RUHR
ECONOMIC PAPERS

Johannes Ludwig

Is there Heterogeneity in the Response of Consumption to Income Shocks?

Ruhr Graduate School in Economics
RUHR UNIVERSITÄT BOCHUM
RUB

#595

Imprint

Ruhr Economic Papers

Published by

Ruhr-Universität Bochum (RUB), Department of Economics
Universitätsstr. 150, 44801 Bochum, Germany

Technische Universität Dortmund, Department of Economic and Social Sciences
Vogelpothsweg 87, 44227 Dortmund, Germany

Universität Duisburg-Essen, Department of Economics
Universitätsstr. 12, 45117 Essen, Germany

Rheinisch-Westfälisches Institut für Wirtschaftsforschung (RWI)
Hohenzollernstr. 1-3, 45128 Essen, Germany

Editors

Editorial Office

Ruhr Economic Papers #595

Responsible Editor: Thomas Bauer

ISSN 1864-4872 (online) – ISBN 978-3-86788-691-8

Ruhr Economic Papers #595

Johannes Ludwig

Is there Heterogeneity in the Response of Consumption to Income Shocks?

Ruhr Graduate School in Economics

RUHR UNIVERSITÄT BOCHUM RUB

Bibliografische Informationen der Deutschen Nationalbibliothek

Die Deutsche Bibliothek verzeichnet diese Publikation in der deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über: *http://dnb.d-nb.de* abrufbar.

http://dx.doi.org/10.4419/86788691
ISSN 1864-4872 (online)
ISBN 978-3-86788-691-8

Johannes Ludwig[1]

Is there Heterogeneity in the Response of Consumption to Income Shocks?

Abstract

While recently more and more research has focused on the aggregate response of consumption to income shocks, little is known about how this response differs for households at different ends of the income distribution. This paper investigates how consumption reacts to transitory and permanent shocks to disposable income for households with an income above or below the median. Panel data on income and consumption from the PSID between 1998 and 2012 is used to estimate consumption insurance parameters. Although households below the median are found to be exposed to larger transitory and permanent income shocks, they can buffer permanent shocks to income significantly better compared to households above the median. The latter, though, are better insured against transitory income shocks. In general, the poorer households are, the more similarly they react to the two kinds of income shocks.

JEL Classification: D12, D31, E21

Keywords: Consumption response to income shocks; consumption insurance

December 2015

1 Johannes Ludwig, RGS Econ and RUB. *– The author is grateful to Anna Klabunde and Christoph Kaufmann for helpful comments and suggestions.**– All correspondence to: Johannes Ludwig, Chair of Macroeconomics, Building GC 2/149, Ruhr-Universität Bochum, 44780 Bochum, Germany, e-Mail: johannes.ludwig@rgs-econ.de*

1 Introduction

Economic theory has come up with a wide variety of models that explain how consumption and income are related: the permanent income hypothesis, buffer-stock saving, hand-to-mouth consumption, etc. Although economists have also spent much effort on testing the implications of these models, no real consensus has emerged so far on the question how strongly consumption reacts to changes in income empirically. This is unfortunate since knowing how household consumption responds to income shocks yields important information on how well agents in the economy can handle (income) risk. If consumption is very vulnerable to shocks, governments will be able to enhance welfare by policies that help to smooth consumption over time.

A huge body of microeconometric literature has focused on analyzing how specific income shocks affect consumption (see Jappelli & Pistaferri 2010 for a review), while the more macroeconomic-oriented literature has tried to estimate parameters that indicate how well households in general are insured against shocks, differentiating only between transitory and permanent shocks to income which are believed to affect consumption differently (Blundell et al. 2008, Jappelli & Pistaferri 2011). These more general measures are not only of great value for policy advice, they are also used as reference values for computational macroeconomic models (Kaplan & Violante 2010).

Most of the literature has concentrated on estimating parameters of average insurance. However, there is striking evidence that the response of consumption to income shocks differs for households at different ends of the income distribution: Deaton & Paxson (1994) argue in their classic paper that the relationship of income and consumption inequality is an indicator for how well households can absorb risk. Krueger et al. (2010) show that for households below the median the distance between the levels of inequality of disposable income and inequality of consumption is substantially smaller than for households above the median. Meyer & Sullivan (2013b) find similar evidence. In the perspective of Deaton & Paxson (1994) this means that poorer households are able to reduce risk to a greater extent than richer households. This counterintuitive finding has received relatively little attention so far, though. Krueger et al. (2010) hypothesize that either the fraction of transitory shocks that hit households is larger for those below the median or that those above the median are less well able to insure consumption against income shocks. Unfortunately, they cannot answer this question due to the cross-sectional nature of their data.

This paper, therefore, seeks to answer the question of Krueger et al. (2010) whether the response of consumption to transitory and permanent income shocks differs for households in the lower and upper part of the income distribution. I use data from the Panel Study of Income Dynamics (PSID) which after a redesign in 1999 includes longitudinal information on both income and consumption. The descriptive analysis of this data shows that for the years 1998-2012 households below the median do not face a higher level of consumption inequality compared to households above the median although they are exposed to a much larger level of income inequality.

The magnitude of transitory and permanent income shocks and the parameters of consumption insurance are identified by a minimum distance estimation that makes use of the covariance structure of income and consumption growth. The estimation results show that at the aggregate level households are well insured against permanent shocks to their net income, but not perfectly insured against transitory shocks. While 37.5% of a permanent income shock translates into consumption, 11.7% of a transitory shock is transmitted.

When the aggregate sample is split up into households with a net income below the median and those above the median, households in the lower half of the income distribution are found to be subject to a substantially higher income risk compared to the upper half. Both the variances of transitory and of permanent income shocks are considerably larger. However, measured relative to the size of the permanent shocks, transitory shocks are found to be less important for poorer households than for those in the upper part of the income distribution. Moreover, households below the median also react rather similarly to both kinds of income shocks, the difference between the estimated consumption insurance parameters is smaller compared to the full sample benchmark. Households above the median, in contrast, differentiate considerably between permanent and transitory shocks to their net income. They are better insured against transitory shocks, but have to adjust consumption more strongly in face of permanent shocks compared to the full sample. Thus, although they face higher risk, households below the median can insure their consumption against permanent shocks to a significantly higher degree. This explains why relatively poor households do not face a higher consumption inequality although income inequality is larger compared to rich households.

This paper contributes to the literature by showing that the response of consumption to income shocks differs significantly for households at different positions in the income distribution. This is not only important for the design of welfare policies intended to stabilize the consumption of poor households, but also for the evaluation of macroeconomic heterogeneous agent models as demanded by Kaplan & Violante (2010). The remainder of the paper is organized as follows: Section 2 summarizes the most important results of the literature. Section 3 describes details about the data and variable construction. Descriptive results are presented in Section 4. The parametric model for income and consumption is introduced in Section 5 whereas the identification of the model parameters is discussed in Section 6. The results of the minimum distance estimation of model parameters is presented in Section 7. Section 8 concludes.

2 Related literature

Many important contributions to the literature on the empirical relationship of income and consumption are based on data from repeated cross-sections (Cutler & Katz 1991, Deaton & Paxson 1994, Blundell & Preston 1998, Krueger & Perri 2006). Repeated cross-sections allow to estimate variances and contemporaneous covariances of income and consumption so that the dispersion of both variables can be tracked over time. Unfortunately, due to the lack of autocovariances, the exact size of transitory and permanent shocks to income cannot be inferred although consumption is expected to react very differently to them. Thus, repeated cross-sections enable to hypothesize why income and consumption distributions comove or not, but definitive causes - e.g. an increase in the size of some kind of income shock - cannot be identified (or only under very strong assumptions).

