

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/143759/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Rondinelli, Concetta, Kovacs, Agnes and Trucchi, Serena 2021. Permanent versus transitory income shocks over the business cycle. *European Economic Review* 139 , 103873. 10.1016/j.euroecorev.2021.103873

Publishers page: <https://doi.org/10.1016/j.euroecorev.2021.103873>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Permanent versus Transitory Income Shocks over the Business Cycle

Agnes Kovacs*

Concetta Rondinelli [†]

Serena Trucchi [‡]

August 16, 2021

Abstract

This paper investigates how income shocks shape consumption dynamics over the business cycle. First, we break new ground and create a unique panel dataset of transitory and permanent income shocks by combining household-level income expectations with realizations from the DNB Household Survey for the Netherlands in 2006-2018. We then use the first and second moments of the identified income shocks in a structural life-cycle framework and show that the model matches the observed consumption patterns well. Finally, using counterfactual model simulations, we assess the importance of the nature of income shocks (permanent income hypothesis), future income uncertainties (precautionary saving motive), and cohort effects, and show how they have individually shaped consumption dynamics over that period in the Netherlands.

Keywords: subjective expectations; income shocks; consumption; business cycle.

JEL classification: C13; D12; D91; E21.

*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). Corresponding author.

[‡]Cardiff University, Cardiff Business School, Aberconway Building, Colum Road, CF103EU, Cardiff, United Kingdom (trucchis@cardiff.ac.uk).

1 Introduction

How household consumption reacts to transitory and permanent income shocks is a long-standing question in macroeconomics, which is crucial both to understanding consumption behaviour and in evaluating policy change. The identification of the level of these income shocks, however, is challenging for many reasons. We normally observe total income changes, rather than transitory and permanent income changes separately. Moreover, there exists an information asymmetry between individuals and the econometrician that could lead to misclassification problems of income changes. As a consequence, the prevalent strategy to measure the transmission of income shocks to consumption is that proposed in the seminal paper of Blundell, Pistaferri, and Preston (2008). Their approach does not require direct identification of the income shocks but imposes strong covariance restrictions between the income and consumption processes in order to measure the variances of transitory and permanent income shocks, and ultimately, the transmission of income to consumption.

In this paper, we use a completely different strategy to study how income shocks affect households' consumption dynamics by looking directly at the level of the income shocks. In doing so, we rely on the approach of Pistaferri (2001) and a rich micro-dataset for the Netherlands to 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 2018. We then use the first and second moments of the identified income shocks in a standard, structural life-cycle framework to evaluate the importance of the nature of income shocks, future income uncertainties, and the cohort effects in shaping consumption dynamics in the Netherlands over this period.

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 revisions in income expectations, while transitory shocks are differences between income realizations and the expectation of future income, once the predictable life-cycle components are removed. Using this theoretical result, we exploit the joint availability of subjective income expectations and realizations in a micro panel dataset, the Dutch National Bank Household Survey (DHS), to compute the level of permanent and transitory income shocks. To the best of our knowledge, ours is the first attempt to use subjective income expectations for an extended period to decompose

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

income shocks into transitory and permanent components.² The direct observation of income shocks has two important advantages. First, we can differentiate shocks by sign and size and, therefore, look at the asymmetries in consumption response along these dimensions. While most papers in the literature focus on the consumption response to an income increase,³ we shed light on the impact of both positive and negative income shocks. Second, in contrast to previous literature (such as Blundell, Pistaferri, and Preston, 2008), our method does not require shocks to be uncorrelated across households and, therefore, allows us to accommodate aggregate income shocks. Disregarding aggregate shocks to income is problematic over economic recessions. For this reason, our methodology is well suited to analyse consumption behaviour over an extended time span, when household income is likely to be affected by aggregate factors.

The period we analyse covers both expansionary and recessive phases of the business cycle. We observe two recessions - the 2008-2009 Global Financial Crisis and the 2011-2013 Sovereign Debt Crisis - which are characterised by a drop in output and in aggregate consumption. Using the described identification strategy, we find that Dutch households face both significant shocks to their income and to their income uncertainties. These shocks are most significant during the two crisis periods. Negative income shocks are small and transitory during the Global Financial Crisis, albeit the increase in income uncertainties is substantial. The precautionary saving motive plays a key role in driving the fall in consumption in 2008-2009: facing higher income uncertainties, households wish to save more and consume less. During the Sovereign Debt Crisis, negative income shocks are permanent and large, however the increase in income uncertainties is also sizeable. We show that the consumption drop during 2011-2013 is triggered both by the level of income shocks and the precautionary saving motive. When we consider different cohorts, we find that the 2008-2009 crisis hits all cohorts in a similar manner, while the 2011-2013 crisis affects the income of the younger cohorts more.

After analysing the dynamics of the income shocks, in the second part of the paper we take full advantage of a structural life-cycle model in order to validate our identification strategy of transitory and permanent income shocks. We use the first and second moments of the identified income shocks in the model to simulate the consumption-savings behaviour of 10 different cohorts of households. We show the model-implied consumption dynamics between 2006 and 2018, both at the aggregate and the cohort levels, and compare these consumption profiles to their empirical counterparts.

Our relatively simple life-cycle model with identified income shocks, precautionary saving motive and many different cohorts is able to match the observed consumption

²However, the idea of using income expectations and their realizations together to circumvent the income shocks misspecification problem dates back to Hayashi (1985).

³Notable exceptions are Baugh et al. (2021) and Christellis et al. (2019).

dynamics well in the Netherlands between 2006 and 2018. In line with the aggregate data, our model generates two significant contractions in consumption over the sample period: one for 2008-2009 and another for 2011-2013, reflecting the Global Financial Crisis and the Sovereign Debt Crisis, respectively. We also show the simulated trajectories of consumption for different cohorts between 2006 and 2018 and compare them to consumption data from the DHS to show that they are broadly consistent with each other and are in line with the dynamics of cohort-level shocks.

To more fully understand the mechanism of income transmission to consumption, we compare different counterfactual scenarios. Income transmission to consumption within our theoretical framework is affected by the nature of income shocks (permanent income hypothesis effect), future income uncertainties (the precautionary saving effect) and the age of the households (cohort effect). Using the model, we can isolate the effect of these different channels from each other, and from other potential factors (e.g., real interest rates, wealth shocks) that might simultaneously affect consumption dynamics. We evaluate the relevance of each channel by comparing consumption profiles induced by different model variants to those observed in the data. These counterfactual simulations show that all three highlighted channels play a key role in determining consumption patterns in the period of our analysis. Turning off any of these channels would result in a simulated consumption profile that does not fit the consumption patterns observed in the data.

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), Carroll (2009), Kaplan and Violante (2010), Guvenen and Smith (2014), or Baugh et al. (2021). 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 ours are Pistaferri (2001) and Attanasio, Kovacs, and Molnar (2020). 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 restrictions, 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 (2020) 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 DNB 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. In Section 2, we introduce the DNB Households Survey, which uniquely collects information on both expected and realized income. Following the approach proposed by Pistaferri (2001), in Section 3, we identify transitory and permanent income shocks and analyse their behaviour over 2006-2018. We then take the first and second moments of the identified shocks and use them in a life-cycle framework in Section 4 and show that the model can successfully match the observed consumption patterns. Section 5 presents different model counterfactuals to gauge the importance of the permanent vs transitory nature of income shocks, income uncertainties and cohorts, and show how they have individually shaped aggregate consumption dynamics. Finally, Section 6 concludes our paper.

2 Income and Consumption in the DNB Household Survey

In our analysis, we use data from the Dutch National Bank Household Survey (DHS) administered by CentERdata (Tilburg University, The Netherlands), which is a longitudinal survey representative of the Dutch-speaking population, which is collected annually on behalf of the Dutch National Bank via an online survey. The survey is designed for gathering information about the psychological and the economic determinants of households' financial behaviour. The dataset includes responses to six questionnaires seeking information on the general household, work, health and income, accommodation and mortgages, assets and liabilities, and psychological data.

The unique feature of the 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 panel dataset which collects subjective income expectations together with their realizations covering a period of more than 10 years, including upturns and downturns in the economic business cycle.

2.1 Sample Selection.

Our initial sample consists of 18,856 household heads and partners aged 21-65, interviewed in the period 2006-2018 and who are asked questions in the income module in the DHS questionnaire. As we later explain in detail, we use the panel-dimension of the

survey in order to identify income shocks; therefore, we restrict our sample to individuals who are observed at least twice, which gives us 17,520 observations. Further, we also make sure that individuals we observe understand expected-income-related questions. For this purpose, we use a simple rule and exclude individuals whose maximum expected future income is below their minimum expected income or/and whose subjective probabilities attached to future events are inconsistent. We end up with 13,412 observations. Finally, to deal with outliers, we simultaneously trim the top and bottom 5% of observed and expected income and are left with a sample of 10,670 individuals.⁴ More details on the sample selection and descriptive statistics can be found in Appendix B.1.

