

Excess Smoothness of Consumption in an Estimated Life Cycle Model

Dmytro Hryshko*
University of Alberta

Abstract

In the literature, econometricians typically assume that household income is the sum of a random walk permanent component and a transitory component, with uncorrelated permanent and transitory shocks. Using U.S. data on household wealth, consumption, and income I estimate a life cycle model where households smooth permanent and transitory income shocks by means of self-insurance, and find that household consumption is excessively smooth. That is, in the data consumption responds to income shocks to a lesser extent than in the model. To reconcile the model with the data, I explore the possibility that households have more information about components of income, transitory and permanent, than econometricians. I find that income shocks are negatively correlated and the model fits the data better but consumption is still excessively smooth. The model replicates the patterns in the data well when household information about components of income and partial risk sharing against permanent income shocks are allowed for.

Keywords: Buffer stock model of savings; Method of simulated moments; Consumption dynamics; Life cycle; Income processes.

JEL Classifications: C15, C61, D91, E21.

*University of Alberta, Department of Economics, 8-14 HM Tory Building. Edmonton, Alberta, Canada, T6G 2H4. E-mail: dhryshko@ualberta.ca. Phone: 780-4922544. Fax: 780-4923300.

1 Introduction

Since Friedman (1957), household income is typically assumed to be well represented by the sum of a permanent random walk component and a short-lived transitory component, with no correlation between transitory and permanent income shocks.¹

Models of household consumption over the life cycle that allow for self-insurance and liquidity constraints predict that households insure against transitory shocks almost perfectly but achieve limited insurance of permanent shocks. Using simulations of a buffer stock model of savings Carroll (2001) finds, for a plausible set of model parameters, that households smooth between 5 to 20 percent of permanent shocks to income. However, Blundell, Pistaferri, and Preston (2008) and Attanasio and Pavoni (2007) recently showed, using U.S. and U.K. data respectively, that households achieve substantial insurance against permanent income shocks. Following the literature on consumption dynamics in macro data, household consumption is said to be “excessively smooth.”²

In this paper, using U.S. data I study the extent of consumption smoothness reflected in the empirically observed sensitivity of household consumption growth to income growth at one to four-year horizons.³ In an estimated life cycle model with uncorrelated permanent and transitory income shocks and self-insurance, I confirm that household consumption in the U.S. is excessively smooth, that is, the model predicts that households should be more sensitive to income shocks than what is found in the data. To reconcile the model with the data I allow for a seemingly innocuous possibility—that, in an otherwise standard decomposition of idiosyncratic income into permanent and transitory components, permanent and transitory shocks are potentially correlated. Households may have better information about components of income, and therefore about the stochastic processes that govern the dynamics of each component and make their consumption and savings choices utilizing that knowledge. I use a life cycle model of consumption to identify the parameters of the household idiosyncratic income process, the volatility of permanent and transitory shocks, and the correlation between them, along with the prefer-

¹Notable examples are Carroll and Samwick (1997) and Meghir and Pistaferri (2004). They split income changes into permanent and transitory parts, and, under the assumption of orthogonality between permanent and transitory shocks, estimate household or group-specific volatility of permanent and transitory shocks.

²If income is non-stationary and income growth exhibits positive serial correlation—as supported by aggregate data—the Permanent Income Hypothesis (PIH) predicts that consumption should change by an amount greater than the value of the current income shock. Consequently, consumption growth should be more volatile than income growth. Consumption growth in aggregate data, however, is much less volatile than income growth. Therefore consumption growth is said to be “excessively smooth” relative to income growth. See, e.g., Deaton (1992).

³Using four-year growth rates allows me to explore the reaction of consumption growth to income growth over a longer horizon, when permanent shocks become relatively more important.

ence parameters. Using Friedman’s words (1957, p.23), “the precise line to be drawn between permanent and transitory components is best left to be determined by the data themselves, to be whatever seems to correspond to consumer behavior.”

Correct identification of permanent versus transitory shocks is important for the prediction of economic behavior and was shown to be important for understanding the “excess smoothness” puzzle in the aggregate data. For different reduced form models of aggregate income, Quah (1990) shows that there exists a decomposition of income into permanent and transitory components that helps solve the PIH “excess smoothness” puzzle in the aggregate data. This decomposition of income into its components, which can be reasonably assumed to be known to households, may or may not coincide with the decomposition done by econometricians using income data alone.

