

Euler Equations, Subjective Expectations and Income Shocks*

Orazio Attanasio[†]

Agnes Kovacs[‡]

Krisztina Molnar[§]

October 25, 2018

Abstract

In this paper, we make three substantive contributions: first, we use elicited subjective income expectations to identify the levels of permanent and transitory income shocks in a life-cycle framework; second, we use these shocks to assess whether households' consumption is insulated from them; third, we use the shock data to estimate an Euler equation for consumption. We find that households are able to smooth transitory shocks, but adjust their consumption in response to permanent shocks, albeit not fully. The estimates of the Euler equation parameters with and without expectational errors are similar, which is consistent with rational expectations. We break new ground by combining data on subjective expectations about future income from the Michigan Survey with micro data on actual income from the Consumer Expenditure Survey.

Keywords: life cycle models; estimating Euler Equations; survey expectations

JEL classification: C13; D12; D84; D91; E21

*We thank for Felicia Ionescu, Hamish Low, Peter Neary, Morten Ravn, Guglielmo Weber and the participants of the Vienna Macroeconomics Workshop for helpful comments.

[†]University College London and Institute for Fiscal Studies, Email: o.attanasio@ucl.ac.uk

[‡]The University of Manchester, Email: agnes.kovacs@manchester.ac.uk

[§]Norwegian School of Economics, Email: krisztina.molnar@nhh.no

1 Introduction

In recent years and a number of contributions, starting with Manski (2004), have stressed that data on subjective expectations can be very useful. The availability of direct data on subjective expectations has many advantages. In some contexts, it is possible to avoid strong assumptions such as that of rational expectations, and to disentangle uncertainty from heterogeneity. However, despite being more common, these data have rarely been used in the context of a structural model of individual behaviour.

In this paper, we use data on subjective income expectations from the Michigan Survey (MS) to study the life cycle model of consumption and, in particular, how transitory and permanent shocks to income are reflected in consumption. In order to do that we combine data from the MS with data from the Consumer Expenditure Survey (CEX), to construct a quasi-panel that has information on both expected and realized income. This approach allows us to improve our understanding of the nature of income shocks and their effects on households' consumption behaviour.

First, we decompose income shocks into their permanent and transitory components in a life-cycle framework. We find that the standard deviation of the permanent component is 30% larger than that of the transitory component. Second, we use these shocks to establish the extent to which households' consumption reflects them or is isolated from them. We find evidence that households are able to smooth transitory shocks, but adjust their consumption in response to permanent shocks, albeit not completely. Third, by estimating the Euler equation of our model with and without expectational errors, we show that our estimates are consistent with rational expectations.

We start with a standard life-cycle model and assume that household income can be decomposed into a permanent and a transitory component (in addition to a deterministic life cycle component). For the empirical implementation, we combine data on subjective income expectations from the MS and data on income realisations from the CEX. Since these surveys interview different households in each period we combine the two datasets by creating a synthetic panel. We then show that using the approach of Pistaferri (2001), it is possible to combine income expectations and realisations in order to identify permanent and transitory income shocks separately. Once we remove predictable life-cycle effects, permanent income shocks are identified by the change in the subjective expectations of income, while transitory income shocks are identified by the difference between income realisations and their subjective expectations.

Having constructed income shock measures, we make use of one of the optimality conditions of the life cycle model, the consumption Euler equation, which can be seen as a conditional expectation of a function of data and parameters. To express it in terms of observables it is useful to re-write the Euler equation as the difference between a data

equivalent of such a function and its theoretical expectation. As we discuss below, such a ‘residual’ includes several components: expectational errors, unobserved heterogeneity or ‘taste shocks’, measurement error and, when working with a log-linearised version of the equation, innovations to the conditional second and higher moments.

Typically, expectational errors are not observed and, given rational expectations, identification is achieved by assuming that they are uncorrelated with lagged information available to the consumers. In our exercise, we construct estimates of expectational errors of the Euler equation. We use the approximation developed by Blundell, Pistaferri, and Preston (2008) to map income shocks into expectational errors of consumption growth. This approach, therefore, allows us to use our estimates of permanent and transitory income shocks directly in the Euler equation. The coefficients we obtain on these shocks have an interesting interpretation as they represent the fraction of each shock that is reflected in consumption innovations. They are therefore analogous to the parameters estimated - with a completely different methodology by Blundell, Pistaferri, and Preston (2008).

In a standard permanent income model, consumption growth should react one-to-one to permanent income shocks, while it should not be affected much by transitory shocks. We find that the coefficient on transitory income innovations is statistically not different from zero, indicating that temporary income shocks are effectively insured. We estimate the coefficient on permanent innovations at 0.29, indicating that there is a substantial amount of insurance of permanent income shocks. Blundell, Pistaferri, and Preston (2008) report estimates between 0.2 and 0.6, depending on the definition of income they use. Our results, therefore are at the lower end of the estimates obtained by Blundell, Pistaferri, and Preston (2008). We discuss why that could be the case in Section 5, after presenting our results.

Using expectation data directly in the Euler equation has several other justifications, apart from testing the empirical significance of the income shocks in affecting consumption. First, we can use subjective expectations as useful instruments when estimating the Euler equation, which imply a potential gain in the efficiency of the estimates. Second, comparing estimation results with and without expectational errors in the Euler equation can be informative about the validity of the model and about the rationality of expectations. Finally, the availability of expectations can change the nature of the identification strategies available for the estimation of the Euler equation. The estimation of Euler equation, as typically done in the literature, needs long time series data, since the orthogonality conditions only hold in expectations. As our Euler equation directly accounts for the expectational errors, the estimates are consistent even when estimated on short time series.

Our results show that accounting for expectational errors improves the efficiency of our estimates but does not lead to statistically different point estimates. This result is consistent with the hypothesis that individual behaviour is reflected in an Euler equation and, when that is estimated with appropriate instruments, it is informative about individual preferences. It is also consistent with the hypothesis that individual expectations are rational, although the tests of this hypothesis might be, in our context, where we can only use cohort averages, not very powerful.

In the last part of the paper, we simulate an artificial panel of household income and consumption in a life-cycle model. We then estimate the model counterpart of our Euler equation. By comparing the coefficients of the Euler equation estimated on real data and on the simulated data, we are able to tell whether saving through a risk-free asset can generate similar effects of permanent and transitory shocks on consumption growth as observed in the data. Our model delivers qualitatively similar results to our estimates on U.S. data. Households are able to smooth transitory shocks, while permanent shocks are reflected in consumption. However, the size of the coefficient on the permanent shocks we get using the simulated data is substantially higher than what we get in our empirical exercise. This ‘excess smoothness’ of consumption has been observed, in a different context, by Campbell and Deaton (1989). Attanasio and Pavoni (2011) interpret it as an indication that individual households can smooth consumption more than in a simple Bewley model where the only asset available for intertemporal transactions is a bond. It is possible that implicit or explicit state contingent contracts provide additional insurance possibilities.

There are several papers in the literature analysing the relationship between income shocks and consumption growth in different contexts, but only a few make use of the available data on subjective expectations. The closest papers to the present one are Pistaferri (2001) and Blundell, Pistaferri, and Preston (2008). Pistaferri (2001) uses a unique dataset, the Survey of Italian Households (SHIW), that contains both income expectations and realisations at the individual level to disentangle income shocks and examine savings behaviour. The drawback of this dataset is that expectations are only observed for two years, hence it is impossible to derive a time-series for the income shocks or to estimate an Euler equation as we do. Blundell, Pistaferri, and Preston (2008) estimate the fraction of permanent and transitory shocks reflected in consumption, just as we do. However, they use an approach that is completely different from ours: they use relative evolution of consumption and income inequality, rather than the ‘augmented’ Euler equation we use.

The rest of the paper is organised as follows. In Section 2, we describe the model taking into account expectational errors. In Section 3, we show how to identify perma-

ment and transitory income shocks separately. In Section 4, we describe the Consumer Expenditure Survey and the Michigan Survey in detail. In Section 5 we discuss the econometric issues that arise in estimating the Euler equation and present our estimation results. In Section 6, we report the results of our simulations. In Section 7, we discuss the implications of our analysis and conclude the paper.

2 Life-cycle consumption and expectation errors

We use a simple model of life-cycle consumption and savings in a dynamic stochastic framework. We make a number of stark assumptions to focus on the main points we want to make. Some of these assumptions (such as deterministic life length or the absence of bequests), can be easily relaxed and would not affect the nature of the empirical exercise we present below. After sketching the basic life cycle set up and deriving specifications that can be estimated empirically, we focus on the nature of the residuals of such equations and discuss how information on subjective expectations and expectation errors could be incorporated in them.

2.1 The life cycle problem

Household h maximises lifetime expected utility, given available resources, by choosing (non-durable) consumption $C_{h,t}$. Utility is assumed to be intertemporally separable and the future is discounted geometrically at a rate β . We assume that preferences are of the Constant Relative Risk Aversion (CRRA) form. Life is assumed to be finite and of known length T . Households do not have a bequest motive, so that they are assumed to consume all their resources by age T . We follow Attanasio and Weber (1995) and assume that utility is shifted by a number of variables. Some of them, denoted by $Z_{h,t}$ are observable to the econometrician, while others, which we denote with $v_{h,t}$ are unobservable. These variables can be thought of as reflecting changing needs over the life cycle that modify the relationship between the amount of consumption and utility enjoyed by the households. We assume that the Z variables are exogenous and deterministic from the point of view of the household. Households are assumed to be able to move resources over time using a risk-free asset.¹ We denote with $A_{h,t+j}$ the stock of asset in period $t + j$ with risk free

¹In reality, of course, households have the possibility to invest in different (riskier) assets. We keep our model simple, because for our exercise, to describe optimal intertemporal choices, it is sufficient to consider only the Euler equation for a single asset on which consumers are not at a corner. We choose the riskless asset as, for the age group we are considering, most consumer will not be at a corner. The Euler equation we consider would also hold in the presence of multiple assets and even in the presence of frictions in the markets for these other assets. Obviously it would be interesting to study changes in the marginal rate of substitution for consumers who do participate to the market for risky assets, as in Attanasio, Banks, and Tanner (2002), but that is beyond the scope of the current paper.

interest rate of r_{t+j} between periods $t+j$ and $t+j+1$. The interest rate is the same across households. Given these assumptions, the consumer problem is given as follows:

$$\max_{\{C_{h,t}\}_t^T} \mathbb{E}_t \sum_{j=0}^{T-t} \beta^j \frac{C_{h,t+j}^{1-\gamma}}{1-\gamma} e^{\theta' Z_{h,t+j} + v_{h,t+j}} \quad (1)$$

subject to the intertemporal budget constraints:

$$\begin{aligned} A_{h,t+j+1} &= (1 + r_{t+j})(A_{h,t+j} + Y_{h,t+j} - C_{h,t+j}), \quad j = 1, T-t. \\ A_{h,T} &= 0 \end{aligned} \quad (2)$$

$Y_{h,t+j}$ is the labor income at period $t+j$, which is assumed to be exogenous and is assumed to be a combination of deterministic and random components. The latter, in turn, is made of a permanent and transitory component. In particular we assume the following decomposition of log income:

$$\log Y_{h,t} = \pi' B_{h,t} + p_{h,t} + \varepsilon_{h,t}, \quad (3)$$

where $B_{h,t}$ is the vector of deterministic time-varying income components, $p_{h,t}$ is the permanent component and $\varepsilon_{h,t}$ is the transitory component, which is assumed to be normally distributed, $\varepsilon_{h,t} \sim N(-0.5\sigma_\varepsilon^2, \sigma_\varepsilon^2)$. Furthermore, in line with many previous empirical studies (MaCurdy (1982), Moffitt and Gottschalk (2011), Meghir and Pistaferri (2004), Blundell, Pistaferri, and Preston (2008)) we assume that the permanent component follows a martingale process of the form

$$p_{h,t} = p_{h,t-1} + \zeta_{h,t}, \quad (4)$$

with ζ_{ht} being the serially uncorrelated innovation to the permanent income, with normal distribution, $\zeta_{h,t} \sim N(-0.5\sigma_\zeta^2, \sigma_\zeta^2)$. The transitory and permanent income shocks, $\varepsilon_{h,t}$ and $\zeta_{h,t}$ are uncorrelated with each other. This is an income process that is widely used in labor economics, and has been shown to fit income data well (Carroll (2001)), although some of the contributions in labour economics that estimate this process (such as MaCurdy (1982) and Abdowd and Card (1989)) model the temporary shocks as an MA(1). We work with an i.i.d. process for the temporary component because of the nature of our data, as we discuss below. In Appendix 4 we discuss how our results would change in the presence of an MA(1) process.² Labor income at any time after retirement

²Recent work on income dynamics, that allow general heterogeneous lifetime income profiles (Guevenen (2007)) allow less overall persistence in the permanent component. It would be a very useful exercise to extend the model in this direction. We leave this for further research given the data limitations we have.

is assumed to be zero.