In what follows, we describe the main variables we use in our analysis: households' income realizations, their subjective income expectations, and consumption. Variables are expressed in 2010 euros by using annual consumer price indices from Statistics Netherlands.

2.2 Income Measures

Income Realizations.

The measure of household income that we use in the empirical analysis is gathered through the following question:

“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 combined. 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 that is used to elicit income expectations, namely total net household income. Even though the questions on income refer to household income, we use the answers related to both the household head and the spouse. We exploit other information collected by DHS and find that the majority of net income comes from labour earnings. On average, financial revenues represent about 18% of net income for all the respondents, and about 31% if we consider owners of financial assets only.⁵ To assess 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 labour income.⁶ The

⁴The main patterns illustrated in the paper are confirmed after a 1% trimming.

⁵Less than 2% of households declare income from housing wealth.

⁶Gross labour income is obtained as the sum of earnings of all household's members. Net labor income is not available.

two variables turn out to be strongly correlated, with a regression line close to the 45-degree line.⁷ 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:

“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:⁸

“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. More precisely, the expected value of household income is calculated multiplying the central value of each interval by the self-reported probability that future income will be in that interval. Income values below the lower bound and above the upper bound are given zero probability.

Dynamics of expected and observed income.


Figure 1 plots the average income expectations (the dashed line), together with actual income data (the solid line). The year on the horizontal axis is the year of interview; that is, when the information was elicited. Shaded areas in the figures indicate the two crisis periods that took place in the period of the analysis: the Global Financial Crisis (2008-2009) and the Sovereign Debt Crisis (2011-2013).

⁷The 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 B.2 in the Appendix B.3. The estimated regression is $\ln y = 0.922 + 1.065 \ln x$, where the coefficient for $\ln x$ is significant at the 1% level.

⁸Heterogeneity in the way income expectations are elicited over time is discussed in Appendix B.1.

There are two episodes of sudden drop both in expected and in observed income. The first contraction occurs around the time of the Global Financial Crisis (2008-2009), when these two variables fall by similar magnitudes. The second drop is during the Sovereign Debt Crisis, when subjective income expectations fell much more than observed income. Moreover, expectations about future income remain below income realization until 2017. In 2018 we observe an increase in income expectations, not followed by a rise in their realization. The increased pessimism that we detect in the DHS dataset is also observed in the Consumer Confidence Indicator for the Netherlands, which is plotted in Figure A.2. In a similar manner to the DHS data, the Consumer Confidence Indicator exhibits a first decline starting from the second half of 2008, followed by a second more sizeable and prolonged drop, lasting from the end of 2011 until the beginning of 2014.

Figure 1: Observed and expected income



Notes: Our calculations from DHS data for the period 2006-2018. Weighted average computed using sample weights. Real values (euros 2010) are calculated using annual consumer price indices from Statistics Netherlands (CBS). Observed income refers to calendar year, while expected income refers to one year ahead. Shaded area indicates crisis periods.

The Sovereign Debt Crisis follows the Global Financial Crisis with only a brief upturn occurring between them. The revisions in expectations that occurred in 2011-2013 may, thus, be related to the unique timing of these two episodes. Previous literature shows that large macroeconomic shocks affect individuals' beliefs and preferences, such as risk attitudes and expectations (Malmendier and Nagel, 2011). Notably, Malmendier and Nagel (2011) show that individuals exposed to periods of low stock returns are more pessimistic about future returns and, more generally, that experienced macroeconomic shocks are important in affecting households' present behaviour. This argument might

be relevant in our case, as well for understanding the behaviour of income expectations over the two crises. Having experienced unpredicted, large, aggregate income shocks over the 2008- 2009 Global Financial Crisis, households display more caution in terms of their income expectations. As a result, facing large aggregate shocks again in 2011, households used their past experiences and adjusted their expectations downwards, accordingly. In addition, the worsening of labour market conditions is more dramatic during the Sovereign Debt Crisis. The unemployment rate rises by less than one percentage point (from 3.7% to 4.4%) between 2008 and 2009, while it steadily increases during the Sovereign Debt Crisis and touches historically high levels in the Netherlands in 2013 (7.3%). This trend reflects in perceived job loss probabilities, which remain steadily high during the years 2011-2013 (see Figure A.3), suggesting a relevant impact of labour market conditions on the downward revision in income expectations. Finally, differences in primary drivers of the two crises may affect the perception of their impact on future households' income. The Global Financial Crisis is prominently an "imported crisis", with a prolonged fall in international trade dragging the Dutch economic activity down, as shown by the dramatic fall in exports. The Sovereign Debt Crisis is, instead, a "Euro-zone crisis", related to the collapse of financial institutions, high government debt, and rapidly rising bond yield spreads in government securities.⁹

Reliability of Income Measures.

The 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 information value and the accuracy of subjective expectations elicited by the DHS survey. 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, we show that the distribution of subjective income expectations has a regular shape and shadows that of income realizations.¹⁰ This evidence is reassuring in terms of the limited diffusion of random or inaccurate responses, which points to the reliability of expected income variables. Second, we illustrate the internal coherency between subjective expectations regarding income and job status. Working and not-working respondents are asked, respectively, about the probability of losing or finding a job in the next 12 months. We test the conditional correlation between expected job status and


⁹Similar evidence has been shown for other European countries which were also affected by two severe recessions (e.g., Caivano, Rodano, and Siviero (2011) and Buseti and Cova (2013) for Italy).

¹⁰The density function of income expectations and realizations for the pooled cross-section dataset is plotted in Figure B.3 in the Appendix B.3. 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.

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 B.5 in the Appendix B.3, show a correlation heading in the expected direction. Working respondents who report 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 of the small number of unemployed respondents in the sample). Overall, 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 households form their income expectations rationally (i.e., using their full information set), we must detect a strong *ex-post* correlation between subjective income expectations and their realization. We exploit the longitudinal component of the dataset, and we examine the link between income realization and subjective expectations elicited one period ahead. We start with a simple scatter plot, shown in Figure 2, that shows observations of (logarithm of) actual income (y-axis) as a function of (logarithm of) expected income (x-axis), together with the 45-degree line and a regression line that is predicted by a linear regression of observed income on expected income.

Figure 2: Expected and Realized Income



Notes: Our calculations from DHS data for the period 2006-2018. Real values (euros 2010) are calculated using annual consumer price indices from Statistics Netherlands (CBS). The estimated regression is $\ln \ln y_t = 1.491 + 0.857 \ln x_t$.

The majority of the observations are clustered around the 45-degree line and the linear regression line is close to the 45-degree line, indicating high correlation between expectations and future realizations. To examine the reason why we observe a slight deviation from the 45-degree line, we compute forecast errors and analyse their behaviour. We define the forecast error (μ_{it-1}) of household i as the difference between the household's expected (log) income at $t-1$ and its realization at time t : $\mu_{it-1} = \mathbb{E}[y_{it}|\Omega_{t-1}] - y_{it}$. We find that the average forecast error is mostly negative, as can also be seen in Figure 1. This finding is in line with the results presented by Rozsypal and Schlafmann (2017), who use the Michigan Survey to document households' systematic pessimism (see also Appendix B.2 for more details). However, following Kaufmann and Pistaferri (2009), we can also interpret these negative forecast errors as the result of persistent measurement errors in subjective reports of future income. Among income groups, low-income households underestimate their income growth, high-income households are too optimistic and overestimate their income growth, in line with Rozsypal and Schlafmann (2017).