In this paper, I explore an idea similar to that in Quah (1990) in the context of the buffer stock model of savings. I first simulate life cycle buffer stock models that only differ in terms of decompositions of the same reduced form income process, and analyze the simulated economies at the household level. I find that models with more negatively correlated permanent and transitory shocks, but the same reduced form income dynamics, result in a significantly lower marginal propensity to consume (MPC) out of shocks to current income, and a lower MPC out of shocks to income cumulated over the four-year horizon. Thus, there may exist the decomposition of household income into permanent and transitory parts, consistent with household data, that explains the excess smoothness of household consumption. The intuition behind these results is the following. Households react to the newly arrived permanent and transitory innovations, that comprise a portion of the observable income growth. When permanent and transitory shocks are negatively correlated, the sum of innovations is smoother compared with income models that feature uncorrelated or positively correlated shocks: positive permanent shocks, in the case of a negative correlation, come together, on average, with negative transitory shocks. If the unpredictable part of the observable income growth is smoother, consumption is also smoother.

I further use the MPCs estimated from empirical micro data to identify parameters of the income process, including the correlation between permanent and transitory shocks. Importantly, this correlation cannot be identified from the univariate dynamics of integrated moving average processes. I estimate parameters of the income process by the Method of Simulated Moments (MSM). Using a life cycle buffer stock model, I simulate the MPCs, the variance and persistence of income, and wealth-to-income ratio over the life cycle, and match them to the same moments constructed from the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX) data. I find significantly negative contemporaneous correlation be-

tween transitory and permanent income shocks of about -0.60 ; and a low degree of patience and high degree of risk aversion as in Cagetti (2003).⁴ While the model with negatively correlated permanent and transitory income shocks fits the reaction of consumption to current income shocks, it still falls short of explaining the MPC out of shocks cumulated over longer horizons; that is, consumption is still excessively smooth in the data. Deaton (1992), in a summary of the literature on consumption volatility in aggregate data, defines excess smoothness as an insufficient responsiveness of consumption to the *current* income shock. The model with negatively correlated permanent shocks is, therefore, capable of explaining excess smoothness in household data as defined in Deaton (1992) but my results highlight that excess smoothness should be evaluated—in macro and household data—not only against the adjustment of consumption to current income shocks, but also to the shocks cumulated over longer horizons.

In reality, households may have access to a wide array of assets and risk-sharing mechanisms that allow for consumption smoothing over time and across states of nature. I further enrich the model with partial risk sharing of permanent income shocks, as in Attanasio and Pavoni (2007). The model with negatively correlated income shocks and partial risk sharing of permanent income shocks replicates well the patterns in household consumption, income and wealth data. I estimate substantial risk sharing of permanent income shocks: to be consistent with the data, the model requires smoothing of 52 percent of permanent income shocks before households make their savings decisions.

The results in this paper relate to Quah (1990), Pischke (1995), Ludvigson and Michaelides (2001), which show that a focus on households' information about the dynamics of income components may shed some new light on the fit of consumption theory to the data. The paper is also related to Gourinchas and Parker (2002) and Cagetti (2003) who estimate the preference parameters in a structural model using average consumption and wealth holdings over the life cycle, respectively. In this paper, I estimate both the preference and income process parameters utilizing information not only on the amount of wealth households choose to hold at different stages of the life cycle, but also on the amount of income risk they face, and the resulting sensitivity of consumption to income shocks.

In a recent paper, Blundell, Pistaferri, and Preston (2008) find substantial insurance of

⁴Friedman (1963), in an attempt to clarify the controversial points in his book on the consumption function, pointed out that the correlation between permanent and transitory shocks may be of any sign and, if present, should be allowed for in analysis of the consumption function. An example of a negative correlation between permanent and transitory income shocks can be found in Belzil and Bognanno (2008). Using earnings data for American executives in U.S. firms, they find that promotions (these events result in an increase of the base pay and, if unpredictable, can be thought of as positive permanent income shocks) come together with bonus cuts (negative transitory income shocks). They interpret the negative comovement between changes in the base pay and bonuses as a result of a compensation smoothing strategy adopted by firms.

household consumption against permanent income shocks in the U.S. data. For the whole sample, they find that about 64 percent of permanent shocks translate into consumption, while the rest are insured away. Kaplan and Violante (2010) calibrated a life cycle model with self-insurance in order to match the degree of insurance against permanent and transitory shocks estimated in Blundell, Pistaferri, and Preston (2008). They find that self-insurance provides less smoothing of permanent shocks relative to the amount of smoothing found in the data. While my paper is very closely related to those papers, there are some important differences as well. Differently from Blundell, Pistaferri, and Preston (2008), I match the mean of wealth-to-income ratio over the life cycle, besides the moments that describe the dynamics of household idiosyncratic income and the comovement of consumption and income. As emphasized in Kaplan and Violante (2010), it is important to match the amount of wealth that households hold over the life cycle in order to assess the amount of insurance that households can do on their own by accumulating assets. Further, I estimate rather than calibrate a model similar to that in Kaplan and Violante (2010), and infer the degree of consumption smoothing from the empirical sensitivity of consumption to income growth over one to four-year horizons, while matching the income moments (the amount of income risk and the persistence of that risk) and the average wealth-to-income ratio over the pre-retirement stage of the life cycle. My paper is also related to Guvenen and Smith (2010), which structurally estimates the income process using information on household consumption and income over the life cycle.