When the effect of transitory and permanent income changes on consumption is to be estimated, panel data both on earnings and consumption expenditure is needed. One of the main problems that the literature has faced so far is that this kind of data is not available for most countries and for most spans of time. Two ways how to overcome this problem have been proposed:

Hall & Mishkin (1982) use data from the PSID for the years 1969-1975. At this time, the study included only information on household expenditure for food. Hence, Hall & Mishkin use food consumption as a proxy for nondurable consumption. Moreover, they concentrate on the effect of transitory shocks and require permanent shocks to be fully transmitted into consumption. Their analyses indicate that transitory income shocks translate into food consumption to a substantial degree and that this effect is larger than predicted by the permanent income hypothesis.

Blundell et al. (2008) circumvent the problem of missing panel data by generating an imputed consumption variable for PSID data. They use information on consumption from the cross-sectional Consumer Expenditure (CE) Survey to construct a measure of nondurable consumption within the PSID for the years 1979-1992. With this longitudinal information on household income and consumption they are able to identify consumption insurance parameters. Their main result is that 64% of a permanent shock to disposable income is transmitted into consumption whereas the reaction to a transitory shock cannot be distinguished from being zero.

It is clear that both studies have severe limitations. Using food consumption as a proxy for nondurable consumption can be misleading since food is a necessity and should react less strongly to income shocks compared to overall nondurable consumption. The result of studies that make use of imputed variables depends very much on the quality of the imputation procedure. Besides that, imputed variables usually have a lower variance compared to survey data which possibly biases consumption insurance parameters (Ludwig 2015).

This is why two datasets that include panel data on both income and consumption have recently attracted the interest of researchers analyzing the interplay of the two variables: the Italian Survey of Household Income and Wealth (SHIW) (Kaufmann & Pistaferri 2009, Jappelli & Pistaferri 2011, Krueger & Perri 2011, Hryshko 2014) and the redesigned PSID for the years from 1999 on (Attanasio & Pistaferri 2014, Kaplan et al. 2014, Blundell et al. 2015). The SHIW has been used by Jappelli & Pistaferri (2011) to estimate aggregate consumption insurance parameters for the years 1987-2006. They find that households in Italy fully translate permanent income shocks into consumption expenditure while transitory shocks only affect consumption at 28% of their size. However, Jappelli & Pistaferri do not analyze whether these values are different for the upper and lower part of the income distribution. The redesigned PSID has not been used for these purposes so far, to the best of my knowledge. The empirical framework used by Blundell et al. (2015) bears some similarities with the one used in this paper, but the focus of their study is rather on the intra-family labor supply responses following income shocks.

3 The data

Income and consumption data used in this paper stem from the Michigan Panel Study of Income Dynamics. This dataset - the world's longest running household panel dataset - started to be collected in 1968 and is renowned for very detailed information on household income. In the late 1990s the dataset was subject to substantial changes in design and in the method of data collection. While information was collected annually up to 1997, it has been collected only every other year since 1999. Moreover, new variables were added in that year, among others more detailed information on household consumption.

This redesign represents a major step forward for researchers interested in consumption. Before 1999 the only variables on consumption expenditure that were continuously available in the PSID were food expenditure (food at home, restaurant meals and value of food stamps) and rent. In 1999 the list of expenditure categories was extended significantly. Information on expenditure for household utilities, gasoline, car maintenance, transportation, health, education and child care are now continuously available. If these variables are compared to the aggregate numbers of the National Income and Product Accounts (NIPA), the new variables combined with food expenditure and rent account for around 70% of total household nondurable consumption (Blundell et al. 2015). In 2005 another set of expenditure variables (clothing, entertainment, vacation, telephone, household equipment) was added so that nearly all of the largest NIPA expenditure categories are represented in the PSID.[1] Since panel data

[1] The biggest expenditure categories that are missing in the PSID are alcohol and tobacco which are

on income and consumption is needed for the analyses of this paper, only the years where comprehensive consumption data is available can be used. My observation period, therefore, starts in 1998 and ends in 2012, the last year for which data is publicly available.[2]

When households decide how much to spend on consumption and how much to save, the decisive variable is net income. As the goal of this paper is to analyze the empirical response of consumption to changes in income, the income variable used in the following has to be household net income. The PSID only includes information on household gross income (including public transfers) so that tax payments still need to be subtracted. NBER TAXSIM is used to simulate federal and state taxes for the sample households (for more information on TAXSIM see Feenberg & Coutts 1993). "Income" in the following, thus, always means household gross income minus (simulated) federal and state taxes.

The consumption variable used below is the sum of all consumption expenditure categories that are continuously available in the PSID since 1999. Hence, it includes food expenditure, rent and the categories added in 1999. Therefore, "consumption" in the following means 70% of total nondurable consumption. For some robustness checks a second variable is used that includes the expenditure categories added in 2005. However, this limits the observation period to 2004-2012 and observations are lost. All income and expenditure variables are deflated by the Consumer Price Index (CPI) of the Bureau of Labor Statistics with the base year 2000 and equivalence-weighted by the OECD equivalence scale. The basic sample that is used for all following analyses consists of 5,882 households with 36,968 observations on income and 35,303 observations on consumption (details on sample selection can be found in Appendix A1).

4 Descriptive results

Figure 1 shows the trends of average real equivalence-weighted income and consumption between 1998 and 2012 for the households of my sample. Average net income increases until 2008, but then decreases slightly after the financial crisis has hit the economy. The trend of consumption looks very similar: Until 2006 it increases (even more strongly than the respective income values) and decreases thereafter. The figure also shows that the trend of consumption is not different when the consumption variables that were added to the PSID in 2005 are included in the variable of consumption. The variable that includes them is characterized, of course, by a higher level of aggregate consumption, but this level moves in

usually not measured very well in survey data.

[2] Income is reported retrospectively in the PSID. Thus, in the interviews of e.g. the 1999-wave of the study interviewers ask for the income in 1998. The reference period of consumption variables varies for the different categories of household consumption. To be able to identify model parameters I assume that consumption values also refer to the previous year as it is done in the majority of the literature.

Figure 1: Average household net income and consumption, 2000$, aggregate level.

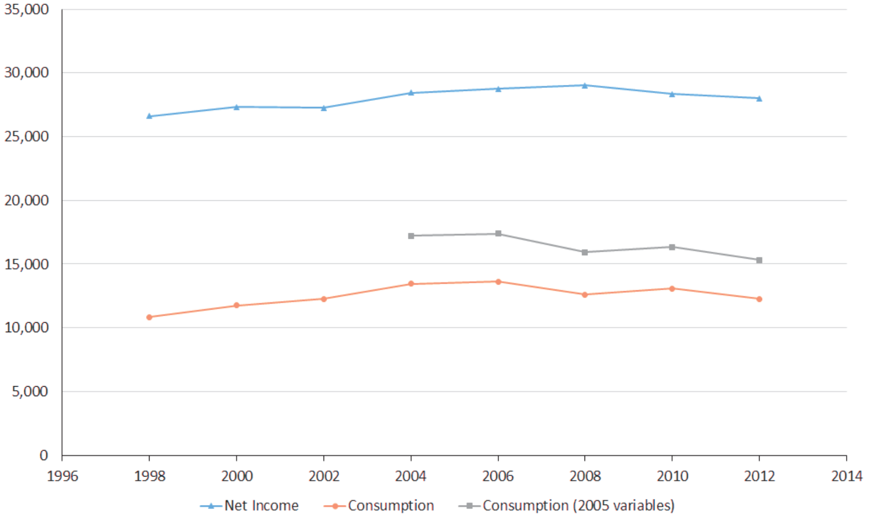

parallel with the benchmark consumption variable.

