

Permanent versus Transitory Income Shocks over the Business Cycle

IFS Working Paper W19/29

Agnes Kovacs
Concetta Rondinelli
Serena Trucchi

Permanent versus Transitory Income Shocks over the Business Cycle*

Agnes Kovacs[†]

Concetta Rondinelli[‡]

Serena Trucchi[§]

October 31, 2019

Abstract

This paper investigates how different income shocks shape consumption dynamics over the business cycle. First, we break new ground by creating a unique, panel dataset of transitory and permanent income shocks, using subjective income expectations from the Dutch Household Survey. Second, we evaluate whether these observed income shocks help to explain contractions in aggregate consumption over the two most recent crises. We find that the income shocks experienced during the 2008-2009 Global Financial Crisis are of a different nature than the shocks experienced during the 2011-2012 Sovereign Debt Crisis, with the 2011-2012 shocks being perceived as more permanent. This helps explain why consumption falls less during the Global Financial Crisis, despite the fact that income declines more than during the Sovereign Debt Crisis.

Keywords: subjective expectations; income shocks; consumption; financial crisis

JEL classification: D12; E21.

*We would like to thank Orazio Attanasio, Alessandro Bucciol, Mauro Mastrogiamomo, Patrick Moran, Stefano Neri, Roberta Zizza and Francesco Zollino for helpful comments. We also thanks participants in seminar at Leiden University, Ca' Foscari University, and Mannheim University; Netspar International Pension Workshop 2018; RES Annual Conference 2018, CESifo Workshop on Subjective Expectations and Probabilities in Economics 2018; Bergamo Workshop on Household Consumption 2019; SIEP Annual Conference 2019. This project has also received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 655770 (Trucchi).

[†]Department of Economics, The University of Manchester and IFS, Arthur Lewis Building, Oxford Road, Manchester M13 9PL, United Kingdom (agnes.kovacs@manchester.ac.uk).

[‡]Bank of Italy, Economic Outlook and Monetary Policy Directorate, via Nazionale, 91, Roma, 00184, Italy (concetta.rondinelli@bancaditalia.it).

[§]Department of Economics, Ca' Foscari University of Venice, Cannaregio 873, Fondamenta San Giobbe, 30121 Venezia (serena.trucchi@unive.it).

1 Introduction

In most European countries, households' consumption response to the two most recent economic crises was remarkably different. During the 2008-2009 Global Financial Crisis, households' consumption dropped much less than their income. By contrast, during the 2011-2012 Sovereign Debt Crisis the drop in households' consumption was much sharper and more prolonged than the decline in their income. Why does consumption react to changes in income so differently across the two periods?

In this paper, we study whether the nature of income shocks can explain households' consumption dynamics over the last decade, with a special focus on the Global Financial Crisis and the Sovereign Debt Crisis. By doing so, we make two substantive contributions. First, using the approach of Pistaferri (2001) and data from the Dutch Household Survey (DHS), we identify permanent and transitory income shocks at the household level.¹ This allows us to build a unique panel dataset of these shocks for the period between 2006 and 2016 that covers the last two business cycles, the 2008-2009 and the 2011-2012 crises. Second, feeding the identified income shocks into a structural life-cycle model, we isolate the impact of income shocks from other potential factors (e.g. real interest rate) that might simultaneously affect consumption dynamics. We evaluate the importance of income shocks in shaping consumption comparing consumption profiles induced by the model to those observed in the data.

In the first part of the paper we consider the most widely used income process that assumes both permanent and transitory income shocks. Within this framework, we show that income shocks can be identified as different combinations of subjective income expectations and their realizations, following Pistaferri (2001). More specifically, permanent shocks are the changes in income expectations, while transitory shocks are differences between income realizations and their expectations, once we remove predictable life-cycle components. Using this theoretical result, we exploit the joint availability of subjective income expectations and realizations in a micro panel dataset, the Dutch Household Survey, to compute the level of income shocks between 2006 and 2016. To the best of our knowledge ours is the first attempt to identify household-level income shocks for an extended period.²

In terms of average income shocks, the two crisis are rather different. Our results show that before 2010 households mainly experience transitory income shocks, while after 2010 they primarily experience permanent income shocks. Based on this evidence,

¹Earlier empirical studies show that income changes are best described by the combination of permanent and transitory income shocks. See for example MaCurdy (1982) and Blundell and Preston (1998).

²Note that Pistaferri (2001), given data limitations, can only compute shocks for one specific year, while the paper by Attanasio, Kovacs, and Molnar (2017) identifies cohort averages of income shocks.

we argue that the Financial Crisis and the succeeding Sovereign Debt crisis are fundamentally different because households perceived income shocks differently. Looking at the distribution of the identified income shocks, instead, we find that the 2008-2009 and 2011-2012 crises behave quite similarly: variance of income shocks does not show counter-cyclicality, while the left-skewness of the shock distribution is countercyclical, which is line with findings from Guvenen, Ozkan, and Song (2014).

Different types of income shocks could lead to substantial differences in consumption responses. According to the simplest version of the permanent income hypothesis, only unanticipated permanent income shocks should induce substantial changes in consumption. On the contrary, expected or temporary income shocks should not alter consumption significantly. In this context, we expect consumption to change substantially only over the 2011-2012 crisis, when income shocks are perceived as permanent, but not over the 2008-2009 crisis, when shocks are mainly perceived as transitory. While this simple prediction is qualitatively in line with what we observe in aggregate data (i.e. a fall in income generates proportionally smaller drop in aggregate consumption over the Global Financial Crisis compared to what is observed over the Sovereign Debt Crisis), it is crucial to quantify the transmission of income shocks to consumption at the household level. Lack of suitable information on consumption expenditure in the Dutch Household Survey though prevents us from directly estimating this relationship. However, to overcome this limitation in the data we next make use of a structural model.

In the second part of the paper, we develop a dynamic life-cycle model of consumption and simulate households behavior between 2006 and 2016 assuming that they face the permanent and transitory income shocks we identify in the first part of the paper. In this way, we are not only able to measure the effect of income shocks on household-level consumption, but also to separate the effect of income shocks from other possible factors that affected consumption between 2006 and 2016. We find that in the model income shocks alone cannot generate the significant fall in consumption for 2008-2009 that is observed in the data during the Financial Crisis (1.6%). On the contrary, income shocks imply a large aggregate consumption drop for 2011-2012 in the model (2.9%), which is similar in magnitude to the consumption drop experienced by the Dutch economy during the Sovereign Debt Crisis (2.8%). Moreover, the model predicts that young households face the most dramatic consumption drop in 2011-2012.

Our analysis relies on a life-cycle model of consumption, which incorporates important real-life features. Households are liquidity constrained, hence they accumulate wealth not only for supporting their retirement years but also for precautionary reasons. They face income uncertainty, which can be either permanent or transitory in nature, as observed in the data. Households belong to different cohorts based on their date of

birth, which is crucial to take into account for at least two reasons. First, households in different cohorts have experienced different income shocks, hence their income and consumption trajectories can be very different. Second, households in different cohorts are by definition at different stages of their life-cycle and as a result their consumption reacts differently to similar income shocks. In turn, we define aggregate consumption in the model as the weighted sum of household consumption over different cohorts.

How household consumption reacts to transitory and permanent income shocks is a longstanding question in macroeconomics, which is crucial both to understand consumption behavior and to evaluate policy changes.³ Identification of the level of these income shocks, however, is challenging for many reasons. Even if panel data on income is available, we only observe total income changes, rather than transitory and permanent income changes separately. For this reason, a large empirical literature tries to proxy the different type of income shocks by identifying episodes of income changes using for example weather shocks in developing countries, unemployment or lottery winnings (Wolpin, 1982; Paxson, 1993; Browning and Crossley, 2001; Fagereng, Holm, and Natvik, 2019). Moreover, even if we are able to identify income shocks, often we can only do that from the econometrician’s point of view. This information asymmetry between individuals and econometrician may introduce a misclassification problem of income changes as we often cannot distinguish between shocks to income versus expected changes in income. The literature based on subjective income expectations attempts to come around this problem by directly asking individuals about their subjective income expectations. This is the approach exploited by Pistaferri (2001), endorsed by Manski (2004) and also explored in this paper.

Our paper contributes to a large literature on measuring the transmission of different income shocks to consumption. Important examples include Pistaferri (2001), Blundell, Pistaferri, and Preston (2008), Kaufmann and Pistaferri (2009), Kaplan and Violante (2010), and Guvenen and Smith (2014). It is also closely related to empirical studies that examine how idiosyncratic income shocks are affected by business cycle movements, see for example Storesletten, Telmer, and Yaron (2004), Guvenen, Ozkan, and Song (2014). The closest papers to the present one are Pistaferri (2001) and Attanasio, Kovacs, and Molnar (2017). They both identify income shocks using data on subjective income expectations. Pistaferri (2001) uses the Italian Survey on Household Income and Wealth (SHIW), which collects information on subjective income expectations and realizations in two specific waves (1989 and 1991). Because of data restriction, this paper can only provide a snapshot of transitory and permanent shocks under strong assumptions about individuals’ information set. The paper by Attanasio, Kovacs, and Molnar (2017)

³For example, the effectiveness of fiscal stimulus, such as tax rebates, in a recession depends on how much households spend out of rebates.

combines two data sources to construct a synthetic panel: one for income realization from the Consumer Expenditure Survey and one for subjective expectations from the Michigan Survey. Given the synthetic panel structure of their data, they can only identify cohort-level income shocks. Our analysis differs from theirs in that the joint availability of subjective income expectations and realizations in the Dutch Household Survey allows us to construct a household-level panel dataset of permanent and transitory shocks.