To control for predictable life-cycle effects, in the empirical analysis we also assume that the deterministic time-varying income component of income can be well approximated by a quadratic polynomial in age (see also Pistaferri (2001)) and therefore³

$$\pi' B_{h,t} = \pi_0 + \pi_1 age_{h,t} + \pi_2 age_{h,t}^2. \quad (5)$$

Substituting equation (5) in equation (3), the log of labor income can be written as follows

$$\log Y_{h,t} = \pi_0 + \pi_1 age_{h,t} + \pi_2 age_{h,t}^2 + p_{h,t} + \varepsilon_{h,t}. \quad (6)$$

2.2 Euler Equations and the Expectational Error

Given the problem above, the household chooses consumption paths that satisfy a number of first order conditions: the Euler equations. Focusing on the Euler equation is particularly useful because, even in the simple set up we have sketched, it is impossible to obtain closed form solutions for consumption. In our context, the Euler equation for optimal consumption is such that the discounted expected marginal utility is kept constant over time.

$$C_{h,t}^{-\gamma} = \mathbb{E}_t [\beta(1 + r_t) e^{\theta' \Delta Z_{h,t+1} + \Delta v_{h,t+1}} C_{h,t+1}^{-\gamma}]. \quad (7)$$

where \mathbb{E}_t is the expectation operator, that takes expectations of variables conditional on the information available the household h at time t . These Euler equations are equilibrium conditions that can be used to derive orthogonality conditions in order to estimate parameters and test the validity of some model assumptions. In particular, if we define the expectational error for the Euler equation as:

$$\tilde{u}_{h,t+1} \equiv \beta(1 + r_t) e^{\theta' \Delta Z_{h,t+1} + \Delta v_{h,t+1}} \left(\frac{C_{h,t+1}}{C_{h,t}} \right)^{-\gamma} - 1 \quad (8)$$

assuming rational expectations implies that such an error is orthogonal to any information available to the consumer:

$$\mathbb{E}_t [\tilde{u}_{h,t+1} \mathbf{W}_{h,t}] = 0 \quad (9)$$

³At individual level, one could control for other components of predictable income, like occupation, education, industry, household demographic variables (see Carroll and Samwick (1997)). As we discuss below, the empirical analysis will be done at the level of year of birth cohort, and, at this level, these changes depend on the cohort composition and would be complicated to keep track of.

where $\mathbf{W}_{h,t}$ is a vector of variables available to the individual household h at time t . Equation 9 can be used to obtain estimates of the structural preference parameters and, if the dimension of the vector $\mathbf{W}_{h,t}$ is larger than the number of parameters to estimate, to test the validity of the model.

When taking the model to the data, for a variety of reasons discussed, for instance, in Attanasio and Low (2004), it is useful to log-linearize the Euler equation (8). Log-linearizing is particularly useful when considering an income process which is linear in logs, such as the one considered above. Following, for example, Hansen and Singleton (1983), log-linearizing the Euler equation (8) yields an expression of the following form:

$$\Delta \log C_{h,t+1} = \alpha + \frac{1}{\gamma} \log(1 + r_{t+1}) + \theta' \Delta Z_{h,t+1} + u_{h,t+1} \quad (10)$$

where the parameter $1/\gamma$ is the elasticity of intertemporal substitution, whilst α contains constants and the unconditional means of second and higher moments of consumption growth and real interest rate.

The residual term $u_{h,t+1}$ is made of several components: it contains the expectational errors $u_{h,t+1}^{\text{exp,c}} \equiv [\Delta \log C_{h,t+1} - \mathbb{E}_t[\Delta \log C_{h,t+1}]]$ and $u_{h,t+1}^{\text{exp,r}} \equiv \frac{1}{\gamma} [\log(1 + r_{t+1}) - \mathbb{E}_t[\log(1 + r_{t+1})]]$, the unobserved heterogeneity term $\Delta v_{h,t+1}$, possibly measurement error in consumption and the deviations of conditional second and higher moments of consumption growth and real interest rate from their unconditional means. We denote with $\eta_{h,t+1}$ all the components of $u_{h,t+1}$ except for the expectational error⁴ and write:

$$u_{h,t+1} = u_{h,t+1}^{\text{exp,c}} + u_{h,t+1}^{\text{exp,r}} + \eta_{h,t+1} \quad (11)$$

This paper's focus is on the expectational error part of the residual $u_{h,t+1}$. In particular, we will use information on elicited subjective expectations to obtain measures of these quantities that can be inserted in equation (10) when bringing it to data. There are several reasons to do that. First, it can improve the efficiency of the estimation procedure. Second, comparing the results one obtains when using these measures to those obtained without them can be informative about the validity of the model and, indirectly, about the rationality of expectations. Finally, and more subtly, the availability of subjective expectations (and expectational errors) can change the nature of the identification strategies available for the estimation of equation (10). It is to this last point we turn now.

To use the orthogonality conditions in equation (9) it is necessary, in general, to use T-asymptotics to obtain consistent estimates of the parameters of interest. Whilst the point is not fully appreciated, it has been made in a number of places: Chamberlain

⁴Notice that $\eta_{h,t+1}$ is not necessarily i.i.d.. Its properties will depend on the nature of the taste shocks v , the process by which conditional higher moments evolve over time and measurement error.

(1984) is one of the first references, while Hayashi (1987) and Attanasio (1999) also discuss it extensively. The issue is quite intuitive: to exploit the orthogonality conditions in equation (9) it is not enough to have many observations in the cross-section as expectations errors will not average out to zero in the cross section. Estimating the average error at a point in time (for instance adding a time dummy) is not enough as for every instrument one considers, one would have to add an additional parameter, a point discussed clearly by Altug and Miller (1990). Only when markets are complete (so that idiosyncratic risk is diversified and there is a unique aggregate shock), does such a strategy achieve identification. The implication of this discussion is that unless one is willing to assume complete markets, orthogonality conditions that include expectational errors require a long time series so that, under rational expectations, these errors can average out to zero.

The availability of information about expectational errors can change the empirical and identification strategy of the model considered substantially. In particular, one does not necessarily need a long period to ensure that unobserved components of the residuals average out to zero. And even if the information on expectational errors is not perfect, one can use in the estimation of Euler equation as long as the deviation between actual expectations and the available measure of expectations is uncorrelated with the instruments used in estimating the Euler equation. Finally, one can also use information on subjective expectations as useful instruments when estimating the Euler equation. This, and the fact that the expectational error might account for a fraction of the residual of the Euler equation imply a potential gain in the estimates' precision.

Although some recent papers, such as Crump et al. (2015), use data on subjective expectations on consumption growth, most data with subjective expectations questions refer to income and inflation. We therefore need to relate expectations and innovations to income to consumption innovation. We follow Blundell, Pistaferri, and Preston (2008) and by an approximation we relate the expectational errors on consumption changes to permanent and transitory innovations to income. Given the power utility assumption and the log-linear income process considered above, Blundell, Pistaferri, and Preston (2008) derive the following expression:

$$u_{h,t+1}^{\text{exp,c}} = \phi \zeta_{h,t+1} + \psi \varepsilon_{h,t+1} \quad (12)$$

Permanent income shocks, $\zeta_{h,t+1}^{\text{exp}}$ have an impact on consumption growth innovations with a loading factor ϕ , while transitory income shocks, $\varepsilon_{h,t+1}^{\text{exp}}$ have an impact on that with loading factor ψ .⁵ The parameters ϕ and ψ reflect the ability households have to

⁵ Blundell, Pistaferri, and Preston (2008), allow the coefficients ψ and ϕ to be time-varying and identify them by considering movements in the cross-sectional distributions of income and consumption.

smooth income shocks. They depend on the type of markets households have access to in order to insure idiosyncratic shocks as well as on the nature of the income shocks (aggregate and idiosyncratic) that hit them. Transitory shocks should be considerably easier to insure, while permanent shocks, especially of an aggregate nature, should be reflected into consumption. In a standard Bewley model with an infinite horizon, for instance, $\phi = 1$, while $\psi = 0$. In a more complex model, where individuals have access to some contingent assets that might be allowing to smooth out part of the idiosyncratic permanent shocks, ϕ might be lower than 1 (see, for instance, Attanasio and Pavoni (2011)).

To sum up, we can write the expectational-error-adjusted log-linearised Euler equation in the following form:

$$\Delta \log C_{h,t+1} = \alpha + \frac{1}{\gamma} \log(1 + r_{t+1}) + \theta' \Delta Z_{h,t+1} + \phi \zeta_{h,t+1}^{exp} + \psi \varepsilon_{h,t+1}^{exp} + \kappa u_{h,t+1}^{exp,r} + v_{h,t+1} \quad (13)$$

The main contribution of this paper is the use of direct estimates of $\zeta_{h,t+1}^{exp}$, $\varepsilon_{h,t+1}^{exp}$ and $u_{h,t+1}^{exp,r}$ derived from questions aimed at eliciting subjective expectations of income, interest rates and inflation. In addition to the potential efficiency gains in estimating equation (13) using direct estimates of expectational errors, we are also able to test the empirical significance of the three shocks in affecting actual consumption growth, by identifying the parameters ψ , ϕ and κ separately. Each of these parameters measures the effect of innovations of different components of income and interest rates on consumption growth. In doing so, we are able to test alternative models of consumption smoothing. In this respect, the first two parameters are particularly interesting: as discussed above, a simple Bewley model would imply $\phi = 1$ and $\psi = 0$, in contrast with the evidence on 'excess smoothness' of consumption presented, for instance, by Campbell and Deaton (1989), Blundell, Pistaferri, and Preston (2008) and Attanasio and Pavoni (2011), who estimate ϕ to be significantly less than 1.

3 Identification of Income Shocks

The income process described by equations (3)-(5) has been used extensively in the study of consumption behaviour and, in particular, in models of life cycle consumption. The decomposition of income shocks in 'permanent' and 'transitory' components is particularly useful as the model has, given a certain asset structure, very strong implications about how consumption should react to them: transitory shocks should be smoothed out, while permanent ones should not. In this section, we show how with the

In what follows, we exploit mainly the time-series variation and estimates of the income shocks, so that we cannot allow time-varying loading factors.

parametrization of the income model in equations (3)-(5) and data on subjective expectations on income and data on actual income over time, it is possible to follow Pistaferri (2001) and identify separately transitory and permanent shocks.

We assume that expectations held by individual consumers are rational and that the parameters π_1 and π_2 in equation (5) are already estimated and known by the econometrician. In Section 4.2 of the paper we also show how we estimate these parameters on the dataset available. For now, using the above given income process, we can write the one-period ahead expected income as follows:

$$\begin{aligned}\mathbb{E}\left[\log Y_{h,t}|\Omega_{h,t-1}\right] &= \pi_0 + \pi_1 age_{h,t} + \pi_2 age_{h,t}^2 + p_{h,t-1} \\ \mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] &= \pi_0 + \pi_1 age_{h,t+1} + \pi_2 age_{h,t+1}^2 + p_{h,t}\end{aligned}\quad (14)$$

where $\Omega_{h,t}$ refers to the information set available to the consumer h at time t . Subtracting one equation in expression (14) from the other we obtain:

$$\mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] - \mathbb{E}\left[\log Y_{h,t}|\Omega_{h,t-1}\right] = \pi_1 + \pi_2 + 2\pi_2 age_{h,t} + p_{h,t} - p_{h,t-1}\quad (15)$$

Using this expression and the definition of permanent income in equation (4), permanent income shocks are easily calculated:

$$\zeta_{h,t} = \mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] - \mathbb{E}\left[\log Y_{h,t}|\Omega_{h,t-1}\right] - \pi_1 - \pi_2 - 2\pi_2 age_{h,t}\quad (16)$$

In words, permanent income shocks are identified by the change in the subjective expectations of income, once one removes predictable life-cycle effects. Next, note that the expectational error in income can be written as the sum of the temporary and permanent income shocks:

$$\log Y_{h,t} - \mathbb{E}\left[\log Y_{h,t}|\Omega_{h,t-1}\right] = \zeta_{h,t} + \varepsilon_{h,t}\quad (17)$$

Therefore, it is possible to compute transitory income shocks by subtracting equation (16) from equation(17):

$$\varepsilon_{h,t} = \log Y_{h,t} - \mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] + \pi_1 + \pi_2 + 2\pi_2 age_{h,t}\quad (18)$$

that is, the income innovation between time t and $t + 1$ given the information available at time t and a factor that governs predictable life-cycle income.

We have therefore established that both temporary and permanent income shocks can be easily identified by combining observed and expected income data at hand. As it is detailed in the next section, merging the Michigan Survey with the Consumer Expen-

diture Survey provides all the information which is necessary to implement equations (16) and (18) and to identify the income shocks separately.⁶

Changing the nature of the income process might affect the ability to identify the income shock components. If the transitory shocks, for example, follow an MA(1) process, as in MaCurdy (1982) and Abdowd and Card (1989), we cannot identify permanent and transitory shocks separately. We discuss this issue further in Section A.4 in the Appendix and we show the consequences of a misspecified transitory income shock later in Section 4.3.

4 Data Description

For our estimations we combine three sources of data. The Consumer Expenditure Survey (CEX) is used to obtain the household level data that is needed in estimating Euler equations (10) and (13). We obtain data on subjective expectations, which are not collected in the CEX, from the so-called Michigan Survey of Consumers. To calculate expectational errors of macro variables we use the macro data from the Federal Reserve Economic Data (FRED). In order to calculate expectational errors of household income, we match the Michigan Survey to the CEX data. As we combine two surveys that interview different samples of households, neither of which is followed over time, we use synthetic panel techniques as those pioneered by Deaton (1985) and Browning, Deaton, and Irish (1985). These techniques consists in following groups of households with fixed membership, rather than individual households.

4.1 CEX dataset

The CEX is a survey run by the Bureau of Labor Statics, which, in the first two decades of its existence, interviewed about 5000 households every quarter. The sample is representative of the *U.S.* population. 80 percent of them are then reinterviewed the following quarter, but the remaining 20 percent are replaced by a new, random group. Hence, each household is interviewed at most four times over a period of year. After 1998, the size of the sample increased dramatically to about 7500 interviews per quarter.