The development of aggregate income and consumption inequality is depicted in Figure 2 where inequality is measured as the 90/10 percentile ratio.[3] Inequality of household net income increases relatively steadily between 1998 and 2012. Consumption, on the contrary, not only has a much lower level of inequality, but inequality also increases only until the middle of the observation period and declines thereafter. The fact that income inequality increases after the financial crisis, but consumption inequality decreases has also been found by Meyer & Sullivan (2013a) who use income data from the Current Population Survey (CPS) and consumption data from the Consumer Expenditure Survey. Again, the development of consumption inequality is practically unchanged when the variables of the 2005-wave are included. This reassures that the benchmark variable does not miss any important consumption items. In general, income and consumption inequality as well as average income and consumption do not show extreme trends between 1998 and 2012 since the period under review is rather short.

How do income and consumption develop for households in the upper and lower half of the income distribution? Figure 3 shows the evolution of mean income and consumption and inequality values separately for households above the median and below the median.[4] Average

[3]The development of the gini coefficient and of the standard deviation of log-variables (not depicted here) look very similar. Results are available from the author upon request.

[4]Inequality is measured as 50/10-percentile ratio in panel (c) and as 90/50-percentile ratio in panel (d). To compare the same households with each other in panel (a) and (c) and in panel (b) and (d), I have to

Figure 2: Inequality of household income and consumption, 90/10-percentile ratio, aggregate level.

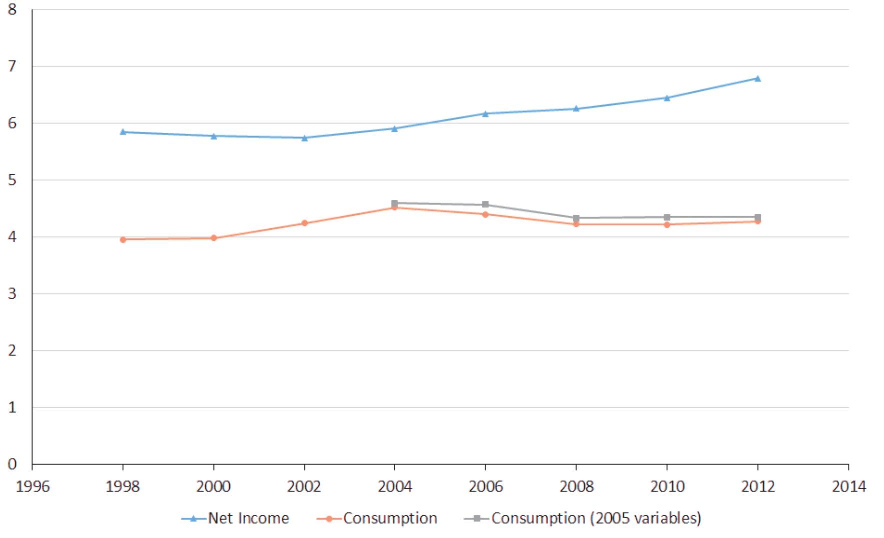

income and consumption is depicted for the households below the median in panel (a) and for households above the median in panel (b). The trends are qualitatively equivalent to the aggregate case, but the levels, of course, differ substantially. The most important difference that can be inferred from panels (a) and (b) is that households above the median spend a smaller fraction of their income on nondurable consumption.

The 50/10-percentile ratio of income and consumption is depicted in panel (c) whereas panel (d) shows the 90/50-percentile ratio. It is immediately visible that households below the median have a substantially higher level of income inequality in all years. However, the level of consumption inequality does not seem to differ significantly between the two panels. This means that for households below the median there is a huge difference between the level of income inequality and the level of consumption inequality while this difference is relatively small for households in the upper part of the income distribution. Krueger et al. (2010) have already described this phenomenon and shown that it not only occurs in the US, but also in a group of other industrialized countries. Using income data from the CPS and consumption data from the CE Survey Meyer & Sullivan (2013b) have found very similar evidence. Thus,

exclude households with an income below the 10th percentile from panel (a) and with an income above the 90th percentile from panel (b). Thus, "below the median" in Figure 3 means households with an income between the 10th percentile and the median. "Above the median" means households with an income between the median and the 90th percentile. Average values without this modification do not differ qualitatively. This restriction applies only to Figure 3, not to the estimation results below.

Figure 3: Income and consumption above and below the median, mean values are in 2000$, inequality is the 50/10- and 90/50-percentile ratio.

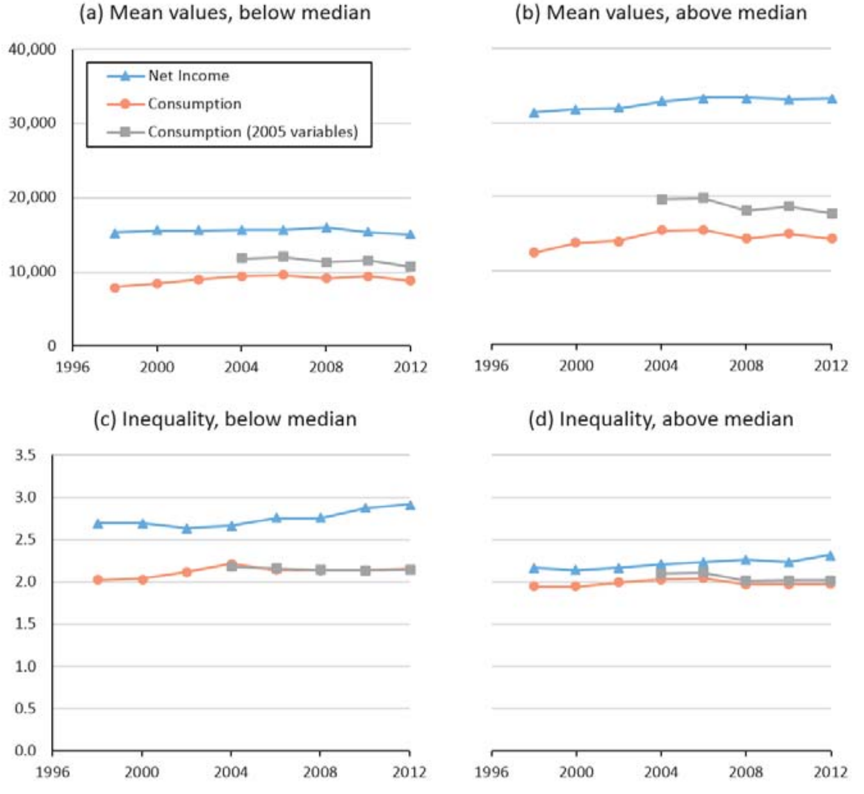

this finding seems to be relatively robust and not driven by PSID data.

5 The income and consumption processes

The parametric model for income and consumption is a version of the seminal model developed by Blundell et al. (2008). It is assumed that household income Y can be separated into a part explained by observable characteristics X and a residual which comprises the effect of all unobserved factors on household income. To be able to estimate the effect of transitory and permanent income shocks on consumption, the residual has to distinguish between unobserved factors that affect household income permanently and those that only have a temporary effect. Thus, the log of income of household i at time t can be expressed as

$$\ln Y_{i,t} = X'_{i,t}\varphi_t + P_{i,t} + \epsilon_{i,t} \quad (1)$$

where $P_{i,t}$ is the permanent component of the residual and $\epsilon_{i,t}$ summarizes the effect of all unobserved factors that have an influence on household income only at time t.

Blundell et al. (2008) use an MA(1)-process for transitory income whereas I use a simple one-period shock that shows no correlation over time. One key difference to the study of Blundell et al. (2008) is the structure of the data. Although we both use PSID data, their data have an annual frequency while my data is biennial since I only use observations from 1999 on. An MA-process in transitory income cannot be identified when only biennial data are available. Thus, I assume transitory income to be a one-period shock. I have shown in earlier work that the resulting income process (1) matches PSID data well (Ludwig 2015). The permanent component of household income is modeled as a random walk with

$$P_{i,t} = P_{i,t-1} + \zeta_{i,t} \quad (2)$$

so that the innovations $\zeta_{i,t}$ have an effect on income that is fully transferred into period $t+1$ and all following periods.