2.3 Consumption Measure.

Since household consumption expenditure is not directly collected by the DHS survey, we need to compute it in an indirect way as the difference between net household income and household saving. Respondents are asked whether they put any money aside in the previous 12 months. In the case of a positive answer, they indicate "about how much money" the household saved in the same period by selecting the appropriate range out of seven possible value bands.¹¹ For each band, we compute the central value of the interval and subtract it from the household net income. Note that, consequently, the measurement error in our consumption variable has two sources: first, income and saving variables might be reported with an error; and second, saving is only collected in brackets. For this reason, we interpret all results based on this measure of consumption with caution and compare them with measures from the national accounts to assess their reliability.¹²

Figure 3 illustrates the average consumption patterns both at the aggregate and the cohort levels. To ease the comparison of different consumption measures, we normalize them using 2008 as the reference year ($= 100$). The left panel of Figure 3 shows aggregate

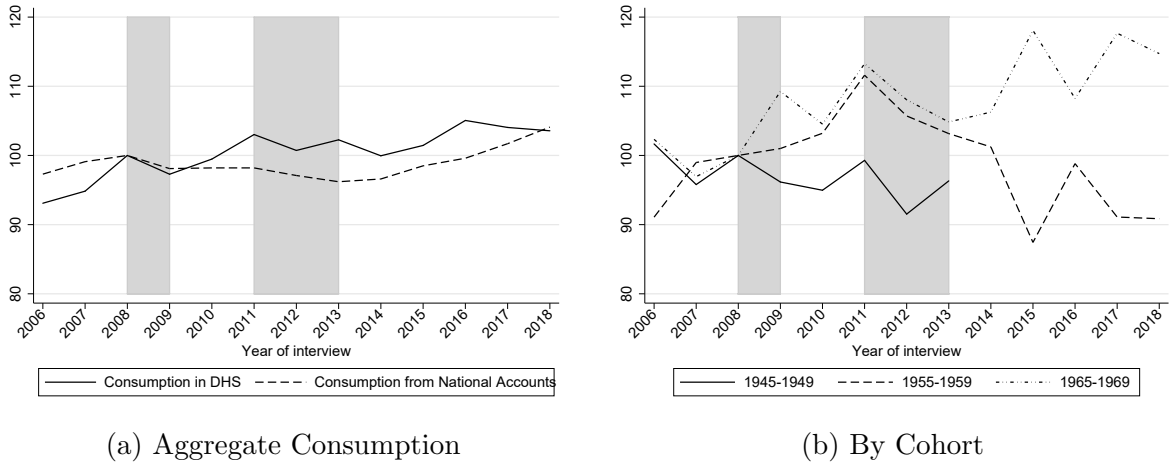
¹¹Value bands are the following: less than 1,500 euros; 1,500-5,000 euros; 5,000-12,500 euros; 12,500-20,000 euros; 20,000-37,500 euros; 37,500-75,000 euros; more than 75,000 euros.

¹²We also constructed a saving and a corresponding consumption measure by using changes in household wealth. This consumption measure, however, prove to be more noisy and deviate more from the observed consumption (from National Accounts) than the consumption measure relying on direct saving information.

consumption dynamics in DHS recovered as described above (the solid line), together with aggregate consumption data from Eurostat, as measured in the national accounts (the dashed line). Consumption in DHS displays an overall increasing trend during the period 2006-2018, which is interrupted by two contractions, which coincide with the two recessions in 2008-2009 and 2011-2013. Consumption recovers to its pre-crisis level in 2010, one year after the Global Financial Crisis, while its fall is more prolonged during the Sovereign Debt Crisis. This pattern mirrors the aggregate measure of consumption retrieved from national accounts relatively well. The only major difference is that in DHS data consumption increases between 2009 and 2011, while it is stable in the aggregate data.

The right panel of Figure 3 plots the consumption dynamics of three different cohorts of individuals, who were born in the following years: 1945-1949 (the solid line), 1955-1959 (the dashed line), and 1965-1969 (the dotted line).¹³ Consumption patterns of the three cohorts diverge during the Global Financial Crisis, when consumption falls for the 1945-49 cohort only. Consumption of the oldest cohort recovers to pre-crisis level in 2011, while it shows an increasing trend for the middle and young cohorts. Consumption substantially falls for all the cohorts during the Sovereign Debt Crisis. It starts increasing after 2013 for the youngest, while it slightly declines for the 1955-1959 cohort.

Figure 3: Consumption: aggregate and by cohorts (indices: 2008=100)



Notes: Our calculations from DHS data for the period 2006-2018. Weighted average computed using sample weights. Real values (euros 2010) are calculated using annual consumer price indices from Statistics Netherlands (CBS). Aggregate consumption from National Accounts (Eurostat) is real Household and NPISH (non-profit institutions serving household) final consumption expenditure. Shaded area indicates crisis periods.

¹³Cohorts are defined by individuals' date of birth. Cohort 1955-1959, for example, includes households born between 1955 and 1959. We only consider cohort-year cells with at least 30 observations.

3 Identification of Income Shocks

In this section, we first describe the strategy we use to separately identify permanent and transitory income shocks. We then present and analyse both the first and higher moments of the identified income shocks between 2006 and 2018. Finally, we show how the dynamics of transitory and permanent income shocks vary by cohorts.

3.1 The Methodology

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 (2020). This method hinges on the relationship between subjective income expectations and the corresponding income realizations. We start with the following, standard decomposition of the (logarithm of) income, as in Blundell, Pistaferri, and Preston (2008):

$$\begin{aligned} y_{it} &= \Pi' Z_{it} + \alpha' V_i + p_{it} + \varepsilon_{it} \\ \Pi' Z_{it} &= \pi_0 + \pi_1 age_{it} + \pi_2 age_{it}^2 \end{aligned} \tag{1}$$

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 . The permanent income component follows a Markov process:

$$p_{it} = p_{it-1} + \zeta_{it} \tag{2}$$

where ζ_{it} is the permanent income shock. Permanent and transitory shocks are assumed to be orthogonal (at all leads and lags), unanticipated and serially uncorrelated and with zero means. Note, that we allow shocks to be correlated across households in order to accommodate aggregate shocks. Consequently, we interpret each shock as a combination of idiosyncratic, cohort-specific, and aggregate shocks.

Combining equations (1) and (2) we obtain the following equation for income growth:

$$\Delta y_{it} = \Pi' \Delta Z_{it} + \zeta_{it} + \Delta \varepsilon_{it}. \tag{3}$$

If we disregard the predictable income component ($\Pi' \Delta Z_{it}$), income changes in response to either permanent income shocks (ζ_{it}) or changes in transitory income shocks ($\Delta \varepsilon_{it}$). 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 detail in Appendix C.1. As a result, transitory and permanent income shocks can be rewritten, respectively, as:

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

and

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

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).¹⁴

In this way, we can offer a straightforward interpretation of the transitory and permanent income shocks based on subjective income expectations and realizations. Apart from a predictable age affect, a transitory income shock, ε_{it} , is identified by the gap between income realization and future subjective income expectation; while a permanent shock, ζ_{it} , is identified as the change in the subjective expectations of income. Therefore, this method allows us to identify transitory and permanent income shocks separately using data only on observed and expected income, as long as shocks are serially uncorrelated.

Identifying Assumption.

As discussed earlier, one of the main advantages of our identification strategy is that it does not require the income shocks to be i.i.d. Relaxing the i.i.d. assumption is of crucial importance for interpreting our results as it allows us to consider aggregate income shocks alongside the household and/or cohort-specific shocks, and represents a key contribution to the literature, which typically require stronger restrictions on the structure of shocks (e.g. Blundell, Pistaferri, and Preston, 2008; Kaplan and Violante, 2010). Instead of the i.i.d. assumption, the method we use needs permanent and transitory income shocks to be serially uncorrelated at the household level.

To test this assumption, we consider an autocorrelation test 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

¹⁴Assuming 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).

different from zero over a time series. The null hypothesis in our case states that there is no autocorrelation in the transitory (permanent) income shocks. To perform this test, we naturally need to restrict our sample to households that are observed at least in two consecutive time periods, which reduces the number of observations and creates a different sample from that used in our analysis. However, results from a smaller sample are still indicative of how shocks behave at the household level.

For the transitory shock, the p-value is greater than 0.16 for 90% of cases, while for the permanent shock the p-value is greater than 0.11 for 90% of cases. Therefore, in more than 90% of the tests performed we cannot reject the null hypothesis of no autocorrelation at the 10% level of significance. Consequently, we argue that the behaviour of our calculated income shocks is not inconsistent with our assumption that they are serially uncorrelated at the household level.

Predictable Income Component.

In order to use equations (4)-(5) to calculate the income shocks, we need to determine the coefficients of the deterministic income component, γ_0 and γ_1 . Having data both on income realizations and subjective income expectations makes it easy to calculate one-year-ahead income growth expectations. Then simply by regressing reported expected income growth on a constant and on age, we can obtain estimates for γ_0 and γ_1 . The estimated coefficients are $\hat{\gamma}_0 = .0082$ and $\hat{\gamma}_1 = -.0016$. The combinations of income realizations, subjective income expectations, and predictable income components over the life-cycle identify 5,490 transitory and permanent shocks, as expressed in equations (4) and (5).


3.2 Identified Income Shocks

After discussing the method to identify income shocks separately, we next present our results both for permanent and transitory income shocks. We analyse first moment and distribution of the identified income shocks over time and their heterogeneity across cohorts.

First Moments of Income Shocks.

In Figure 4, we illustrate the dynamics of average permanent (the solid line) and transitory (the dashed line) income shocks between 2006 and 2018. In addition, we also present the changes in average transitory income shocks (the dotted line) that, besides the permanent income shocks, drive income changes, as shown in equation (3).

Figure 4: Permanent and transitory shocks



Notes: Our calculations from DHS data for the period 2006-2018. Weighted average computed using sample weights. Real values (euros 2010) are calculated using annual consumer price indices from Statistics Netherlands (CBS). Permanent and transitory shocks are calculated following the method described in Section 3. Shaded area indicates crisis periods.