Given the rotating panel nature of the survey, it is not possible to follow individual households for more than the four quarters over which it is observed. For the purpose of studying life cycle behaviour we therefore use synthetic panel techniques and, naturally, define groups by the year of birth of the household head, or cohorts. Cohorts are defined

⁶Since we work with quarterly data, but expectations are collected every quarter for one year ahead, we have to be careful when applying equations (16)-(17). Moreover, we have to make a further assumption that cohorts do not update their information set within a year. See details in the appendix.

over five year bands, as reported in Table 1. The head is defined as the male in the male-female couple and as the reference person otherwise. We examine quarterly cohort averages instead of individual data. This way we have sufficient time dimension for our analysis and we can follow more or less homogeneous groups over time. It is important to construct cohorts with a big cell size (number of observations per quarter per cohort) to minimize the impact of unobserved household heterogeneity on the cohort averages.

<i>Cohort</i>	<i>Year of Birth</i>	<i>Age in 1994</i>	<i>Average Cell Size</i>	
			<i>in CEX</i>	<i>in MS</i>
1	1970-74	20-24	442	108
2	1965-69	25-29	496	137
3	1960-64	30-34	567	165
4	1955-59	35-39	554	175
5	1950-54	40-44	491	156
6	1945-49	45-49	381	129
7	1940-44	50-54	261	92

Table 1: Cohort Definition

During the interviews, a number of questions are asked concerning household characteristics and detailed expenditures over the three month prior to the interview. We make use of the following household characteristics: age of the household head, family size, and the number of children who are younger than 2 years of age. We use non-durable consumption expenditure data, which is available on monthly basis for each household. We create quarterly consumption by aggregating monthly expenditures. To avoid the complicated error structure that the timing of the interviews would imply on quarterly data, we take the spending in the month closest to the interview and multiply it by three (see also Attanasio and Weber (1995)).

We exclude non-urban households⁷ and those households who have incomplete income information. Furthermore, we only keep households of which the head is at least 21 and no more than 60.⁸ We ended up with 233,443 observations (interviews), for around 85,880 households for the sample period 1994q1-2012q4. We work with real data, hence we deflate all variables by the consumer price index.

⁷Most papers in the literature that use CEX data drop rural households, who are, obviously, very different from urban ones. Moreover, rural households were not interviewed in some years.

⁸For a more detailed explanation about the exclusions see section 5.

4.2 Survey of Consumers and Aggregate Data

The Survey of Consumers is a monthly survey conducted by the Survey Research Centre at the University of Michigan. Each month around 500 interviews are conducted by telephone and the respondents answer approximately 50 questions. Each of these questions tracks a different aspect of consumer attitudes and expectations. The Survey focuses on three areas: how consumers view prospects for their own financial situation, how they view prospects for general economy on the short and long term. In our estimations we make use of elicited expectations on four variables: household income, inflation, interest rate and unemployment rate. We have altogether 72,809 observations on a quarterly basis on the same sample as the CEX, 1994q1 to 2012q4. From these we generate the same cohorts as in the CEX dataset (see table 1).⁹

Our main survey variable in our Euler equation estimations is household income, because, as we have shown in the previous section, the expectational error of this variable affects the consumption path. Consumers are surveyed about the expected change in their family income both qualitatively and quantitatively. Since most of the households answered both questions, we opt to use the quantitative answers in our analysis:¹⁰

“By about what percent do you expect your (family) income to increase/decrease during the next 12 months?”

It is not clear from the wording of this question whether households have before or after tax income in mind when replying. In our analysis we use before tax income, however the results do not change if we use after tax income.

We merge the Michigan Survey data with the CEX data at the cohort level to calculate expectational errors of household income.¹¹ We calculate a cohort’s income expectations with multiplying their actual income from the CEX dataset with the cohort’s average expected percentage change of family income from the Michigan Survey.

The reported one-year-ahead income growth expectations also let us identify π_1 and π_2 , the coefficients of the deterministic income component in equation (6). The fact that $\mathbb{E}\left[\log Y_{h,t+1} - \log Y_{h,t} | \Omega_{h,t}\right] = (\pi_1 - \pi_2) + 2\pi_2 age_{h,t+1} - \varepsilon_{h,t}$ implies that simply regressing reported expected income growth on a constant and on age, we can obtain estimates for π_1 and π_2 ($\hat{\pi}_1$ and $\hat{\pi}_2$).

In addition to household level income expectations, we use data on subjective expectations on three macro variables that may be relevant for the household’s dynamic

⁹For completeness, we note that similarly to the CEX, the Michigan survey also has a rotating panel component, a fraction of households are re-interviewed in half a year.

¹⁰We also estimated Euler equations using qualitative expectations on household income and the results remain unchanged.

¹¹For an alternative matching of the two datasets see Souleles (2004), who uses imputation to match at the individual level.

consumption choice: inflation, interest rates and unemployment rates. Inflation and interest rate expectations enter the Euler equation, and it’s expectational errors will show up in the error term. Unemployment rate expectations might impact the household’s outlook on their own employment status and future earnings. The expectation questions on these variables in the Michigan Survey, however, are of a ‘qualitative’ nature.¹² For example consumers are asked:

“No one can say for sure, but what do you think will happen to interest rates for borrowing money during the next 12 months will they go up, stay the same, or go down?”

We quantify these ‘qualitative’ expectations on the three macro variables by a method, detailed in Appendix A.2, and due to Carlson and Parkin (1975). This approach has three crucial assumptions, which make it possible to recover quantitative expectations from qualitative survey answers. First, the distribution of the expected change of each economic variable is assumed to be known. Second, it assumes that a respondent of the survey has an indifference interval around zero: her qualitative answer will only be different from ‘no change’, if her quantitative expectation of the change in that economic variable is greater/smaller than some cutoff value c . We assume that this cutoff value is symmetric around zero and the same for all respondents.

We compute expectational errors on inflation, interest rates and aggregate unemployment rates by subtracting the subjective expectations on these variables from actual data, taken from Federal Reserve Economic Data (FRED), St. Louis Fed.

	1994		2003		2012	
	CEX	MS	CEX	MS	CEX	MS
Age	39.73	39.09	44.07	43.83	49.29	50.21
Family size	2.84	2.90	2.91	2.82	2.85	2.74
No. of children	0.94	0.93	0.95	0.95	0.80	0.77
White	0.83	0.85	0.82	0.86	0.79	0.85
HS graduate	28.38	31.02	25.92	25.10	24.20	21.22
College dropout	28.01	23.83	18.82	22.97	18.16	29.69
At least College	30.21	39.34	42.80	47.46	45.94	45.15

Table 2: Comparison of Means: CEX and MS

¹²Quantitative questions are also available on inflation expectation, but we decided to use the qualitative answers for two reasons. First, using quantitative inflation measure did not change our regression results significantly. Second, much more households answer the qualitative question than the quantitative one.

4.3 Descriptive Statistics

In Table 2, we compare the average demographic and socioeconomic characteristics of households observed in the two different dataset for selected years: 1994, 2003 and 2012. There is basically no difference in the age of respondents between the CEX and the Michigan Survey and a slight difference only in terms of other demographic variables. The only visible difference between the two datasets is in the distribution of households by schooling levels. The Michigan Survey tends to overrepresent higher educated households in the sample.¹³

Figure 1 plots the quantitative (for income changes) and quantified (for inflation, changes in unemployment rates and changes in interest rates) one year ahead average survey expectations, together with the 12-month growth in actual data. In line the wording of the survey question, which asks consumers about the expected direction of change one year ahead, interest rate and unemployment rate are shown as annual percentage point change while annual percentage change in the Consumer Price Index (CPI) and family income is used.

Comparing survey questions to actual data is hindered by two issues. One, the method of quantification affects the level of the of the the expected relevant variable, and second, wording of the Michigan survey is not explicit about the variable respondents are asked to forecast. Therefore comparing survey responses to actual data requires some assumptions.

To be more precise, with the Carlson and Parkin (1975) method we choose a symmetric cutoff value c , which is the cutoff over which individuals are assumed to answer the qualitative question as ‘increase’ or ‘decrease’. The level of the expected relevant variable is only identified up to a proportional constant, given by this cutoff value c . We choose this constant arbitrarily at 1%. This implies that the comparison between the actual and expected series should be done with caution: for the expectations derived from the qualitative answers, the changes over time (rather than the level) of these expectations should be compared to actual data.¹⁴

To deal with the fact that interest rate expectations pertain to “interest rates for borrowing money,” and do not specify the measure it refers to, we assume that the answers to an analogous question about government bonds with would be the same.¹⁵

¹³In Section 5 we also show estimates for different different education groups. This way we can gauge whether household choices differ with schooling, and we can also make the households matched from the CEX to the Michigan survey more similar in their schooling.

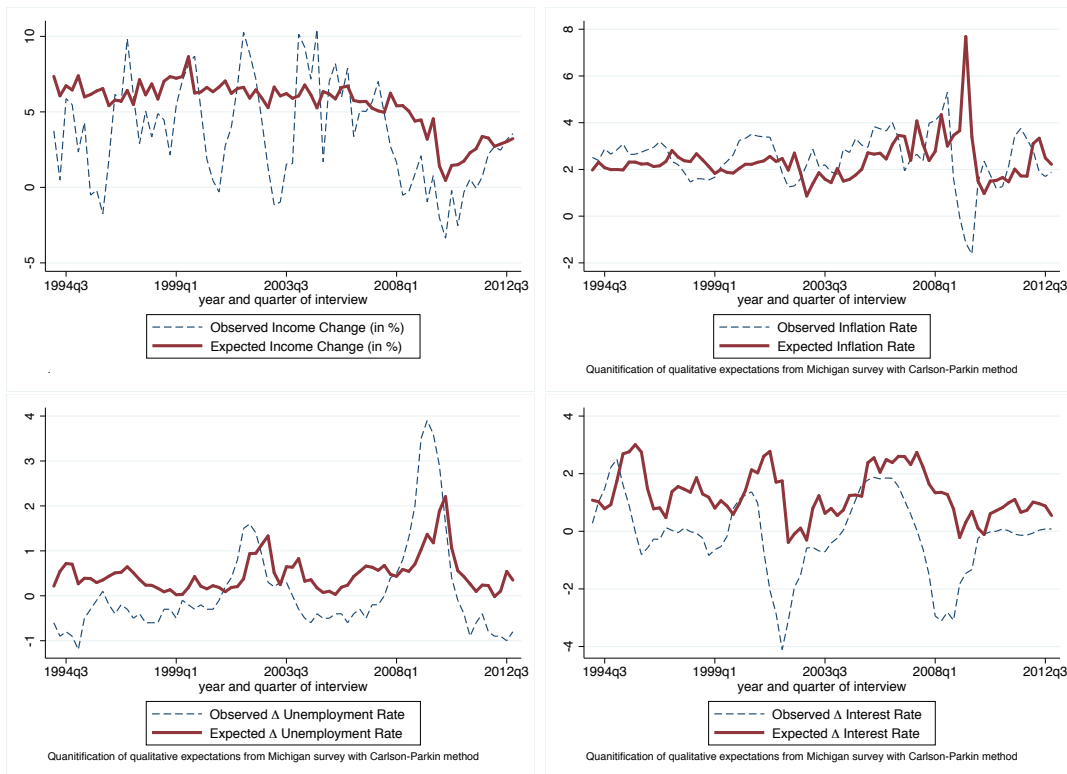
¹⁴In our alternative quantification, we used cutoff values such that the mean of the survey is equal to the mean of the actual data. This does not impact the main results, but the variability of the survey responses depend very much on the actual data chosen, see next paragraphs.

¹⁵This is a good assumption as long as the spread between the borrowing rates that a household has in mind and the corresponding government bond rate do not vary too much.

As the survey asks about one year ahead interest rate changes, in Figure 1 we compare them to the year on year change of Treasury Bill rates with one year maturity. The level of interest rate expectations is generally higher than the actual interest rate because of the quantification, it only makes sense to compare changes in the level.

Similarly, the survey question about prices refer to the general level of inflation expected for the next 12 months, and does not specify a particular measure. We assume this refers to the headline inflation, measured by year-on-year change of the the CPI. This is sensible assumption, given that the CPI measures the price of the consumption basket of a typical household.

Figure 1: Expectations and Actual Variables

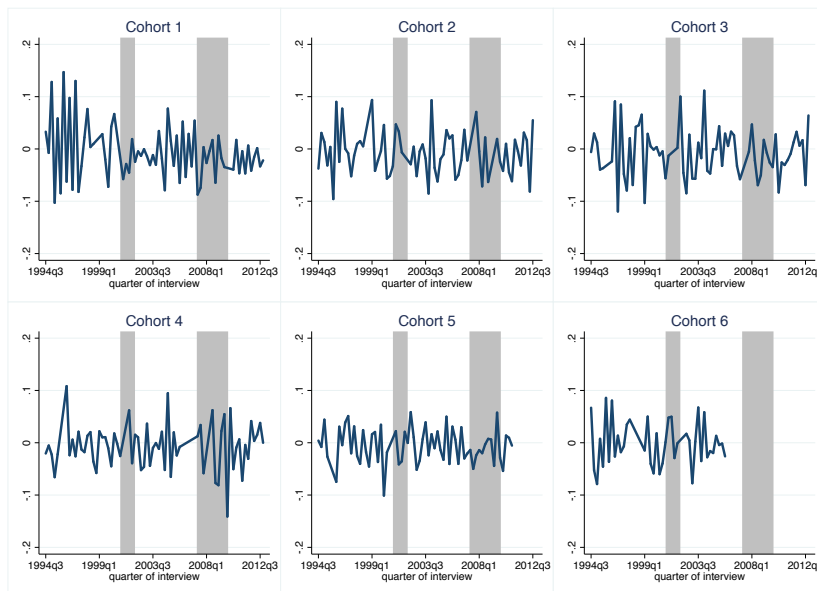


To deal with the fact the survey asks about “family income” and does not specify whether this is before or after tax income, we simply try both measures. Figure 1 and our baseline estimation uses before tax income, but the main results are the same using after tax income. Recall, that to calculate expected income we chain the expected change from the survey to the actual income from the CEX at cohort level. Next, we calculate income shocks using the level of income expectations and the level of actual income. The choice of before or after tax income impacts mainly the time pattern of transitory shocks, and not the permanent shocks. The top-left panel of Figure 1, reports actual and expected nominal income changes. Note that expectations are naturally much smoother

than actual income movements, which follows from the fact that temporary shocks do not change income expectations, but do impact actual income. In our estimations we use real expected income, and deflate income expectations with the CPI data. We opted against deflating with individual inflation expectations, as this would introduce an additional noise to the data.