To abstract from changes in income that can be explained by changes in observed household characteristics, I will work with the income residual $y_{i,t} = \ln Y_{i,t} - X'_{i,t}\varphi_t = P_{i,t} + \epsilon_{i,t}$ throughout the analysis. Changes in characteristics like age and household size are often foreseen and do not qualify as shocks to household income. To control for the effect of observable characteristics on income and consumption, an OLS-regression is run that uses the following variables as explanatory factors: employment status of head and wife (if present), education of head and wife, number of children in the household, presence of other income earners than head and wife in the household, state of residence, ethnicity of head and wife, sex of the household head and age of the household head. These variables explain 52.4% of the variation in income and 48.0% of the variation in consumption.[5]

The model is estimated in differences, but the biennial structure of the data needs to be taken into account. To make it explicit that the observed changes over time are not yearly first differences $\Delta y_{i,t}$, I use the second (seasonal) difference of the residual for the estimation that is defined as follows:

$$\Delta_2 y_{i,t} = y_{i,t} - y_{i,t-2} = \zeta_{i,t} + \zeta_{i,t-1} + \epsilon_{i,t} - \epsilon_{i,t-2}. \quad (3)$$

Blundell et al. (2008) assume that there are three possible sources for a change in (residual) consumption between two periods: as a reaction to a shock to permanent income $\zeta_{i,t}$,

[5] The trends of residual income and consumption inequality do not differ from the trends of the overall variables depicted in Figures 2 and 3 above.

as a reaction to a transitory income shock $\epsilon_{i,t}$ or due to some event that is unrelated to household income. One of their main contributions consists in not defining by how much consumption ought to change in response to a specific income shock. Instead, they let the data determine this amount freely. Therefore, they introduce two parameters that measure how much of an income shock is translated into household consumption. The parameters $\phi_{i,t}$ and $\psi_{i,t}$ are insurance parameters that express how well a household can insure itself against permanent and transitory shocks to income, respectively. Hence, the first difference of residual consumption is defined by Blundell et al. (2008) as

$$\Delta c_{i,t} = \phi_{i,t}\zeta_{i,t} + \psi_{i,t}\epsilon_{i,t} + \xi_{i,t} \tag{4}$$

where $\xi_{i,t}$ summarizes factors that affect consumption but are not related to income. The insurance parameters $\phi_{i,t}$ and $\psi_{i,t}$ should lie between zero and one for the vast majority of households. If they are found to be zero, the income shock has no effect on consumption, i.e. consumption is fully insured. If they are equal to one, the shock is fully translated into consumption. Finally, if they lie somewhere in between, households are only partially insured against the respective income shock.

Blundell et al. also allow for measurement error in consumption since their consumption variable is imputed and very likely to be measured with error. Equation (4) then becomes

$$\Delta c_{i,t} = \phi_{i,t}\zeta_{i,t} + \psi_{i,t}\epsilon_{i,t} + \xi_{i,t} + u^c_{i,t} - u^c_{i,t-1} \tag{5}$$

where $u^c_{i,t}$ represents measurement error in consumption in period t. As the consumption variable used in my analysis is not imputed, it is not clear whether it is measured with substantial error. However, in the following I will use the specification that allows for measurement error in consumption (Appendix A2 discusses the role of measurement error in more detail). Again, equation (5) needs to be adjusted to the biennial data structure of this paper. If it is assumed that the insurance parameters both do not differ between t and $t-1$, the modified equation becomes

$$\Delta_2 c_{i,t} = \phi_{i,t}\left(\zeta_{i,t} + \zeta_{i,t-1}\right) + \psi_{i,t}\left(\epsilon_{i,t} + \epsilon_{i,t-1}\right) + \xi_{i,t} + \xi_{i,t-1} + u^c_{i,t} - u^c_{i,t-2}. \tag{6}$$

6 Identification

The main problem when trying to identify the parameters of the model is that single transitory and permanent income shocks that hit households cannot be observed directly from the data. If only the residual $y_{i,t}$ is known, it is impossible to use equation (6) and simply regress the shocks on the change in consumption by e.g. an OLS-regression. Thus, to identify aggregate consumption insurance parameters ϕ_t and ψ_t it has to be relied on the variances, covariances and autocovariances of $\Delta_2 y_{i,t}$ and $\Delta_2 c_{i,t}$. Insurance parameters and the variances of transitory and permanent income shocks can then be identified by minimizing the distance between the empirical second moments and the predictions of the parametric model described above.

Equation (3) shows how the change in residual household income depends on transitory and permanent shocks. The structure of the variance and the autocovariances of income change can be derived from this equation:

$$Cov\left(\Delta_2 y_{i,t}, \Delta_2 y_{i,t+s}\right) = \begin{cases} 2Var\left(\zeta_{i,t}\right) + Var\left(\epsilon_{i,t}\right) + Var\left(\epsilon_{i,t-2}\right) & \text{for } s = 0, \\ -Var\left(\epsilon_{i,t}\right) & \text{for } s = 2, \\ 0 & \text{for } s = 4, 6, 8, 10, 12, \end{cases}$$

where it has to be assumed that $\zeta_{i,t}$ and $\epsilon_{i,t}$ are mutually and serially uncorrelated. Moreover, I have to assume that the size of the permanent income shocks does not differ between period t and $t-1$. Thus, the magnitude of permanent shocks cannot be identified for every single year, but the variances represent an average value for t and $t-1$. However, this should be sufficient to detect general trends over time.

Note that due to the biennial frequency of the data $\Delta_2 y_{i,t}$ can only be observed in t, $t+2$, $t+4$, etc. The autocovariance between income change in period t and $t+2$ identifies the variance of the transitory income shocks $\epsilon_{i,t}$. If the size of transitory shocks is known, the variance of income change identifies the variance of permanent shocks $\zeta_{i,t}$. Hence, $Var\left(\Delta_2 y_{i,t}\right)$ and $Cov\left(\Delta_2 y_{i,t}, \Delta_2 y_{i,t+2}\right)$ together suffice to identify the variances of both income shocks. All other autocovariances of income change are equal to zero.

Adding consumption data allows consumption insurance parameters to be identified. The central moments that determine the size of ϕ_t and ψ_t are the covariances between the change in income and the change in consumption. Equation (3) and (6) predict the following:

$$Cov\left(\Delta_2 c_{i,t}, \Delta_2 y_{i,t+s}\right) = \begin{cases} 2\phi_t Var\left(\zeta_{i,t}\right) + \psi_t Var\left(\epsilon_{i,t}\right) & \text{for } s = 0, \\ -\psi_t Var\left(\epsilon_{i,t}\right) & \text{for } s = 2, \\ 0 & \forall s > 2 \text{ and } s < 0. \end{cases}$$

Assuming the variances of the income shocks are known, the covariance between $\Delta_2 c_{i,t}$ and $\Delta_2 y_{i,t+2}$ identifies the insurance against transitory income shocks. With ψ_t, $Var\left(\epsilon_{i,t}\right)$ and $Var\left(\zeta_{i,t}\right)$ already identified, the contemporaneous covariance of income and consumption growth determines the parameter for insurance against permanent shocks ϕ_t.

Finally, the autocovariances of consumption change provide additional moments for the estimation and allow to identify the remaining parameters of the model, $Var\left(\xi_{i,t}\right)$ and $Var\left(u_t^c\right)$.