Transitory shocks over the period of observation are almost always positive, except for 2009-2011 and 2018. As seen in equation (4), transitory income shocks are positive when subjective income expectations for the future ($\mathbb{E}[y_{t+1}|\Omega_t]$) are below their observed value in the present (y_t). We can interpret the systematic (yet small) discrepancy between future income expectations and today's income realizations as a measure of general pessimism. Looking at changes in transitory income shocks, which are relevant for income changes, we only document positive changes for the years between 2010 and 2013 and for 2016. Permanent income shocks over the same period show higher volatility than transitory income shocks. Between 2006-2011 and after 2015 permanent income shocks are positive, while between 2011 and 2015 they are significantly negative. As seen in equation (5), positive (negative) permanent income shocks imply upward (downward) revisions in subjective income expectations.

Focusing now on the two crisis periods over the sample, there are visible differences between income shocks during the Global Financial Crisis and the Sovereign Debt Crisis. In 2009, the average permanent income shock is zero, while the average transitory income shock is only slightly negative but decreasing. Consequently, the observed income drop in Figure 1 in 2009 is triggered by a negative change in the transitory income shock. By contrast, in 2012 and 2013 the average permanent income shocks are large and negative, while the average transitory income shocks are positive throughout. As a


result, the observed income drop in Figure 1 in 2012 is driven by negative permanent income shocks.

It is also worth noting here that transitory shocks and permanent income shocks are negatively correlated via the expected future income, as seen in equations (4) and (5). *Ceteris paribus*, a decrease in expectation on future income reflects in a negative permanent shock and a positive transitory shock. The opposite movement of the two shocks during the Sovereign Debt Crisis is a clear example of a decrease in future income expectation.

Higher Moments of Income Shocks.

In Figure 5, we look at higher moments of transitory and permanent income shocks by presenting the kernel densities for each of them. On the left-hand side, we plot densities for permanent income shocks, while on the right-hand side we plot that for transitory income shocks. In general, the kernel densities are well-behaved: the centre of the income shocks distributions are at zero with variances that are larger for the permanent income shocks than for transitory income shocks.

Figure 5: Kernel densities of permanent and transitory income shocks



Notes: Our calculations from DHS data for the period 2006-2018. Real values (euros 2010) are calculated using annual consumer price indices from Statistics Netherlands (CBS). Permanent and transitory shocks are calculated following the method described in Section 3.

A relevant aspect of idiosyncratic income shock heterogeneities is whether and to what extent these heterogeneities are affected by the business cycle. Storesletten, Telmer, and Yaron (2004), for example, 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 centre of the income shock distribution. This results in a countercyclical left-skewness.

To examine the business cycle effects on shocks' distribution in our sample, we separate our identified income shocks according to when they are observed, and plot their kernel densities in Figure 5, separately. We differentiate between three particular periods: no- recession periods, which we call our benchmark (the solid grey line), the Global Financial Crisis period between 2008-2009 (the solid black line), and the Sovereign Debt Crisis period between 2001-2013 (the dashed line).

The centre of the income shock distributions does not move much during recessions, compared to the no-recession period, while the tails of the shock distributions move quite asymmetrically. Considering first the distributions of permanent income shocks, we observe no clear pattern in the shift of shock distribution during the Global Financial Crisis: exceptionally large negative income shocks become less likely, while small negative shocks and positive income shocks become more likely. The shift in distribution is more clearly seen during the Sovereign Debt crisis, when permanent income shocks are substantial: large negative income shocks become more likely, whereas the probability of experiencing large positive income shocks decreases. These findings reinforce the results of Guvenen, Ozkan, and Song (2014), who state that idiosyncratic permanent shocks are not countercyclical, instead their left-skewness is countercyclical.

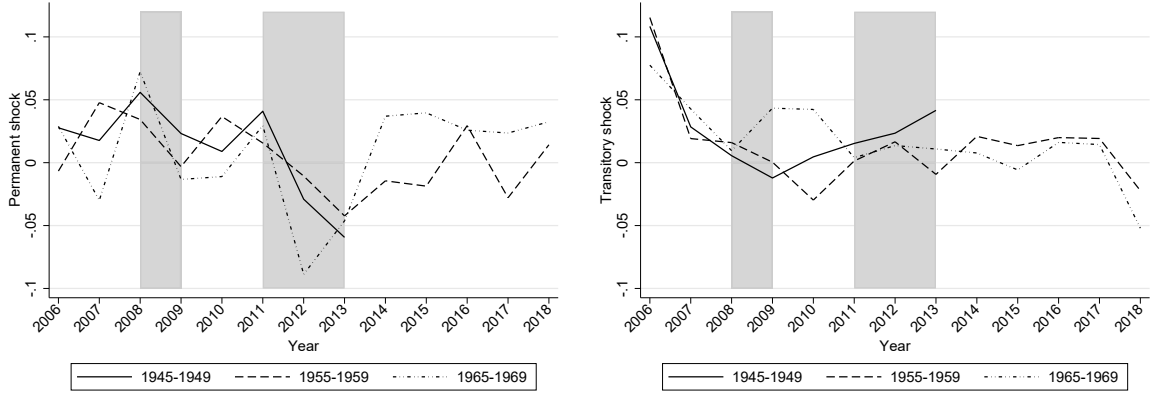
Considering next the distributions of transitory income shocks, we observe a significant shift in distribution during the Global Financial Crisis: small negative income shocks become more likely, while the probability of experiencing large positive income shocks decreases. The transitory income distribution does not change significantly during the Sovereign Debt Crisis.

Heterogeneities by Cohort.

Average income shocks potentially mask heterogeneities across households, which might shed light on the channels driving the dynamics of aggregate variables. For this reason, in Figure 6 we illustrate the time trend of income shocks for cohorts described earlier in Figure 3: for households born between 1945-1949 (the solid line), 1955-1959 (the dashed line), and 1965-1969 (the dotted line).¹⁵

¹⁵Tables D.2 and D.3 report the evolution of permanent and transitory shocks for all cohorts.

Figure 6: Permanent and transitory shocks by cohort



Notes: Our calculations from DHS data for the period 2006-2018. Weighted average computed using sample weights. Real values (euros 2010) are calculated using annual consumer price indices from Statistics Netherlands (CBS). Permanent and transitory shocks are calculated following the method described in Section 3. Shaded area indicates crisis periods.

The left panel in Figure 6 highlights two important facts. First, only the 1965-69 cohort face slightly negative permanent income shocks (a downward revision of income expectations) over the 2008-2009 Global Financial Crisis. Second, all of the cohort face large and negative permanent income shocks over the 2011-2013 Sovereign Debt Crisis, however the youngest cohort suffers the most negative shocks. After 2013, the 1955-59 cohort experiences relatively small permanent shocks, which fluctuate around zero, while the youngest cohort continuously faces positive permanent income shocks.

The right panel in Figure 6 plots the dynamics of transitory income shocks for the same cohorts. We observe a decline in transitory shocks for all the cohorts in the initial period, until the 2008-2009 crisis, while only the 1945-49 cohort is hit by a negative transitory shock during the Global Financial Crisis. All the cohorts face positive transitory shocks in 2011-2017, including the Sovereign Debt Crisis, and experience a large negative shock in 2018.

4 A Life-Cycle Model of Consumption

In this section, we use a structural model to map income shocks identified in Section 3 into consumption dynamics. We consider a standard incomplete market life-cycle model, which has become the workhorse framework for quantitative analysis in macroeconomics over the last decades (see for instance Deaton, 1991; Carroll, 1997; Blundell, Pistaferri, and Preston, 2008; Kaplan and Violante, 2010).¹⁶ In this framework, households do not

¹⁶Our model is similar to that used by Kaplan and Violante (2010), who evaluate the precision of the insurance coefficient derived and estimated by Blundell, Pistaferri, and Preston (2008). They find

only save for expected future income drops (say for retirement) but also for precautionary reasons motivated by uncertain future income and liquidity constraints. We use the model to simulate household consumption using the income shocks identified in Section 3. We then compare aggregate consumption dynamics from the model to that traced in the Dutch data.

4.1 The Model

We build a standard incomplete market model of life-cycle consumption and savings, where households face permanent and transitory income uncertainty. We assume that households live for T periods as adults: they work for W periods and retire afterwards. footnoteIn our model, we make a number of stark assumptions to focus on the main points we want to make. Most of these assumptions (such as deterministic length of life, the absence of bequests or the absence of different assets), can be easily relaxed and would not affect the nature of the exercise we present below. 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 for life-cycle and precautionary purposes. The only uncertainty households face in the model comes from different income shocks of a transitory and permanent nature. We consider different cohorts of households who differ in terms of their expectations and realization of the income shocks and income uncertainties they face.