One feature that emerges from these graphs is a well known pattern of expectation surveys: households often revise their one year ahead expectations in line with changes in the current data. For example when unemployment rate grows more than before, households forecast this to happen one year ahead as well. (More on this see for example Ang, Bekaert, and Wei (2007), Coibion and Gorodnichenko (2012), Long (1997), Dotsey and DeVaro (1995).) Nevertheless, average surveys are still very good forecasters; Ang, Bekaert, and Wei (2007) shows that the Michigan inflation survey is largely unbiased and it forecasts better than state of the art forecasting methods¹⁶

Figure 2: Permanent Income Shocks by Cohorts

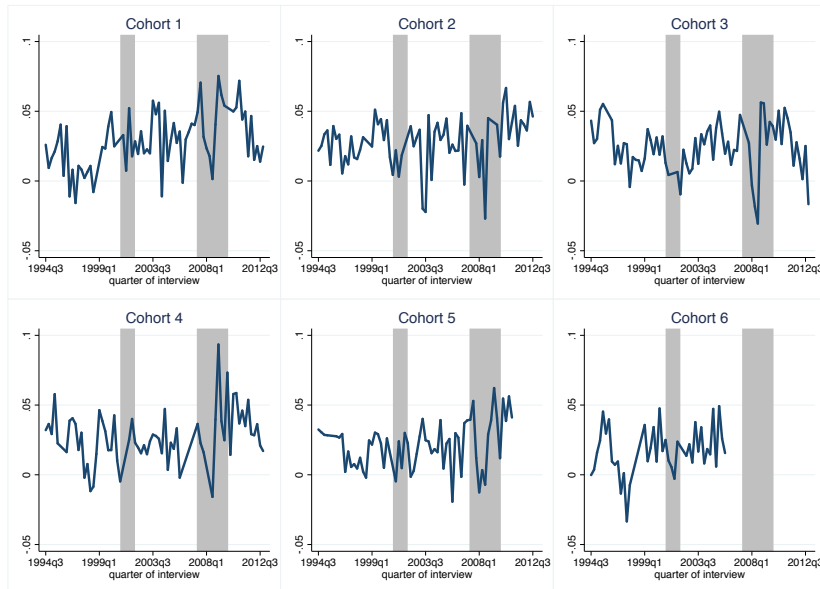


The impact of the great recession, which started in December 2007 (US National Bureau of Economic Research definition) is clearly visible in Figure 1. There was a remarkable decline in household income and income expectations as well. After the 2nd quarter of 2008 average household income kept declining and income growth stayed low throughout our sample. Households' income growth expectations followed suit, yet with a delay: one-year-ahead income growth expectations decreased in the 4th quarter of

¹⁶Ang, Bekaert, and Wei (2007) compares survey forecasts to time series, term structure and model based methods, including forecast combinations.

2008. This pessimism in households' income growth expectation was long lasting, after 2010 average income growth expectations dropped on average by 6 percentage points. Unemployment rate and its survey expectations were increasing at the beginning of the crises. Unemployment rate peaked at the end of 2010, then started declining; this was forecasted remarkably well by households. The monetary policy response to the crises is visible on the second graph in Figure 1. The treasury bill rate and its survey expectations declined because of the monetary easing: the Federal Reserve repeatedly decreased its leading interest rate in 2008-9 and implemented a large scale asset purchase program. Interestingly, during the great recession the largest deviation between expected and actual data is for the figures on inflation. While actual inflation declined dramatically and even became negative, the Michigan survey suggests that households seemed to have believed that the monetary stimulus will be effective and raise inflation.

Figure 3: Transitory Income Shocks by Cohorts



Having observations on actual household income from the CEX and expected households income growth from the Michigan Survey, we can apply the method summarised by equations (16) and (18) in Section 3, to compute the levels of the permanent and the transitory income shocks.

Figure 2 and 3 plot the log levels of the estimated permanent and transitory income shocks (ζ and ε) for the period 1994q1 to 2012q4, and for the cohorts 1-6 as defined in Table 1. The grey shaded areas in each of these graphs indicate recession periods in the U.S. as reported by the NBER. The youngest 3 cohorts seem to be subject to larger permanent income shocks than the oldest three.

As for the temporary shocks, we notice the large size, for most cohorts, of the 2008 recession shocks.¹⁷

Figure 4: Cumulative Periodogram White-Noise Test for the Permanent Income Shock

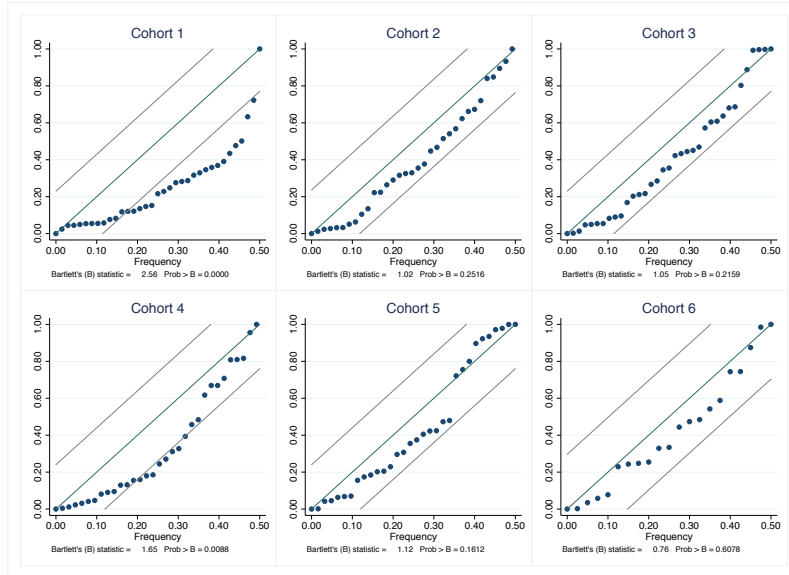
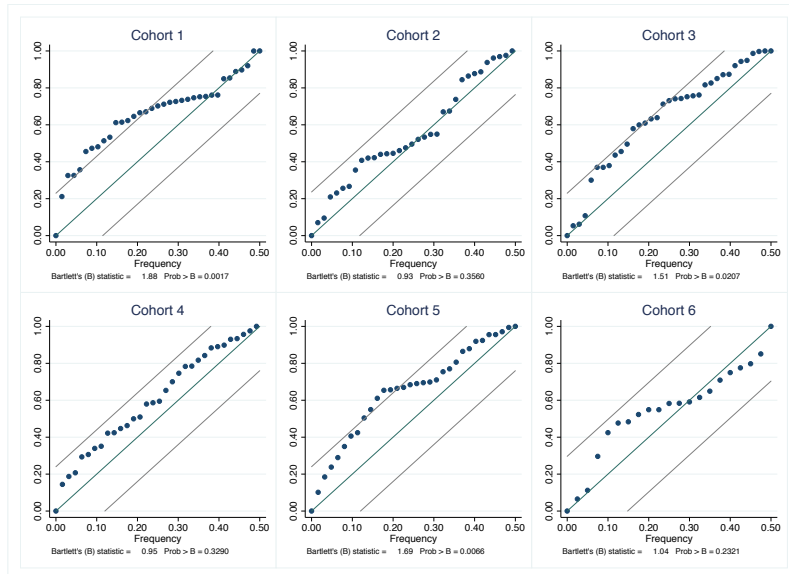


Figure 5: Cumulative Periodogram White-Noise Test for the Transitory Income Shock



¹⁷Some blips also seem to coincide with large stock market movements. While it is difficult to give a structural interpretation to the large changes observed in the Figures, it is possible that large innovations in financial markets might generate changes in expectations. This, together with recent evidence on irrationality in return expectations (See Martin (2017), Greenwood and Shleifer (2014)) suggest that it would be an interesting avenue for future research to exploring links to expectations of risky returns. Unfortunately, our data limitations do not allow to examine individual expectations and decisions.

As we mentioned above, the identification of the income shocks crucially depends on whether this process is specified correctly. To show whether our cohort-level data is consistent with the assumptions we made on the income process, we take a look at the nature of the calculated income shocks. More specifically, we use the Bartlett test, which applies the Kolmogorov-Smirnov test to the cumulative periodogram to assess whether any group of autocorrelations of the income shock time series are different from zero. The results are plotted in Figures 4 and 5.

The interpretation of these graphs is fairly straightforward. When the test values are within the 95% confidence band (which is also shown in the graph), we do not reject the null hypothesis that the process we are testing is actually a white noise process. For the permanent income shocks, the test does not reject the hypothesis of i.i.d. innovations for most of the cohorts. The results are very similar for transitory income shocks. Therefore, we conclude that the stochastic part of the income process in fact is well described by the sum of a random walk component and a white noise component.

For the cohorts for which we reject the null of i.i.d transitory shocks, we obtain estimates of first order autocorrelation of the order of 0.2, which is not inconsistent with the MA(1) specifications that have been estimated in the literatures (MaCurdy (1982) and Abdowd and Card (1989)). In order to see how our calculated shocks would be affected by the assumption of an MA(1) transitory income component, in graph 6, we plot transitory and permanent income shocks under two scenarios. First, we assume that the transitory shock is i.i.d., then we assume that it follows an MA(1) process with a persistence parameter of $\rho = 0.2$.¹⁸ As it is seen in the graph, the times series of shocks under the two scenarios are very similar. Therefore, we conclude that the possibility of an MA(1) transitory income shock does not affect substantially the estimates of the income shocks which we use in the Euler equation estimation.

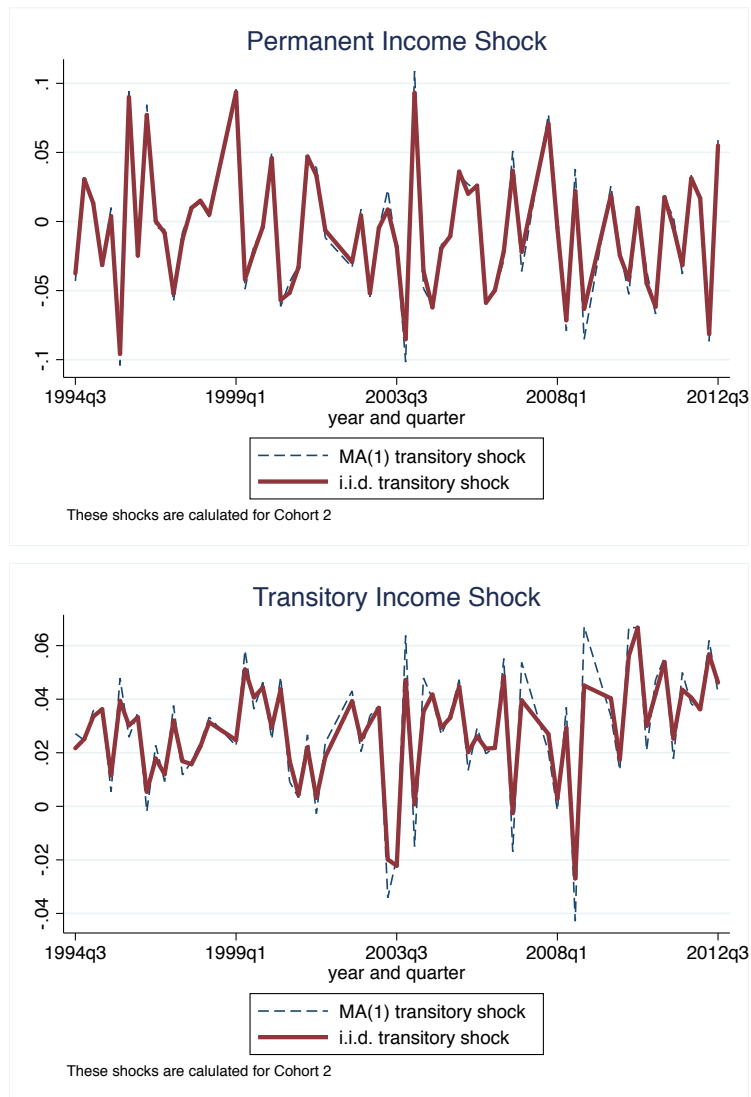
In our sample period 1994q1-2012q4, we estimate the standard deviation of the permanent and transitory shock to be 0.046 and 0.019 respectively. These standard deviations are lower than other estimates in the literature. It should be stressed, however, that others estimate income shock variances at the household (Blundell, Pistaferri, and Preston (2008)) or individual level (Meghir and Pistaferri (2004)), while our estimates are at the cohort level.¹⁹ Given that average income of a cohort may include some form of implicit or explicit insurance, we expect our estimates to be lower.²⁰

¹⁸Note that in Section A.4, we show how to calculate the shocks under special circumstances and assuming MA(1) transitory income shocks.

¹⁹Blundell, Pistaferri, and Preston (2008) estimate the standard deviation of permanent shocks to be between 0.07-0.17, while for the transitory shock it is 0.14-0.28.

