusting e.g. the labor income of the spouse. Educated households on average earn higher incomes and should therefore have a better access to credit as well as a higher stock of savings to buffer shocks. Therefore, we differentiate between four different population subgroups: Highly educated couple households, less educated couple households, highly educated single households and less educated single households.³

Figure 4 depicts the trends of the insurance parameters for these subgroups. Quali-

³All individuals that have spent twelve years or less in education ("only high school") are classified as "less educated". Individuals that have thirteen years of education or more ("at least some college") are classified as "highly educated". Highly educated couple households are only those where both partners have spent more than 12 years in education.

Figure 4: Insurance parameters for demographic subgroups over time



tatively, the development in most years for most groups does not differ from the aggregate case: ϕ_t usually lies above ψ_t , ϕ_t decreases over the whole time horizon whereas ψ_t decreases until 2006 before it jumps upward. However, there are also some rather counterintuitive deviations from the aggregate: Both insurance parameters for educated couples do not seem to have a decreasing trend over time, for the less educated subgroups there are some years where transitory shocks are less well insured than permanent ones and for highly educated singles it is not ψ_t that jumps upward in the final period but ϕ_t . A possible explanation for these deviations is the lower sample size in the subgroups that could lead to less precise estimates.

The differences between the subgroups concerning the insurance against permanent shocks are in line with general intuition: Highly educated couples are the group that is insured best against permanent shocks to income in the 1980s and 1990s. Both less educated couples and singles are less well insured against permanent shocks compared to

their respective highly educated counterparts. The insurance against transitory shocks, on the contrary, does not seem to differ very much between the demographic subgroups.

Quantitatively, it is surprising how small the differences between the four groups are. Especially, the differences between couple and single households are minor. This is stunning since couples can buffer shocks to the income of one partner by adjusting the income of the other partner, a possibility that single households obviously lack. Except for the fact that educated couples are found to be better insured against permanent shocks the differences between educated and less educated households are also found to be negligible.

The variances of transitory and permanent shocks to income are not shown here for the subgroups. However, there are significant differences between groups: the magnitude of both shocks is considerably higher for singles compared to couples and permanent shocks are relatively more important for the highly educated groups. This is in line with earlier research (Ludwig 2014). Hence, the results suggest that the degree of insurance does not differ substantially between demographic groups, rather it is the size of the shocks that matters for differences in consumption inequality.

5.3 Robustness

The robustness of the findings can be assessed by varying (i) the kind of variables used as well as (ii) the observations included in the sample. The first set of robustness tests tries to figure out by how much the results depend upon the imputed consumption variable. Two kinds of alternative consumption variables are logical candidates: the real values for total household consumption of food that are available for the whole sample and the real values for household nondurable consumption which are available from 1998 on.

Using only food consumption allows to check whether the improvement in consumption insurance can also be seen for another consumption variable that is available for the whole sample period. However, an important difference to nondurable consumption arises that has already been pointed out by Blundell et al. (2008): the food consumption variable only captures a part of total nondurable consumption expenditure and the fraction of food consumption to total consumption can change over time. The estimated parameters for consumption insurance will then be a product of the true parameter and the elasticity of food consumption with respect to total consumption expenditure (e.g. $\hat{\phi}_t = \beta_t \phi_t$, where β_t is the mentioned elasticity). If household incomes and the standard of living increase on average since the 1980s, β_t will decrease over time. Then, the estimated parameters $\hat{\phi}_t$ and $\hat{\psi}_t$ will automatically decrease as well even if the true in-

insurance parameters are constant. Combining a decreasing β_t and decreasing insurance parameters yields a stronger decrease than in the benchmark case. Thus, it can be expected that estimated insurance parameters will decrease more strongly over time when food consumption is used instead of the imputed variable for nondurable consumption. Moreover, since $\beta_t < 1$, the parameters found should also be smaller.

Figure A1 in the Appendix shows estimated parameters for the aggregate sample when only food consumption is used. The qualitative development is similar to the benchmark case and quantitatively the parameters show larger decreases as has been expected before. However, the estimated parameters are not smaller than in the benchmark case, but mostly larger. This counterintuitive finding can be explained by the nature of the imputed consumption variable: As the benchmark consumption variable has been imputed, its variance is smaller in level compared to comparative real data. Since one of the main identifying equations for the insurance parameters is $Var(\Delta c_{i,t}) = \phi_t^2 Var(\zeta_{i,t}) + \psi_t^2 Var(\epsilon_{i,t}) + Var(\xi_{i,t}) + Var(u_{i,t}^c) + Var(u_{i,t-1}^c)$, it is clear that $\hat{\phi}_t$ and $\hat{\psi}_t$ will increase when $Var(\Delta c_{i,t})$ increases and the income shock variances are held constant. This effect can explain why the level of the insurance parameters is higher when real data for food consumption is used. However, the finding that the insurance parameters decrease more strongly over time in Figure A1 is exactly in line with expectations and backs the benchmark results.