Equation (6) predicts:

$$Cov\left(\Delta_2 c_{i,t}, \Delta_2 c_{i,t+s}\right) = \begin{cases} 2\phi_t^2 Var\left(\zeta_{i,t}\right) + 2\psi_t^2 Var\left(\epsilon_{i,t}\right) & \\ +2Var\left(\xi_{i,t}\right) + Var\left(u_t^c\right) + Var\left(u_{t-2}^c\right) & \text{for } s=0, \\ & \\ -Var\left(u_t^c\right) & \text{for } s=2, \\ 0 & \text{for } s=4,6,8,10,12, \end{cases}$$

where it needs to be assumed that $Var\left(\xi_{i,t}\right) = Var\left(\xi_{i,t-1}\right)$. Concerning identification, the same logic as above also applies here: The autocovariance of consumption growth in period t and period $t+2$ pins down the size of the measurement error, $Var\left(u_t^c\right)$, and the variance of consumption growth then identifies the size of residual influences on consumption, $Var\left(\xi_{i,t}\right)$.

In principle, the framework allows the size of all parameters to vary over time. However, for the benchmark estimation the variance of the consumption change residual $\xi_{i,t}$ is required to be equal in all periods as there is no theoretical justification why it should change over time. Since the observation period is relatively short, insurance parameters ϕ and ψ will also be treated as stationary for most of the following analysis. After collecting all second moments of the data, the model parameters are estimated by equally-weighted minimum distance (EWMD). More details on the minimum distance estimation can be found in Appendix A3.

7 Results of the minimum distance estimation

7.1 Aggregate level

The development of the variances of transitory and permanent income shocks as estimated by the minimum distance procedure is depicted in Figure 4. Transitory shocks that have an effect on household income only in the respective year are much larger in size, but slightly decreasing over time. The variance of permanent shocks, on the contrary, is smaller, but slightly increasing over time. In 2008 when the financial crisis hit the economy, $Var\left(\zeta_{i,t}\right)$ peaks, but quickly declines to an average level thereafter. Thus, the composition of income shocks changes mildly over time and permanent shocks become more important towards the end of the observation period.

Table 1 summarizes the estimated parameters and standard errors for different specifications of the minimum distance procedure. For brevity, the variances of transitory and permanent shocks are only listed with their average value for all cases and, therefore, no standard errors are reported. The first column of Table 1 summarizes the results of the benchmark model that uses the full sample of households and EWMD as estimation method. It shows that 37.5% of a permanent income shock is translated into household consumption.

Figure 4: Variances of transitory and permanent shocks with 95%-confidence bands, aggregate level.

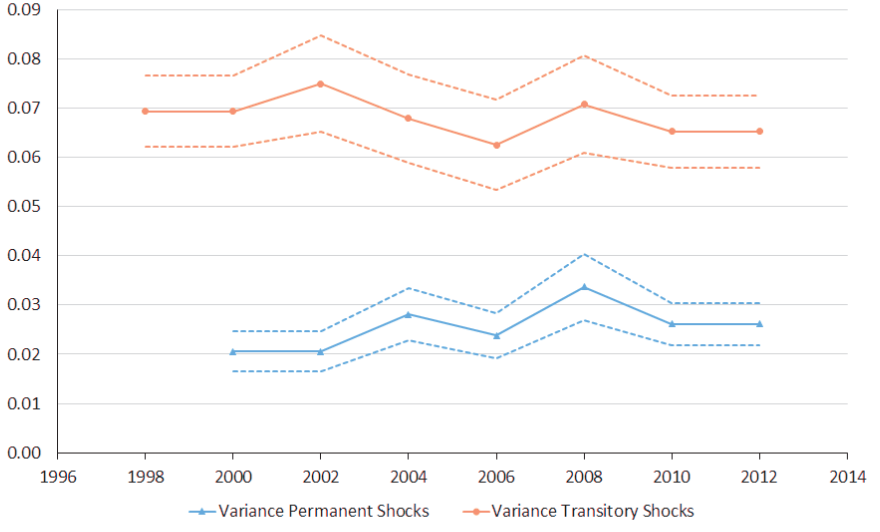

This value of ϕ is much lower than the respective values found by Blundell et al. (2008) and Jappelli & Pistaferri (2011) and indicates that the households in my sample are relatively well insured against shocks that have a long-lasting effect on household income. However, Jappelli & Pistaferri (2011) analyze the Italian economy which differs from the US economy and Blundell et al. (2008) use data from the 1980s. In earlier research I have shown that consumption insurance parameters can change over time when longer time spans are analyzed (Ludwig 2015). Hence, I do not consider my results for ϕ to be in direct contradiction to their values.

The coefficient for ψ indicates that 11.7% of a transitory shock translates into household consumption and this estimate is significantly different from zero. This shows that although the households in my sample are well insured against these shocks which are much larger in size compared to permanent shocks, consumption expenditure is adjusted to a small degree. This is in line with the findings of Hall & Mishkin (1982) and Jappelli & Pistaferri (2011) who both find excess sensitivity, i.e. a significant reaction of consumption to transitory income shocks, but stands in contrast to Blundell et al. (2008) who estimate a lower ψ which is not statistically distinct from zero. The estimate for $Var\left(\xi_{i,t}\right)$ is small and has only negligible effects on consumption growth.

Table 1: Results of the minimum distance estimations at the aggregate level.

	Full sample EWMD	Full sample DWMD	No SEO & Immigrant EWMD	2005- Cons. EWMD
ϕ	0.375	0.396	0.401	0.418
(Insurance perm. shock)	(0.030)	(0.032)	(0.038)	(0.037)
ψ	0.117	0.114	0.108	0.174
(Insurance trans. shock)	(0.020)	(0.020)	(0.023)	(0.028)
Av. $Var\left(\zeta_{i,t}\right)$	0.026	0.025	0.023	0.027
(Av. variance perm. shock)				
Av. $Var\left(\epsilon_{i,t}\right)$	0.068	0.069	0.063	0.064
(Av. variance trans. shock)				
$Var\left(\xi_{i,t}\right)$	0.013	0.013	0.012	0.010
(Variance unobs. heterog.)	(0.001)	(0.001)	(0.001)	(0.001)
N	5,773	5,773	3,709	5,145

Note: standard errors in parentheses, EWMD = Equally-weighted minimum distance, DWMD = Diagonally-weighted minimum distance, SEO = Survey of Economic Opportunities subsample of the PSID.

7.2 Robustness

An obvious question is whether the results are affected by the financial crisis that hit the US economy in 2008 which is roughly the middle of my observation period. The variances of transitory and permanent shocks in Figure 4 that are allowed to vary with t indicate that at least permanent shocks are larger after the crisis. The benchmark case discussed above does not allow that the insurance parameters vary over time. However, the financial crisis led to a drop in the supply of credit which is believed to have an important effect on how income shocks translate into consumption (Krueger & Perri 2006). If credit supply decreases due to the financial crisis and this worsens the ability of households to smooth consumption over time, the values of ϕ and ψ should increase after 2008. Table 2 shows the results of an EWMD estimation that allows the insurance coefficients to take on different values for the time periods 2000-2006 and 2008-2012. Since the other variables are hardly affected by this adjustment, I concentrate on the values of ϕ_t and ψ_t here.

Table 2 shows that as hypothesized both consumption insurance parameters increase after the financial crisis. Thus, from 2008 on, it becomes harder to buffer permanent and transitory shocks to income. However, the increase in ϕ_t is rather small and not statistically significant.

Table 2: Result of the minimum distance estimation with time-dependent insurance parameters.

Insurance permanent shock:		Insurance transitory shock:	
$\phi_{2000/2006}$	$\phi_{2008/2012}$	$\psi_{2000/2006}$	$\psi_{2008/2012}$
0.330	0.373	0.097	0.170
(0.036)	(0.045)	(0.024)	(0.035)

Note: standard errors in parentheses.

The increase in ψ_t, though, is relatively large: before 2008 only about 10% of a transitory shock to household income is transmitted into consumption whereas thereafter it is 17%. Hence, the results in Table 2 are in line with an assumed effect of the financial crisis on consumption insurance, but this effect is stronger on the insurance against transitory income shocks.