The rest of the paper is organized as follows. First, looking at aggregate data, we put forward our hypothesis that understanding the nature of income shocks is more important than the size of income shocks in explaining consumption behavior. As we can't directly observe temporary or permanent income shocks, we show the methodology we use to identify them (Section 2). We demonstrate that the identification of income shocks requires data both on subjective income expectations and on income realizations. We then describe the Dutch Household Survey, which uniquely collects these information in one panel dataset (Section 3). We show that the identified shocks are mainly transitory before 2010, while mostly permanent after 2010 (Section 4). Moreover, using these shocks in a structural life-cycle model, we illustrate how transitory and permanent income shocks are reflected in consumption. More specifically, we demonstrate that the observed income shocks can explain the consumption behavior during the 2011-2012 crisis, but not in 2008-2009 (Section 5). We conclude by discussing next steps for future research (Section 6).

2 Consumption, Income and Income Shocks

Consumption and income are naturally tied together. Income over the life-cycle determines lifetime resources and therefore consumption possibilities today and tomorrow. Consequently, understanding how consumption responds to changes in income/resources is crucial to understand consumer behavior, and also to evaluate policies that affect households' available resources. According to the textbook version of the permanent income hypothesis (PIH), only unanticipated permanent income shocks should induce substantial changes in consumption. Expected or transitory income shocks, instead, should not alter consumption significantly.


In this Section, we first illustrate dynamics of aggregate consumption in light of predictions from the PIH model. We then describe the strategy used to identify income shocks of different nature (permanent versus transitory), which uses combinations of subjective income expectations and income realizations. There exist only a few micro datasets that jointly keeps record of households' income expectations and income realizations. One of them is the Dutch Household Survey, which we rely on throughout this

paper.

2.1 Aggregate Behavior

Over the last decade, most European countries, including the Netherlands, experienced two periods of substantial decline in aggregate consumption. These periods coincide with the two most recent crises: the 2008-2009 Global Financial Crisis and the 2011-2012 Sovereign Debt Crisis.⁴ Figure 1 illustrates disposable income and consumption dynamics of the Dutch economy for the period between 2006 and 2016. During the

Figure 1: Households' disposable income and consumption



Notes: Notes: Our calculations based on Eurostat quarterly data; indices, 2008-Q1=100. Consumption is real Household and NPISH (Non-profit institutions serving household) final consumption expenditure. Disposable income is the sum of compensation per employee, gross operating surplus and mixed income deflated with GDP deflator. Shaded area indicates crisis period.

2008-2009 Global Financial Crisis, households' disposable income in the Netherlands shrinks by 3.6% and simultaneously their consumption drops by 1.6%. In turn, during the 2011-2012 Sovereign Debt Crisis the fall in disposable income is roughly 2.2%, which coincides with a large 2.8% contraction in aggregate consumption. The fact that during the Sovereign Debt Crisis consumption falls significantly and more than households' disposable income suggests a potentially important role of the nature of income shocks, as endorsed by the PIH model. More specifically, Figure 1 suggests that income shocks during 2011-2012 may be perceived more permanent than those occurred during 2008-

⁴Crises years are defined as a contraction in Gross Domestic Product (GDP) for two consecutive quarters or longer.

2009, which in turn triggers larger consumption responses during the Sovereign Debt Crisis.

The most important driver of households' consumption besides their actual income and intertemporal prices is undoubtedly what households expect about their future income. Whenever households face an unexpected income shock, they re-evaluate their optimal resource allocations and adjust their consumption based on what they think about the persistence of the income shock (i.e. how much the shock affects their income in the future). Consumption adjustments are in accordance with households' consumption smoothing motive. For instance, after an income shock which households expect to be permanent, households feel (lifetime) wealthier and might increase their current consumption. By contrast, after an income shock which households expect to be transitory, households only feel temporarily wealthier, which they prefer to smooth over time and they don't change their consumption much.

Having information on how households perceive different income shocks could help us to confirm the premise that income shocks during the 2011-2012 Sovereign Debt Crisis are perceived more persistent than the shocks over the 2008-2009 Global Financial Crisis. It could also help us to better understand consumption dynamics in the last decade. Data on different income shocks though is not available. Even if we have panel data on income, we only observe total income changes, rather than transitory and permanent income changes, separately. For this reason, we create our own household-level panel dataset of transitory and permanent income shocks by using subjective income expectations.

2.2 Identifying Income Shocks

In order to identify the permanent and transitory components of income shock, we follow the approach proposed by Pistaferri (2001) and exploited by Attanasio, Kovacs, and Molnar (2017). This method hinges on the relationship between subjective income expectations and the nature of income shocks.

We start with the following, standard decomposition of the log income, as in Blundell, Pistaferri, and Preston (2008):

$$y_{it} = \Pi' Z_{it} + \alpha' V_i + p_{it} + \varepsilon_{it} \quad (1)$$

$$\Pi' Z_{it} = \pi_0 + \pi_1 age_{it} + \pi_2 age_{it}^2$$

where y_{it} is the log of household income i at time t ; $\Pi' Z_{it}$ is a deterministic time-varying component (second order polynomial of age), and $\alpha' V_i$ is a deterministic time invariant component, which includes gender, education and household fixed effects. p_{it} and ε_{it} are, respectively, the permanent and transitory components of income of household i at

time t .⁵ The transitory component is i.i.d, while the permanent component follows a Markov process:

$$p_{it} = p_{it-1} + \zeta_{it} \quad (2)$$

where ζ_{it} is the permanent income shock. Permanent and transitory shocks are assumed to be orthogonal (at all leads and lags) and to be unanticipated, similarly to Blundell, Pistaferri, and Preston (2008). Combining equations (1) and (2) we obtain the following equation for change in log of income:

$$\Delta y_{it} = \Pi' \Delta Z_{it} + \zeta_{it} + \Delta \varepsilon_{it} \quad (3)$$

Note that Blundell, Pistaferri, and Preston (2008) use covariance restrictions of the income growth characterized by equation (3) to identify the variances of different income shocks. We use a completely different strategy by which we are not only able to derive the variances, but also the levels of the income shocks.

Under the assumption of rational expectations, we can express the two income shocks as a function of income expectations and realizations, which is described in details in Appendix A.1. As a result, transitory and permanent income shocks can be respectively rewritten as:

$$\begin{aligned} \varepsilon_{it} &= -\mathbb{E}[\Delta y_{it+1} | \Omega_t] + (\gamma_0 + \gamma_1 \text{age}_{it+1}) = \\ & y_{it} - \mathbb{E}[y_{it+1} | \Omega_t] + (\gamma_0 + \gamma_1 \text{age}_{it+1}) \end{aligned} \quad (4)$$

and

$$\zeta_{it} = \mathbb{E}[y_{it+1} | \Omega_t] - \mathbb{E}[y_{it} | \Omega_{t-1}] - (\gamma_0 + \gamma_1 \text{age}_{it+1}) \quad (5)$$

where \mathbb{E} is the expectation operator that takes expectations of variables conditional on the information set available to households. Ω_t is the set of information available to household i at time t . Coefficients γ_0 and γ_1 are functions of the parameters π_1 and π_2 , the coefficients on the second-order polynomial of age in equation (1).⁶

In this way, one can give a straightforward interpretation of the temporary and permanent income shocks based on subjective income expectations and realizations. Apart from a predictable age affect, temporary income shock, ε_{it} , is identified by the gap between income realization and its subjective expectation; while permanent shock, ζ_{it} , is identified as the change in the subjective expectations of income. Therefore, this method allows us to identify temporary and permanent income shocks separately using

⁵A formal test of this assumption is discussed in Section 4.

⁶Assuming that individuals only face unanticipated income shocks is crucial for our identification strategy. When we allow for both anticipated and unanticipated income shocks, it is not possible to identify the level of income shocks, but it is possible to compute the variances of the shocks (as shown by Kaufmann and Pistaferri (2009)).

data on observed and expected income only.

3 Data

In this Section, we describe the Dutch Household Survey (DHS), the dataset we use for our analysis in this paper. The key feature of this dataset is the joint availability of expected and realized income at the household-level, which is crucial to separately identify transitory and permanent income shocks. To the best of our knowledge, this is the only dataset which collects subjective income expectations together with their realizations at the household level for a longitudinal sample covering a time period of more than 10 years, including the two crises episodes described earlier.

The Dutch Household Survey is a longitudinal survey representative of the Dutch-speaking population. Data are collected annually on behalf of the Dutch National Bank via an online survey. Everyone aged 16 or over is interviewed within each household. To the purpose of our analysis, we restrict the sample to working respondents aged 21-65 and we focus on the period between 2006 and 2016.⁷ We also drop individuals who have obvious misunderstanding of the subjective expectation questions or who give inconsistent answers to these questions. We end up with a sample of about 450 individuals per year.⁸ Each respondent is observed, on average, 2.8 times.