²⁰Our sample period is also different, it does not include the 1980s, when Blundell, Pistaferri, and Preston (2008) document a dramatic increase in income inequalities (and a corresponding rise in the variance of income shocks). Yet, while Blundell, Pistaferri, and Preston (2008) also document a decline in inequalities at the beginning of our sample period, income inequalities are still widening during our

Figure 6: Shocks under Different Income Processes



5 Euler Equation Estimation

In this section, we first discuss the econometric issues relevant for the estimation of consumption Euler equation on cohort-level data, and then present our estimation results.

5.1 Econometric Issues

In order to estimate the expectation-error-adjusted Euler equation (13), we construct a synthetic panel dataset merging the Michigan Survey and the CEX Survey. Since these surveys interview different groups of households in each period, we cannot follow

sample period.

individual households behaviour over time. However, we can circumvent this problem following Deaton (1985) and Browning, Deaton, and Irish (1985), and constructing synthetic or pseudo panels. That is, rather than following individual households, we identify groups of households that have fixed membership and, using repeated cross sections (or rotating panels) drawn from the same population, we follow the cohort averages for the variable of interests. Given the structure of our surveys, we construct pseudo panels with a quarterly frequency.

The ‘true’ cohort mean of the variables of interest is unobserved. However, using our samples, we can construct estimates of these averages. The sample means will therefore be used as measures of the population means, albeit affected by ‘measurement error’.²¹ To minimise the impact of this type of error, in our estimation we only use cells containing more than 100 observations per quarter. The necessity to work with relatively large cells informs the definition of cohorts: by using wider year of birth intervals we have larger cells, albeit at the cost of including less homogeneous households.

We also impose an age limit on the cohorts and exclude observations for cohorts whose head on average is younger than 21 years or older than 60 years. Young households are more likely to be affected by binding liquidity constraints, so that their consideration might bias the estimation of the coefficients of the Euler equation.²² As for older households, one could argue that their preferences might be undergoing substantial changes, maybe related to health status. Therefore, the Euler equation might be mis-specified for young and old households.

There is an additional reason to exclude households headed by young and old individuals. The synthetic panel approach assumes that group membership is, in the population of reference, constant. Individuals with different socio-economic background might be starting a household at different ages. At the end of the life cycle, on the other hand, differential mortality between affluent and poor consumers might be changing systematically the composition of the cohorts. For these reasons, considering households headed by individuals that are neither too young nor too old makes it more likely to satisfy the assumption of constant group membership when constructing the pseudo panels.

As Chamberlain (1984) highlighted, the estimation of Euler equations needs long time series data since the orthogonality conditions hold in expectations. Using realisations to proxy expectations imply the use of the rational expectations hypothesis to derive orthogonality restrictions: the Euler equation errors include an expectational error that should be uncorrelated with past information. Rational expectations, however,

²¹As we know the size of the cells, we can construct estimates of the variance of measurement errors for each of the variables of interest.

²²Clearly, we are not sure that those above age 21 are not liquidity constrained. Therefore we tried different cutoffs by age. When we use a sample of households between age 30 and 60 for example, who are assumably less likely to be liquidity constrained, our results stay the same.

are correct on average over time, not across individuals, which explains the need for a long time period.

When we estimate the expectational-error-adjusted Euler equation Chamberlain's conditions do not apply, and it is enough to have large cross-sectional dimension to get a consistent estimate of the Euler equation. This is because we explicitly account for the expectational errors. Since we both have a long time-series and cross-sectional dimension, we do not need to worry about the consistency of our estimates (even though we would get consistent estimate even without a long time-series dimension).

We estimate equation (13) for all cohorts simultaneously using instrumental variable techniques. Using the synthetic panel approach for estimation raises some important issues to take into account before the estimation. As mentioned above, we do not observe population means but only somewhat noisy estimates of them. This is equivalent to having measurement error in the level of a variable. The fact that (log) consumption enters in first differences in the log linearized Euler equation creates an MA(1) structure for the residuals in equation (13). Consequently, we cannot use one-period lagged variables as instruments. However instruments lagged two or more periods gives consistent estimates. This is not the case for our macro variables, like interest rate or inflation, which can be used in the one-period lagged form. The instruments we use in our favourite specification are the different lags of consumption growth, nominal interest rates, inflation rate and household characteristics. Household characteristics are the number of family members, number of family members who are younger than 2 and dummy for single households. We also try to use different lags of the expectational errors of three macro variables as instruments: interest rate, inflation and unemployment rate.

Because of the presence of MA(1) residuals for each cohort and because we estimate equation (13) for several cohorts simultaneously, the error structure of the Euler equation is quite complicated. This has to be taken into account in order to construct an efficient estimator. Therefore, residuals for a given cohort are assumed to have an MA(1) structure, while between cohorts we only allow residuals to have contemporaneous correlation.

5.2 Results

In Section 2, we discussed how to incorporate data on subjective expectations within the estimation of an Euler equation. As clear from equation (13), if one observes $\zeta_{h,t}^{exp}$ and $\varepsilon_{h,t}^{exp}$, one could add them to the equation to be estimated and, by doing so, improve the efficiency of the estimates and, at the same time, obtain estimates of ϕ and θ .

There is another way, however, in which subjective expectations data can be used. The orthogonality conditions derived from the Euler equation imply that any variable

in the consumer’s information set at time t is a valid instrument. Such an instrument would be a useful instrument if it predicts the variables to be instrumented, in our case consumption growth and interest rates. The subjective expectations data, therefore, appropriately lagged can also be used as instrument and, as such, could also improve the efficiency of the estimates.

In Table 3, we report estimates of the Euler equation parameters with and without the subjective expectations data. Standard errors are in brackets and they are robust to the presence of the MA(1)-structured residuals. In the first column of the Table, we report estimates of the elasticity of intertemporal substitution (EIS) obtained from the CEX synthetic panel without using subjective expectation data. The EIS is estimated at 0.54, which is not substantially different from other estimates of the EIS in the literature (see for example Attanasio and Weber (1993), Blundell, Browning, and Meghir (1994)). We also report the estimates of the coefficients on taste shifters (family size and the number of children less than 2). Our results confirm earlier estimates, both family size and number of children are significant, suggesting that changing family needs impact consumption growth. The coefficients on the demographic variables are sensible: a growing family raises consumption, but younger children are less costly (see also Attanasio and Weber (1995), Browning and Ejrnaes (2009)).

Column 2 presents estimates of the same parameters obtained using the appropriately lagged subjective expectations data as additional instruments. The point estimates of the EIS and of the other parameters do not change much. The EIS increases slightly from 0.54 to 0.60.

In column 3 of the table, we add the estimates of expectational errors to the Euler equation. In particular, as specified in equation (13), we add innovations to the interest rate and to transitory and permanent components of income. Whilst the coefficient on the interest rate innovation is small and insignificantly different from zero, the coefficient on the permanent innovation to income is equal to 0.29, reflecting the extent to which these innovation to permanent income are reflected into consumption growth. The coefficient on the temporary innovations to income, instead, is small and not statistically different from zero.

In column 4, we remove the expectational error on the interest rate, which does not have a significant impact on consumption.²³ The other coefficients do not change much relative to column 3. In the last column we change the cohort definition: instead of having 5-year brackets for the birth of the household head, we use 10-year brackets.

²³Similar results were found by a few papers, that analyze the relationship between survey inflation expectations and spending (see Bachmann, Berg, and Sims (2015), Burke and Ozdagli (2013)). Their evidence suggests that high inflation expectations that decrease real interest rate expectations might not boost consumption.

$\Delta \log C$					
VARIABLES	(1)	(2)	(3)	(4)	(5)
r	0.546** (0.276)	0.602** (0.276)	0.739*** (0.268)	0.676*** (0.254)	0.609** (0.303)
Δ family size	0.471*** (0.073)	0.471*** (0.073)	0.321*** (0.070)	0.322*** (0.069)	0.453*** (0.084)
Δ #(children < 2)	-0.582*** (0.137)	-0.582*** (0.136)	-0.522*** (0.127)	-0.513*** (0.126)	-0.554*** (0.142)
ζ^{exp}			0.293*** (0.042)	0.292*** (0.041)	0.159** (0.069)
ε^{exp}			0.120 (0.087)	0.130 (0.086)	0.036 (0.089)
$u^{\text{exp},r}$			-0.001 (0.001)		
EE instruments	<i>N</i>	<i>Y</i>	<i>Y</i>	<i>Y</i>	<i>Y</i>
Observations	377	377	377	377	220
Sargan	0.70	0.82	0.79	0.90	0.56
R-squared	0.51	0.51	0.56	0.56	0.56

Standard errors are in parentheses, which are corrected for the MA(1) structure of the error term. All specification include a constant and three seasonal dummies. EE instruments: lags of expectational errors used as instruments. ζ^{exp} is the permanent shock, ε^{exp} the transitory shock and $u^{\text{exp},r}$ the interest rate expectation error. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Euler Equations

As discussed above, having larger cells reduces the sampling error in estimating cohort means (at the cost of having cohorts that are less homogenous than with a narrower definition). This approach seems to decrease the precision of the estimates without changing considerably their point value. We also note that, when including direct estimates of expectational errors, the estimated EIS increases from 0.60 in column 2 to 0.67. Moreover, and consistently with our expectations data capturing at least part of the variability of the Euler equation errors, the precision of the estimates improves.

The coefficients on the two components of income innovations can be compared to the estimates obtained by Blundell, Pistaferri, and Preston (2008). As these authors, we find that the coefficient on transitory income innovations is not statistically different from zero, indicating that temporary income shocks are effectively insured. We estimate the coefficient on permanent innovations at 0.29 with a standard error of 0.04. Blundell, Pistaferri, and Preston (2008) report estimates between 0.2 and 0.6, depending on the definition of income they use. Our results, therefore are not too far from those in

Blundell, Pistaferri, and Preston (2008) that were obtained with a completely different approach. The evidence, therefore, is that, consistently with standard versions of the life cycle model, households seem to be able to smooth transitory shocks. Persistent shocks, however are reflected in consumption. However, the loading factor of these shocks is considerably below 1. A coefficient less than unity is consistent with the ‘excess smoothness’ of consumption some authors have identified and with access to more sophisticated asset markets.

As mentioned above, the parameters on the subjective expectations ζ^{exp} and ε^{exp} (ϕ and θ) can be interpreted as reflecting the extent to which permanent and transitory shocks to income are reflected into consumption. It is therefore instructive to examine whether these parameters change when we estimate the Euler equation on different education groups. As suggested in Blundell, Pistaferri, and Preston (2008), better educated

	Baseline	No College	College
ζ^{exp}	0.16** (0.06)	0.29*** (0.08)	0.16*** (0.06)
ε^{exp}	0.04 (0.11)	0.25** (0.12)	0.02 (0.10)

Standard errors are in parentheses, which are corrected for the MA(1) structure of the error term. All specification include a constant and three seasonal dummies. EE instruments: lags of expectational errors used as instruments. ζ^{exp} is the permanent shock, ε^{exp} the transitory shock and $u^{exp,r}$ the interest rate expectation error. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: Euler equation by education groups

individuals might have better insurance possibilities. In Table 4, we present the results of such an exercise, in which the Euler equation is estimated on households headed by a college graduate and households headed by somebody without a college degree separately. In the Table, we report only the two insurance parameters, that is the coefficients on permanent income shocks ζ^{exp} and transitory income shocks ε^{exp} . Consistently with the evidence in Blundell, Pistaferri, and Preston (2008), the point estimates of these coefficients indicate that households headed by better educated individuals have better insurance possibilities. It should be stressed, however, that the low precision of these estimates implies that they are not statistically different from each other.

An interesting feature of Table 4 is that the point estimates of the coefficient on the permanent shock ζ^{exp} is at the lower end of the interval of estimates reported by Blundell, Pistaferri, and Preston (2008) and smaller than their favourite estimates. Although this difference is unlikely to be statistically significant, an interesting question is why would one get smaller coefficients on the innovations identified from the subjec-

tive expectations than in the Blundell, Pistaferri, and Preston (2008) procedure. There are several plausible hypotheses. One is that our estimates of permanent shocks are affected by measurement error, possibly induced by noise in the way the expectations questions are asked, and that induces an attenuation bias which reduces the size of the relevant coefficients. Second, it is possible that there are two types of news, idiosyncratic and cohort level news and that, for some reason, cohort level shocks (which is what we measure) are better insured than individual level shocks. Finally, it is possible that individuals cannot distinguish between aggregate and idiosyncratic components of income, as in Pischke (1995). If the aggregate component is persistent and the idiosyncratic is temporary, the ‘innovations’ to the individual income process will be less persistent than the aggregate process; individuals will interpret a permanent shock as partly temporary and, therefore, will react less to it.