The rest of the paper is organized in the following way. In Section 2, I present the model and the income process. In Section 3, I discuss results from simulations of the model. In Section 4, I discuss the moment matching method used for estimation of the model parameters, and construction of the empirical moments used in matching. Section 5 presents the main results. Section 6 concludes.

2 The Model

In this section, I set up a model of household consumption over the life cycle, and discuss the potential importance of different income models with the same autocovariance structure for consumption dynamics and consumption smoothness.

Assume that households value consumption, supply labor inelastically, face income uncertainty over the working part of the life cycle, and are subject to liquidity constraints. Households start their life cycle at period 0, retire at period R , face age-dependent mortality risk until period T when they die with certainty. Thus, a household's problem is:

$$\max_{\{C_{it}\}_{t=0}^T} E_{i0} \sum_{t=0}^T \beta^t s_t U(C_{it}), \quad (1)$$

subject to the accumulation (cash-on-hand) constraint,

$$X_{it+1} = R_{t+1}(X_{it} - C_{it}) + Y_{it+1}, \quad (2)$$

and the liquidity constraint:

$$C_{it} \leq X_{it}, \text{ for } t = 0, \dots, T. \quad (3)$$

Cash-on-hand available to household i in period $t + 1$, $X_{it+1} = W_{it+1} + Y_{it+1}$, consists of labor income realized in period $t + 1$, Y_{it+1} , and household wealth at time $t + 1$, W_{it+1} ; R_{t+1} is a gross interest rate on a risk-free asset held between periods t and $t + 1$. β is the common pure time discount factor, s_t is the unconditional probability of surviving up to age t , C_{it+1} is household i 's consumption in period $t + 1$, and E_{i0} denotes household i 's expectation about future resources based on the information available at time 0. I assume that utility is CRRA, $U(C_{it}) = n_{it} \frac{(C_{it})^{1-\rho}}{1-\rho}$, where n_{it} stands for household i 's effective family size when the head is of age t . Households are subject to liquidity constraints so that their total consumption is constrained to be below their total cash-on-hand in each period—equation (3).

The Income Process

A popular and empirically justifiable income model decomposes household income, Y_{it+1} , into a random walk permanent component, P_{it+1} , and a transitory component, ϵ_{it+1}^T :⁵

$$Y_{it+1} = P_{it+1} \epsilon_{it+1}^T \text{ for } t = 0, \dots, R - 1 \quad (4)$$

$$P_{it+1} = G_{t+1} P_{it} \epsilon_{it+1}^P \text{ for } t = 0, \dots, R - 1, \quad (5)$$

where ϵ_{it+1}^P is an innovation to the permanent component, and G_{t+1} is the gross growth rate of income between ages t and $t + 1$ common to all households of age t .

After retirement, household income is assumed to be proportional to the permanent compo-

⁵In the context of computational consumption models, this model was first used by Zeldes (1989) and Carroll (1992).

ment of income received at age R :

$$Y_{it} = \kappa P_{iR} \text{ for } t = R + 1, \dots, T,$$

where κ is the replacement rate.

Taking natural logs, the first difference of household income during the working part of the life cycle is:⁶

$$\Delta \log Y_{it+1} = g_{t+1} + u_{it+1}^P + \Delta u_{it+1}^T, \quad (6)$$

where $\log Y_{it+1}$ is household i 's log-income at age $t + 1$; g_{t+1} is the log of its gross growth rate at age $t + 1$; u_{it+1}^P is the log of ϵ_{it+1}^P ; and u_{it+1}^T is the log of ϵ_{it+1}^T . g_{t+1} is composed of the aggregate productivity growth and the growth in the predictable component of income over the life cycle (which accounts, e.g., for the growth in income due to experience). After removing g_{t+1} from equation (6), the growth in income is affected solely by idiosyncratic shocks. Specifically, it is composed of the current value of the permanent shock, u_{it+1}^P , and the first difference in transitory shocks, u_{it+1}^T and u_{it}^T .