The Value Function

Households have time-separable expected utility given by:

$$\mathbb{E}_0 \sum_{t=1}^T \beta^{t-1} U(C_{i,c,t}) \quad (6)$$

hence, we can formulate households value function in a recursive form as follows:

$$V_{i,c,t}(X_{i,c,t}, P_{i,c,t}) = \max_{\{C_{i,c,t}\}} U(C_{i,c,t}) + \beta \mathbb{E}_{i,c,t} V_{i,c,t+1}(X_{i,c,t+1}, P_{i,c,t+1}), \quad (7)$$

subject to:

$$X_{i,c,t+1} = R^X(X_{i,c,t} - C_{i,c,t}) + Y_{i,c,t+1} \quad (8)$$

that the estimated insurance coefficients by Blundell, Pistaferri, and Preston (2008) are very similar to those predicted by the structural model, however they are, in general, downward biased.

where $V_{i,c,t}$ is the value function for household i belonging to cohort c at time t . $C_{i,c,t}$ is non-durable consumption, $Y_{i,c,t}$ is labour income, and $P_{i,c,t}$ is the permanent part of the labour income, to be defined later in this section. $X_{i,c,t}$ is cash-on-hand, defined as the sum of savings and labour income in period t .¹⁷ Finally, parameter β is the discount factor.

Sources of Uncertainty.

In our framework, the only source of uncertainty households face 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. The (log) labour income, $y_{i,c,t}$, for household i belonging to cohort c at time t is defined as:

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

with G_t being a deterministic function of age only; $p_{i,c,t}$ is the permanent income component for household i belonging to cohort c at time t , while $\varepsilon_{i,c,t}$ is the transitory income shock for the same household. The permanent income component follows a martingale process:

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

where $\zeta_{i,c,t}$ is the shock to permanent income. We assume that both the transitory and permanent income shocks are normally distributed over individuals in a given cohort, with cohort-specific distributional parameters.

Income at any time after retirement is a constant, a , a fraction of the last working year's permanent labour income, such as a pension that is wholly provided by the employer and/or the state.

Utility Function.

We assume CRRA (Constant Relative Risk Aversion) utilities:

¹⁷Cash-on-hand in period t the sum of the assets carried over from time $t-1$ to t ($A_{i,c,t-1}$) augmented with the constant interest rate (R^X) and labour income in period t ($Y_{i,c,t}$):

$$X_{i,c,t} = R^X A_{i,c,t-1} + Y_{i,c,t}.$$

$$U(C_{i,c,t}) = \frac{C_{i,c,t}^{1-\rho}}{1-\rho} \quad (11)$$

where the curvature parameter $\rho \geq 0$ represents the risk aversion parameter that equals the inverse of the elasticity of intertemporal substitution. CRRA utility functions are able to capture the precautionary motive of households, the motive to “save for a rainy day”, which might be of crucial importance in understanding the consumption behaviour of households,¹⁸ and which we analyse in Section 5.

4.2 Solution and Simulation

In this section, we first show details of our calibration, and then discuss the steps of the model’s 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.

Calibration.

Time Preference. Papers estimating time preference parameter β (see for instance Gourinchas and Parker, 2002) find that the estimates vary around the value of 0.95 (at annual frequency). As a result, the most widely used value for calibrating β is 0.95; we use $\beta = 0.95$ in our model.

Risk Aversion Parameter. The existing literature reports estimated values for the risk aversion parameter, ρ , that vary roughly between 0 and 2 (see for instance Attanasio and Weber, 1993; Blundell, Browning, and Meghir, 1994; Gruber, 2013; Kovacs, Low, and Moran, 2021). We experiment with four different values of $\rho = \{0.3, 0.5, 1.0, 2.0\}$ in the model, and choose to match the aggregate consumption profile presented in Figure 3. As a result, we set parameter ρ to 0.5 in our baseline model.¹⁹

Income. The deterministic component of income (G_t in equation (9)) is approximated by a second-order polynomial of age on observed (log) income from the DHS. The estimated coefficients of this polynomial are listed in Table D.6 in Appendix D.5.

Cohorts. Taking full advantage of the cohort-level panel dataset of the income shocks we constructed in Section 3.1, we assume that households in our structural model also

¹⁸See, for example, Zeldes (1989), Kimball (1990), Deaton (1991) or Carroll (1997), who have all emphasized the importance of precautionary motives for savings.

¹⁹Note that we would have tried values lower than $\rho = 0.3$ or higher than $\rho = 2.0$ if we saw improvement in the fit of the model as we decrease/increase the risk aversion parameter.

belong to different cohorts. It is crucial to take into account the existence of different cohorts at least for two reasons. First, households in different cohorts have experienced different income shocks and different levels of uncertainty, hence their income and consumption trajectories can differ significantly. 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. We consider 10 different cohorts of households in our model. Households are grouped by their age in 2006, using five-year age intervals between the ages of 20 and 65.

All the parameter values that we use to solve and simulate the model are listed in Table D.6 in Appendix D.5.

Solution.

We use backward induction over the normalized value function of the households to obtain the optimal policy functions.²⁰ 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 for each cohort, we are able to use a two-dimensional Gauss-Hermite quadrature to approximate the expectations. See more details in Appendix D.2.

In order to take into account that different groups of households might expect and experience different income shocks, we solve each cohort's problem separately. In doing so, we consider income shock expectations and their variances to be different for different cohorts. In particular, we assume that the expected income shock for a household in a given cohort is the average of all the historically observed income shocks within that particular cohort, and therefore can be calculated as:

$$\mu_{\varepsilon,c} = \frac{\sum_i \sum_t \varepsilon_{i,c,t}}{N_c}, \quad \mu_{\zeta,c} = \frac{\sum_i \sum_t \zeta_{i,c,t}}{N_c}$$

where $\mu_{\varepsilon,c}$ ($\mu_{\zeta,c}$) is the mean value of transitory (permanent) income shock for cohort c , while N_c is the number of observations in a given cohort. The corresponding expected income shock variances can be easily computed at the cohort level as:

$$\sigma_{\varepsilon,c}^2 = \frac{\sum_i \sum_t (\varepsilon_{i,c,t} - \mu_{\varepsilon,c})^2}{N_c}, \quad \sigma_{\zeta,c}^2 = \frac{\sum_i \sum_t (\zeta_{i,c,t} - \mu_{\zeta,c})^2}{N_c}$$

where $\sigma_{\varepsilon,c}^2$ ($\sigma_{\zeta,c}^2$) is the variance of transitory (permanent) income shocks for cohort c . To calculate these statistics, we rely on the identified transitory and permanent income

²⁰Following Carroll (1992), variables are normalised by permanent income for ease of computation. In Appendix D.1, we show the detailed derivation of the standardized model.

shocks $(\varepsilon_{i,c,t}, \zeta_{i,c,t})$ from Section 3.²¹

Simulation.

Once the decision rules/policy functions are obtained via our solution algorithm, we simulate the behaviour of 50,000 households for 10 different cohorts (5,000 households per cohort). When simulating the model, we consider income shocks to be normally distributed, with cohort and time-specific distributional parameters. In particular, we assume that the mean of the distribution is the average of the observed income shocks in a particular cohort and year, and therefore can be calculated as:

$$\mu_{\varepsilon,c,t} = \frac{\sum_i \varepsilon_{i,c,t}}{N_{c,t}}, \quad \mu_{\zeta,c,t} = \frac{\sum_i \zeta_{i,c,t}}{N_{c,t}}$$

where $\mu_{\varepsilon,c,t}$ ($\mu_{\zeta,c,t}$) is the mean value of transitory (permanent) income shock for cohort c at time t , while $N_{c,t}$ is the number of observation in a given cohort c at time t . The corresponding variance of the distribution can be computed as:

$$\sigma_{\varepsilon,c,t}^2 = \frac{\sum_i (\varepsilon_{i,c,t} - \mu_{\varepsilon,c,t})^2}{N_{c,t}}, \quad \sigma_{\zeta,c,t}^2 = \frac{\sum_i (\zeta_{i,c,t} - \mu_{\zeta,c,t})^2}{N_{c,t}}.$$

where $\sigma_{\varepsilon,c,t}^2$ ($\sigma_{\zeta,c,t}^2$) is the variance of transitory (permanent) income shock for cohort c at time t . To calculate these statistics, we rely on the identified transitory and permanent income shocks $(\varepsilon_{i,c,t}, \zeta_{i,c,t})$ from Section 3.²²

In each individual simulation, we draw realizations for the two income shocks $(\varepsilon_{i,c,t}, \zeta_{i,c,t})$ from normal distributions characterized by parameters $(\mu_{\varepsilon,c,t}, \sigma_{\varepsilon,c,t}^2)$ and $(\mu_{\zeta,c,t}, \sigma_{\zeta,c,t}^2)$, respectively. We assume that each household starts its life with zero wealth, and only receives labour income; therefore, early in life households are liquidity constrained. When aggregating variables, we use cohort weights, which are representative weights of the Dutch population.²³

5 Simulation Results

In this section, we present results from our structural model. First, we discuss the simulation results from our baseline model with the identified transitory and permanent income shocks. We then show several counterfactual simulations to gauge the importance of the nature of income shocks, future income uncertainties, and cohort-effects, in

²¹Income shocks and variances by cohort are reported in Table D.1 in Appendix D.3.