The real values for nondurable consumption that are available for the period 1998-2010 can be used to assess whether the estimated insurance coefficients for this period differ from the benchmark results. Again, the variances for the real consumption variable are larger compared to imputed data. Hence, the level of the estimated coefficients should be higher compared to the benchmark case, but the change over time should be comparable. Figure A2 depicts the development of ϕ_t and ψ_t for the periods 2000/02, 2004/06, 2008/10. For ψ_t the results are in line with expectations: the level of insurance against transitory shocks is a bit lower (i.e. the estimated parameters are higher), but the change over time is parallel to the benchmark results. For ϕ_t the result is at least qualitatively similar to the benchmark case: the insurance parameter first increases in 2004 and then decreases back in 2008. Quantitatively, though, the level of ϕ_t in the initial period (2000/02) seems to be a bit too low to be in line with the benchmark case. This deviation could be due to problems related to the initial period and the fact that only a short spell of observations is available for the real consumption variable compared to the imputed data. In general, this robustness test also shows that the level of consumption insurance varies over time and thereby supports the previous findings.

I have also experimented with generating a new imputed consumption variable by

using the additional consumption expenditure variables that the PSID added in 2005 (clothing, vacation and entertainment). This broadens the "net consumption" variable, but requires to drop all observations from the years 1999-2003 to rerun the imputation procedure. The results are not shown here since the outcome of the imputation regression does not differ much from the benchmark results. Thus, the results for ϕ_t and ψ_t are also very close to the benchmark case.

The second set of robustness tests varies the amount of observations included in the dataset to check whether the results depend on sample selection. Many comparable studies exclude the SEO-subsample of the PSID that was introduced to collect information on poorer households in the US. Therefore, I will test how the exclusion of this subsample affects the results. Moreover, the minimum amount of observations required for a household to be included in the dataset is an arbitrary choice that needs to be controlled. The benchmark sample requires 4 observations on income for a household to be included in the dataset. This leads to a dataset consisting of a relatively high number of households with only few observations on income and consumption. To control that the results are not driven by households that are only part of the sample for few periods, I use three alternative samples that include only households with at least 30, 35 and 40 observations on income and consumption.

Figure A3 compares estimated insurance parameters for the benchmark and for the sample that excludes the SEO-subsample. Figure A4 shows the benchmark case and the sample that excludes households with less than 30 observations on income and consumption. The respective figures for a minimum of 35 and 40 observations are not depicted here, but look very similar. Both variations in sample selection do not lead to substantially different insurance parameters. In most cases the estimates are very close to each other. These results show that the improvement in consumption insurance is not a product of the sample selection used here. In general, the key findings of the paper seem to be very robust with respect to changes in the consumption variable and the sample selection.

6 Discussion

Figure 2 has shown that transitory shocks to household net income increase in size over time while permanent shocks are roughly constant. Figure 3, though, made clear that a changing composition of income shocks - from less insurable permanent to transitory shocks - is not enough to explain the development of consumption inequality in the last three decades. The ability of households to insure their consumption expenditures

against shocks to income changes as well. Both, insurance against transitory and against permanent shocks improve significantly over time (at least until 2006). By combining the evidence from the two figures it becomes obvious why the dramatic increase in income inequality is not mirrored in consumption inequality. The two explanations offered in the literature are present at the same time: Transitory shocks become more important as suggested by Blundell et al. (2008), but the insurability of all income shocks also improves over time as suggested by Krueger & Perri (2006).

The robustness tests have revealed that the exact level of consumption insurance is very sensitive to the kind of consumption variable used. Imputed variables have a lower variance compared to real data and therefore yield smaller insurance coefficients. Hence, it is surprising that the quantities estimated by Blundell et al. (2008) who also use an imputed consumption variable, have received so much attention (e.g. Kaplan & Violante 2010). This paper has focused rather on the trends of consumption insurance parameters over time which seem to be relatively robust to changes in the consumption variable. Contrary to the findings of Blundell et al., both insurance coefficients fluctuate substantially over the course of the thirty years studied.