Descriptive statistics of the final sample, for three representative waves, are reported in Table 1. The average age of the respondents is about 50 and households are composed of less than three members of which about 2 are adults. Roughly 70% of respondents work and 10% are retired. There is also significant heterogeneity in terms of education: middle educated respondents and those having attended vocational schools represent about 40% and 20% of the sample, respectively. Hereafter, we describe the two main variables of interest: income realizations and subjective expectations.

Income Realizations.

The measure of household income that we use in the empirical analysis is gathered through the following question, which is asked to household heads and their spouses:

“What is the total net income for your household in [year]? The total net income for your household is the net income of all household members com-

⁷The latter sample restriction is due to data limitations. Hence, over this period, DHS collects self-reported income amounts; if the value is not reported respondents choose among income bracket. Until early 2000s, instead, income was collected in brackets only, providing a less precise measure of observed income.

⁸For each year, we drop from the sample the top and bottom 5% of observed and expected income and of permanent and transitory shocks.

Table 1: Descriptive statistics

	2006		2011		2015	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Age	47.488	10.875	51.130	10.638	48.537	11.091
No. of hous. members	2.657	1.358	2.468	1.245	2.636	1.429
No. of adults	1.770	0.445	1.762	0.468	1.773	0.514
Work	0.656	0.476	0.698	0.460	0.770	0.421
Retired	0.085	0.279	0.114	0.318	0.047	0.211
Unemployed	0.024	0.154	0.019	0.135	0.034	0.180
No education	0.036	0.187	0.024	0.153	0.018	0.133
Low education	0.268	0.443	0.241	0.428	0.158	0.365
Middle education	0.412	0.493	0.415	0.493	0.432	0.496
Vocational education	0.199	0.400	0.177	0.382	0.202	0.402
University education	0.076	0.265	0.140	0.348	0.183	0.388
Observed income	30,389	10,998	33,512	11,379	33,177	12,692
Expected income	30,609	12,241	32,675	11,296	30,027	11,285
No. obs	578		378		387	

Notes: Our calculations from DHS; real values (euros 2010).

bined. Net income means the income after deduction of taxes and social security benefits.”

This question is particularly well-suited to our purpose, since it refers to the same income measure which is used to elicit income expectations, namely total net household income. We also find that the majority of net income comes from labor earnings. On average, financial revenues represent less than 4% of net income for all the respondents, and less than 10% if we consider owners of financial assets only.⁹ To gauge the contribution of labour earnings to total household resources, we also examine the correlation between self-assessed total net income - our measure of interest - and gross labor income.¹⁰ The two variables turn out to be strongly correlated, with a regression line close to the 45 degrees line.¹¹ This result further supports the key role of labour income, which represents the main determinant of total household income.

Subjective Income Expectations.

Subjective income expectations are collected through two sets of questions. Respondents start reporting the lower and upper bounds for expected income, respectively:

⁹Less than 2% of households declare income from housing wealth.

¹⁰Gross labour income is obtained as the sum of earnings of all household’s members. Net labor income is not available.

¹¹The plot of the joint distribution of logarithm of net total income and the logarithm of gross labour earnings, along with the regression line, is shown in Figure A.1 in the Appendix A.4. The estimated regression is $\ln y = 1.864 + 0.967 \ln x$, where the coefficient for $\ln x$ is significant at the 1% level. The sign of the intercept is due to the fact that labour income is gross, while total income is net.

“We would like to know a little bit more about what you expect will happen to the net income of your household in the next 12 months. What do you expect to be the lowest (highest) total net yearly income your household may realize in the next 12 months?”

The interval between the lower (l) and upper (h) bounds is divided into equal intervals:

$$l + (h - l)x, \text{ with } x = \frac{2}{10}, \frac{4}{10}, \frac{6}{10}, \frac{8}{10}.$$

Respondents declare, then, the probability that future income will be lower than the threshold $l + (h - l)x$. More precisely, for each threshold, they are asked

“What do you think is the probability (in percent) that the net yearly income of your household will be less than euro [threshold] in the next 12 months?”

We exploit this information to compute the expected value of net household income.¹²

Figure 2 plots the one one-year-ahead average income expectations, together with the actual income data. There are two episodes of sudden drop both in the expected and in the observed income. The first drop happens around the Global Financial Crisis, when these two variables fall by similar magnitudes. The second drop is at the time of the Sovereign Debt Crisis, when subjective income expectations fall much more than observed income itself. Moreover, expectations about future income remain low till the end of our observational period, till 2016.


Identification of income shocks and interpretation of our results hinges on the reliability of expected income measure. For this reason, we provide evidence to support the informative value and the accuracy of subjective expectations elicited by the DHS survey, which has never been done for this specific dataset. Hereafter, we document their well-behaved distribution, the internal coherency between different questions about the future, and the predictive power of subjective income expectations, as suggested by Manski (2004).

First, the distribution of subjective income expectations has a regular shape and shadows the one of income realizations.¹³ This evidence is reassuring about the limited diffusion of random or inaccurate responses, which points to the reliability of expected income variable. Second, we test internal coherency between subjective expectations on income and on job status. Working and not-working respondents are asked, respectively, about the probability of loosing or finding a job in the next 12 months. We test

¹²Heterogeneity in the way income expectations are elicited over time is discussed in Appendix A.2.

¹³The density function of income expectations and realizations for the pooled cross-section dataset is plotted in Figure A.2 in the Appendix A.4. The distribution of expectations is more left-skewed and presents a mass for very low annual income (close to 0), consistent with pessimistic expectations over the period.

Figure 2: Observed and expected income



Notes: Real values (euros 2010). Mean values refer to weighted means. Shaded area indicates crisis period.

the conditional correlation between expected job status and income by regressing the latter on the probability of job loss - or job finding -, controlling for an unemployment indicator and a set of covariates. Estimation results, which are reported in Table A.2 in the Appendix A.4, show a correlation which goes in the expected direction. Working respondents reporting higher probability of losing their job are also significantly more pessimistic about future income. On the contrary, the effect of self-reported probability of finding a job on expected income is positive, although not statistically significant (possibly also because the small number of unemployed respondents in the sample). Also, actual income turns out to be highly correlated with the expected one, with conditional correlation of 0.85 and significance of 1%, and unemployed respondents reports lower expected income on average. All in all, these results support the internal coherence among questions eliciting subjective expectations, corroborating the informative power of expected income.

If declared income predictions are accurate and exploit the respondent's complete information set, we must detect an *ex-post* correlation between subjective income expectations and realizations. Therefore, even if we cannot observe the information set used to make predictions on future income, we can evaluate their accuracy by testing whether they have some "predictive power" to explain future income realization (Manski, 2004). To this purpose, we exploit the longitudinal component of the dataset and we regress income realization on subjective expectations elicited one period ahead. Results are

shown in Table A.3 in the Appendix A.4. As it is seen in the first column, unconditional correlation is higher than 0.8. Even if the inclusion of other controls reduces its magnitude, the correlation stays as high as 0.5 in all the specifications. As shown in columns 4 and 5, the correlation between expectations and future income realizations is lower for households, which are more uncertain about their future income (proxied by the square root of income expectations). These findings suggest that income expectations have a high predictive power for actual income realization, further sustaining their reliability.

Timing.

An important aspect to be discussed is the time period which our collected survey information refers to. First of all, we consider a time span of one year, since both questions on observed and expected income refer to a 12 months period. Identification of transitory shocks requires computing the difference $y_{it} - E[y_{it+1}|\Omega_t]$, as shown in equation 4. Since the DHS questionnaire measures y_{it} as observed household income earned in the previous calendar year, expected income should be ideally elicited on January 1st (and referring to the coming calendar year). The gap between the date of the interview and the beginning of the year is, thus, a source of time discrepancy. In our sample, this issue is mitigated by the fact that more than two thirds of interviews are run between weeks 10 and 18, and only 8% of respondents reply after week 30. In our baseline measure of income shock, we implicitly assume that no shock has occurred within this time span (January 1st and time of the interview), but we also measure income shocks i) including only respondents with a time discrepancy lower than 18 weeks and ii) using a ‘corrected’ measure of observed income, which is meant to be consistent with expected income by measuring observed income realizations during the 12 months preceding the interview.¹⁴ Figure A.3 shows that the distribution of observed income (referred to the previous year) and the distribution of ‘corrected’ household income (referring to the 12 months before the month of the interview) are broadly comparable. Figure A.4 shows that the paths of transitory shock based on either observed income or ‘corrected’ income are very similar, supporting our assumption that time discrepancies do not play a big role in this context. Permanent shocks, instead, hinge on a measure of change in subjective expectation, e.g. $E[y_{it+1}|\Omega_t] - E[y_{it}|\Omega_{t-1}]$ (see equation 5). Time discrepancy, in this case, refers to the moment when subjective expectations are retrieved, in two subsequent waves. This discrepancy is less than one week in one third of cases, while it is lower than four weeks in the large majority of interviews (almost 60%).