²²Income shocks and their variances by cohort and year are reported in Tables D.2-D.5.

²³Cohort weights are reported in Table D.1 in Appendix D.3.

individually shaping aggregate consumption dynamics.

For the ease of comparison to Figure 3, which shows the observed evolution of aggregate variables, we normalize all the simulated consumption profiles using 2008 as the base year (=100).

5.1 Baseline Simulations.

We first describe the performance of our calibrated baseline model with the first and second moments of the identified transitory and permanent income shocks. Figure 7 shows the simulated aggregate life-cycle profile of consumption²⁴ (the solid line), its empirical counterpart from Eurostat data (the dashed line), and the consumption measure calculated from DHS data (the dotted line). As discussed in detail in Section 2.3, the latter needs to be interpreted with caution as it is subject to large measurement errors. For this reason, we focus on consumption data from the Eurostat when comparing our model to the data. Overall, the simulated model obtains a good fit with the empirical data. Matching the observed aggregate consumption is not only successful over “normal times”, but also over crisis periods. Our model generates two significant contractions in consumption over the sample period: one for 2008-2009, and another for 2011-2013, reflecting the Global Financial Crisis and the Sovereign Debt Crisis, respectively.


Consumption dynamics in our model can only be linked to three possible driving mechanisms: the nature of the income shocks, future income uncertainties, and age. Therefore, the model-implied consumption path shown in Figure 7 is a combination of these three forces, which we analyse one by one in Section 5.2. Consumption dynamics in reality, however, might be linked to other channels which are not analysed in this paper. For instance, Slacalek (2009), Mian, Rao, and Sufi (2013) and Bottazzi, Wakefield, and Trucchi (2020) show that shocks in financial and housing wealth affect consumption decline during the Great Financial crisis. Nevertheless, the three channels we examine play a key role in determining the aggregate consumption path between 2006 and 2018, as shown in Figure 7.

The Global Financial Crisis and the Sovereign Debt Crisis.

As shown in Figure 7, our model generates two substantial falls in consumption over the period of interest. Consumption in the first contraction period, between 2008 and 2009, drops by 1.9%, which corresponds exactly to what is experienced in the Netherlands

²⁴The simulated model is based on calibrated parameters that are reported in Table D.6 and risk aversion parameter of $\rho = 0.5$. For different values of ρ , results can be found in Figure E.1 in Appendix E.1.

Figure 7: Consumption dynamics: the model, aggregate and DHS data (indices: 2008=100)



Notes: The model is simulated under baseline parameter settings listed in Table D.6 and risk aversion parameter of $\rho = 0.5$. We calculate aggregate consumption in the baseline model by simulating cohort-level consumption profiles and we aggregate them using appropriate cohort weights (representative of the Dutch population). We then create a consumption index by using 2008 as the base year (2008=100). Aggregate consumption (from Eurostat) is real Household and NPISH (non-profit institutions serving household) final consumption expenditure. Data from the DHS is obtained as illustrated in Section 2.3. Shaded area indicates crisis periods.

during the Global Financial Crisis.²⁵ Consumption in the second contraction period, between 2011 and 2013 decreases by 2.5% in the model, similar to the actual 2.1% drop observed that during the Sovereign Debt crisis (Figure A.1 in Appendix A). To better understand the driving forces behind these contraction periods in our model, it is worth revisiting the income shocks that households face over the two episodes.

Starting with the Global Financial Crisis in 2008-2009, Figure 4 highlights that, on average, neither the level of permanent nor transitory income shocks are significantly negative in 2009. As a result, the 2009 consumption drop depicted in Figure 7 cannot be driven by the level of the income shocks. However, the distribution of income shocks in Figure 5 show a significant increase in the uncertainty of the transitory income shocks over the Global Financial Crisis. Compared to ‘normal’ times, in 2008-2009 large upward movements in income, triggered by transitory income shocks, become less likely, while large downward movements become more likely. This change in income uncertainty induces a stronger precautionary saving motive in our model: households, facing

²⁵See Appendix A for more details.

an increased probability of a negative income shock, wish to save more and therefore consume less.

Turning to the Sovereign Debt Crisis in 2011-2013, Figure 4 shows that permanent income shocks are significantly negative both in 2012 and in 2013, which reflect in the 2012-2013 consumption drop predicted by the model and shown in Figure 7. Furthermore, Figure 5 highlights the importance of the change in permanent income shock uncertainties: in 2011-2013 large downward movements in income, triggered by permanent income shocks, become more likely. As a result, the precautionary saving motive becomes stronger in the model: facing an increased probability of a negative income shock, households wish to save more and therefore consume less. Finally, cohort heterogeneity in income shocks affects the consumption drop in 2011-2013. As shown in Figure 6, permanent shocks during the Sovereign Debt Crisis are larger for younger cohorts, who respond more to worsening lifetime resources due to their longer time horizon.

Calculated income shocks and the results of the structural model suggest that the drivers of the contractions in consumption over the two crises are different. The increase in income uncertainties (predominantly transitory shocks) is the main determinant of the consumption fall in 2008-2009. In contrast, the negative income shocks and the increase in income uncertainties (predominantly permanent shocks) play a major role in explaining consumption drop during the Sovereign Debt Crisis.

To further examine the contribution of alternative channels in explaining consumption dynamics within our framework, we implement counterfactual simulations illustrated below.

5.2 Counterfactual Simulations.


In this section, we consider the three main channels that potentially drive consumption movements in our framework: first, the nature of the income shocks; second, the precautionary saving motive; and third, the cohort effects. We run counterfactual simulations in order to analyse the impact of these three different channels on consumption behaviour individually.

The Nature of the Shocks.

According to the textbook version of the permanent income hypothesis, only unanticipated permanent income shocks should induce substantial changes in consumption. Expected or transitory income shocks, instead, should not alter consumption significantly. When we consider a finite-period version of the permanent income hypothesis, however, the predictions are not as straightforward, given that the age of the individuals

also plays a role in the responses to different shocks. This is particularly relevant in our case, where many different cohorts of households are present at the same time.

Figure 8: Simulated consumption profile: observed shocks perceived as permanent



Notes: Both the baseline and the counterfactual model is simulated under parameter settings listed in Table D.6 and risk aversion parameter of $\rho = 0.5$. The counterfactual model assumes that all the shocks are perceived by households as permanent. We calculate aggregate consumption in both models by simulating cohort-level consumption profiles and aggregate them using appropriate cohort weights (representative of the Dutch population). We then create a consumption index by using 2008 as the base year (2008=100). Aggregate consumption (from Eurostat) is real Household and NPISH (non-profit institutions serving household) final consumption expenditure.

We illustrate how the nature of the income shocks shapes aggregate consumption profile by simulating a model variant where households perceive all the income shocks (which we have identified) as permanent. In this hypothetical scenario, households cannot differentiate between transitory and permanent income shocks, and they mistakenly consider all shocks as permanent.

Figure 8 presents simulated consumption profiles from this counterfactual model (the dotted line) together with our baseline results (the solid line) and the data (the dashed line), with the latter two being identical to those seen in Figure 7. It is evident that consumption from the counterfactual model and the data are completely different: until 2008, consumption from the model is below the observed consumption data, while after 2008 it is above.

Recalling the trajectories of permanent and transitory income shocks, shown in Figure 4 and their distributions in Figure 5, we can interpret the results in Figure 8. We

have to consider the sum of the two income shocks that, in this counterfactual model, represent the perceived permanent income shocks households face. The sum of the two shocks is only negative between 2011 and 2013, implying no significant consumption drop in our model, except for the period of the Sovereign Debt Crisis. Note also, that by perceiving permanent income shocks only, households also face higher future income uncertainties (compared to our baseline model with both transitory and permanent income shocks), and hence a stronger motive for precautionary saving. As a result, households in the model have lower consumption till 2008, compared to the data, and they do not react to the Global Financial Crisis as much as when the shocks are perceived correctly (as in our baseline model).

From this counterfactual model simulation, we conclude that the nature of income shocks is crucial in shaping aggregate consumption: by assuming that all the income shocks are perceived as permanent, the model cannot explain the dynamics of the Dutch economy between 2006-2018. In particular, the model predicts a 1.2% increase in aggregate consumption between the 2008-2009 Global Financial Crisis, as opposed to the observed drop of 1.9%. Moreover, assuming that all the income shocks are perceived as permanent, the model generates a 0.3% increase in aggregate consumption between the 2011-2013 Sovereign Debt Crisis, which is far from the observed drop of 2.1%. Note that assuming that households perceive all the shocks as transitory instead also implies unreasonable aggregate consumption profiles, as shown in Figure E.2 in the Appendix.