Can the results potentially also be explained by superior information? As described above, it is hard to differentiate between consumption insurance and superior information of the household with the given data. The level of insurance found for both kind of shocks is relatively high and a part of what is identified as “insurance” is probably information. However, it is hard to argue that there is a trend in superior information in the last thirty years. It seems implausible that households have significantly more information on their stochastic income processes today compared to the 1980s or that the income shocks itself have become more predictable. Thus, the main finding of the paper that consumption insurance improves over time, cannot be explained by superior information.

Krueger & Perri (2006) hypothesized that an expansion of credit supply is the reason for the trend in consumption inequality. The empirical framework of this paper, though, did not allow to identify the exact channel through which consumption insurance has improved. Hence, we are not able to verify the credit supply-hypothesis within our analysis, but there are some empirical findings which are very well in line with it. First, the insurance against transitory shocks worsens dramatically in the period 2007-2010. This finding can be explained by a drying up of credit induced by the financial crisis that makes it harder to buffer temporary income losses. Second, it is well-known that credit to poorer households has been expanded since the 1980s, especially during the 2000s. The subgroup analysis shows that highly educated couples are insured better against

permanent shocks compared to all other groups in the 1980s and early 1990s. The other groups, though, converge to the level of highly educated couples over time. This could be explained by a convergence in credit supply for all demographic subgroups. Thus, empirical evidence supports the credit supply-hypothesis which offers the best explanation for the increase in consumption insurance so far. Future research, however, should be concerned with ways to identify the precise channels that affect consumption insurance parameters.

7 Conclusion

This paper has analyzed the joint development of income and consumption inequality in the US over the last 30 years. By using a new method for imputing consumption into PSID data developed by Attanasio & Pistaferri (2014) a panel dataset that contains income and consumption could be created. This data has been used to analyze by how much consumption reacts to transitory and permanent shocks to income and whether this reaction has changed over the course of the past 30 years. Many other papers before have found that consumption inequality does not seem to react very much to increases in income inequality and the reasons for this phenomenon are still debated. While Krueger & Perri (2006) believe that the ability of households to insure consumption against income shocks has improved due to a better credit supply, Blundell et al. (2008) argue that it is the changing composition of income shocks that determines consumption inequality.

The results of this paper show that inequality of imputed consumption is also characterized by a relatively small increase since the 1980s whereas income inequality is steadily rising over the same time span. The minimum distance estimation of the model parameters shows that both explanations contribute to the development of consumption inequality. On the one hand the composition of shocks to income changes over time. The variance of transitory shocks increases constantly while the variance of permanent shocks is rather flat. On the other hand the capability of households to insure their consumption against income shocks improves over time. Both the ability to buffer permanent and transitory shocks increases so that shocks to household income are to a lesser extent transmitted into consumption. Taken together, these two developments explain why consumption inequality is relatively flat even in face of increasing income inequality.

Appendix A1: Sample Selection and Variable Construction

After merging all observations of the PSID waves between 1978 and 2011, the Latino and the Immigrant subsample is dropped. Since they are both only available for some years of the sample, its inclusion would change the sample composition for these years substantially and make comparisons to other years more difficult. Especially, the imputation procedure could yield erroneous results if the Immigrant sample was not left out in the regression and the results were projected onto years before 1996.

As the main focus of the work is to identify the households response to income shocks stemming from labor market risk, households with a head who is not in working age (i.e. where the head is younger than 25 or older than 65) are excluded from the sample. Moreover, observations are dropped when income or consumption values are drastic outliers. Income outliers are those values where household net income increases by more than 500% or by less than -80% compared to the period before. Observations where household net income, household food expenditure or the sum of all expenditure variables that were added in 1999 (household utilities, gasoline, car maintenance, transportation, health, education and child care) is less than 100\$ a year are also identified as outliers and excluded. Households with a missing value for the educational attainment of the household head are dropped as well.

Finally, households that have less than four observations on household net income are excluded. This is done to prevent increases in income variation stemming from households that are only part of the dataset for one or two periods and for which there are no observations on continuous first differences. The decision to exclude those with exactly less than four observations is, of course, somewhat arbitrary. The reason is that four periods (i.e. three first differences) suffice to identify the key parameters of the study. However, robustness tests are conducted to see by how much the results depend on this decision. Table A1 summarizes how many observations are dropped in the process of sample selection.