To calibrate the parameters of the household income process researchers use micro data, or rely on other studies of household income processes like Abowd and Card (1989) or MaCurdy (1982). What are the informational assumptions behind the income model in equations (4)–(6)? It is implicitly assumed that households can differentiate between permanent and transitory shocks, and that both econometricians and households know the joint distribution function of permanent and transitory shocks, usually assumed to be uncorrelated at all leads and lags. Thus, if the growth rate of income and interest rate are non-stochastic, the time- t (income) information set of household i is $\Omega_{it}^h = \{\epsilon_{it}^P, \epsilon_{it}^T, \epsilon_{it-1}^P, \epsilon_{it-1}^T, \epsilon_{it-2}^P, \epsilon_{it-2}^T, \dots, Y_{i0}\}$ while the econometrician's information set is $\Omega_{it}^e = \{Y_{it}, Y_{it-1}, Y_{it-2}, \dots, Y_{i0}\}$, where superscripts h and e stand for the household and econometrician, respectively. How important is the distinction of the

⁶For some evidence that idiosyncratic household log income is a difference stationary process see, e.g., Meghir and Pistaferri (2004) and Guiso, Pistaferri, and Schivardi (2005). Another model of idiosyncratic household income advanced in the literature is the heterogeneous growth-rate model (see, e.g., Baker 1997, Guvenen 2009) where idiosyncratic household log-income is a person-specific function of experience or age. Meghir and Pistaferri (2004) tested the null hypothesis that idiosyncratic household income is a difference stationary process against the growth-rate heterogeneity alternative and could not reject it. In a recent paper, Hryshko (2008) finds that male earnings data in the PSID are best represented by the model that contains a permanent random walk component and no deterministic growth-rate heterogeneity.

informational sets of econometricians and households? Assume a household knows that the shocks to its permanent and transitory income are negatively correlated. For example, when the head gets promoted, he expects his bonuses to be cut off. This (negative) correlation helps the household sharpen its predictions on the smoothness of the unpredictable part of the income growth, and adjust consumption appropriately. Econometricians, in turn, do not differentiate between income news known to households, but can decompose them into orthogonal permanent and transitory components. Consequently, econometricians make spurious conclusions about the joint distribution of permanent and transitory components, and this may lead to their wrong predictions of household reactions to income growth.⁷

Within the PIH, the correct identification of permanent versus transitory component of income has been proven to be important. Quah (1990) showed that if econometricians observe income news different from the news households observe, they may falsely reject the PIH, even though households behave exactly in accordance with it. This is the main point made by Quah (1990) that provides one of the solutions to the excess smoothness puzzle. Quah (1990) constructs different representations of several reduced form models of the aggregate US income, and finds that there always exists an income model consistent with the relative pattern of variances of consumption and income observed in the aggregate US data, and consistent with the PIH. Thus, the excess smoothness puzzle in macro data can be solved if the importance of the permanent component is “reduced.” It is possible to suppress the permanent component within an income model without distortion of the properties of the reduced form process.

I will now briefly outline this idea in the context of the PIH. If the reduced form income process follows an ARIMA(0,1,q) process, the PIH consumption rule for a dynastic household implies the following relation of consumption changes to income news (see, e.g., Deaton 1992):

$$\Delta C_{it} = \frac{r}{1+r} \frac{\theta_q \left(\frac{1}{1+r} \right)}{\left(1 - \frac{1}{1+r} \right)} \epsilon_{it} = \theta_q \left(\frac{1}{1+r} \right) \epsilon_{it},$$

⁷Throughout the paper, I assume that households know the joint distribution function of distinct income components. Other views on household versus econometrician’s (income) information have been explored in the literature. Pischke (1995), for example, assumes that household income consists of idiosyncratic and aggregate components and that a household cannot decompose the shock to its income into aggregate and idiosyncratic parts. For example, a household differentiates with a lag whether the head’s unemployment spell is due to an economy-wide shock, or whether it is the idiosyncratic shock. This assumption enables Pischke to provide micro-foundations for the “excess sensitivity” puzzle in macro data without violating the orthogonality condition of Hall (1978) at the micro level. Wang (2004) assumes that income consists of two potentially correlated processes of different persistence. He theoretically shows that a precautionary savings motive strengthens if an individual imperfectly observes innovations to each component compared to the case of the perfect knowledge about each component.

where $\theta_q(\cdot)$ is the lag polynomial of order q in L evaluated at $\frac{1}{1+r}$, and ϵ_{it} is a reduced form income shock. If, for example, $q = 1$ so that $\theta(L) = 1 + \theta L$ and, consistent with empirical micro data, θ is negative, consumption should change by $1 + \frac{\theta}{1+r}$. Parameter θ controls the mean reversion in income, and, along with the standard deviation of income shocks, determines the volatility of consumption changes. If θ is zero, income is a random walk and consumption should change by the full amount of the (permanent) income shock. The closer θ to -1.0 is, the less persistent is the income process, the smaller is the response of consumption to a permanent shock, and the smaller is the volatility of consumption changes for a given volatility of income shocks.