The Precautionary Saving Motive.

If we extend the textbook version of the permanent income hypothesis model with labour uncertainties (or liquidity constraints) and prudent preferences²⁶, the model can accommodate the so-called precautionary saving motive that triggers households to save more and consume less if the downside risk to their future income increases. Given the CRRA preferences and the presence of both labour uncertainties and liquidity constraints in our model, the precautionary saving motive plays a role in households' consumption/saving behaviour in our baseline simulations.

The consumption literature has adopted the view that precautionary saving is an important aspect of households' savings behaviour, however there are some studies suggesting that wealth accumulation for precautionary reasons is relatively small.²⁷ In what


²⁶Kimball (1990) defines relative prudence using the second and third derivatives of the utility function:

$$-\frac{c \cdot u'''}{u''}.$$

Prudence measures the strength of the precautionary saving motive as a function of risk.

²⁷Jappelli, Padula, and Pistaferri (2008), Hurst et al. (2010) or Fulford (2015) find a small effect of precautionary saving, while Carroll and Samwick (1997), Gruber and Yelowitz (1999), Dynan, Skinner,

Figure 9: Simulated consumption profile: no precautionary motive



Notes: Both the baseline and the counterfactual model is simulated under parameter settings listed in Table D.6 and risk aversion parameter of $\rho = 0.5$. The counterfactual model is the perfect certainty version of our baseline model. We calculate aggregate consumption in both models by simulating cohort-level consumption profiles and aggregate them using appropriate cohort weights (representative of the Dutch population). We then create a consumption index by using 2008 as the base year (2008=100). Aggregate consumption (from Eurostat) is real Household and NPISH (non-profit institutions serving household) final consumption expenditure.

follows, we evaluate the importance of the precautionary channel within our framework. In doing so, we simulate another model variant, assuming that households do not face income uncertainties and hence they do not accumulate wealth for precautionary reasons. Following Carroll (1997), we consider the perfect certainty version of our baseline CRRA model described above. By assuming that transitory and permanent income shocks are known in advance and that households do not expect these shocks to deviate from their expectations, we are able to cease the precautionary saving effects.

Figure 9 shows the consumption profile from the perfect certainty version of our CRRA model (the dotted line) together with our baseline consumption profile (the solid line) and the data (the dashed line). In comparison with the baseline model, the alternative model without the precautionary saving motive implies higher volatility in consumption. The reason is straightforward: if households do not expect any downside risk to their future income, they do not accumulate wealth for precautionary purposes. These households consume more and save less than those in the baseline model with the and Zeldes (2004) find that a significant fraction of savings is due to the presence of income uncertainty.

precautionary motive.²⁸ The lack of the precautionary saving motive makes households less insured against negative income shocks.

Ignoring the precautionary saving motive causes the counterfactual model to fail to match the empirical consumption profile over the period of interest. In particular, the model predicts a 0.8% decrease in aggregate consumption between the 2008-2009 Global Financial Crisis, as opposed to the observed drop of 1.9%. Moreover, in the absence of the precautionary saving motive, the model generates a 4.4% decrease in aggregate consumption between the 2011-2013 Sovereign Debt Crisis, which is more than twice as much as the observed drop of 2.1%. These results suggest that the precautionary saving motive plays a crucial role in households' consumption and savings behaviour in our framework. Therefore, our baseline model is more suitable for task of investigating consumption sensitivity to income shocks than a model without the precautionary saving motive.

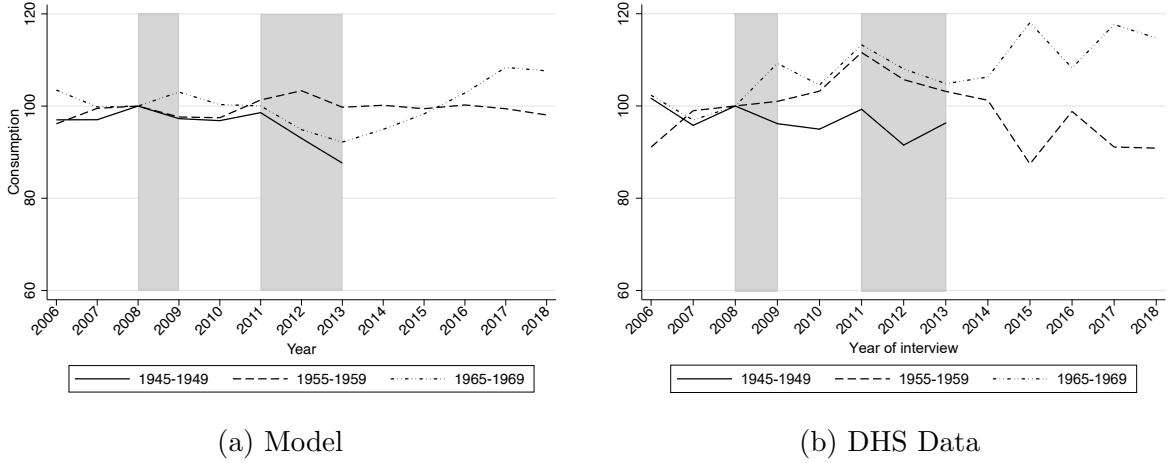
Different Cohorts.

Our structural framework is based on cohort-level differences. As we point out in Section 4.2, taking into account the existence of various cohorts is crucial at least for two reasons. First, households belonging to different cohorts might experience quite different shocks to their income. Second, households belonging to different cohorts might react differently to similar income shocks as predicted by a finite version of the permanent income hypothesis. A same-sized transitory income shock, for instance, triggers larger consumption responses for older cohorts, as they face a shorter time horizon ahead to smooth over income shocks. In contrast, a same-sized permanent income shock causes larger consumption responses for younger cohorts, as the effect of the shock on lifetime resources is greater for those with longer time horizon ahead. Next, we analyse these cohort level differences and their impact on the aggregate consumption dynamics between 2006 and 2018 in the Netherlands.

Figure 10a presents the simulated consumption paths for three different cohorts born between 1945-1949 (the solid line), 1955-1959 (the dashed line), and 1965-1969 (the dotted line), while Figure 10b shows their empirical counterparts (already seen in Figure 3). The cohort-level consumption profiles from the DNB Household Survey data exhibit higher volatility than those from our simulated model, due to large measurement errors in imputed consumption. Consequently, we interpret differences in the model and the data with caution and focus mainly on comparing general trends of consumption dynamics rather than levels of consumption.

²⁸In Figure E.3 in Appendix E.1, we compare households' savings behaviour over the life-cycle in the baseline model and in the model without the precautionary saving motive.

Figure 10: Simulated cohort consumption profiles



Notes: The model is simulated under baseline parameter settings listed in Table D.6 and risk aversion parameter of $\rho = 0.5$. We then create a consumption index for each cohort by using 2008 as the base year (2008=100). Data from the DHS is obtained as illustrated in Section 2.3. Shaded area indicates crisis periods.

During the 2008-2009 crisis, the oldest cohort suffers the largest drop in consumption both in our simulated model and in the Dutch data. This highlights the importance of the negative transitory income shock the oldest cohort faces in 2009, as seen earlier in Figure 6. These transitory income shocks can trigger considerable changes in consumption for older households, as they face a short time horizon ahead to smooth over income shocks. Moreover, as the uncertainties over these transitory income shocks also increase during the Global Financial Crisis, the impact of these shocks on consumption is further amplified. Note, however, that the reason the consumption of the oldest cohort drops most significantly is also related to life-cycle effects (i.e. households in cohort 1945-1949 become retired around 2011, which may be associated with a consumption drop; see for instance Battistin et al., 2009).

Having established that the simulated cohort-level consumption profiles during the 2008-2009 crisis are in line with those from the DNB Household Survey, we next consider how aggregate and cohort-level consumption profiles compare. As documented earlier, aggregate consumption drops by 1.9% between 2008-2009 both in our simulated baseline model and in the data. When we only consider the oldest cohort in our sample, cohort 1945-1949 for example, the drop in consumption is significantly higher with a 2.7% decrease.

During the 2011-2013 crisis, we find that consumption of all cohorts decreases substantially, both in our simulated model and in the Dutch data. Permanent income shocks trigger large declines in consumption for young households, as the effect of permanent

shocks on their lifetime resources is substantial.²⁹ Moreover, as the uncertainties over these permanent income shocks also increase over the 2011-2013 crisis, the impact of these shocks on consumption is further amplified. Comparing our cohort-level results to the observed aggregate consumption, we find that consumption would have dropped by almost five times as much as aggregate consumption if we only considered the oldest cohort, cohort 1945-1949. Consumption for this cohort drops as much as 11.1% between 2011-2013.