One of the two central variables of the paper is household net income. To construct an accurate measure of annual household net income all top-coded income values in the PSID are set to missing. All income values that report weekly or monthly earnings are adjusted to annual values. Some of these adjusted values are unrealistically high so that they have to be dropped them from the sample. Household net income is then defined as total family money minus the federal and state tax payments of head and wife and the federal tax payments of other family unit members.

Since federal tax payments are only available in the PSID up to 1991 and state taxes are not available at all, the variables have to be simulated. The simulation is

Table A1: Number of total observations in the dataset by step of sample selection

Reason for exclusion	# dropped	# remain
Initial Sample		205,860
Latino & Immigrant Subsample	16,272	189,588
Head not in Working Age	40,652	148,936
Income & Consumption Outliers	4,852	144,084
No information on head's education	3,497	140,587
Too few income observations	9,895	130,692

realized with the help of NBER TaxSim (Feenberg & Coutts 1993). Unfortunately, it is not possible to simulate tax payments for other family unit members since there is too few information available in the PSID. Thus, federal tax payments of other family unit members can only be deducted up to 1991. However, as this variable is usually very small or zero, household net income should not be affected too much by this change in definition.

Household net income and consumption expenditure are then deflated using the corresponding Consumer Price Index (CPI) variables taken from the U.S. Bureau of Labor Statistics and, thus, are referring to prices of the year 1979. All variables are also equivalence weighted by the modified OECD scale to control for differences in family composition. All income variables of the PSID are referring to the previous year. I follow most of the literature by also treating consumption expenditure variables as referring to the previous year. Hence, yearly income data is available for the years 1978-1996 and biennial data for the years 1998, 2000, 2002, 2004, 2006, 2008 and 2010. (Imputed) Consumption data is available for the same years except for 1987 and 1988 where food consumption variables have not been collected.

Appendix A2: Results of the Imputation Regression

Imputation quality of the procedure by Attanasio & Pistaferri (2014) depends on the validity of the regression equation (1). I have tried to mirror the imputation procedure of Attanasio & Pistaferri as close as possible. However, due to slight differences in sample selection and the treatments of outliers and top-coded income values minor differences in the regression results can be expected. Table A2 summarizes and compares the regression results of Attanasio & Pistaferri and of this paper.

Table A2: Comparison of Regression Results

	A&P (2014)		This paper	
Hours worked/100	0.0057	***	0.0073	***
Family size	0.1250	***	0.0878	***
Education dummy 1	-0.4343	***	-0.4657	***
Education dummy 2	-0.3252	***	-0.3369	***
Education dummy 3	-0.1815	***	-0.1676	***
White	0.0373		0.0559	**
Black	0.0059		-0.0553	**
CPI	0.0088	***	0.0060	***
CPI_Food at home	-0.0092	***	-0.0069	***
CPI_Food away	-0.0057	***	-0.0052	***
CPI_Rent	0.0086	***	0.0077	***
Food consumption	0.0859	***	0.0629	***
Food consumption ²	-0.0026	***	-0.0015	***
Food consumption ³	0.0254	***	0.0142	***
Home owner dummy	0.3271	***	0.2047	***
N	26,815		28,846	
R ²	0.5058		0.5392	

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

From a qualitative perspective nearly all point estimates of the two regressions share the same algebraic sign. The only important difference is that the ethnicity dummies (*White*, *Black*) are not statistically significant in the A&P-sample, but in my sample. The dummy *Black* has a negative impact on net consumption in this paper while it is very small, but positive in the sample of Attanasio & Pistaferri. Quantitatively, most of the values are rather close to the results of Attanasio & Pistaferri. Besides the ethnicity dummies only the variable for the family size and the dummy for home ownership show larger discrepancies. However, the comparison shows that there are no systematic differences between the regression results of Attanasio & Pistaferri (2014) and the results of this paper. The robustness test using the real consumption values for the years 1998-2010 also shows that the imputed values are reasonably close to the true values (see Figure A2).