¹⁴For instance, if the survey is run during week 10 of year 2010, we construct ‘corrected’ income as a weighted average of observed income in 2010 and 2009, where the weight for the first component is given by the incidence of income 2010 in the calculation of income in the previous 12 months (i.e. 10 weeks out of 52). In this case, $y_{corr} = (10 * y_{2010} + (52 - 10) * y_{2009})/52$.

Predictable Income Component.

Having data both on income realizations and subjective income expectations makes it easy to calculate one-year-ahead income growth expectations. They let us identify γ_0 and γ_1 , the coefficients of the deterministic income component in equations (4)-(5). Simply regressing reported expected income growth on a constant and on age, we can obtain estimates for γ_0 and γ_1 .¹⁵ The estimated coefficients are $\hat{\gamma}_0 = .0063$ and $\hat{\gamma}_1 = -.0002$. The combinations of income realizations, subjective income expectations, and predictable income components over the life-cycle identify transitory and permanent shocks as expressed in equations (4) and (5).

4 Identified Income Shocks

In this Section we first confirm that the assumptions we make in order to identify income shocks are satisfied in the data, we then analyze the behavior of the transitory and permanent income shocks. We illustrate the evolution of average income shocks between 2006 and 2016 and show how their dynamics vary by cohorts, by the level of income and over the business cycle.

Identifying Assumptions.


The most crucial assumption we make for the identification of the income shocks in Section 2.2 is that at the household level, both transitory and permanent income shocks have to be i.i.d. In order to test these assumptions, we restrict our sample to households observed at least twice in the period we consider, and we test the null hypothesis of no autocorrelation of shocks with the Q-statistics suggested by Ljung and Box (1978). Instead of testing autocorrelation at different lags separately, the Ljung-Box statistics tests whether any group of autocorrelations are different from zero over a time series. The null hypothesis of no autocorrelation in the transitory (permanent) income shock cannot be rejected for more than 95% (90%) of households. Therefore, the behavior of our calculated income shocks are not inconsistent with the income process that is widely used in labor economics, and assumed in Section 2.2.

¹⁵To avoid possible biases related to the two recessions, we enlarge the time span and we include in the sample all the household heads interviewed from 1997 to 2016.

Averages.

The dynamics of average transitory and permanent income shocks between 2006 and 2016 are illustrated in Figure 3.¹⁶ During the Global Financial Crisis, households are exposed to negative transitory shocks, which only become positive after the middle of 2010. At the same time they face a small, one-period negative permanent income shock in 2009. Our unexpected finding is that this permanent income shock does not mirror the severity of the 2008-2009 financial crisis. During the Sovereign Debt crisis, in turn, households are exposed to large and positive transitory shocks. At the same time they face a prolonged period of time with large, negative permanent income shocks.

Figure 3: Permanent and transitory shocks



Notes: Real values (euros 2010). Mean values refer to weighted means. Shaded area indicates crisis period.

According to our identification, differences in income shocks over the two crisis periods reflect differences in households' subjective income expectations. Based on equation (4), negative transitory shocks, which we witness up until 2010, indicate that subjective income expectations are systematically higher than income realizations, i.e. households are in general optimistic. After 2010, the sign of transitory shocks changes to positive, showing that income expectations are below their realizations, i.e. households are in general pessimistic. According to equation (5), small permanent income shocks, which we observe up until 2010, are due to limited revision in subjective income expectations

¹⁶The evolution of median transitory and permanent shocks is plotted in Figure A.5 in Appendix A.4. The pattern is robust to the normalisation with respect to total income in previous year (Figure A.6).

over this period. Between 2010 and 2015, permanent income shocks are much larger and more negative, indicating a constant downward revision of households' income expectations. The upward revision of expectations only started in the last year of our sample, in 2016, where the permanent income shock became positive.

These patterns confirm the premise that the nature of the income shocks over the two crises are very different and that triggered different consumption responses. Shocks over the 2008-2009 Global Financial Crisis are mainly perceived as transitory, which leads to only a modest response of consumption. On the contrary, shocks over the 2011-2012 Sovereign Debt Crisis are primarily perceived as permanent, therefore consumption drops much more significantly.

Heterogeneities by Age and Income.

Average income shocks potentially mask some heterogeneity across households, which can contribute to shed light on the channels behind the dynamic of aggregate variables. For this reason, we next illustrate the time trend of income shocks according to cohort and households' financial position.

We start by analyzing average values of income shocks for three different cohorts of households, who were born in the following years: 1950-1954, 1960-1964, and 1970-1974.¹⁷ Figure 4 shows the dynamics of permanent income shocks for the three different cohorts. The figure highlights two important facts. First, only the 1960-64 cohort faced significant negative permanent income shocks over the 2008-2009 crisis (i.e. there was only significant revision in households' income expectations for cohort 1960-64). Second, all the cohort faced large and negative permanent income shocks over the Sovereign Debt Crisis, but the oldest cohort suffered the most negative shocks (see also Table A.4, which shows the evolution of the shocks for all the cohorts).

Figure 5, in turn, shows the dynamics of transitory income shocks for the three different cohorts. It highlights that transitory income shocks are negative for all the cohorts during the 2008-2009 crisis, while they are positive for all the cohorts during the 2011-2012 crisis. Young cohorts face the most negative and least positive transitory income shocks over the 2006-2016 period.


¹⁷Cohorts are defined by households' date of birth. Cohort 1950-1954, for example, includes households born between 1950 and 1954. We only consider cohort-year cells with more than 55 observations.

Figure 4: Permanent shocks by cohort



Notes: Sample size for each cohort: > 55 observations. Real values (euros 2010). Mean values refer to weighted means. Shaded area indicates crisis period.

Figure 5: Transitory shocks by cohort



Notes: Sample size for each cohort: > 55 observations. Real values (euros 2010). Mean values refer to weighted means. Shaded area indicates crisis period.


Another relevant dimension to gauge income shock heterogeneity is household income, therefore next we group households by income quartile and examine how their

income shock differ. Results are shown in Figure A.7 in the Appendix A.4.¹⁸ While the level of transitory income shocks are very similar across households in different income quartiles, the same is not true for permanent shocks. The 2011-2012 crisis affected households in higher income quartiles the most. In fact, the permanent income shocks faced by the top two quartiles over this period was on average twice as large as the shocks faced by the bottom two quartiles.

Variances.

Another relevant aspect related to income shock heterogeneity is whether and to what extent variances of idiosyncratic income shock are affected by the business cycle. Storesletten, Telmer, and Yaron (2004) show that idiosyncratic permanent shock variances are countercyclical, which result in higher income uncertainty during recessions: households can receive both larger positive and larger negative permanent income shocks. In contrast, Guvenen, Ozkan, and Song (2014) document greater uncertainty in recessions without an increasing chance of upward movements in income. They show that during recessions large upward income movements become less likely, without a change in the center of income shock distribution. This results in a countercyclical left-skewness.

Figure 6: Kernel densities of permanent income shocks



To test these two theories on our data, we plot the kernel densities of the calculated permanent income shocks for the non-crisis periods in our sample (2006-2007, 2010, 2013-2015), and for two crisis periods (2008-2009 and 2011-2012). Results are shown in Figure

¹⁸Results are robust when we focus on permanent and transitory shock as a fraction of income in $t - 1$, as shown in Figure A.8.

6. The center of the income shock distribution does not move much during recessions, compared to the non-recession period, while the tails of the shock distribution move quite asymmetrically. During the Sovereign Debt crisis, when permanent income shocks were substantial, large positive income shocks become less likely, whereas the probability of experiencing large negative income shocks increases. These findings reinforce the results from Guvenen, Ozkan, and Song (2014), that idiosyncratic permanent shocks are not countercyclical, instead their left-skewness is countercyclical.¹⁹

5 Life-Cycle Simulations

The lack of suitable information on consumption expenditure in the Dutch Household Survey prevents us from assessing the transmission of income shock transmission to consumption household-level using our dataset. In this Section, however, we make use of a structural model to get an insight of the importance of income shocks.

We solve a full life-cycle model and simulate households between 2006 and 2016 assuming that they face the permanent and transitory income shocks we identify in Section 4. In this way, we are not only able to see the effect of income shocks on household-level consumption, but also to separate the effect of these shocks from other possible factors that affected consumption between 2006 and 2016. Finally, we can aggregate simulated, household-level consumption and compare the resulting consumption dynamics to aggregate consumption profile observed in the data.

5.1 Model Structure

We build a dynamic single-asset model of life-cycle consumption and savings, where households face permanent and transitory income uncertainty. They live for T periods as adults: they work for W periods and retire afterwards. Households maximize their present discounted lifetime utility, which only depends on their non-durable consumption. To reallocate resources between periods, households have access to one-period bond, which yields a gross interest rate of R^X . There is no credit market in the model, hence households are liquidity constrained at the beginning of their life and accumulate wealth both for life-cycle and precautionary purposes. Households also belong to different cohorts based on their date of birth, which is crucial to take into account for at least two reasons. First, households in different cohorts have experienced different income shocks, hence their income and consumption trajectories can be very different. Second, households in different cohorts are, by definition, at different stages of their

¹⁹In Figure A.9 in the Appendix A.4, we also present the distributional changes for the transitory income shocks.

life-cycle and, as a result, their consumption reacts differently to similar income shocks. Later, we define aggregate consumption in the model as the weighted sum of household consumption over different cohorts.