Given that consumption seems to react on income innovations slowly, see Malloy, Moskowitz, and Vissing-Jorgensen (2009) for example, an interesting avenue for future research would be to focus on consumption responses over a longer time horizon. Instead of using the Euler equation between two consecutive periods, one can formulate an optimality condition between periods further away from each other, as in Parker and Julliard (2005) for example, and then estimate the corresponding insurance parameters. The lack of more detailed information on income expectations in the Michigan Survey unfortunately prevents us from following this avenue in the current paper.²⁴

6 Simulation

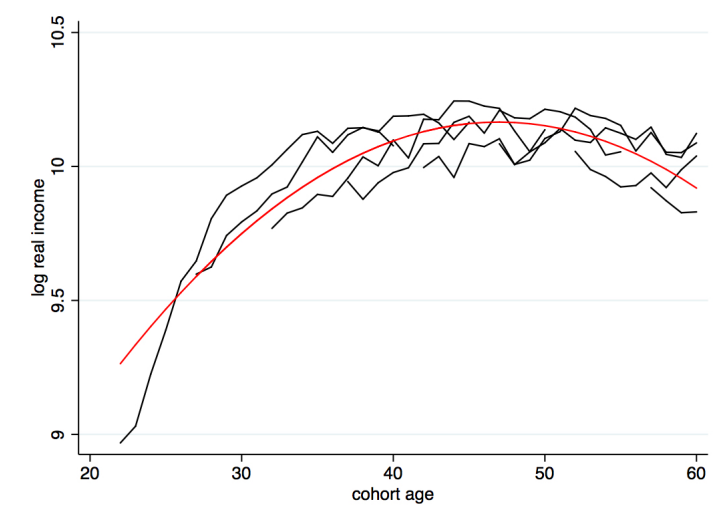
To get a sense of whether a basic life-cycle model predicts similar insurance possibilities as we observed in the data, we simulate an artificial panel of household income and consumption. Using this artificial panel, we then estimate the model counterpart of our Euler equation (13), to calculate the effect of transitory and permanent income shocks on consumption growth. By comparing coefficients, we are able to tell whether borrowing and saving through a risk-free asset over the life cycle can generate similar self-insurance

²⁴The Euler equation we are currently estimating is based on the optimality condition $u'(c_t) = E_t [\beta R_{t+1} u'(c_{t+1})]$. In case we want to apply the formulation by Parker and Julliard (2005), we need to estimate $u'(c_t) = E_t [\beta R_{t+S} u'(c_{t+S})]$, where S is the horizon over which the consumption response is studied. Estimating this version of the Euler equation including the income shocks as we do, would practically mean that we need more information on income expectations, which we don't have in the Michigan Survey. Households are only asked about their income expectations for next year, but not for later ahead. Having expectations on income for both one year and two years ahead, for example, would let us identify the income shocks in the first and second year, hence we would be able to estimate an Euler equation on the two-year horizon. Without this additional information on income expectations, unfortunately, we are unable to test whether our approach is robust to a slow reaction of consumption to income shocks.

of permanent and transitory income shocks as observed in the data.

The households problem is characterised by Equations (1)-(5). With CRRA preferences, households have an incentive to smooth consumption. In the absence of perfect insurance markets households undertake precautionary saving.

Figure 7: Log Income



Predictable Income Component. We use the household income data that is available in the CEX in order to estimate the deterministic time-varying income component (yearly) of labor income, which we use in our simulation. In Figure 7, we plot log disposable income for different cohorts against age (black lines). Continuous lines for cohorts overlap because we defined cohorts in five year intervals. Income shows the usual hump-shaped profile, peaking before retirement (see for example Attanasio et al. (1999))

	log y	Std. dev.
Age	0.1422***	(0.0042)
Age ²	-0.0015***	(0.0000)
Constant	6.8492***	(0.0885)
Observations	122	
R-squared	0.764	

Standard errors are in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: Income Process

We approximate the deterministic, time-variant income component ($B_{h,t}$) by a second-order polynomial in age. Focusing on cohort level observations, the parameters for the

labor income process is approximated by the following regression

$$\ln(y_t)^c = \beta_0 + \beta_1 age_{c,t} + \beta_2 age_{c,t}^2 + u_{c,t}^y \quad (19)$$

where superscripts and subscripts c stand for cohort averages. The age of the cohort $age_{c,t}$ in a given period is calculated by taking average age over those household heads who belong to the same cohort. Our regression results are presented in Table 5. Figure 7 plots the predicted average log income profile (red line), which gives a good approximation to cohort incomes and shows a similar hump-shaped profile.

Real Interest Rate. The real interest rate, which is the only aggregate uncertainty in the model is assumed to follow an AR(1) process.

$$r_t = c + \rho_r r_{t-1} + \xi_t \quad \xi_t \sim N(0, \sigma_\xi^2) \quad (20)$$

The interest rate process is estimated on U.S. 3-month Treasury Bill data²⁵ between 1984 and 2012, which we adjust for inflation. The estimation results are reported in Table 6. The persistence parameter of the real interest rate, ρ_r , is estimated to be 0.72, while the standard deviation of the interest rate shock, σ_ξ , is 0.014.

	r	
c	-0.002 (0.003)	0 <i>constrained</i>
$r(-1)$	0.688*** (0.132)	0.725*** (0.108)
R-squared	0.51	0.63
σ_ξ		0.014

Standard errors are in parenthesis. ***

$p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Real Interest Rate

All the parameters used for simulations are listed in Table 8 in Appendix A.5. For the artificial panel, we simulate the behaviour of 10,000 households over random realisations of the idiosyncratic permanent and temporary labor income shocks. After simulating the optimal life cycle paths for consumption, we are able to replicate the same regression as the one that is used on actual data in Table 3. The only difference between the

²⁵Taken from the Federal Reserve Bank of St.Louis (Fred)

Euler equation estimated on actual and simulated data is that the latter do not include demographic variation, as our simple life cycle model does not take into account changes in family composition. Therefore the regression we run on the simulated data is a simplified version of equation (13) given by:

$$\Delta \log C_{t+1}^m = \alpha^m + \beta_1^m \log(1 + r_t^m) + \beta_2^m \zeta_t^m + \beta_3^m \varepsilon_t^m + v_{t+1}^m \quad (21)$$

where similarly to our previous notations, C^m is the simulated level of consumption, r^m is the real interest rate, while ζ^m and ε^m are the permanent and transitory income shocks in the model, respectively.

We exclude retirement period of households from the regression as there is no uncertainty around income after age 65. Instruments are the first, second and third lags of the real interest rate, and the expectational error on the real interest rate. We end up using 420,000 observations for 10,000 households. Table 7 presents the results of the regression on the simulated data.

VARIABLES	$\Delta \log C^m$
r^m	0.630*** (0.011)
ζ^m	0.933*** (0.000)
ε^m	0.255*** (0.000)
Observations	420,000

Standard errors are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 7: Regression from Simulation

Our life-cycle model delivers qualitatively similar results to our estimates on U.S. data. Households are able to smooth transitory shocks, while permanent income innovations are reflected in consumption (compare Table 3 and Table 7). However, the loading factor of these shocks are very different to what we found in the data.

Households increase consumption by 9% after a 10% positive permanent income shock in our simulated life-cycle model, while, as our estimates show, they are able to insure themselves against a big fraction of it in reality.²⁶Therefore, a permanent income

²⁶In theoretical models, similarly high consumption responses to permanent shocks were found by other authors. Carroll (2001) found that in steady state consumption responds between 0.85 and 0.95 in a simulated buffer-stock model. Blundell, Low, and Preston (2004) found an estimate of 0.8 with a simulated model with CRRA preferences and a similar income process as ours. More recently, Kaplan and Violante (2010) found a estimates of 0.77 in a calibrated life-cycle version of a standard incomplete

model with self-insurance through a simple asset structure provide too little insurance possibilities compared to the data. This “excess smoothness” of consumption in the data was identified by several authors, starting with Campbell and Deaton (1989). In the 80’s and 90’s this was interpreted as a failure of the permanent income life-cycle model. Recently several authors provided evidence that individuals have access to several other markets than savings and borrowing to insure themselves against shocks.²⁷ The role of family networks was suggested in Attanasio and Rios-Rull (2000). Browning and Crossley (2009) suggest that in the short-run households can insure their non-durable consumption against income fluctuations with cutting back on total expenditures on durables. Other authors emphasize the insurance role of government tax programs (Kimball and Mankiw (1989), Auerbach and Feenberg (2000)) and different government public policy programs, like the unemployment insurance (Engen and Gruber (2001)) and food stamps (Blundell and Pistaferri (2003)). Attanasio and Pavoni (2011) show that in a more complex model than the Bewley economy, in which there are informational problems, individuals can enter insurance contracts that provide better insurance possibilities than self-insurance.²⁸

About 70 % of the transitory shocks are insured away by consumers on average in the theoretical model. In contrast, our empirical analysis on U.S. data shows that transitory shocks have no effect on consumption. Apart from the simple asset structure in the model, this discrepancy can be explained by the finite horizon setup: close to the “end of life” transitory shocks are like permanent shocks, they become a bigger component of lifetime income and become more difficult to insure against. In an infinite horizon setup, as in the Bewley model, this end of life effect is not present and transitory shocks are insured away.

In Table 7 we also present our estimate for the elasticity of intertemporal substitution parameter (the coefficient on the real interest rate). The EIS is estimated to be around 0.63, which corresponds to a risk aversion parameter of around 1.59. This result is not surprising, given that we have calibrated the risk aversion parameter in the model to be 1.66.

markets model with borrowing constraints. For other numerical simulations with precautionary savings, see Carroll (2001).

²⁷For an alternative interpretation see Primiceri and van Rens (2009).

²⁸In Attanasio and Pavoni (2011) the extra smoothness of consumption depends on the severity of the information problems. Their economy is characterized by hidden assets, (for example individuals have a hidden access to a credit market), and moral hazard.

7 Concluding Remarks

In this paper, we break new ground by showing how data on subjective expectations can be used to estimate a structural model. In particular, we use data on subjective income expectations within the estimation of an Euler equation for consumption. Expectations obviously play an important role in dynamic models of consumer choice. Under the assumption of rational expectations, if a long panel (or pseudo panel) of observations on individual consumption is available, one can estimate structural parameters even in the absence of data on subjective income expectations. However, data on subjective expectations can be used either to improve these estimates, to relax some of the assumptions made, or to reduce the data requirements to obtain them. In particular, as we discussed earlier, with data on subjective expectations, one does not necessarily need a long time period to get consistent estimates of the structural parameters of interest.

We have shown that the use of subjective expectations, while not changing substantially the point estimates of the preference parameters of the Euler equation, delivers more precise estimates. Besides the estimation of structural parameters, within the context of the life cycle consumption model, data on subjective expectations are also useful for different reasons. One can study, for instance, the extent to which new information on income (or income shocks) is translated (or not) into changes in consumption. Under commonly used specifications of the income process that decompose it into a permanent and transitory component, following Pistaferri (2001) it is possible to use subjective expectations data to identify permanent and transitory shocks. These shocks can then be added to an Euler equation and in order to estimate the extent to which they are translated into consumption changes. Such an exercise is informative about the degree of insurance against income fluctuations that is available to individuals.

We use data from the Michigan Survey on consumer confidence and the Consumer Expenditure Survey to perform such an exercise. We obtain estimates for the elasticity of intertemporal substitution that are not too different from those reported in the literature. Indeed, estimating the parameters of the Euler equation with or without the subjective expectations data - yields similar results, which is consistent with rational expectations.

We also show that transitory shocks are not transmitted to consumption. Interestingly, we find that only a relatively small fraction of permanent shocks is reflected in consumption. This result is consistent with the evidence on 'excess smoothness' reported by Campbell and Deaton (1989) and others, and discussed by Attanasio and Pavoni (2011), using a model of imperfect risk sharing. The parameters on the extent of insurance that we obtain are also not too different from those obtained, with a completely different approach by Blundell Pistaferri and Preston (2008).

Our exercise is not without limitations, some of which we have already mentioned.

Our measures are not a direct measure of subjective expectations, but we need to use the procedure described in Appendix A.2 to translate the answers to a specific survey question into quantitative measures of expectations. This approach, necessarily gives rise to measurement error in the expectations data we do not deal explicitly with, and that could also give rise to attenuation bias in the estimated coefficients on the shocks. Furthermore, as the expectation data do not have a longitudinal dimension, we are forced to use a synthetic cohort approach to estimate a dynamic relationship such as the Euler Equation. Moreover, our cohort averages necessarily lose any idiosyncratic variation in subjective expectations.

Further work is needed to go beyond our approach. The most important direction for extension is to address the fact that the lack of actual consumption data and of a longitudinal dimension in the Michigan data forces us to use synthetic cohort techniques. Because of that we cannot fully exploit the potential of the subjective expectations data. For instance, we cannot give up on a relatively long time period as we lose the cross sectional dimension. An important and substantive consequence of this limitation in the data is that, when studying the extent to which shocks are reflected in consumption, we lose all the individual shocks, as they are averaged out in the construction of the synthetic panel, and are forced to focus only on the cohort-level shocks. In future work, given the right individual level data we hope to explore possible extensions. For example strong rejections of rationality in return expectations (See Martin (2017), Greenwood and Shleifer (2014)), suggest extending our analysis to these would be an interesting avenue for future research, as consumption-savings choices of households crucially depend on the menu of financial and real assets available for them and their expected returns.