Appendix A3: Minimum Distance Estimation

Treatment of biennial data

For the years after 1996 yearly observations cease to be available and the dataset changes to a biennial frequency. Thus, first differences cannot be constructed and it has to be relied on second (seasonal) differences. The second difference of residual income is defined as follows:

$$\begin{aligned}\Delta_2 y_{i,t} &= y_{i,t} - y_{i,t-2} = P_{i,t} + \epsilon_{i,t} - (P_{i,t-2} + \epsilon_{i,t-2}) \\ &= \zeta_{i,t} + \zeta_{i,t-1} + \epsilon_{i,t} - \epsilon_{i,t-2}.\end{aligned}$$

Using this definition it is again possible to set up covariance relationships. Note that due to the biennial structure $Cov(\Delta_2 y_{i,t}, \Delta_2 y_{i,t+s})$ is only available for $s = 0, 2, 4, \dots$. The structure of the resulting covariances is very similar to the case of yearly data:

$$\begin{aligned}Var(\Delta_2 y_{i,t}) &= Var(\zeta_{i,t}) + Var(\zeta_{i,t-1}) + Var(\epsilon_{i,t}) + Var(\epsilon_{i,t-2}), \\ Cov(\Delta_2 y_{i,t}, \Delta_2 y_{i,t+2}) &= -Var(\epsilon_{i,t}), \\ Cov(\Delta_2 y_{i,t}, \Delta_2 y_{i,t+s}) &= 0 \quad \forall s > 2.\end{aligned}$$

Unfortunately, it is not possible to identify the single yearly variance of permanent innovations $Var(\zeta_{i,t})$, only the sum of the variance in year t and the previous year ($Var(\zeta_{i,t}) + Var(\zeta_{i,t-1})$) can be determined. Thus, it has to be assumed that the variances are equal for the respective years. The variance of growth in residual income therefore becomes:

$$Var(\Delta_2 y_{i,t}) = 2Var(\zeta_{i,t}) + Var(\epsilon_{i,t}) + Var(\epsilon_{i,t-2})$$

Hence, the depicted variances in the years after 1996 always represent arithmetic means of two years. However, this should suffice to identify a broad trend in the permanent component.

The second difference of residual consumption is defined as

$$\begin{aligned}\Delta_2 c_{i,t} &= c_{i,t} - c_{i,t-2} = (c_{i,t} - c_{i,t-1}) + (c_{i,t-1} - c_{i,t-2}) = \Delta c_{i,t} + \Delta c_{i,t-1} \\ &= \phi_{i,t} \zeta_{i,t} + \phi_{i,t-1} \zeta_{i,t-1} + \psi_{i,t} \epsilon_{i,t} + \psi_{i,t-1} \epsilon_{i,t-1} + \xi_{i,t} + \xi_{i,t-1} + u_t^c - u_{t-2}^c.\end{aligned}$$

If it is assumed that the insurance coefficients do not differ for t and $t - 1$, we receive

the following expression:

$$\Delta_2 c_{i,t} = \phi_{i,t} (\zeta_{i,t} + \zeta_{i,t-1}) + \psi_{i,t} (\epsilon_{i,t} + \epsilon_{i,t-1}) + \xi_{i,t} + \xi_{i,t-1} + u_t^c - u_{t-2}^c.$$

For the variances and covariances this implies (again assuming also that $Var(\zeta_t) = Var(\zeta_{t-1})$, $Var(\epsilon_t) = Var(\epsilon_{t-1})$ and $Var(\xi_t) = Var(\xi_{t-1})$):

$$\begin{aligned} Var(\Delta_2 c_{i,t}) &= \phi_t^2 (2Var(\zeta_t)) + \psi_t^2 (2Var(\epsilon_t)) + 2Var(\xi_t) + Var(u_t^c) + Var(u_{t-2}^c), \\ Cov(\Delta_2 c_{i,t}, \Delta_2 c_{i,t+2}) &= -Var(u_t^c), \\ Cov(\Delta_2 c_{i,t}, \Delta_2 c_{i,t+s}) &= 0 \quad \forall s > 2. \end{aligned}$$

as well as

$$\begin{aligned} Cov(\Delta_2 c_{i,t}, \Delta_2 y_{i,t}) &= \phi_t (2Var(\zeta_t)) + \psi_t Var(\epsilon_t), \\ Cov(\Delta_2 c_{i,t}, \Delta_2 y_{i,t+2}) &= -\psi_t Var(\epsilon_t), \\ Cov(\Delta_2 c_{i,t}, \Delta_2 y_{i,t+s}) &= 0 \quad \forall s > 2 \text{ and } s < 0. \end{aligned}$$

The structure of the covariance relationships is analogous to the case of annual data. Thus, all important parameters can also be identified when only biennial data is available.