We formulate household i 's value function in period t in a recursive form as follows:

$$V_{i,t}(X_{i,t}, P_{i,t}) = \max_{\{C_{i,t}\}} U(C_{i,t}) + \beta \mathbb{E}_t V_{i,t+1}(X_{i,t+1}, P_{i,t+1}), \quad (6)$$

subject to

$$X_{i,t+1} = R^X(X_{i,t} - C_{i,t}) + Y_{i,t+1} \quad (7)$$

where β is the discount factor, $C_{i,t}$ is non-durable consumption, $Y_{i,t}$ is labor income, and $P_{i,t}$ is the permanent part of the labor income, to be defined later in this section. $X_{i,t}$ is the cash-on-hand, defined as the sum of savings and labor income in period t .

Utility Function.

The period utility function is a CRRA (Constant Relative Risk Aversion) function in nondurable consumption.

$$U(C_{i,t}) = \frac{C_{i,t}^{1-\rho}}{1-\rho} \quad (8)$$

where $\rho \geq 0$ is a curvature parameter, which equals to the relative risk aversion parameter and to the inverse of the elasticity of intertemporal substitution.

Sources of Uncertainty and Cohorts.

In this framework, the only source of uncertainty is idiosyncratic labor income. In line with the income process described by equations (1) and (2), we assume that (log) labor income is exogenously described by a combination of deterministic and random components at any time before retirement. In addition, we assume that income shocks have a cohort-specific component. Here, we index the income process by c to distinguish cohort-specific and idiosyncratic components, but for clarity we drop the superscript c from the rest of the exposition. The (log) labor income, $y_{i,t}^c$, for household i belonging to cohort c at time t is defined as:

$$y_{i,t}^c = G_t + p_{i,t}^c + \varepsilon_{i,t}^c \quad (9)$$

with G_t being a deterministic function of age only; $p_{i,t}^c$ is the permanent income component, and $\varepsilon_{i,t}^c$ is the transitory income shock for the same household. The permanent

income component follows a martingale process:

$$p_{i,t}^c = p_{i,t-1}^c + \zeta_{i,t}^c \quad (10)$$

where $\zeta_{i,t}^c$ is the shock on the permanent income. We assume that both of the innovations on log income can be decomposed into a cohort-specific and a household-specific part. Therefore, we write them as follows

$$\varepsilon_{i,t}^c = \varepsilon_t^c + \varepsilon_{i,t} \quad (11)$$

$$\zeta_{i,t}^c = \zeta_t^c + \zeta_{i,t} \quad (12)$$

where ε_t^c and ζ_t^c are the cohort-specific income shocks. $\varepsilon_{i,t}$ and $\zeta_{i,t}$ are household-specific income shocks, which are assumed to be normally distributed, serially uncorrelated, and independent.²⁰

$$\varepsilon_{i,t} \sim N(-0.5\sigma_\varepsilon^2, \sigma_\varepsilon^2) \quad (13)$$

$$\zeta_{i,t} \sim N(-0.5\sigma_\zeta^2, \sigma_\zeta^2) \quad (14)$$

Labor income at any time after retirement is a constant, a , fraction of the last working year's permanent labor income. One can think of this as a pension that is wholly provided by the employer and/or the state.

5.2 Solution and Simulation

Our life-cycle problem cannot be solved analytically, so we apply numerical techniques. Given the finite nature of the problem, a solution exists and can be obtained by approximating optimal policy functions by backward induction.

Solution.

We use the backward induction technique over the normalized value function of the households to obtain the optimal policy functions.²¹ Expectations in the model refer to uncertain incomes, while they are evaluated using the Gauss-Hermite approximation. Since the innovations of income are log-normally distributed random variables in each period, we are able to use a two-dimensional Gauss-Hermite quadrature to approximate

²⁰The assumption of log-normality of the income shocks with given parameters is a simplification. In this case the mean values of the level of the income shocks equal 1.

²¹Following Carroll (1992), variables are normalised by permanent income for ease of computation. In Appendix A.3, we show the detailed derivation of the standardized model.

the expectations as follows

$$\begin{aligned}
\mathbb{E}_t V_{t+1}(x_{t+1}) &= \int V_{t+1}(x_{t+1}(Z, N)) dF(Z) dF(N) \\
&= \int_{-\infty}^{\infty} \frac{1}{\pi} V_{t+1}(x_{t+1}(\sqrt{2}\sigma_Z Z, \sqrt{2}\sigma_N N)) e^{-(Z^2+N^2)} \\
&\approx \sum_{i \otimes j \otimes k} \frac{1}{\pi} w_i^{GH} w_j^{GH} V_{t+1}(x_{t+1}(\sqrt{2}\sigma_Z Z_i^{GH}, \sqrt{2}\sigma_N N_j^{GH}))
\end{aligned} \tag{15}$$

where Z_i^{GH} and N_j^{GH} are the Gauss-Hermite nodes, while w_i^{GH} and w_j^{GH} are the corresponding weights.

Cohorts.

In order to take into account that different groups of households might have experienced different income shocks, we define cohorts and simulate their behavior separately. Given that our focus period is between 2006 and 2016, we define ten cohorts based on their age in 2006. We use 5-year age intervals between ages 20 and 65.

Table 2: Parameters for the benchmark model

Parameter		Value	Source
T	<i>Number of years as adult</i>	60	
W	<i>Number of years as worker</i>	45	
β	<i>Discount factor</i>	0.95	
ρ	<i>Risk aversion parameter</i>	1.5	Blundell, Browning, and Meghir (1994)
Constant	<i>Age-spec income, constant</i>	8.668	Own calculations, DHS
Age	<i>Age-spec income, linear trend</i>	0.058	Own calculations, DHS
$Age^2/10$	<i>Age-spec income, quadratic trend</i>	-0.001	Own calculations, DHS
a	<i>Replacement rate</i>	0.75	Own calculations, DHS
σ_ζ^2	<i>Var.permanent income shock</i>	0.035	Own calculations, DHS
σ_ε^2	<i>Var.transitory income shock</i>	0.044	Own calculations, DHS
R^X	<i>Liquid asset return</i>	1.02	Gourinchas and Parker (2002)

Simulation.

When simulating the model, we use information we extracted from the Dutch Household Survey data: the deterministic component of income, the level of the cohort-specific permanent and transitory income shocks by year, and the variance of the household-specific income shocks by year. The deterministic component of income (G_t) is approximated by a second-order polynomial of age. The cohort-specific transitory and permanent income

shocks $(\varepsilon_t^c, \zeta_t^c)$, and the variance of the household-specific income shocks $(\sigma_\varepsilon^2, \sigma_\zeta^2)$ are calculated in Section 4 and the results are reported in Table A.4. The parameter values we use for the simulation are also listed in Table 2.

For each simulation, we draw realizations for the two unknown, household-specific income shocks $(\varepsilon_{i,t}, \zeta_{i,t})$ according to equations (13)-(14). We assume that each household starts its life with zero wealth, and only receives labor income. Altogether we run 100,000 simulations, 10,000 households in each cohort. When aggregating variables, we use cohort weights, which are representative weights of the Dutch population.

5.3 Simulation Results

In this Section, we present results from our structural model, imposing different assumption on which income shocks are allowed in the simulation. In our baseline simulation, we solve the model both with the observed transitory and permanent income shocks. We then contrast these baseline results with simulations obtained when we turn off either the temporary or the permanent income shocks. This strategy allows us to gauge the importance of each shock, separately. Finally, we also evaluate the simulation results at the cohort level in order to see their individual effects on consumption dynamics.


For the ease of comparison to Figure 1, which shows the observed evolution of aggregate variables, we normalize the simulated consumption profiles using 2008 as the base year (=100).

Baseline Simulations.

Figure 7 shows the life-cycle profile of consumption when we simulate the model with both the observed transitory and permanent income shocks. This figure highlights two important points. First, income shocks alone cannot replicate the consumption drop that we observe during the 2008-2009 Financial Crisis. Consumption in the simulated model does not change significantly over this period, which is in contrast with aggregate empirical evidence, shown in Figure 1. Second, income shocks generate a significant downturn in the simulated economy, replicating well the dynamics of the economy observed during the 2011-2012 Sovereign Debt Crisis. Consumption in the simulated model shows a prolonged decrease between 2011 and 2012, which is in line with the actual patterns observed in the Dutch economy. Between 2011 and 2012, consumption falls in the simulated economy by 2.9%. The corresponding drop in the data is 2.8%, as shown in Figure 1.

To understand these results better, it is useful to recall Figure 3 in Section 4, which plots our measures of the income shocks. As indicated in this figure, during the 2008-2009 crisis, the combination of permanent and transitory income shocks are not significant,

Figure 7: Simulated aggregate consumption profile



Notes: We calculate aggregate consumption by using simulated, cohort-level consumption profiles and appropriate cohort weights (representative of the Dutch population). Then we create a consumption index by using 2008 as the base year (2008=100). Shaded area indicates crisis period.

therefore aggregate consumption does not change significantly either. In contrast, after 2011, the sizable negative permanent income shocks offset the impact of the positive transitory income shocks, which leads to a sharp drop in aggregate consumption.