A Appendix

A.1 Relating Consumption Growth to Income Shocks

Following Blundell, Low, and Preston (2004) and Blundell, Pistaferri, and Preston (2008) we relate consumption to income shocks by approximating the budget constraint. Let us start by rewriting the period budget constraint given by equation (2) to the lifetime budget constraint assuming constant interest rate:

$$\sum_{j=0}^{T-t} \frac{C_{h,t+j}}{(1+r)^j} = \sum_{j=0}^{T-t} \frac{Y_{h,t+j}}{(1+r)^j} + A_{h,t} \quad (\text{A.1})$$

Now we log-linearize this budget constraint using Taylor approximation. Along the derivation of Blundell, Low, and Preston (2013), we define a function $F(\xi) = \ln \sum_j \exp \xi_j$ and take the first-order Taylor expansion around a point ξ^0 :

$$F(\xi) \simeq \ln \sum_{j=0}^N \exp \xi_j^0 + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{i=0}^N \exp \xi_i^0} (\xi_j - \xi_j^0) \quad (\text{A.2})$$

Next, let us take the expectations of (A.2) subject to information set \mathcal{I}

$$\mathbb{E}_{\mathcal{I}}[F(\xi)] \simeq \ln \sum_{j=0}^N \exp \xi_j^0 + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{i=0}^N \exp \xi_i^0} (\mathbb{E}_{\mathcal{I}} \xi_j - \xi_j^0) \quad (\text{A.3})$$

We use this relationship to approximate the lifetime budget constraint, (A.1). We start with the expected present value of consumption and set,

$$\begin{aligned} N &= T - t \\ \xi_j &= \ln C_{h,t+j} - j \ln(1+r) \\ \xi_j^0 &= \mathbb{E}_{t-1} \ln C_{h,t+j} - j \ln(1+r), \quad j = 0, \dots, N \end{aligned} \quad (\text{A.4})$$

hence we get

$$\begin{aligned} \mathbb{E}_{\mathcal{I}} \left[\ln \sum_{j=0}^{T-t} \frac{C_{h,t+j}}{(1+r)^j} \right] &\simeq \ln \sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} \ln C_{h,t+j} - j \ln(1+r)] \\ &\quad + \sum_{j=0}^{T-t} \theta_{t+j} [\mathbb{E}_{\mathcal{I}} \ln C_{h,t+j} - \mathbb{E}_{t-1} \ln C_{h,t+j}] \end{aligned} \quad (\text{A.5})$$

where

$$\theta_{t+j} = \frac{\exp \xi_j^0}{\sum_{i=0}^{T-t} \exp \xi_i^0} = \frac{\exp[\mathbb{E}_{t-1} \ln C_{h,t+j} - j \ln(1+r)]}{\sum_{i=0}^N \exp[\mathbb{E}_{t-1} \ln C_{h,t+i} - i \ln(1+r)]}$$

Now we turn to the expected present value of resources and set,

$$\begin{aligned} N &= T - t + 1 \\ \xi_j &= \ln Y_{h,t+j} - j \ln(1+r) \\ \xi_j^0 &= \mathbb{E}_{t-1} \ln Y_{h,t+j} - j \ln(1+r), \quad j = 0, \dots, N-1 \\ \xi_N &= \ln A_{h,t} \\ \xi_N^0 &= \mathbb{E}_{t-1} \ln A_{h,t} \end{aligned} \tag{A.6}$$

hence we get

$$\begin{aligned} \mathbb{E}_{\mathcal{I}} \left[\ln \sum_{j=0}^{T-t} \frac{Y_{h,t+j}}{(1+r)^j} + A_{h,t} \right] &\simeq \ln \left[\sum_{j=0}^{T-t} \exp[\mathbb{E}_{t-1} \ln Y_{h,t+j} - j \ln(1+r)] + \exp \mathbb{E}_{t-1} \ln A_{h,t} \right] \\ &\quad + \delta_t \sum_{j=0}^{T-t} \alpha_{t+j} [\mathbb{E}_{\mathcal{I}} \ln Y_{h,t+j} - \mathbb{E}_{t-1} \ln Y_{h,t+j}] \\ &\quad + (1 - \delta_t) [\mathbb{E}_{\mathcal{I}} \ln A_{h,t} - \mathbb{E}_{t-1} \ln A_{h,t}] \end{aligned} \tag{A.7}$$

where

$$\alpha_{t+j} = \frac{\exp[\mathbb{E}_{t-1} \ln Y_{h,t+j} - j \ln(1+r)]}{\sum_{i=0}^{T-t} \exp[\mathbb{E}_{t-1} \ln Y_{h,t+i} - i \ln(1+r)]} \tag{A.8}$$

and

$$\delta_t = 1 - \frac{\exp \xi_N^0}{\sum_{i=0}^N \exp \xi_j^0} \tag{A.9}$$

We use (A.7) and take the difference between expectations at period $t-1$ and period t . By doing so, we recall equations (3)-(4) and use that

$$\log Y_{h,t} = \pi' B_{h,t} + p_{h,t-1} + \zeta_{h,t} + \varepsilon_{h,t} \tag{A.10}$$

We can write the expectational errors as:

$$[\mathbb{E}_t - \mathbb{E}_{t-1}] \left[\ln \sum_{j=0}^{T-t} \frac{C_{h,t+j}}{(1+r)^j} \right] = [\mathbb{E}_t - \mathbb{E}_{t-1}] \left[\ln \sum_{j=0}^{T-t} \frac{Y_{h,t+j}}{(1+r)^j} + A_{h,t} \right] \quad (\text{A.11})$$

which simplifies to

$$u_{h,t}^{exp,c} \simeq \delta_t (\zeta_{h,t} + \alpha_t \varepsilon_{h,t}) \quad (\text{A.12})$$

So the innovation to consumption, u_t^c is related to the idiosyncratic income shocks, ζ_t and ε_t . (A.12) is equation (12) in the text, with $\phi_t = \delta_t$ and $\psi_t = \delta_t \alpha_t$.

A.2 Quantification of qualitative survey data

We will use the traditional probability approach of Carlson and Parkin (1975) to obtain quantitative measures for qualitative survey expectations. We can for example think of a question on how the interest rate changes from one year to the other, with possible answers of “increases”, “stays the same” and “decreases”. The Carlson and Parkin (1975) method assumes that a respondent answer decrease, if her mean expected value of the change in the certain economic variable (k_t) by the end of next period $t+1$, $E_t \Delta k_{i,t+1}$ is smaller than some value l_i . If we assume for example $l_i = 2$, the respondent’s answer is “decreases” if she expects the interest rate to decrease more than 2 percentage points. Here $E_t \Delta k_{i,t+1}$ is defined as the individual expectation on the variable for tomorrow net of the currently observed value of the same variable, $E_t k_{i,t+1} - k_t$. Similarly, the individual answers increase, if $E_t \Delta k_{i,t+1}$ is larger than some value h_i . Finally, the respondent answers remain the same if $E_t \Delta k_{i,t+1}$ is within the lower and upper bound, l_i and h_i . We will assume that lower and upper bounds are the same for all respondents in the survey (i.e. $l_i = l$ and $h_i = h$).

Let’s denote the percentage of respondents who expect the economic variable to increase/decrease by INC_t and DEC_t , then we can think of the survey as sampling from some aggregate distribution (Φ), consequently

$$\begin{aligned} INC_t &= P(E_t \Delta k_{t+1} \geq h) = 1 - P(E_t \Delta k_{t+1} < h) \\ DEC_t &= P(E_t \Delta k_{t+1} \leq l) \end{aligned}$$

We will assume that the distribution for the change of the expected economic variable

is normal. In this manner, we can write down the relationships in standardised forms

$$1 - INC_t = P\left(0 < \frac{h - E_t\Delta k_{t+1}}{sd_{t+1}}\right) = \Phi\left(\frac{h - E_t\Delta k_{t+1}}{sd_{t+1}}\right)$$

$$DEC_t = P\left(0 < \frac{l - E_t\Delta k_{t+1}}{sd_{t+1}}\right) = \Phi\left(\frac{l - E_t\Delta k_{t+1}}{sd_{t+1}}\right)$$

where $E_t\Delta k_{t+1}$ is the mean and sd_{t+1} is the standard deviation of the aggregate distribution of the economic variable expectations. The quantiles can be simply calculated from observed percentages of respondents as follows:

$$r_t = \Phi^{-1}(1 - INC_t) = \frac{h - E_t\Delta k_{t+1}}{sd_{t+1}}$$

$$f_t = \Phi^{-1}(DEC_t) = \frac{l - E_t\Delta k_{t+1}}{sd_{t+1}}$$

Using this two equations, we can solve for $E_t\Delta k_{t+1}$

$$E_t\Delta k_{t+1} = \frac{hf_t - lr_t}{f_t - r_t}$$

We follow Carlson and Parkin (1975) and assume that the upper and lower bounds of the so called indifference interval are symmetric, $c = -l = h$. In our calculation we chose $c = 1\%$ for all the variables, thus we assume households respond “decrease” when they expect the relevant variable to decrease by more than 1 percent.

$$E_t\Delta k_{t+1} = \frac{c(f_t + r_t)}{f_t - r_t} \tag{A.13}$$

Now in case we have long time-series observation for the given economic variable, we can even estimate c by assuming that the average value of past realisations and the average value of expectations are equal (unbiasedness of expectations):

$$\frac{1}{T} \sum_{t=1}^T E_t\Delta k_{t+1} = \frac{1}{T} \sum_{t=1}^T (k_t - k_{t-1})$$

which is

$$\sum_{t=1}^T \frac{c(f_t + r_t)}{f_t - r_t} = \sum_{t=1}^T (k_t - k_{t-1})$$

And the estimate for c is:

$$\hat{c} = \frac{\sum_{t=1}^T (k_t - k_{t-1})}{\sum_{t=1}^T \left(\frac{f_t + r_t}{f_t - r_t}\right)}$$

A.3 Identifying Permanent and Transitory Income Shocks

Households expected income can be calculated based on their reported expectations on their income growth (from the Michigan Survey) and their realized income (from the CEX). Also note that the Surveys does not constitute a panel, hence we need to rely on the so-called synthetic panel technique, once we merge the two datasets. Consequently, what we have available in period 1995q1 for example is a cohort's income expectation for period 1996q1 based on the information it has in period 1995q1:

$$\mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right]$$

where $y_{i,96q1} = \log Y_{i,96q1}$ denotes the log of the labor income for a given cohort i , and $\Omega_{i,95q1}$ is the information set in 1995q1. For our derivations the important assumption is that expectations of each cohort are formed rationally and that these cohorts do not update their expectations about any future income within a year. Under the assumed income process in equations (3)-(5) the following holds

$$\begin{aligned}\mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] &= \pi' B_{i,96q1} + p_{i,95q1} \\ &= \pi_0 + \pi_1 age_{i,96q1} + \pi_2 age_{i,96q1}^2 + p_{i,95q1}\end{aligned}$$

Now let us calculate the sum of the permanent and transitory shocks for period 1996q1 as follows

$$\begin{aligned}y_{i,96q1} - \mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] &= \pi' B_{i,96q1} + p_{i,95q1} + \zeta_{i,96q1} + \varepsilon_{i,96q1} - \pi' B_{i,96q1} - p_{i,95q1} \\ \zeta_{i,96q1} + \varepsilon_{i,96q1} &= y_{i,96q1} - \mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] + \pi' B_{i,96q1} - \pi' B_{i,96q1} \\ \zeta_{i,96q1} + \varepsilon_{i,96q1} &= y_{i,96q1} - \mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] \\ &\quad + \pi_1 (age_{i,96q1} - age_{i,95q1}) + \pi_2 (age_{i,96q1}^2 - age_{i,95q1}^2)\end{aligned}$$

In our next step, we use a linear function for age in order to let age in different quarters differ.²⁹ If we do so, then we can get the following expression for the sum of shocks in 1996q1:

$$\begin{aligned}\zeta_{i,96q1} + \varepsilon_{i,96q1} &= y_{i,96q1} - \mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] + 0.75\pi_1 + 0.75\pi_2 (age_{i,96q1} + age_{i,95q1}) \\ &= y_{i,96q1} - \mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] + 0.75\pi_1 + 0.75\pi_2 (2age_{i,96q1} + 0.75) \\ &= y_{i,96q1} - \mathbb{E}\left[y_{i,96q1}|\Omega_{i,95q1}\right] + 0.75(\pi_1 + 0.75\pi_2) + 1.5\pi_2 age_{i,96q1}\end{aligned}$$

²⁹If someone reports to be aged 55 in 1996q1, we assume that he is 55.25 in 1996q2 etc.

where each element on the right side is observable in the surveys, while π_1 and π_2 can be estimated.

The permanent part of the income shocks can also be calculated from survey data for period 1996q1 for example, as follows:

$$\mathbb{E}\left[y_{i,97q1}|\Omega_{i,96q1}\right] - \mathbb{E}\left[y_{i,96q4}|\Omega_{i,95q4}\right] = \pi' B_{i,97q1} + p_{i,96q1} - \pi' B_{i,96q4} - p_{i,95q4}$$

therefore

$$\zeta_{i,96q1} = \mathbb{E}\left[y_{i,97q1}|\Omega_{i,96q1}\right] - \mathbb{E}\left[y_{i,96q4}|\Omega_{i,95q4}\right] - (\pi' B_{i,97q1} - \pi' B_{i,96q4}).$$

Now using the definition of function B and the smooth approximation for age, we can derive a simple relationship for the permanent income shock in 1996q1:

$$\begin{aligned} \zeta_{i,96q1} &= \mathbb{E}\left[y_{i,97q1}|\Omega_{i,96q1}\right] - \mathbb{E}\left[y_{i,96q4}|\Omega_{i,95q4}\right] \\ &\quad - \left[\pi_1(\text{age}_{i,97q1} - \text{age}_{i,96q4}) + \pi_2(\text{age}_{i,97q1}^2 - \text{age}_{i,96q4}^2)\right] \\ &= \mathbb{E}\left[y_{i,97q1}|\Omega_{i,96q1}\right] - \mathbb{E}\left[y_{i,96q4}|\Omega_{i,95q4}\right] - \left[0.25\pi_1 + 0.25\pi_2(\text{age}_{i,97q1} + \text{age}_{i,96q4})\right] \\ &= \mathbb{E}\left[y_{i,97q1}|\Omega_{i,96q1}\right] - \mathbb{E}\left[y_{i,96q4}|\Omega_{i,95q4}\right] - \left[0.25\pi_1 + 0.25\pi_2(2\text{age}_{i,96q1} + 1.75)\right] \\ &= \mathbb{E}\left[y_{i,97q1}|\Omega_{i,96q1}\right] - \mathbb{E}\left[y_{i,96q4}|\Omega_{i,95q4}\right] - 0.25(\pi_1 + 1.75\pi_2) - 0.5\pi_2\text{age}_{i,96q1} \end{aligned}$$

A.4 MA(1) Transitory Income Shocks

Let's alter the income process we use in our model, described by equations (3)-(5), in order to incorporate the possibility of a more flexible transitory income shock component. More specifically, let's assume that the transitory income shock follows an MA(1) process:

$$\begin{aligned} \log Y_{h,t} &= \pi' B_{h,t} + p_{h,t} + \varepsilon_{h,t} + \rho\varepsilon_{h,t-1}, \\ p_{h,t} &= p_{h,t-1} + \zeta_{h,t} \end{aligned}$$

As we showed in equations (14)-(16), the crucial part of our identification was to calculate permanent income shocks as the change in the subjective expectations of income (once predictable life-cycle effects were removed). In what follows, we show that by assuming an MA(1) transitory income component, instead of an i.i.d., our identification fails.