Estimation Procedure

For every household in the dataset the following column-vector of observations containing the first and second differences of residual consumption and income is collected:

$$x_i = \begin{pmatrix} \Delta c_{i,1979}, \Delta c_{i,1980}, \dots, \Delta c_{i,1986}, \Delta c_{i,1990}, \dots, \Delta c_{i,1996}, \Delta_2 c_{i,1998}, \dots, \Delta_2 c_{i,2010}, \\ \Delta y_{i,1979}, \Delta y_{i,1980}, \dots, \Delta y_{i,1996}, \Delta_2 y_{i,1998}, \dots, \Delta_2 y_{i,2010} \end{pmatrix}'$$

For households that are present in the dataset throughout the whole sample period without missing observations, this yields 22 observations on consumption growth and 25 observations on income growth. Hence, $dim(x_i) = 47$. Missing values of income or consumption growth are set equal to zero. To control for missing values another vector d_i of the same dimension is constructed for every household in the dataset. d_i contains ones in years where the observations are non-missing and zeros for years with missing values, respectively.

With the help of x_i and d_i a symmetric 47x47-matrix M can be set up:

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

where \otimes represents an elementwise division. M contains estimates of all variances, covariances and autocovariances of income and consumption growth:

$$M = \begin{pmatrix} \text{Var}(\Delta c_{i,1979}) \\ \text{Cov}(\Delta c_{i,1979}, \Delta c_{i,1980}) \\ \dots \\ \text{Cov}(\Delta c_{i,1979}, \Delta_2 c_{i,2008}) & \dots & \text{Var}(\Delta_2 c_{i,2008}) \\ \text{Cov}(\Delta c_{i,1979}, \Delta_2 c_{i,2010}) & \dots & \text{Cov}(\Delta_2 c_{i,2008}, \Delta_2 c_{i,2010}) & \text{Var}(\Delta_2 c_{i,2010}) & \dots \\ \text{Cov}(\Delta c_{i,1979}, \Delta y_{i,1979}) & \dots & \text{Cov}(\Delta_2 c_{i,2008}, \Delta y_{i,1979}) & \text{Cov}(\Delta_2 c_{i,2010}, \Delta y_{i,1979}) & \dots \\ \text{Cov}(\Delta c_{i,1979}, \Delta y_{i,1980}) & \dots & \text{Cov}(\Delta_2 c_{i,2008}, \Delta y_{i,1980}) & \text{Cov}(\Delta_2 c_{i,2010}, \Delta y_{i,1980}) & \dots \\ \dots & \dots & \dots & \dots & \dots \\ \text{Cov}(\Delta c_{i,1979}, \Delta_2 y_{i,2008}) & \dots & \text{Cov}(\Delta_2 c_{i,2008}, \Delta_2 y_{i,2008}) & \text{Cov}(\Delta_2 c_{i,2010}, \Delta_2 y_{i,2008}) & \dots \\ \text{Cov}(\Delta c_{i,1979}, \Delta_2 y_{i,2010}) & \dots & \text{Cov}(\Delta_2 c_{i,2008}, \Delta_2 y_{i,2010}) & \text{Cov}(\Delta_2 c_{i,2010}, \Delta_2 y_{i,2010}) & \dots \\ \dots \\ \dots & \text{Var}(\Delta y_{i,1979}) \\ \text{Cov}(\Delta y_{i,1979}, \Delta y_{i,1980}) & \text{Var}(\Delta y_{i,1980}) \\ \dots & \dots \\ \text{Cov}(\Delta y_{i,1979}, \Delta_2 y_{i,2008}) & \text{Cov}(\Delta y_{i,1980}, \Delta_2 y_{i,2008}) & \dots & \text{Var}(\Delta_2 y_{i,2008}) \\ \text{Cov}(\Delta y_{i,1979}, \Delta_2 y_{i,2010}) & \text{Cov}(\Delta y_{i,1980}, \Delta_2 y_{i,2010}) & \dots & \text{Cov}(\Delta_2 y_{i,2008}, \Delta_2 y_{i,2010}) & \text{Var}(\Delta_2 y_{i,2010}) \end{pmatrix}$$

For the minimum distance estimation the matrix M is half-vectorized according to $m = \text{vech}(M)$ so that the resulting vector m is of dimension $\frac{47 \cdot 48}{2} = 1128$ and contains all unique second moments of the data.

$$m = \begin{pmatrix} \text{Var}(\Delta c_{i,1979}) \\ \text{Cov}(\Delta c_{i,1979}, \Delta c_{i,1980}) \\ \dots \\ \text{Cov}(\Delta c_{i,1979}, \Delta_2 y_{i,2010}) \\ \dots \\ \text{Var}(\Delta y_{i,1979}) \\ \text{Cov}(\Delta y_{i,1979}, \Delta y_{i,1980}) \\ \dots \\ \text{Var}(\Delta_2 y_{i,2010}) \end{pmatrix}.$$