The findings from our baseline simulations suggest that the triggering factors of the Global Financial Crisis and of the Sovereign Debt Crisis are fundamentally different. Using income shocks alone, we are only able to rationalize consumption dynamics over the Sovereign Debt Crisis.

Simulations for Different Shocks.

To investigate the importance of the permanent and transitory income shocks separately, in explaining the observed consumption patterns, we next take advantage of our structural model and simulate different counterfactuals. In the first counterfactual, we assume that households only face the earlier identified transitory income shocks, but no permanent income shocks, by setting $\zeta_{i,t}^c = 0$ and $\varepsilon_{i,t}^c \neq 0$. In the second scenario, in turn, we assume that households only face the identified permanent income shocks, but no transitory income shocks, by setting $\varepsilon_{i,t}^c = 0$ and $\zeta_{i,t}^c \neq 0$. Using these counterfactuals allows us to examine how consumption dynamic varies under different hypothetical situations, and to evaluate the relevance of each type of shock independently. Figure 8 shows aggregate results from the two counterfactual simulations for consumption.

The dashed line in Figure 8 presents results from our first counterfactual simulation, where household only experience transitory shocks. Under this scenario, consumption increases smoothly and is not affected significantly by the income shocks over the years, which stems from the transitory nature of these shocks. As a consequence, the identified transitory income shocks alone cannot generate crisis periods as observed in the data, and cannot explain the dynamics of the Dutch economy between 2006-2016.

Figure 8: Disentangling the effect of income shocks on aggregate consumption



Notes: We simulate our model once with the observed cohort-level permanent shocks only ($\varepsilon_t^c = 0, \forall t$), and once with the observed cohort-level transitory shocks only ($\zeta_t^c = 0, \forall t$). We calculate aggregate consumption by using simulated, cohort-level consumption profiles and appropriate cohort weights (representative of the Dutch population). Then we create a consumption index by using 2008 as the base year (2008=100). Shaded area indicates crisis period.

The solid line in Figure 8 shows results from the second counterfactual simulation, where household only experience permanent shocks. Given the permanent nature of the income shocks, under this scenario income shocks induce significant changes in consumption, which shrinks substantially starting from 2011. Between 2011 and 2012, consumption fall by around 4%. While permanent income shocks alone cannot generate the dynamics of the Dutch economy over the 2008-2009 crisis quantitatively, they are clearly the main source of the consumption fall during the 2011-2012 Sovereign Debt Crisis. Without transitory shocks though (which are positive over this period) the simulated drop in consumption is much higher than observed in the data. These results are in line with the permanent income hypotheses: consumption tracks income more closely when the income shocks are perceived to be permanent rather than transitory.


Simulations for Different Cohorts.

When we simulate our baseline model, we aggregate variables over different cohorts by using appropriate cohort weights, which are representative for the Dutch population. Yet, this aggregation clearly masks important heterogeneities across cohorts.

As discussed earlier, cohorts differ in at least two dimensions. First, by definition, cohorts differ by their age. Age determines how households' consumption reacts to different income shocks. Temporary income shocks trigger larger consumption responses of older cohorts, as they face a shorter time horizon ahead to smooth income shocks over. In contrast, permanent income shocks cause larger consumption responses of younger cohorts, as the effect of permanent income shocks on lifetime resources is greater for those with longer time horizon ahead. Second, evidently, cohorts differ by the shocks they face. As seen in Figure 4 and Figure 5 in Section 4, the nature and magnitude of income shocks that households face vary across different cohorts.

To investigate cohort heterogeneity in consumption behavior, we next focus on three different cohorts separately: cohorts that include households born between 1950-1954, 1960-1964, and 1970-1974. By using these particular cohorts, we are able to compare their consumption profiles to the shocks they face, which are shown in Figure 4 and Figure 5. In Figure 9, we plot the simulated consumption paths for the three different cohorts.

Figure 9: Simulated consumption profiles for different cohorts



Notes: We create consumption indices for different cohorts by using 2008 as the base year (2008=100). Shaded area indicates crisis period.

This figure highlights two important findings. First, over the 2008-2009 crisis only

one cohort's consumption is affected significantly by the observed income shocks: the consumption of the mid-aged cohort (1960-1964). The difference between cohorts' reaction in the simulated model can be understood by looking at the shocks these cohorts face over the same period, shown in Figure 4 and Figure 5. Over the 2008-2009 crisis, the youngest (1970-1974) and oldest (1950-1954) cohorts face no significant, negative permanent income shocks, while the mid cohort (1960-1964) faces a large negative permanent income shock. Second, over the 2011-2012 crisis all the cohorts' consumption is affected by the observed income shocks, while the consumption responses vary substantially across these cohorts. The youngest cohort (1970-1974) faces the most dramatic reduction in consumption, while the oldest cohort (1950-1954) only experiences a slight decrease in consumption.²²

Discussion.

Using our simulated structural model with income shocks identified by household-level data from the Dutch Households Survey, we highlight three crucial aspects of income and consumption dynamics over the last decade. First, income shocks are only relevant in explaining consumption behavior over the Sovereign Debt Crisis, while they do not help us in understanding the Global Financial Crisis. Second, using both transitory and permanent income shocks in the model imply a drop in consumption in the simulated economy for 2012-2013 (2.9%), which is similar in magnitude to the consumption drop observed by the Dutch economy during the Sovereign Debt Crisis (2.8%). Third, the income shocks over the Sovereign Debt Crisis affect the younger cohorts the most and they suffered from the largest consumption drop on average.

These observations lead us to conclude that the two recent crises are fundamentally different. Income shocks experienced during the 2011-2012 Sovereign Debt Crisis are of a different nature than the shocks experienced during the 2008-2009 Global Financial Crisis, with the 2011-2012 shocks being perceived as more permanent. This helps explain why consumption declines more than income during the Sovereign Debt Crisis.

6 Conclusions

Consumption theories embody the idea that household consumption respond differently to income shocks of different persistence. Identifying the level of income shocks is, however, difficult for many reasons. The most widespread approach to overcome this

²²Even though the oldest cohort faces the largest negative permanent income shocks over 2011-2012, their response is the smallest in terms of consumption. The reason is that permanent income changes later in life have less impact on lifetime earnings, than income changes early in life.

identification problem was proposed by Blundell, Pistaferri, and Preston (2008). In this paper, we depart from the Blundell, Pistaferri, and Preston (2008) approach and analyze the transmission of income shocks to consumption by looking directly at the level of the shocks.

We exploit the Dutch Household Survey to identify the levels of transitory and permanent income shocks, following the approach by Pistaferri (2001) and to build a unique panel dataset of income shocks. We find large variations of the shocks both over time and between cohorts. Up until 2010, households mainly experienced transitory income shocks, while after 2010 they primarily experienced permanent income shocks. This evidence suggests that the Financial Crisis and the succeeding Sovereign Debt Crisis are fundamentally different because households perceived income shocks differently. This helps explain why consumption falls less than income during the Global Financial Crisis, while it declines more than income during the Sovereign Debt Crisis.

The Dutch Household Survey lacks reliable information on consumption expenditure. Therefore we use a structural model to understand income shock transmissions to consumption. We use the time series of the estimated transitory and permanent income shocks in the structural model and simulate a large number of households. Our results show that income shocks are crucial in explaining consumption behavior over the 2011-2012 Sovereign Debt Crisis, while they do not contribute to our understanding of the 2008-2009 Global Financial crisis. In particular, our model implies no significant consumption change between 2008-2009, while a 2.9% consumption drop between 2011 and 2012, which is similar to the observed patterns in the data. Therefore, the two episodes of economic downturn are very different in nature: income shocks are perceived to be important drivers of consumption dynamics during the Sovereign Debt Crisis.

The results of this paper are very suggestive and lead us to conclude that the nature of income shocks are crucial for understanding consumption dynamics over the two most recent economic recessions. They also point towards further considerations of our approach that for reasons of space could not be explored in this paper: in particular the inclusion of the role of credit markets. While we think that the impact of permanent income shocks on consumption would not be altered significantly by the inclusion of credit markets, the impact of transitory income shocks might be affected. It would be worthwhile to incorporate this issue in future work that extends the model presented here.

References

- Attanasio, Orazio, Agnes Kovacs, and Krisztina Molnar. 2017. “Euler Equations, Subjective Expectations and Income Shocks.” Discussion Paper 820, University of Oxford, Department of Economics.
- Blundell, Richard, Martin Browning, and Costas Meghir. 1994. “Consumer Demand and the Life-Cycle Allocation of Household Expenditures.” *Review of Economic Studies* (61):57–80.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. “Consumption Inequality and Partial Insurance.” *American Economic Review* 98 (5):1887–1921.
- Blundell, Richard and Ian Preston. 1998. “Consumption Inequality and Income Uncertainty.” *The Quarterly Journal of Economics* 113 (2):603–640.
- Browning, Martin and Thomas F. Crossley. 2001. “Unemployment Insurance Benefit Levels and Consumption Changes.” *Journal of Public Economics* 80 (1):1–23.
- Carroll, Christopher D. 1992. “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence.” *Brookings Papers on Economic Activity* 1992(2):61–156.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2019. “MPC Heterogeneity and Household Balance Sheets.” *CESifo Working Paper* (7134).
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. “Consumption over the Life Cycle.” *Econometrica* 70(1):47–89.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. “The Nature of Countercyclical Income Risk.” *Journal of Political Economy* 122 (3):621–660.
- Guvenen, Fatih and Anthony A. Smith. 2014. “Inferring Labor Income Risk And Partial Insurance From Economic Choices.” *Econometrica* 82 (6):2085–2129.
- Kaplan, Greg and Giovanni L. Violante. 2010. “How Much Consumption Insurance Beyond Self-Insurance.” *American Economic Journal: Macroeconomics* 2:53–87.
- Kaufmann, Katja and Luigi Pistaferri. 2009. “Disentangling Insurance and Information in Intertemporal Consumption Choices.” *American Economic Review* 99 (2):387–92.
- Little, Roderick Joseph A. and Donald B. Rubin. 2002. “Statistical Analysis with Missing Data.” 2nd ed. Hoboken, NJ: Wiley .