Columns (2)-(4) of Table 1 depict the results of variations of the benchmark estimation with respect to estimation procedure, sample selection and variable specification. The second column shows the coefficients estimated by a diagonally-weighted minimum distance (DWMD) procedure with the same sample as in the benchmark case, i.e. only the weighting matrix of the objective function is exchanged. DWMD is a slightly more general version of the benchmark EWMD procedure and allows for heteroskedasticity. The results show that the change in the estimation procedure does not have a huge effect: While the coefficient of ϕ is slightly larger than in the benchmark case, all other parameters hardly change.

The benchmark sample includes all subsamples of the PSID. The third column of Table 1 shows how the results change when only households are considered that belong to the original nationally representative Survey Research Center (SRC) subsample of the PSID. All households belonging to the Survey of Economic Opportunity (SEO) subsample of poor households or to the Immigrant sample are dropped. The results show that SRC households are subject to a lower level of income risk, i.e. the average variances of both transitory and permanent shocks are smaller than for the full sample. However, although they are subject to less risk, insurance coefficients do not differ significantly from the whole sample: ϕ increases and ψ decreases, but the changes are small and not significant.

Finally, the fourth column of Table 1 lists the coefficients of an estimation that makes use of the consumption variable which includes the expenditure categories the PSID added in 2005. This variable is only available for the years 2004-2012 and, therefore, shrinks the observation period. The change in the time period can explain why the magnitude of permanent shocks slightly increases and transitory shocks slightly decrease in size (see Figure 4). Both insurance coefficients increase, i.e. households can buffer shocks less well than in the benchmark case. While the change in ϕ is not significant, this is the case for the insurance against transitory shocks. The value found for ψ, though, is similar to the value found in Table 2 for the post-financial crisis period. Thus, it is more likely that this increase is related to the observation period and not driven by the change in the variable for consumption. In general, the results of the benchmark estimation in the first column of Table 1 seem relatively robust to changes in estimation procedure, sample selection or variable specification.

7.3 Upper and lower half of the income distribution

Figure 3 in Section 4 shows that consumption inequality has roughly the same level for households above and below the median although poorer households have a much higher level of income inequality. It remains an open question why those below the median are able to reduce the amount of income risk far more strongly than households that are actually richer. Krueger et al. (2010) mention two potential explanations for this phenomenon: (i) households below the median are characterized by a higher relative importance of transitory income shocks compared to permanent ones (e.g. due to a higher probability of short-term unemployment), i.e. $\frac{Var(\epsilon_{below})}{Var(\zeta_{below})} > \frac{Var(\epsilon_{above})}{Var(\zeta_{above})}$, and (ii) households in the lower half of the income distribution are characterized by a higher ability to buffer shocks to their income, i.e. $\phi^{below} < \phi^{above}$ and/or $\psi^{below} < \psi^{above}$. Krueger et al. (2010) suggest that "transfers from informal networks of family and friends" could be a source of consumption insurance which is more substantial for households below the median.

The first two columns of Table 3 list estimation results for samples that only include households with a net income above or below the median, respectively. Households qualify for each group if their income values are above or below the median in at least 80% of the observed cases.[6] Table 3 reveals that the results for the upper and the lower part of the income distribution are, in fact, very different, not only between each other, but also compared to the full sample benchmark.

Households above the median are characterized by a much lower level of income risk. Both the variances of transitory and of permanent income shocks are substantially lower compared to the benchmark results and also compared to the households below the median. This explains why income inequality is relatively lower for households above the median compared to those below the median. As households above the median have a higher income, one would assume that they also have a higher stock of savings and are better able to buffer shocks to their income. This is true for transitory income shocks: for households above the median 8.4% of a transitory income shock is translated into consumption compared to 11.7% for the full sample. The difference is not statistically significant, though. Surprisingly,

[6] Results do not differ qualitatively when all income observations of a certain household are required to be above or below the median. Sample sizes become very small, though.

Table 3: Results of the minimum distance estimations for different parts of the income distribution.

	≥80% Above median EWMD	≥80% Below median EWMD	≥80% Above p75 EWMD	≥80% Below p25 EWMD
ϕ	0.512	0.312	0.570	0.278
(Insurance perm. shock)	(0.060)	(0.054)	(0.111)	(0.096)
ψ	0.084	0.161	0.059	0.174
(Insurance trans. shock)	(0.030)	(0.039)	(0.043)	(0.061)
Av. $Var\left(\zeta_{i,t}\right)$	0.016	0.031	0.016	0.035
(Av. variance perm. shock)				
Av. $Var\left(\epsilon_{i,t}\right)$	0.053	0.081	0.066	0.107
(Av. variance trans. shock)				
$Var\left(\xi_{i,t}\right)$	0.011	0.014	0.012	0.014
(Variance unobs. heterog.)	(0.001)	(0.002)	(0.002)	(0.005)
N	2,020	2,006	777	849

Note: standard errors in parentheses, EWMD = Equally-weighted minimum distance.

households above the median are significantly less well able to buffer permanent shocks to their income than the full sample. In the benchmark case 37.5% of a permanent income shock is transmitted into consumption, but for the upper half of the income distribution it is 51.2%.

Households below the median, on the contrary, face a higher level of income risk. The average magnitude of both transitory and permanent shocks to household net income is higher compared to the benchmark and to households above the median. The difference of insurance parameters to the benchmark case is contrary to the sample of households above the median. Households in the lower part of the income distribution are less well insured against transitory shocks (16.1% vs. 11.7% for the benchmark sample), but slightly better insured against permanent shocks (31.2% vs. 37.5%). Both parameters do not differ significantly from the benchmark, but differ significantly from the sample of households above the median.

What do these results imply for the relationship of income and consumption inequality for the two parts of the income distribution? The different magnitude of transitory and permanent income shocks explains why income inequality is relatively larger for households below the median. Concerning the similar level of consumption inequality for the two groups, Krueger et al. (2010) hypothesize that this phenomenon could be explained by transitory income shocks being more important for households below the median. The estimation results clearly reject this hypothesis. Although $Var\left(\epsilon_{i,t}\right)$ is absolutely higher for households below the median, the fraction of transitory to permanent shocks is lower. $\frac{Var(\epsilon_{i,t})}{Var(\zeta_{i,t})}$ is 2.613 for households below the median whereas it is 3.313 for those above the median. Thus, transitory shocks to income are relatively more important for households with a net income above the median.

The second potential explanation offered by Krueger et al. (2010) receives more support from the estimation results. While households above the median are better able to buffer transitory shocks to income compared to those below the median, they are significantly less well insured against income shocks that are permanent. The latter translate at 31.2% into the consumption of a household below the median, but at 51.2% into the consumption of a household above the median. Thus, regarding the hypotheses of Krueger et al. (2010), the only explanation for the development of income and consumption inequality above and below the median that is consistent with the estimation results is the low level of insurance against permanent income shocks for households in the upper part of the income distribution.

Why are richer households less well able to buffer permanent income shocks? Of course, networks of family and friends as conjectured by Krueger et al. (2010) could explain the results to some extent. However, it is striking that households above the median react very differently to the two kinds of income shocks whereas for households below the median the difference between ϕ and ψ is smaller. An explanation for this finding could be that the results for households in the lower part of the income distribution are driven by very poor households that cannot reduce their consumption further. These households cannot afford to differentiate their reaction to a negative income shock between permanent and transitory shocks. The result for the upper part of the income distribution, on the contrary, could be driven by households with a high amount of savings that can afford to react "optimally" to any kind of income shock, i.e. they do not adjust consumption in face of a transitory shock, but translate the majority of a permanent shock to income into consumption. This behavior corresponds to the theoretically optimal reaction in a simple permanent income setting.

Alternatively, it could also be argued that it is harder for households in the lower part of the income distribution to distinguish a priori between income shocks that are transitory and those that are permanent. In an experimental study Carbone & Hey (2004) find that agents differ considerably in their ability to plan ahead and optimize dynamically. The ability to solve dynamic optimization problems should be correlated positively with the position in the income distribution. If poor households are less well able to understand the long-run consequences of income shocks, they will not differentiate in their reaction to transitory and permanent income shocks.