The estimation procedure minimizes the distance between the empirical second moments in the vector m and a vector $f(\theta)$ that contains the predictions of the parametric income model for these moments. The vector $f(\theta)$ is a function of the vector θ that contains the parameters of the income model that we seek to estimate (ϕ_t , ψ_t , $\text{Var}(\zeta_{i,t})$, $\text{Var}(\epsilon_{i,t})$, $\text{Var}(u_{i,t}^c)$, $\text{Var}(\xi_i)$):

$$f(\theta) = \begin{pmatrix} \phi_{1979}^2 \text{Var}(\zeta_{i,1979}) + \psi_{1979}^2 \text{Var}(\epsilon_{i,1979}) + \text{Var}(\xi_i) + \text{Var}(u_{i,1979}^c) + \text{Var}(u_{i,1978}^c) \\ -\text{Var}(u_{1979}^c) \\ \dots \\ 0 \\ \dots \\ \text{Var}(\zeta_{1979}) + \text{Var}(\epsilon_{1979}) + \text{Var}(\epsilon_{1978}) \\ -\text{Var}(\epsilon_{1979}) \\ \dots \\ 2\text{Var}(\zeta_{2009/10}) + \text{Var}(\epsilon_{2010}) + \text{Var}(\epsilon_{2008}) \end{pmatrix}.$$

Equally weighted minimum distance (EWMD) estimates the parameter vector θ by

$$\min_{\theta} (m - f(\theta))' (m - f(\theta)).$$

Standard errors for θ can be estimated as shown by Chamberlain (1984):

$$\widehat{\text{Var}}(\hat{\theta}) = (G'G)^{-1} G'VG(G'G)^{-1}$$

where $G = \frac{\partial f(\theta)}{\partial \theta'} \Big|_{\theta=\hat{\theta}}$ and V is the variance-covariance matrix of m .

Estimation Results

Table A3: Results of the Minimum Distance Estimation - Variances of Permanent and Transitory Innovations (Standard Errors in Parantheses)

$Var(\zeta_t)$				$Var(\epsilon_t)$			
		92	0.0412	$t = 78$	0.0563	92	0.0829
			(0.0063)		(0.0063)		(0.0067)
		93	0.0504	79	0.0438	93	0.0836
			(0.0071)		(0.0039)		(0.0063)
$t = 79/80$	0.0316	94	0.0490	80	0.0505	94	0.0720
	(0.0041)		(0.0069)		(0.0043)		(0.0065)
81	0.0321	95	0.0549	81	0.0527	95	0.0691
	(0.0043)		(0.0075)		(0.0045)		(0.0058)
82	0.0359	96	0.0576	82	0.0526	96	0.0769
	(0.0046)		(0.0066)		(0.0041)		(0.0062)
83	0.0415	97/98	0.0293	83	0.0520	98	0.0691
	(0.0055)		(0.0031)		(0.0043)		(0.0055)
84	0.0355	99/00	0.0221	84	0.0543	00	0.0670
	(0.0046)		(0.0030)		(0.0047)		(0.0053)
85	0.0501	01/02	0.0274	85	0.0570	02	0.0722
	(0.0055)		(0.0032)		(0.0047)		(0.0053)
86	0.0376	03/04	0.0295	86	0.0560	04	0.0693
	(0.0047)		(0.0057)		(0.0047)		(0.0053)
87	0.0338	05/06	0.0263	87	0.0568	06	0.0660
	(0.0055)		(0.0032)		(0.0049)		(0.0056)
88	0.0343	07/10	0.0262	88	0.0629	08	0.0722
	(0.0056)		(0.0033)		(0.0051)		(0.0060)
89	0.0328			89	0.0544	10	0.0958
	(0.0052)				(0.0050)		(0.0092)
90	0.0384			90	0.0537		
	(0.0053)				(0.0049)		
91	0.0387			91	0.0586		
	(0.0054)				(0.0054)		

Table A4: Results of the Minimum Distance Estimation - Insurance Parameters (Standard Errors in Parantheses)