- Ljung, Greta M. and George E.P. Box. 1978. "On a Measure of a Lack of Fit in Time Series Models." *Biometrika* 65:297–303.
- MaCurdy, Thomas E. 1982. "The use of time series processes to model the error structure of earnings in a longitudinal data analysis." *Journal of Econometrics* 18 (1):83–114.
- Manski, Charles F. 2004. "Measuring Expectations." *Econometrica* 72 (5):1329–1376.
- Paxson, Christina H. 1993. "Consumption and Income Seasonality in Thailand." *Journal of Political Economy* 101 (1):39–72.
- Pistaferri, Luigi. 2001. "Superior Information, Income Shocks, And The Permanent Income Hypothesis." *The Review of Economics and Statistics* 83 (3):465–476.
- Storesletten, Kjetil, Chris I Telmer, and Amir Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor Market Risk." *Journal of Political Economy* 112 (3):695–717.
- Wolpin, Kenneth I. 1982. "A New Test of the Permanent Income Hypothesis: the Impact of Weather on the Income and Consumption of Farm Households in India." *International Economic Review* 22 (3):583–594.

A Appendix

A.1 Identification of Income Shocks

We assume the following standard decomposition of the log of income process (Pistaferri, 2001; Blundell, Pistaferri, and Preston, 2008):

$$y_{it} = \Pi' Z_{it} + \alpha' V_i + p_{it} + \varepsilon_{it} \quad (\text{A.1})$$

where Z_{it} is a deterministic time variant component of income and $\alpha' V_i$ is a deterministic time invariant one (e.g. it includes gender, education and household fixed effect). p_{it} and ε_{it} are, respectively, the permanent and transitory component of income of household i at time t . The transitory component (ε_{it}) is independently distributed σ_ε^2 , while the permanent component is a Markov process:

$$p_{it} = p_{it-1} + \zeta_{it} \quad (\text{A.2})$$

where ζ_{it} is the permanent shock and it is assumed to be i.i.d. with constant variance σ_ζ^2 . It is orthogonal to the transitory shock, at all lags and leads.

Combining equations (A.1) and (A.2) we obtain the following equation for the change in income:

$$\Delta y_{it} = \Pi' \Delta Z_{it} + \zeta_{it} + \Delta \varepsilon_{it} \quad (\text{A.3})$$

Under the assumption that the deterministic component of the evolution of income is a second order polynomial of age. i.e. $\Pi' Z_{it} = \pi_0 + \pi_1 \text{age}_{it} + \pi_2 \text{age}_{it}^2$, equation (A.3) can be rewritten as:

$$\Delta y_{it} = (\gamma_0 + \gamma_1 \text{age}_{it}) + \zeta_{it} + \Delta \varepsilon_{it} \quad (\text{A.4})$$

where $\gamma_0 = (\pi_1 - \pi_2)$ and $\gamma_1 = 2\pi_2$.

Rewriting equation (A.4) and exploiting the assumption of rational expectations, we can derive the following expression for the transitory shock:

$$\begin{aligned} \varepsilon_{it} &= -E[\Delta y_{it+1} | \Omega_t] + (\gamma_0 + \gamma_1 \text{age}_{it+1}) = \\ &= y_{it} - E[y_{it+1} | \Omega_t] + (\gamma_0 + \gamma_1 \text{age}_{it+1}) \end{aligned} \quad (\text{A.5})$$

Substituting this expression in equation (A.4), we identify the permanent income shock as:

$$\zeta = E[y_{it+1} | \Omega_t] - E[y_{it} | \Omega_{t-1}] - (\gamma_0 + \gamma_1 \text{age}_{it+1}) \quad (\text{A.6})$$

where Ω_t is the set of information available to household i at time t , and coefficients γ_0 and γ_1 are function of parameters π_1 and π_2 . We can interpret the temporary shock

ε_{it} as the gap between income realization and its expected value, given the information available at time t . The permanent shock p_{it} is measured by the revision in subjective income expectations with respect to the previous period ($t - 1$).

A.2 Data: Monthly and Annual Values of Expected Income

Wording of questions in the DNB Household Survey are, unfortunately, not homogeneous across waves. To our purpose, a relevant variation concerns questions eliciting subjective income expectations. While after 2007, they explicitly refer to ‘*annual*’ income, the time frame they refer to is more ambiguous for years 2003-2007. The exact wording of questions since 2008 is: ‘*We would like to know a little bit more about what you expect will happen to the net income of your household in the next 12 months. What do you expect to be the lowest total net yearly income your household may realize in the next 12 months? What do you expect to be the highest total net yearly income your household may realize in the next 12 months?*’.

In waves 2003-2007, the questions are: ‘*We would like to know a little bit more about what you expect will happen to the net income of your household in the next 12 months. What do you expect to be the lowest total net monthly income your household may realize in the next 12 months? What do you expect to be the highest total net income your household may realize in the next 12 months?*’. The first question refers to a time span of 12 months; the second and third refer to monthly income when eliciting the lower bound of the distribution and to any time frame when asking about the upper bound, respectively. In this sense, responses to those questions could be expressed either in annual or monthly terms.

To tackle this issue, we derive information on the relevant time frame for responses in period 2002-2007 by exploiting responses in waves when the reference to annual income is unambiguous. This approach is in the same spirit of imputation methods to tackle missing values described by Little and Rubin (2002), and exploit the panel structure of the sample to derive additional information for the period 2002-2007. We proceed by steps, as described hereafter.

1. For each respondent, we calculate subjective expectations referring to the lower and the upper bounds of annual income, i.e. year 1998-2002 and 2008-2015, to compute their average expected values for this period. This household specific ‘average lower/upper bound for annual income’ may depend on observable variables (family composition, education, etc.) and unobservables (ability of household members, optimism/pessimism of the respondent, information available to the respondent but not to the econometrician, etc).

2. We, then, estimate the lower/upper bound for expected income in each specific year. We use as regressors the household specific mean of subjective expectations described in 1., aimed to capture household specific information and expectations, along with other individual and household characteristics, aimed to capture both heterogeneity of expected income over the life-cycle and time-specific events which may affect expectations. More precisely, we use the pooled sample for periods 1998-2002 and 2008-2015 and we regress the logarithm of expected income on the ‘average lower/upper bound for annual income’, observed net household income, age, the number of workers in the couple, whether the respondent is working, and two dummies derived from a qualitative question about subjective expectations and capturing, respectively, whether the respondent does not expect any significant change in income or whether she expect an income increase.²³ Estimate results are shown in Table A.1. The lower (upper) bound is positively associated with the log mean lower (upper) bound and the log of observed income.
3. We use the estimated ‘typical lower/upper bound for expected income’ to identify respondents who report the upper and lower bound of expected *monthly* income in waves 2002-2007. More precisely, we assume that the upper/lower bound refer to monthly income when the reported value is lower than 20% of predicted annual values.

²³More precisely, we exploit the following question: “As a consequence of what changes (listed below) do you expect the total net yearly income of your household to change in the next 12 months? (More than one answer possible). a) A member of the household who currently has a job, will stop working, b) a member of the household who is currently out of work, will start working, c) a member of the household will change jobs, d) a member of the household will get a promotion e) social security (welfare) benefits (if any) that the household now receives will significantly go up f) social security (welfare) benefits (if any) that the household now receives will significantly go down/ other changes g) I don’t expect any significant changes in the next 12 months h) none of the above.