If these explanations for the results of the upper and lower part of the income distribution hold, results should become even more extreme when moving further up or down the income distribution. The final two columns of Table 3 show the results for samples of households with an income that is for at least 80% of the observations above the 75-percentile or below the 25-percentile, respectively. Thus, these samples include the upper half of the households above the median and the lower half of those below the median. Note that due to the rather small sample size standard errors are relatively large for both samples.

Although most differences to the first two columns of Table 3 are not statistically significant, the tendency of the point estimates is in line with expectations. The results for households above the 75-percentile in the third column show that these households act a bit more like the optimal household in a permanent income setting compared to all households above the median. ϕ increases further to 0.570 whereas ψ decreases further and now does not differ significantly from zero anymore. For households below the 25-percentile the difference between ϕ and ψ narrows further and the two parameters are no longer significantly different from each other. Thus, the richer households are, the more differently they react to transitory and permanent income shocks and the poorer they are, the more alike their reactions to the two kinds of shocks are.

8 Conclusion

Households in the lower half of the income distribution have been found to be able to reduce income risk to a greater extent than those in the upper half (Krueger et al. 2010, Meyer & Sullivan 2013b). This paper is the first attempt to investigate whether the response to transitory and permanent income shocks differs for households above and below the median. Using panel data on income and consumption from the redesigned PSID for the years 1998-2012 I find that households below the median are exposed to larger transitory and permanent shocks to their net income. This explains why the level of income inequality is larger for households in the lower part of the income distribution.

The response to transitory and permanent income shocks is very different for rich and poor households. Households in the bottom 25% of the income distribution are found to react relatively similarly to both kinds of income shock, they translate 27.8% of a permanent income shock into consumption and 17.4% of a transitory shock. This relatively high degree of consumption insurance could be explained by the behavior of very poor households who cannot reduce consumption further. Households in the top 25% of the income distribution, on the contrary, clearly differentiate between transitory and permanent shocks and behave similarly to agents in a permanent income model, i.e. they are found to be perfectly insured against transitory income shocks, but transmit 57% of a permanent income shock into consumption. The significantly stronger response to permanent income shocks of households above the median compared to those below the median explains why the level of consumption inequality does not differ between the two halves of the income distribution although the level of income inequality differs.

Kaplan et al. (2014) have recently shown that there is a group of rich households that invest most of their portfolio in rather illiquid assets. Although their paper is concerned with the response of these households to transitory income shocks, it is conceivable that holding wealth in an illiquid way (e.g. housing) also reduces the ability to buffer shocks to household income that are permanent. This could be an explanation for the finding of this paper that households above the median insure against permanent income shocks to a substantially lower degree than relatively poorer households. Further research in this area is clearly warranted.

Appendix A1: Data and sample selection

Initially, all waves of the PSID from 1999 to 2013 are merged. Thereafter, income and expenditure values that are top-coded are set to missing. Some of the values for particular income or expenditure categories are not reported on an annual basis, but time-units range from weekly to biennial. Therefore, I adjust all of them to an annual value. Some of the adjusted values are obvious outliers and, hence, also set to missing.

The variable for expenditure on health insurance is of particular concern here. It reports biennial expenditure in the PSID waves of 1999-2011, whereas in the most recent wave of 2013 it reports annual expenditure. Even if the biennial values are adjusted to an annual value, average expenditure on health insurance in the 2013-wave is dramatically lower than in the years before. This raises concern whether biennial and annual reportings for this variable are comparable. Although expenditure for health insurance is a rather large expenditure category in the PSID, I have, therefore, decided to exclude it from the variable of aggregate consumption.

The measure of consumption used in the analysis also includes rent payments as a measure for consumption of housing services. To establish comparability between renters and home owners, I follow the example of Attanasio & Pistaferri (2014) and Blundell et al. (2015) who use a rent equivalent for home owners of 6% of the value of the house owned. This value goes back to the study of Flavin & Yamashita (2002).

Table A1: No. of households and observations in the dataset by step of sample selection

	No. of observations		No. of households	
Reason for exclusion	dropped	remain	dropped	remain
Initial sample		65,176		14,329
Head not in working age	12,962	52,214	2,318	12,011
Income & consumption outliers	2,507	49,707	175	11,836
No information on head's education	2,184	47,523	326	11,510
Too few income observations	10,555	36,968	5,628	5,882

As the study focuses on the working population, households where the head is not in working age are dropped from the sample. This exclusion affects household heads who are younger than 25 or older than 65. Moreover, I exclude observations on income and consumption that are very likely to be outliers. Following Blundell et al. (2008), income observations that are more than 500% or less than -80% of the previous value qualify as outliers. In addition, if food consumption or net income of a household is less than 100$ a year the observation is also treated as an outlier. Observations without an information on the educational status of the household head are dropped from the sample as well. Finally, we exclude households that have less than four observations on income. Table A1 summarizes how many observations and households are lost in each step.

Appendix A2: Measurement error in consumption

In a parametric model that abstracts from measurement error in consumption, equation (6) changes to

$$\Delta_2 c_{i,t} = \phi_{i,t}\left(\zeta_{i,t} + \zeta_{i,t-1}\right) + \psi_{i,t}\left(\epsilon_{i,t} + \epsilon_{i,t-1}\right) + \xi_{i,t} + \xi_{i,t-1}.$$

It is straightforward that this change has no effect on the autocovariances of income change or on the covariances between income and consumption. The only moments that are affected are the autocovariances of consumption growth that become

$$Cov\left(\Delta_2 c_{i,t}, \Delta_2 c_{i,t+s}\right) = \begin{cases} 2\phi_t^2 Var\left(\zeta_t\right) + 2\psi_t^2 Var\left(\epsilon_t\right) + 2Var\left(\xi_t\right) & \text{for } s = 0, \\ 0 & \text{for } s = 2,4,6,8,10,12. \end{cases}$$

Thus, the main difference to the benchmark model of Section 5 and 6 is that the covariance between $\Delta_2 c_{i,t}$ and $\Delta_2 c_{i,t+2}$ which is equal to $-Var\left(u_t^c\right)$ above, is now required to be zero.

Table A2: Covariances of consumption growth for t and t+2.

$t = 2000$	$t = 2002$	$t = 2004$	$t = 2006$	$t = 2008$	$t = 2010$
-0.042	-0.043	-0.054	-0.052	-0.050	-0.052
(0.003)	(0.003)	(0.005)	(0.004)	(0.003)	(0.004)

Note: standard errors in parentheses.

Table A2 summarizes the estimates and standard errors of the covariances between $\Delta_2 c_{i,t}$ and $\Delta_2 c_{i,t+2}$ for the years in our sample. All covariances are negative and strongly significantly different from zero as required by a consumption specification allowing for measurement error u_t^c. This shows that the benchmark consumption specification fits the data better than the alternative specification outlined here. Hence, although the consumption data used is survey data that does not need to be generated by an imputation procedure, it suffers from considerable measurement error. Therefore, the specification including u_t^c is used in all analyses of this paper.

Moreover, it would be desirable to also control for measurement error in income. Unfortunately, measurement error in income and transitory income shocks cannot be separately identified. It has to be kept in mind that some fraction of the variance in transitory income could be due to measurement error. This could be an additional explanation why consumption is found to be relatively well insured against transitory shocks to income.

Appendix A3: Minimum distance estimation

The starting point of the empirical investigation is a vector x_i that can be set up for every household in the dataset and that contains all second differences of residual consumption and income:

$$x_i = (\Delta_2 c_{i,00}, \Delta_2 c_{i,02}, ..., \Delta_2 c_{i,12}, \Delta_2 y_{i,00}, \Delta_2 y_{i,02}, ..., \Delta_2 y_{i,12})'$$

where "00" stands for the year 2000, "02" for 2002, etc. As there are 7 observations both on income and consumption growth in the dataset, the dimension of x_i is 14.