Assume that the reduced form income process, ARIMA(0,1, q), can be decomposed into a permanent IMA(1, q_1) component, and a transitory MA component of order q_0 , such that $\max(q_1, q_0 + 1)$ is equal to q , and permanent and transitory shocks are not correlated. It can be shown (see Quah 1990) that an income model that agrees with the reduced form ARIMA(0,1, q) income process implies the following response of consumption changes to transitory and permanent income shocks:⁸

$$\Delta C_{it} = \frac{r}{1+r} \theta_{q_0} \left(\frac{1}{1+r} \right) \epsilon_{it}^T + \theta_{q_1} \left(\frac{1}{1+r} \right) \epsilon_{it}^P. \quad (7)$$

Take $q_1 = 0$ and $q_0 = 0$, so that the order of autocovariance of the structural income process is the same as in the example above. In this case the implied consumption change should equal to the sum of the annuity value of the transitory income shock, and the entire permanent income shock. The response of consumption will be stronger if a permanent shock is larger. Similarly, the volatility of consumption changes will be larger if, within a structural income model, the volatility of permanent income shocks dominates the volatility of transitory income shocks. In general, the volatility of consumption changes, as implied by the PIH, depends on the relative importance of the permanent component. The weight of the permanent component in the income series is governed by polynomials $\theta_{q_1}(L)$, $\theta_{q_0}(L)$, and the relative variances of ϵ_{it}^T and ϵ_{it}^P under the constraint that autocovariance functions of reduced and structural form processes are identical. Since households have better information on the sequences of permanent and transitory shocks,

⁸Note that Quah (1990) considers linear difference stationary processes, while equation (6) features the log-linear income process. Campbell and Deaton (1989), however, show in a study of the PIH excess smoothness puzzle that this distinction is of little empirical importance. Furthermore, equation (7), derived using an UC representation of difference stationary linear income processes, serves only as a motivation for the main analysis of this paper. Thus, to avoid notational complications, for now, I interpret ϵ_{it}^T and ϵ_{it}^P as transitory and permanent innovations to the level of income within linear income processes.

one may conclude, provided the PIH is true, that the “correct” decomposition of income is the one that matches the ratio of the variances of consumption and income growth observed in the aggregate data with the ratio predicted by the PIH, which is not necessarily the one identified by econometricians.

This intuition can be summarized as follows. The relative dynamics of income components is best known to households and this unique knowledge should be reflected in household consumption choices. Econometricians, in turn, make inferences on income components from identified models of the income process which may or may not coincide with the model households “observe.” Ultimately, the importance of the income information sets should be judged by their effect on household choices of consumption. In the next section, I provide some evidence on the importance of this issue within a simulated buffer stock model of savings.

The autocovariance function of the reduced form process modeled as an ARIMA(0,1, q) has $q + 1$ non-zero autocovariances, which is sufficient to estimate q moving average parameters, along with the variance of the reduced form income shock. An estimable model of income may allow at most $q + 1$ non-zero parameters, two of which are the variances of structural shocks and the rest determine the dynamics of each unobserved component of income, $\theta_{q_1}(L)$ and $\theta_{q_0}(L)$. Thus, if the permanent component of income is a random walk and the transitory component is a moving average process of order $q - 1$, one can identify the variances of transitory and permanent shocks, and $q - 1$ moving average parameters; the correlation between the shocks is not identifiable from the sole dynamics of household income.

3 Simulations of the Model

In this section, I use the PSID to estimate a reduced form ARIMA(0,1,1) income model. I then construct several models of income that imply different permanent and transitory components but have the same autocovariance function as the reduced form. I assume that consumers make their consumption and savings choices in accordance with a life cycle buffer stock model, taking into account the knowledge of the joint distribution of permanent and transitory shocks. I further examine the effect of different income decompositions on consumption dynamics in the buffer stock model.