Table A.1: Estimates of (log of) lower/upper bound of expected income

	Lower bound	Upper bound
Ln(mean lower bound)	1.064*** (0.012)	
Ln(mean upper bound)		0.944*** (0.011)
Ln(income observed)	0.241*** (0.012)	0.249*** (0.011)
Age	0.005*** (0.001)	0.004*** (0.001)
No. workers in couple	-0.115*** (0.030)	-0.103*** (0.027)
Working	0.010 (0.029)	0.041 (0.026)
No significant changes in income expected	0.071*** (0.026)	0.068*** (0.023)
Positive reasons for change in income	0.108** (0.052)	0.109** (0.047)
Year 1998	0.143*** (0.055)	0.119** (0.050)
Year 1999	0.138** (0.059)	0.066 (0.053)
Year 2000	0.213*** (0.082)	0.202*** (0.075)
Year 2001	0.047 (0.058)	0.042 (0.053)
Year 2002	0.007 (0.057)	-0.029 (0.052)
Year 2009	-0.140*** (0.051)	-0.161*** (0.046)
Year 2010	-0.141*** (0.050)	-0.107** (0.045)
Year 2011	-0.054 (0.051)	-0.059 (0.046)
Year 2012	-0.215*** (0.050)	-0.167*** (0.045)
Year 2013	-0.201*** (0.050)	-0.223*** (0.046)
Year 2014	-0.287*** (0.048)	-0.245*** (0.044)
Year 2015	-0.049 (0.049)	-0.078* (0.044)
Year 2016	-0.068 (0.050)	-0.040 (0.046)
Constant	-3.482*** (0.153)	-2.294*** (0.141)
Observations	22033	22046

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Coefficients, standard error in parenthesis. Real values (euros 2010).

A.3 Standardization of the Model

The number of state variables of in the problem can be reduced from two $(X_{i,t}, P_{i,t})$ to one $\left(\frac{X_{i,t}}{P_{i,t}}\right)$. At terminal age $t = T$ the value function becomes

$$V_{i,T}(X_{i,T}, P_{i,T}) = \frac{C_{i,T}^{1-\rho}}{1-\rho},$$

With standardized variables, using notation $x_{i,T} = \frac{X_{i,T}}{P_{i,T}}$ and $c_{i,T} = \frac{C_{i,T}}{P_{i,T}}$, the value function can be written as

$$V_{i,T}(x_{i,T}) = U(c_{i,T}) = U\left(\frac{C_{i,T}}{P_{i,T}}\right) = \frac{\left(\frac{C_{i,T}}{P_{i,T}}\right)^{1-\rho}}{1-\rho}$$

Hence the value function with standardized variables can be rewritten as

$$V_{i,T}(x_{i,T}) = \frac{1}{(P_{i,T})^{1-\rho}} \left[\frac{C_{i,T}^{1-\rho}}{1-\rho} \right]$$

Therefore the relationship between the original and standardized value functions is:

$$V_{i,T}(X_{i,T}, P_{i,T}) = P_{i,T}^{1-\rho} V_{i,T}(x_{i,T})$$

Now considering the value function at age $t = T - 1$:

$$\begin{aligned} V_{i,T-1}(X_{i,T-1}, P_{i,T-1}) &= \max_{C_{i,T-1}} \{U(C_{i,T-1}) + E_{i,T-1}\beta V_{i,T}(X_{i,T}, P_{i,T})\} \\ &= (P_{i,T-1})^{1-\rho} \max_{c_{i,T-1}} \left\{ U(c_{i,T-1}) + E_{i,T-1} \left[\beta \left(\frac{P_{i,T}}{P_{i,T-1}} \right)^{1-\rho} V_{i,T}(x_{i,T}) \right] \right\} \end{aligned}$$

And similarly to the previous result, the simple relationship we get is

$$V_{i,T-1}(X_{i,T-1}, P_{i,T-1}) = P_{i,T-1}^{1-\rho} V_{i,T-1}(x_{i,T-1})$$

It can be shown that this relationship holds at a generic time t , hence the value function and the standardized value function at any point in time only differ by a scale factor. It is equivalent to maximize either function.

A.4 Additional Tables and Figures

Table A.2: Income and job expectations

	Expected hh income
Prob. unempl*work	-27.075*** (4.210)
Prob. find job*unempl.	4.167 (12.290)
Unemployed	-2690.982*** (619.580)
Hh income	0.854*** (0.009)
Constant	-1495.573 (2067.976)
Other controls	Yes
Obs.	3,630

Notes: Other control variables are: age, age squared, no. members, no. children, education and year dummies. Mean expected household income in the sample is 32,913.

Table A.3: Predictive power of subjective expectations

<i>Dep. var.</i>	Income in $t + 1$				
	(1)	(2)	(3)	(4)	(5)
Expected income	0.878*** (0.013)	0.485*** (0.034)	0.502*** (0.037)	0.538*** (0.036)	0.552*** (0.039)
Sq. root income expect. (/1000)				1779.370*** (428.597)	1461.108*** (387.556)
Exp. income* sq. root				-0.046*** (0.011)	-0.041*** (0.010)
Hh income		0.423*** (0.034)	0.367*** (0.036)	0.412*** (0.034)	0.363*** (0.036)
Age			116.263 (110.246)		62.617 (110.150)
Age sq.			-1.556 (1.165)		-0.966 (1.157)
N.hh members			1608.688*** (334.255)		1504.943*** (322.620)
N. children			-1695.199*** (417.859)		-1575.491*** (400.416)
Constant	4946.258*** (435.129)	3659.880*** (367.156)	-981.261 (2643.486)	2069.493*** (446.543)	-991.091 (2652.261)
Education dummies	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes
N	2377	2377	2377	2377	2377


Notes: OLS estimate. Errors clustered at the household level.

Table A.4: Permanent and transitory shocks, by cohort (weighted mean values)

Permanent shock								
Year	Coh 3	Coh 4	Coh 5	Coh 6	Coh 7	Coh 8	Coh 9	Coh 10
2006	-0.016	0.013	-0.012	-0.004	0.033	0.016	0.028	0.005
2007	-0.007	-0.010	0.006	0.048	-0.011	-0.031	0.029	0.005
2008	0.024	0.025	0.054	0.045	0.054	0.090	0.055	0.099
2009	-0.100	0.009	-0.012	-0.022	-0.033	-0.011	-0.010	-0.033
2010	.	-0.006	0.011	0.030	-0.009	0.020	0.016	0.010
2011	.	0.010	-0.030	-0.022	0.010	0.020	-0.002	0.064
2012	.	-0.077	-0.072	-0.055	-0.053	-0.106	-0.026	-0.080
2013	.	-0.097	-0.057	-0.058	-0.049	-0.047	-0.050	-0.029
2014	.	-0.045	-0.045	-0.017	0.019	0.028	0.015	-0.006
2015	.	.	-0.012	-0.018	-0.024	0.044	-0.001	0.015
2016	.	.	-0.043	0.027	0.043	0.032	0.049	0.071
Variance of permanent shock variance (pooled data)								
$\sigma_{\zeta}^2 = 0.035$								
Temporary shock								
Year	Coh 3	Coh 4	Coh 5	Coh 6	Coh 7	Coh 8	Coh 9	Coh 10
2006	-0.002	0.073	-0.054	0.044	0.047	-0.010	0.022	0.143
2007	-0.016	-0.030	-0.007	-0.042	0.027	-0.010	0.040	0.013
2008	-0.003	-0.022	-0.025	-0.008	-0.042	-0.031	-0.062	-0.084
2009	0.060	-0.024	-0.017	-0.003	-0.027	0.013	-0.017	-0.043
2010	.	0.009	0.012	-0.031	0.005	0.017	-0.015	-0.053
2011	.	0.047	0.021	0.046	0.015	0.020	0.001	-0.003
2012	.	0.116	0.106	0.086	0.073	0.057	0.022	0.020
2013	.	0.148	0.142	0.067	0.103	0.078	0.134	0.101
2014	.	0.131	0.090	0.093	0.086	0.082	0.096	0.090
2015	.	.	0.124	0.100	0.093	0.072	0.054	0.081
2016	.	.	0.142	0.114	0.103	0.098	0.071	0.104
Variance of temporary shock variance (pooled data)								
$\sigma_{\varepsilon}^2 = 0.044$								


Notes: Coh 1: year of birth lower than 1935; Coh 2: year of birth 1935-39; Coh 3: year of birth 1940-44; Coh 4: year of birth 1945-49; Coh 5: year of birth 1950-54; Coh 6: year of birth 1955-59; Coh 7: year of birth 1960-64; Coh 8: year of birth 1965-69; Coh 9: year of birth 1970-74; Coh 10: year of birth 1975+. Sample size for each cell is > 55 observations. Real values (euros 2010).

Figure A.1: Correlation between self-reported net household income and salary




Notes: One observation per household-year. The estimated regression line is $\ln y = 1.864 + 0.967 \ln x$; the coefficients are significant at the 1% level.

Figure A.2: Kernel densities of logarithm of observed and expected income




Notes: Pooled cross-section.

Figure A.3: Kernel densities of logarithm of observed and corrected income




Notes: Pooled cross-section.

Figure A.4: Transitory shock: based on calendar year and past 12 months




Notes: Corrected shock: Observed income is the average of the last 12 months and not referred to calendar year. As for permanent shocks, the difference between the week of the interview is larger than 12 weeks for only 15% of the sample; it is less than 4 weeks for 50% of the sample. Weighted mean.

Figure A.5: Permanent and transitory shocks




Notes: Real values (euros 2010). Weighted median.

Figure A.6: Permanent and transitory shocks/ income in $t - 1$




Notes: Real values (euros 2010). Weighted mean.

Figure A.7: Permanent and transitory shocks: Heterogeneity across income quartiles in $t-1$



Notes: Weighted mean.

Figure A.8: Permanent and transitory shocks/ income in $t - 1$: Heterogeneity across income quartiles



Notes: Weighted mean.

Figure A.9: Kernel densities of transitory income shocks