Ignoring the fact the households consist of different age groups is not only unrealistic but also has important implications on aggregate consumption dynamics in the model. Different cohorts experience different shocks to their income, and they respond to those shocks differently. By aggregating up the various cohorts, however, our model can successfully match the aggregate consumption profile observed in the data between 2006 and 2018, as shown in our baseline simulation in Figure 7.

6 Conclusions

In this paper, we present a unique panel dataset of transitory and permanent income shocks in the Netherlands between 2006 and 2018. As our first contribution, we identify the level of these income shocks by following the method of Pistaferri (2001) that combines subjective income expectations with income realizations. Our results show that Dutch households faced significant income shocks over the observational period, and especially during the Global Financial Crisis (2008-2009) and the Sovereign Debt Crisis (2011-2013). We find that the income shocks experienced during the 2008-2009 crisis are of a different nature than the shocks experienced during the 2011-2013 crisis, with the 2011-2013 shocks being perceived as more permanent. Looking at the income uncertainties, instead, we highlight that the two crises are somewhat similar, as the distributions of the income shocks exhibited countercyclical left-skewness. Finally, we show that the 2008-2009 crisis hit all the cohorts in a similar manner, while the 2011-2013 crisis affects the income of younger cohorts the most.

As a second contribution, we used the first and second moments of the identified shocks in a structural model to address important points of income transmission to consumption that would not be possible in a reduced-form setting. Our baseline model with the income shocks generates very similar consumption patterns to those observed in the data, both at the aggregate and the cohort level. Further, using counterfactual model simulations we demonstrate the importance of the three channels that drive consumption

²⁹Note that the consumption drop for the oldest cohort is also significant, however this drop is a combination of the negative income shocks and the predicted life-cycle dynamics of income, as households in cohort 1945-1949 are beyond retirement by this time.

movements in our theoretical framework: the nature of income shocks, the precautionary saving motive, and the cohort effects.

The strategy used in this paper is not without limitations. Probably the most important of all is the fact that the identification relies on the assumption of serially uncorrelated income shocks. A less restrictive assumption about the income process would be an MA(1) transitory income shock, in line with labour economics literature (such as MaCurdy, 1982; Abdowd and Card, 1989). In order to gain identification under serially correlated transitory income shocks, we would need further information about the transitory income shocks (the known persistence parameter and the initial level of the transitory income shock), which are not necessarily available. Another limitation is linked to the reliability and information content of the subjective data. This issue, however, is less and less problematic as survey measures improve, partially in response to contributions such as Manski (2004), who stresses the usefulness of using subjective expectations.

Beside its limitations, using income expectations and realizations to identify income shock has obvious advantages. We show that using this method allows us a more in-depth study of income shock transmission to consumption, as it allows for the direct observation of income shocks (level identification) as opposed to other strategies (for instance Blundell, Pistaferri, and Preston, 2008). With this method, we can also avoid making strong assumptions on the information sets of households, and as a result can interpret income shocks more broadly, as the sum of individual/cohort and aggregate shocks. Therefore, we are confident, that empirical papers based on the methodology proposed by Pistaferri (2001) and explored in this paper can provide new evidence in understanding the link between income shocks and consumption.

Acknowledgments

We would like to thank Rob Alessie, Orazio Attanasio, Alessandro Buccioli, Hamish Low, Mauro Mastrogiacomo, Patrick Moran, Stefano Neri, Roberta Zizza, and Francesco Zollino for helpful comments. We also thank participants in seminars at Leiden University, Ca' Foscari University, Mannheim University, and Cardiff 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).

References

- Abdowd, John M. and David Card. 1989. “On The Covariance Structure of Earnings and Hours Changes.” *Econometrica* 57 (2):411–445.
- Attanasio, Orazio, Agnes Kovacs, and Krisztina Molnar. 2020. “Euler Equations, Subjective Expectations and Income Shocks.” *Economica* 87 (346):406–441.
- Attanasio, Orazio P. and Guglielmo Weber. 1993. “Consumption Growth, the Interest Rate and Aggregation.” *Review of Economic Studies* 60(3):631–649.
- Battistini, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber. 2009. “The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach.” *American Economic Review* 99 (5):2209–2226.
- Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2021. “Asymmetric Consumption Smoothing.” *American Economic Review* 111 (1):192–230.
- 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.
- Bottazzi, Renata, Matthew Wakefield, and Serena Trucchi. 2020. “Consumption responses to a large shock to financial wealth: evidence from Italy.” *The Scandinavian Journal of Economics* 122 (2):762–789.
- Busetti, Fabio and Pietro Cova. 2013. “The macroeconomic impact of the sovereign debt crisis: a counterfactual analysis for the Italian economy.” *Questioni di Economia e Finanza (Occasional Papers)* 201, Bank of Italy, Economic Research and International Relations Area. URL <https://ideas.repec.org/p/bdi/opques/qef2013.html>.
- Caivano, Michele, Lisa Rodano, and Stefano Siviero. 2011. “The Transmission of the Global Financial Crisis to the Italian Economy.” *Giornale degli Economisti* 70 (3):1–32.
- Carroll, Christopher D. 1992. “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence.” *Brookings Papers on Economic Activity* 1992(2):61–156.

- . 1997. “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *The Quarterly Journal of Economics* 112(1):1–55.
- . 2009. “Precautionary saving and the marginal propensity to consume out of permanent income.” *Journal of Monetary Economics* 56 (6):780–790.
- Carroll, Christopher D. and Andrew A. Samwick. 1997. “The Nature of Precautionary Wealth.” *Journal of Monetary Economics* 40(1):41–71.
- Christellis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten von Rooij. 2019. “Asymmetric Consumption Effects of Transitory Income Shocks.” *The Economic Journal* 129(129):2322–2341.
- Deaton, Angus. 1991. “Saving and Liquidity Constraints.” *Econometrica* 59 (5):1221–1248.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. “Do the Rich Save More.” *Journal of Political Economy* 112 (2):397–444.
- Fulford, Scott L. 2015. “The Surprisingly Low Importance of Income Uncertainty for Precaution.” *European Economic Review* 79:151–171.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. “Consumption over the Life Cycle.” *Econometrica* 70(1):47–89.
- Gruber, Jonathan. 2013. “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution.” *Quarterly Journal of Finance* 3 (1).
- Gruber, Jonathan and Aaron Yelowitz. 1999. “Public Health Insurance and Private Savings.” *Journal of Political Economy* 107 (6):1249–1274.
- 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.
- Hayashi, Fumio. 1985. “The Permanent Income Hypothesis and Consumption Durability: Analysis Based on Japanese Panel Data.” *The Quarterly Journal of Economics* 100 (4):1083–1113.
- Hurst, Erik, Annamaria Lusardi, Arthur Kennickell, and Francisco Torralba. 2010. “The Importance of Business Owners in Assessing the Size of Precautionary Savings.” *The Review of Economics and Statistics* 92 (1):61–69.

- Jappelli, Tullio, Mario Padula, and Luigi Pistaferri. 2008. “A Direct Test of The Buffer-Stock Model of Saving.” *Journal of the European Economic Association* 6 (6):1186–1210.
- Judd, Kenneth. 1998. “Numerical Methods in Economics.” *The MIT Press* Cambridge, Massachusetts.
- 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.
- Kimball, Miles S. 1990. “Precautionary Savings in the Small and in the Large.” *Econometrica* 58 (1):53–73.
- Kovacs, Agnes, Hamish Low, and Patrick Moran. 2021. “Estimating Temptation and Commitment Over the Life-Cycle.” *International Economic Review* 62 (1):101–139.
- 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.
- Malmendier, Ulrike and Stefan Nagel. 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *The Quarterly Journal of Economics* 126 (1):373–416.
- Manski, Charles F. 2004. “Measuring Expectations.” *Econometrica* 72 (5):1329–1376.
- Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump.” *The Quarterly Journal of Economics* 128 (4):1687–1726.
- Pistaferri, Luigi. 2001. “Superior Information, Income Shocks, And The Permanent Income Hypothesis.” *The Review of Economics and Statistics* 83 (3):465–476.
- Rozsypal, Filip and Kathrin Schlafmann. 2017. “Overpersistence Bias in Individual Income Expectations and its Aggregate Implications.” CEPR

Discussion Papers UPDATE 12028, C.E.P.R. Discussion Papers. URL <https://ideas.repec.org/p/cpr/ceprdp/12028.html>.

Slacalek, Jiri. 2009. “What Drives Personal Consumption? The Role of Housing and Financial Wealth.” *The B.E. Journal of Macroeconomics* 9 (1):1–37.

Storesletten, Kjetil, Chris I Telmer, and Amir Yaron. 2004. “Cyclical Dynamics in Idiosyncratic Labor Market Risk.” *Journal of Political Economy* 112 (3):695–717.

Zeldes, Stephen P. 1989. “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence.” *The Quarterly Journal of Economics* 104 (2):275–298.