Univariate Dynamics of Idiosyncratic Household Income

In this section, I present some results on the univariate dynamics of household income in

PSID data.⁹

The income measure I consider is the residuals from the cross-sectional regressions of household disposable log income on a second degree polynomial in the head’s age, and education dummies.¹⁰ In the literature, it is typically labeled idiosyncratic household income. For the cross-sectional regressions, I use information from the 1981–1997 annual family files of the PSID.¹¹ Sample selection is described in Appendix C. Table 1 presents the autocovariance function for the growth in household idiosyncratic income. As can be seen from the table, the autocovariance function is statistically significant up to order two. This is consistent with an integrated moving average process of order two and the findings in Abowd and Card (1989), and Meghir and Pistaferri (2004). To simplify the computations, in the rest of the paper, I will assume that the reduced form income process is an integrated moving average of order one. This is not at odds with the data as the autocovariances of orders 2 and higher are small in magnitude. In Table 2, I present estimates of the reduced form process for idiosyncratic household income. Idiosyncratic household income is highly volatile, with a standard deviation of the reduced form shocks of about 20% per year, and contains a strong mean-reverting component.

Constructing Different Income Models

In this section, I decompose a moving average process estimated in the previous section into permanent and transitory components of different relative volatilities.

Assume that log income in differences, after the deterministic growth rate g_t has been removed, follows a stationary MA(1) process. This is consistent with an income process represented as the sum of a random walk permanent component and a transitory white noise process. This particular income process has become the workhorse in simulations of the buffer stock model of savings and for computational models of asset holdings over the life cycle.¹² The reduced and structural forms of the process for the first differences in income are:

⁹I will describe the data utilized in more detail in the next section and Appendix C.

¹⁰My specification of the predictable component of labor income is quite flexible: it assumes, for example, time-varying returns to experience and education.

¹¹The PSID collected data biennially after 1997. Inclusion of data after 1997 would require a different modeling strategy, e.g., analyzing idiosyncratic income growth over the two-year horizon. Since this strategy will necessarily result in a loss of data, I use the data available at the annual frequency.

¹²Ludvigson and Michaelides (2001) use this process to analyze “excess smoothness” and “excess sensitivity” puzzles on the aggregated data from a simulated buffer stock model; Michaelides (2001)—to investigate the same phenomena but for a buffer stock economy of consumers with habit forming preferences; Luengo-Prado (2007)—to analyze a buffer stock model augmented with durable goods. Luengo-Prado and Sørensen (2008) use a generalization of this process to explore the effects of different types of risk (idiosyncratic and aggregate) on the marginal propensity to consume in the simulated “state”-level data and US state-level data. Gomes and Michaelides (2005) and Cocco, Gomes, and Maenhout (2005) calibrate the parameters of this income process to investigate consumption and portfolio choice over the life cycle.

$$\begin{aligned}\Delta \log Y_{it}^{rf} &= (1 + \theta L)u_{it} \\ \Delta \log Y_{it}^{sf} &= u_{it}^P + (1 - L)u_{it}^T,\end{aligned}$$

where superscripts *rf* and *sf* denote the reduced and structural form, respectively.

Since the reduced form has only two pieces of information, the autocovariances of order zero and one, one can statistically identify only two parameters, the variance of permanent shocks and the variance of transitory shocks. To explore the impact of the structure of income on the consumption process, I allow for a covariance between the permanent and transitory shocks, and then work out the variance of transitory shocks. I match the moments of constructed series to the moments of the reduced form series, thus keeping the stochastic structure of the series intact. I present the full details of the procedure in Appendix A. I take the estimated parameters of an ARIMA(0,1,1) process from Table 2. The grid of covariances considered in simulations implies the following correlations between structural shocks: $-1.0, -0.75, -0.5, -0.25, 0.0, 0.25, 0.5, 0.75,$ and 1.0 . For the estimated income parameters, an estimate of the variance of innovations to the random walk permanent component equals $(1 + \hat{\theta})^2 \hat{\sigma}_u^2 = 0.02$. The variance of transitory innovations can be estimated by $-\hat{\gamma}(1) - \text{cov}(u_{it}^P, u_{it}^T)$, where $\hat{\gamma}(1)$ is the first order autocovariance of the reduced form process and $\text{cov}(u_{it}^P, u_{it}^T)$ is the covariance between permanent and transitory innovations. Thus, for the covariance equal to -0.0109 (and the corresponding correlation between income shocks approximately equal to -0.50), the standard deviation of transitory innovations is 0.152 ; for the covariance equal to 0.00 , the standard deviation of transitory innovations is 0.111 .

The decompositions differ in terms of the relative volatility of permanent and transitory shocks. Thus, the income model with perfect negative correlation between permanent and transitory shocks has the most volatile transitory shocks, while the income model with perfect positive correlation has the least volatile transitory shocks.