To cope with the unbalanced panel structure of the data, missing values in x_i have to be set to zero. Moreover, a second vector d_i has to be compiled that contains ones and zeros for non-missing and missing values in x_i, respectively. Since x_i and d_i are of the same dimension, variances and (auto)covariances of the variables in x_i can be estimated by the following formula:

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

where $\bigotimes$ stands for an elementwise division.

The symmetric matrix M contains all relevant second moments of the data:

$$M = \begin{pmatrix} Var\left(\Delta_2 c_{i,00}\right) & & & & \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 c_{i,02}\right) & & & & \\ ... & & & & \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 c_{i,08}\right) & ... & & & \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 c_{i,12}\right) & ... & Var\left(\Delta_2 c_{i,12}\right) & & \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 y_{i,00}\right) & ... & Cov\left(\Delta_2 c_{i,12}, \Delta_2 y_{i,00}\right) & Var\left(\Delta_2 y_{i,00}\right) & \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 y_{i,02}\right) & ... & Cov\left(\Delta_2 c_{i,12}, \Delta_2 y_{i,02}\right) & Cov\left(\Delta_2 y_{i,00}, \Delta_2 y_{i,02}\right) & \\ ... & & ... & ... & \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 y_{i,10}\right) & ... & Cov\left(\Delta_2 c_{i,12}, \Delta_2 y_{i,10}\right) & Cov\left(\Delta_2 y_{i,00}, \Delta_2 y_{i,10}\right) & ... \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 y_{i,12}\right) & ... & Cov\left(\Delta_2 c_{i,12}, \Delta_2 y_{i,12}\right) & Cov\left(\Delta_2 y_{i,00}, \Delta_2 y_{i,12}\right) & ... \quad Var\left(\Delta_2 y_{i,12}\right) \end{pmatrix}.$$

To collect all unique moments of M into a single vector, M is half-vectorized with $m = vech\left(M\right)$.

The resulting vector

$$m = \begin{pmatrix} Var\left(\Delta_2 c_{i,00}\right) \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 c_{i,02}\right) \\ \dots \\ Cov\left(\Delta_2 c_{i,00}, \Delta_2 y_{i,12}\right) \\ \dots \\ Var\left(\Delta_2 y_{i,00}\right) \\ Cov\left(\Delta_2 y_{i,00}, \Delta_2 y_{i,02}\right) \\ \dots \\ Var\left(\Delta_2 y_{i,12}\right) \end{pmatrix}$$

collects the variances and (auto)covariances needed for the minimum distance estimation and is of dimension $\frac{\dim(x_i)*(\dim(x_i)+1)}{2} = 105$.

For the estimation a corresponding vector f has to be constructed that contains the predictions of the parametric income model as outlined in Chapter 5 and 6. This vector f is dependent on the insurance parameters ϕ and ψ, on the variances of permanent and transitory shocks and on $Var\left(\xi_t\right)$ and $Var\left(u_t^c\right)$, i.e. on all parameters that I seek to estimate. The targeted parameters are collected in a vector θ so that $f\left(\theta\right)$ has the following structure:

$$f\left(\theta\right) = \begin{pmatrix} 2\phi_{00}^2 Var\left(\zeta_{00}\right) + 2\psi_{00}^2 Var\left(\epsilon_{00}\right) + 2Var\left(\xi_{00}\right) + Var\left(u_{00}^c\right) + Var\left(u_{98}^c\right) \\ -Var\left(u_{00}^c\right) \\ \dots \\ 0 \\ \dots \\ 2Var\left(\zeta_{00}\right) + Var\left(\epsilon_{00}\right) + Var\left(\epsilon_{98}\right) \\ -Var\left(\epsilon_{00}\right) \\ \dots \\ 2Var\left(\zeta_{12}\right) + Var\left(\epsilon_{12}\right) + Var\left(\epsilon_{10}\right) \end{pmatrix}.$$

The minimum distance procedure estimates the parameter vector θ by minimizing the difference between m and $f\left(\theta\right)$ according to

$$\min_{\theta}\left(m - f\left(\theta\right)\right)' A \left(m - f\left(\theta\right)\right)$$

where A is a weighting matrix. Three specifications for A are most common: Equally-weighted minimum distance (EWMD) uses the identity matrix I as weighting matrix. This procedure is unbiased, but generally not considered to be efficient. Optimal minimum distance (OMD) sets $A = V^{-1}$ where V is the variance-covariance matrix of m. Although OMD is considered efficient, Altonji & Segal (1996) have raised concerns that it is biased in small samples. Diagonally-weighted minimum distance (DWMD) uses a diagonal matrix for A with the values on the main diagonal taken from V^{-1}. DWMD is a more general version of EWMD that allows for heteroskedasticity. For most of the analyses of this paper EWMD is used. However, the robustness section shows that the benchmark results do not change significantly when the slightly more general DWMD is used.

For the computation of the standard errors of θ the formula developed by Chamberlain (1984) is applied:

$$\widehat{Var\left(\widehat{\theta}\right)} = \left(G'AG\right)^{-1} G'AVAG \left(G'AG\right)^{-1}$$

with $G = \frac{\partial f(\theta)}{\partial \theta'}|_{\theta=\widehat{\theta}}$.

References

Altonji, J. G. & Segal, L. M. (1996). Small-sample bias in GMM estimation of covariance structures, *Journal of Business & Economic Statistics* **14**(3): 353–366.

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. (2015). Consumption inequality and family labor supply, *American Economic Review (forthcoming)* .

Blundell, R. & Preston, I. (1998). Consumption inequality and income uncertainty, *Quarterly Journal of Economics* **113**(2): 603–640.

Carbone, E. & Hey, J. D. (2004). The effect of unemployment on consumption: an experimental analysis, *The Economic Journal* **114**(497): 660–683.

Chamberlain, G. (1984). Panel data, *Handbook of econometrics*, Vol. 2, Elsevier, pp. 1247–1318.

Cutler, D. M. & Katz, L. F. (1991). Macroeconomic Performance and the Disadvantaged, *Brookings Papers on Economic Activity* **1991**(2): pp. 1–74.

Deaton, A. & Paxson, C. (1994). Intertemporal Choice and Inequality, *Journal of Political Economy* **102**(3): 437–467.

Feenberg, D. & Coutts, E. (1993). An introduction to the TAXSIM model, *Journal of Policy Analysis and Management* **12**(1): 189–194.

Flavin, M. & Yamashita, T. (2002). Owner-occupied housing and the composition of the household portfolio, *American Economic Review* **92**(1): 345–362.

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.

Hryshko, D. (2014). Correlated income shocks and excess smoothness of consumption, *Journal of Economic Dynamics and Control* **48**: 41 – 62.

Jappelli, T. & Pistaferri, L. (2010). The Consumption Response to Income Changes, *Annual Review of Economics* **2**: 479–506.

Jappelli, T. & Pistaferri, L. (2011). Financial Integration and Consumption Smoothing, *The Economic Journal* **121**(553): 678–706.

Kaplan, G. & Violante, G. L. (2010). How Much Consumption Insurance Beyond Self-Insurance?, *American Economic Journal: Macroeconomics* **2**(4): 53–87.

Kaplan, G., Violante, G. L. & Weidner, J. (2014). The Wealthy Hand-to-Mouth, *Brookings Papers on Economic Activity* **48**(1): 77–138.

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. (2011). How Do Households Respond to Income Shocks. Working Paper.

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. (2015). Income Shocks or Insurance-What Determines Consumption Inequality? Ruhr Economic Papers, No. 550.

Meyer, B. D. & Sullivan, J. X. (2013a). Consumption and Income Inequality and the Great Recession, *American Economic Review Papers & Proceedings* **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.