Using the modified income process above, we can write the one-period ahead expected

income as follows:

$$\begin{aligned}\mathbb{E}\left[\log Y_{h,t}|\Omega_{h,t-1}\right] &= \pi' B_{h,t} + p_{h,t-1} + \rho\varepsilon_{h,t-1} \\ \mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] &= \pi' B_{h,t+1} + p_{h,t} + \rho\varepsilon_{h,t}\end{aligned}\tag{A.14}$$

Subtracting one equation in expression (A.14) from the other we obtain:

$$\begin{aligned}\mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] - \mathbb{E}\left[\log Y_{h,t}|\Omega_{h,t-1}\right] &= \pi'(B_{h,t+1} - B_{h,t}) + p_{h,t} - p_{h,t-1} + \rho\varepsilon_{h,t} - \rho\varepsilon_{h,t-1} \\ &= \pi'(B_{h,t+1} - B_{h,t}) + \zeta_{h,t} + \rho(\varepsilon_{h,t} - \varepsilon_{h,t-1})\end{aligned}$$

This expression shows that the permanent income shock, ζ , is only identified by the change in subjective income expectations if the persistence parameter of the MA(1) transitory income component, ρ , is zero.

Now let's subtract income expectation from income in period t :

$$\log Y_{h,t} - \mathbb{E}\left[\log Y_{h,t+1}|\Omega_{h,t}\right] = \pi'(B_{h,t} - B_{h,t+1}) + \varepsilon_{h,t} + \rho(\varepsilon_{h,t-1} - \varepsilon_{h,t})$$

This expression shows that the transitory income shock, ε , is only identified by the difference between income realisations and income expectations if the persistence parameter of the MA(1) transitory income component, ρ , is zero. In case ρ is different from zero and estimable, we still need information on the initial transitory income shock in order to have identification.

A.5 Model Calibration

Parameter		Value	Source
T_start	<i>Age of entering model</i>	20	
T_ret	<i>Age of retirement</i>	65	
β	<i>Discount factor</i>	0.95	
γ	<i>Risk aversion parameter</i>	1.666	Own calculations, CEX
π_0	<i>Age-specific income, constant</i>	6.849	Own calculations, CEX
π_1	<i>Age-specific income, linear trend</i>	0.142	Own calculations, CEX
π_2	<i>Age-specific income, quadratic trend</i>	-0.0015	Own calculations, CEX
ρ_r	<i>Persistence parameter for interest rate</i>	0.725	Own calculations, Fred
σ_ζ	<i>Std.dev.permanent income shock</i>	0.14	Blundell, Pistaferri, and Preston (2008)
σ_ε	<i>Std.dev.transitory income shock</i>	0.20	Blundell, Pistaferri, and Preston (2008)
σ_ξ	<i>Std.dev.interest rate shock</i>	0.014	Own calculations, Fred

Table 8: Parameters for the Benchmark Model

References

- Abdowd, John M. and David Card. 1989. "On The Covariance Structure of Earnings and Hours Changes." *Econometrica* 57 (2):411–445.
- Altonji, Joseph G. and Aloysius Siow. 1987. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *The Quarterly Journal of Economics* 102 (2):pp. 293–328.
- Altug, Sumru and Robert A Miller. 1990. "Household Choices in Equilibrium." *Econometrica* 58 (3):543–70.
- Ang, Andrew, Geert Bekaert, and Min Wei. 2007. "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics* 54 (4):1163–1212.
- Angelini, Viola, Agar Brugiavini, and Guglielmo Weber. 2009. "Ageing and Unused Capacity in Europe: Is There an Early Retirement Trap?" *Economic Policy* 59:453–508.
- Attanasio, Orazio and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104 (6):pp. 1227–1262.
- Attanasio, Orazio and Jose-Victor Rios-Rull. 2000. "Consumption smoothing in island economies: Can public insurance reduce welfare?" *European Economic Review* 44 (7):1225–1258.
- Attanasio, Orazio P. 1999. "Consumption." in Taylor, J.B. and M Woodford (eds) *Handbook of Macroeconomics* .
- Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Life-Time Consumption." *Journal of Business and Economics Statistics* 17(1):22–35.
- Attanasio, Orazio P., James Banks, and Sarah Tanner. 2002. "Asset Holding and Consumption Volatility." *Journal of Political Economy* 110 (4):771–792.
- Attanasio, Orazio P. and Martin Browning. 1995. "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review* 85(5):1118–1137.
- Attanasio, Orazio P. and Hamish Low. 2004. "Estimating Euler Equations." *Review of Economic Dynamics* 7 (2):405–435.

- Attanasio, Orazio P. and Nicola Pavoni. 2011. “Risk Sharing in Private Information Models With Asset Accumulation: Explaining the Excess Smoothness of Consumption.” *Econometrica* 79 (4):1027–1068.
- Attanasio, Orazio P. and Susann Rohwedder. 2003. “Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom.” *American Economic Review* 93(5):1499–1521.
- Attanasio, Orazio P. and Guglielmo Weber. 1989. “Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption.” *The Economic Journal* 99:59–73.
- . 1993. “Consumption Growth, the Interest Rate and Aggregation.” *Review of Economic Studies* 60(3):631–649.
- . 1995. “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey.” *Journal of Political Economy* 103(6):1121–1157.
- . 2010. “Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy.” *Journal of Economic Literature* 48:693–751.
- Auerbach, Alan and Daniel Feenberg. 2000. “The Significance of Federal Taxes as Automatic Stabilizers.” *Journal of Economic Perspectives* 14 (3):37–56.
- Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims. 2015. “Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence.” *American Economic Journal: Economic Policy* 7 (1):1–35.
- Banks, James, Richard Blundell, and Agar Brugiavini. 2001. “Risk Pooling, Precautionary Saving and Consumption Growth.” *Review of Economic Studies* 68 (4):757–779.
- Banks, James, Richard Blundell, and Sarah Tanner. 1998. “Is There a Retirement-Savings Puzzle?” *American Economic Review* 88 (4):769–788.
- Battistin, 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.
- 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, Hamish Low, and Ian Preston. 2004. "Income risk and consumption inequality: a simulation study." IFS Working Papers W04/26, Institute for Fiscal Studies.
- . 2013. "Decomposing Changes in Income Risk Using Consumption Data." *Quantitative Economics* 4:1–37.
- Blundell, Richard and Luigi Pistaferri. 2003. "Income Volatility and Household Consumption: The Impact of Food Assistance Programs." *The Journal of Human Resources* 38:1032–1050.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5):1887–1921.
- Browning, Martin and Thomas Crossley. 2009. "Shocks, Stocks, and Socks: Smoothing Consumption Over a Temporary Income Loss." *Journal of the European Economic Association* 7 (6):1169–1192.
- Browning, Martin, Angus Deaton, and Margaret Irish. 1985. "A Profitable Approach to Labour Supply and Commodity Demands over the Life-Cycle." *Econometrica* 53 (3):503–544.
- Browning, Martin and Mette Ejrnæs. 2009. "Consumption and Children." *The Review of Economics and Statistics* 91 (1):93–111.
- Burke, Mary A. and Ali K. Ozdagli. 2013. "Household inflation expectations and consumer spending: evidence from panel data." Working Papers 13-25, Federal Reserve Bank of Boston.
- Campbell, John Y and Angus Deaton. 1989. "Why Is Consumption So Smooth?" *Review of Economic Studies* 56 (3):357–73.
- Carlson, John A. and Michael Parkin. 1975. "Inflation Expectations." *Economica* 42 (166):123–138.
- Carroll, Christopher D. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 1992(2):61–156.
- . 1997a. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics* 112(1):1–55.
- . 1997b. "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)." NBER Working Papers 6298, National Bureau of Economic Research, Inc.

- . 2001. “Precautionary Saving and the Marginal Propensity to Consume out of Permanent Income.” NBER Working Papers 8233, National Bureau of Economic Research, Inc.
- Carroll, Christopher D. and Andrew A. Samwick. 1997. “The Nature of Precautionary Wealth.” *Journal of Monetary Economics* 40(1):41–71.
- Chamberlain, Gary. 1984. “Panel Data.” *Handbook of Econometrics* 2:1247–1328.
- Cochrane, John H. 1991. “A Simple Test of Consumption Insurance.” *Journal of Political Economy* 99 (5):957–76.
- Coibion, Olivier and Yuriy Gorodnichenko. 2012. “What Can Survey Forecasts Tell Us about Information Rigidities?” *Journal of Political Economy* 120 (1):116–159.
- Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa. 2015. “Subjective Intertemporal Substitution.” *Federal Reserve Bank of New York Staff Reports* 734.
- Danninger, Stephan. 2016. “What’s Up with U.S. Wage Growth and Job Mobility?” Imf working papers, International Monetary Fund.
- Deaton, Angus. 1985. “Panel Data from the Time Series of Cross-Sections.” *Journal of Econometrics* 30:109–126.
- . 1991. “Saving and Liquidity Constraints.” *Econometrica* 59(5):1221–1248.
- Dotsey, Michael and Jed L. DeVaro. 1995. “Was the disinflation of the early 1980’s anticipated?” *Economic Quarterly* (Fall):41–60.
- Dreze, Jacques H. and Franco Modigliani. 1972. “Consumption decisions under uncertainty.” *Journal of Economic Theory* 5 (3):308–335.
- Edelberg, Wendy. 2004. “Risk-based Pricing of Interest Rates in Household Loan Markets.” 2004 Meeting Papers 442, Society for Economic Dynamics.
- Engen, Eric M. and Jonathan Gruber. 2001. “Unemployment insurance and precautionary saving.” *Journal of Monetary Economics* 47 (3):545–579.
- Friedman, Milton. 1957. *A Theory of the Consumption Function*. National Bureau of Economic Research, Inc.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. “Consumption over the Life Cycle.” *Econometrica* 70(1):47–89.

- Greenwood, Robin and Andrei Shleifer. 2014. "Expectations of Returns and Expected Returns." *The Review of Financial Studies* 27 (3):714–746.
- Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97 (3):687–712.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 96:971–987.
- Hansen, Lars Peter and Kenneth J. Singleton. 1983. "Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy* 91 (2):249–265.
- 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.
- . 1987. "Tests For Liquidity Constraints: A Critical Survey and Some New Observatiions." *T.F. Bewley (ed.), Advances in Econometrics: Fifth World Congress II*:91–120.
- 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 (4):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 and N Gregory Mankiw. 1989. "Precautionary Saving and the Timing of Taxes." *Journal of Political Economy* 97 (4):863–79.
- Krueger, Dirk and Fabrizio Perri. 2005. "Understanding Consumption Smoothing: Evidence from the U.S. Consumer Expenditure Data." *Journal of the European Economic Association* 3 (2-3):340–349.
- Long, J. Bradford De. 1997. *America's Peacetime Inflation: The 1970s*. in *Reducing Inflation: Motivation and Strategy*. The University of Chicago Press, Chicago.
- 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.

- Malloy, Christopher J., Tobias J. Moskowitz, and Annette Vissing-Jorgensen. 2009. “Long-Run Stockholder Consumption Risk and Asset Returns.” *The Journal of Finance* LXIV (6):2427–2479.
- Manski, Charles. 2004. “Measuring Expectations.” *Econometrica* 72 (5):1329–1376.
- Martin, Ian. 2017. “What Is the Expected Return on the Market.” *The Quarterly Journal of Economics* 132 (1):367–433.
- Meghir, Costas and Luigi Pistaferri. 2004. “Income Variance Dynamics and Heterogeneity.” *Econometrica* 72 (1):1–32.
- Moffitt, Robert. 1993. “Identification and estimation of dynamic models with a time-series of repeated cross-sections.” *Journal of Econometrics* 59:99–123.
- Moffitt, Robert and Peter Gottschalk. 2011. “Trends in the covariance structure of earnings in the U.S.: 1969–1987.” *The Journal of Economic Inequality* 9 (3):439–459.
- Nardi, Mariacristina De, Eric French, and David Benson. 2012. “Consumption and the Great Recession.” *Economic Perspectives* (Q I):1–16.
- Parker, Jonathan A. and Christian Julliard. 2005. “Consumption Risk and the Cross Section of Expected Returns.” *Journal of Political Economy* 113 (1):185–222.
- Paxson, Christina H. 1992. “Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand.” *American Economic Review* 82 (1):15–33.
- Pischke, Jorn-Steffen. 1995. “Individual Income, Incomplete Information, and Aggregate Consumption.” *Econometrica* 63 (4):805–40.
- Pistaferri, Luigi. 2001. “Superior Information, Income Shocks, And The Permanent Income Hypothesis.” *The Review of Economics and Statistics* 83 (3):465–476.
- Primiceri, Giorgio E. and Thijs van Rens. 2009. “Heterogeneous life-cycle profiles, income risk and consumption inequality.” *Journal of Monetary Economics* 56 (1):20–39.
- Souleles, Nicholas. 2004. “Expectations, heterogenous Forecast Errors and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys.” *Journal of Money, Credit and Banking* 36 (1):39–72.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. “Consumption and risk sharing over the life cycle.” *Journal of Monetary Economics* 51 (3):609–633.

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97:305–346.