Results for Simulated Life Cycle Buffer Stock Economies

I solve the model introduced in the previous section by utilizing the Euler equation linking marginal utility from consumption in adjacent periods. I assume that the gross interest rate R_t is non-stochastic and the joint probability density function of transitory and permanent shocks is time-invariant. In addition, shocks are assumed to be jointly log-normal. I further assume that households start their life cycle at age 26 ($t = 0$ in the model), retire at age 65,

and die with certainty at age 90 ($T = 64$ in the model). Before retirement, the unconditional probability of survival is set to 1; after the retirement, households face an age-dependent risk of dying. The conditional probabilities of surviving up to age t provided the household is alive at age $t - 1$ for all $R < t \leq T - 1$ are taken from Table A.1 in Hubbard, Skinner, and Zeldes (1994).¹³ The replacement rate κ is set to 0.60. This value is similar to an estimate of the replacement rate for U.S. high school graduates in Cocco, Gomes, and Maenhout (2005). The average effective family size over the life cycle, n_t , is estimated using PSID data following Scholz, Seshadri, and Khitatrakun (2006) as $(\text{no. adults}_t + 0.7 \times \text{no. children}_t)^{0.7}$, where “no. adults $_t$ ” (“no. children $_t$ ”) is the “typical” number of adults (number of children) when household head is of age t . The age-dependent deterministic growth rate in household disposable income, G_t , is estimated using CEX data. I discuss construction of those profiles in the empirical section of the paper.

After I find the age-dependent consumption functions, I simulate the economy populated by 5,000 ex ante identical consumers, who are differentiated ex post due to different history of income draws. I assume that households have zero wealth in the beginning of their life cycle, at age 26. Since I am interested in the properties of consumption for different decompositions of a given reduced form model of income, I hold all other parameters of the buffer stock model fixed. Thus, in my first set of simulations, I do not vary the behavioral parameters of the model. I set the gross real interest rate to 1.03, the time discount factor to 0.85, and the coefficient of relative risk aversion to 6.0. I choose low patience and high degree of risk aversion since I will later estimate similar values using PSID data. Those values are also consistent with the estimates in Cagetti (2003) who used wealth data to fit a life cycle model. I take draws from the joint distribution of log-normal transitory and permanent shocks, the parameters of which are derived from the reduced form ARIMA(0,1,1), as already discussed in the previous subsection. The details of the model solution are provided in Appendix B.

I run pooled panel regressions of the growth of household consumption on the growth of household income over one and four-year horizons. The magnitude of the coefficient on the current income growth should depend on the smoothness of income innovations. Long differences in log income will be largely dominated by the permanent shocks, which should be reflected in the long differences in log consumption. In the empirical evaluation of the model, I will match the wealth-to-income ratio over the life cycle. This information is important to identify the time discount factor and the coefficient of relative risk aversion as shown in Cagetti (2003). Matching the wealth-to-income ratio also provides some discipline on the ability of households

¹³This is the mortality data on women for 1982.

to smooth consumption using their own assets before other forms of insurance are allowed for. Thus, in addition to the moments that describe the sensitivity of consumption to income shocks, I tabulate the average wealth-to-income ratio at ages 31–35 and 61–65, and the reduced form income parameters, an autoregressive persistence and the variance of income growth. The results for income models with negative correlation, no correlation, and positive correlation between the shocks are presented in Table 4.¹⁴

In the first three rows of Table 4, I show that consumption is contemporaneously less sensitive to income when the correlation between the shocks is the lowest. Similar results hold for the sensitivity of consumption growth to income growth over the four-year horizon. The average wealth-to-income ratio at early and late stages of the working part of the life cycle is not affected much by the choice of the income process.

The basic intuition behind the results is the following. Absent borrowing restrictions, households react only to the newly arrived permanent and transitory innovations, u_{it}^P and u_{it}^T . The sensitivity of household consumption to income news can be described by the equation

$$\Delta \log C_{it} = \alpha_P u_{it}^P + \alpha_T u_{it}^T, \quad (8)$$

where α_P and α_T are the (“partial insurance”) coefficients that depend on the endogenously accumulated wealth and, therefore, on the relative risk aversion parameter, the time discount factor, the real interest rate, and the volatility of permanent and transitory shocks. While the regression can be estimated using simulated data since permanent and transitory innovations can be observed, in the real data one can only relate $\Delta \log C_{it}$ to the observable income growth, $\Delta \log Y_{it}$, which, for the income process analyzed, equals $u_{it}^P + u_{it}^T - u_{it-1}^T$. Thus, one can evaluate the above equation to make predictions, for simulated economies with households facing different structural income processes, on the coefficient β_1 from an OLS regression $\Delta \log C_{it} = \beta_0 + \beta_1 \Delta \log Y_{it} + \epsilon_{it}$, and β_k from an OLS regression $\Delta_k \log C_{it} = \beta_0 + \beta_k \Delta_k \log Y_{it} + \epsilon_{it}$, where $\Delta_k \log C_{it} = \log C_{it} - \log C_{it-k}$, and similarly for $\Delta_k \log Y_{it}$. Intuitively, if permanent and transitory innovations are negatively correlated, the portion of the unpredictable income growth to which households react, $u_{it}^P + u_{it}^T$, is smoother compared with the case when the structural innovations are uncorrelated or positively correlated. For the case of a negative