ϕ_t		ψ_t		$Var(\xi_t)$	
$t = 79/82$	0.3247	$t = 79/82$	0.1428		0.0038
	(0.0375)		(0.0217)		(0.0004)
83/86	0.3088	83/86	0.1223		
	(0.0311)		(0.0211)		
90/93	0.2854	90/93	0.1092		
	(0.0374)		(0.0191)		
94/98	0.2489	94/98	0.0834		
	(0.0255)		(0.0173)		
99/02	0.1829	99/02	0.0934		
	(0.0362)		(0.0260)		
03/06	0.2146	03/06	0.0644		
	(0.0425)		(0.0273)		
07/10	0.1860	07/10	0.1560		
	(0.0533)		(0.0361)		

Appendix A4: Robustness

Figure A1: Robustness test: Insurance parameters when consumption of food only is used as variable of interest




Figure A2: Robustness test: Insurance parameters when actual nondurable consumption is used for the years 1998-2010




Figure A3: Robustness test: Insurance parameters when SEO-subsample is dropped





Figure A4: Robustness test: Insurance parameters when 30 observations per household are required



References

- Abowd, J. M. & Card, D. (1989). On the Covariance Structure of Earnings and Hours Changes, *Econometrica* **57**(2): 411–445.
- Aguiar, M. & Bils, M. (2013). Has Consumption Inequality Mirrored Income Inequality? Working Paper.
- Attanasio, O., Hurst, E. & Pistaferri, L. (2013). The Evolution of Income, Consumption, and Leisure Inequality in the US, 1980-2010, *Improving the Measurement of Consumer Expenditures*, University of Chicago Press.
- Attanasio, O. & Pistaferri, L. (2014). Consumption Inequality Over the Last Half Century. Some Evidence using the New PSID Consumption Measure, *American Economic Review Papers & Proceedings* **104**(5): 122–126.
- Blundell, R., Pistaferri, L. & Preston, I. (2008). Consumption Inequality and Partial Insurance, *American Economic Review* **98**(5): 1887–1921.
- Blundell, R., Pistaferri, L. & Saporta-Eksten, I. (2012). Consumption inequality and family labor supply. National Bureau of Economic Research Working Paper 18445.
- Chamberlain, G. (1984). Panel data, *Handbook of econometrics*, Vol. 2, Elsevier, pp. 1247–1318.
- Feenberg, D. & Coutts, E. (1993). An introduction to the TAXSIM model, *Journal of Policy Analysis and Management* **12**(1): 189–194.
- Fisher, J. D., Johnson, D. S. & Smeeding, T. M. (2013). Measuring the Trends in Inequality of Individuals and Families: Income and Consumption, *The American Economic Review* **103**(3): pp. 184–188.
- Hall, R. E. & Mishkin, F. S. (1982). The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households, *Econometrica* **50**(2): pp. 461–481.
- Heathcote, J., Perri, F. & Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the United States, 1967-2006, *Review of Economic Dynamics* **13**(1): 15 – 51.
- Kaplan, G. & Violante, G. L. (2010). How Much Consumption Insurance Beyond Self-Insurance?, *American Economic Journal: Macroeconomics* **2**(4): 53–87.

- Kaufmann, K. & Pistaferri, L. (2009). Disentangling insurance and information in intertemporal consumption choices, *American Economic Review Papers & Proceedings* **99**(2): 387–392.
- Krueger, D. & Perri, F. (2006). Does income inequality lead to consumption inequality? Evidence and theory, *The Review of Economic Studies* **73**(1): 163–193.
- Krueger, D., Perri, F., Pistaferri, L. & Violante, G. L. (2010). Cross-sectional facts for macroeconomists, *Review of Economic Dynamics* **13**(1): 1 – 14.
- Ludwig, J. (2014). The role of education and household composition for transitory and permanent income inequality: Evidence from PSID data. Ruhr Economic Papers, No. 478.
- Meghir, C. & Pistaferri, L. (2004). Income Variance Dynamics and Heterogeneity, *Econometrica* **72**(1): 1–32.
- Meghir, C. & Pistaferri, L. (2011). Earnings, Consumption and Life Cycle Choices, *in* D. Card & O. Ashenfelter (eds), *Handbook of Labor Economics*, Vol. 4, Part B, Chapter 9, Elsevier, pp. 773 – 854.
- Meyer, B. D. & Sullivan, J. X. (2013a). Consumption and Income Inequality and the Great Recession, *The American Economic Review* **103**(3): pp. 178–183.
- Meyer, B. D. & Sullivan, J. X. (2013b). Consumption and Income Inequality in the U.S. Since the 1960s. Working Paper.
- Primiceri, G. E. & van Rens, T. (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality, *Journal of Monetary Economics* **56**(1): 20 – 39.