¹⁴The MPCs out of shocks to current income, and the shocks cumulated over the four-year horizon are larger for models with a higher correlation between the shocks. Thus, without losing valuable information, I chose to report only the results for the income models with the correlation between the shocks equal to -0.50 , 0.0 , and 0.50 .

correlation, a positive permanent shock is, on average, accompanied by a negative transitory shock, smoothing out the sum of income innovations. Hence, income becomes smoother and this is reflected in lower coefficients measuring the sensitivity of current consumption to current income growth (β_1), and cumulative consumption growth to cumulative income growth over the four-year horizon (β_4). For the case of a positive correlation, positive (negative) permanent shocks arrive, on average, together with positive (negative) transitory shocks, making the sum of innovations less smooth and this is consequently reflected in higher coefficients measuring the sensitivity of consumption to income growth at different horizons (β_1 and β_4). In statistical terms,

$$\hat{\beta}_1 = \frac{\text{cov}(\Delta \log C_{it}, \Delta \log Y_{it})}{\text{var}(\Delta \log Y_{it})} = \frac{\alpha_P \sigma_{u^P}^2 + \alpha_T \sigma_{u^T}^2 + (\alpha_P + \alpha_T) \text{cov}(u_{it}^P, u_{it}^T)}{\text{var}(\Delta \log Y_{it})}.$$

The denominator is the same for all structural decompositions of the reduced form income model, the “smoothing” term is measured by $(\alpha_P + \alpha_T) \text{cov}(u_{it}^P, u_{it}^T)$ in the numerator. It follows that the sensitivity of current consumption to current income growth is lower for structural income models with more negatively correlated shocks. The sensitivity of cumulative consumption growth to cumulative income growth over k periods is measured by

$$\hat{\beta}_k = \frac{k\alpha_P \sigma_{u^P}^2 + \alpha_T \sigma_{u^T}^2 + (\alpha_P + k\alpha_T) \text{cov}(u_{it}^P, u_{it}^T)}{k\sigma_{u^P}^2 + 2\sigma_{u^T}^2 + 2\text{cov}(u_{it}^P, u_{it}^T)}.$$

Again, the denominator is the same for different structural income processes while the numerator contains the “smoothing” term $(\alpha_P + k\alpha_T) \text{cov}(u_{it}^P, u_{it}^T)$, which is larger, in absolute value, for the processes with more negatively correlated permanent and transitory shocks.

In rows (4) and (5), I explore the sensitivity of the moments in the benchmark model in row (1) to different values of the risk aversion parameter. Lower risk aversion results in a much lower accumulation of assets over the life cycle—households arrive with virtually no assets at retirement when the coefficient of relative risk aversion is set to 2 and the degree of impatience is kept at a high level. As a result, households are very sensitive to income shocks, as reflected in high values of $\hat{\beta}_1$ and $\hat{\beta}_4$. The reverse is true for a higher degree of risk aversion. In rows (5) and (6), I examine the sensitivity of the model moments to variations in the time discount factor, holding other parameters at their values in row (1). The results are intuitive: more patient consumers accumulate larger amounts of wealth and are able to better smooth consumption

over the life cycle.

Lastly, in rows (8)–(10) I explore the effect of introducing partial risk sharing against permanent and/or transitory income shocks on the model moments. I do not model risk sharing in a structural way. Rather, I follow Attanasio and Pavoni (2007) who show, for a model with hidden access to asset markets, that the bond (self-insurance) Euler equation holds for household resources that have been smoothed by state-contingent transfers or other mechanisms before households make their decisions on savings. The sensitivity of consumption to income shocks at one and four-year horizons is halved when 50 percent risk sharing of permanent and transitory shocks are allowed in the model. The results are similar when partial risk sharing of permanent income shocks only is introduced into the model—row (9). Households appear to substantially smooth transitory shocks using accumulated assets (for a similar result, see Kaplan and Violante 2010). Consumption reaction to income shocks is similar to the no-insurance case when households do not have access to partial risk sharing of permanent income shocks but 50 percent of transitory shocks are smoothed away before self-insurance—rows (1) and (10). Thus, the model moments are not affected much by the availability of partial risk sharing of transitory shocks. It can be concluded that partial risk sharing of transitory income shocks beyond self-insurance is not likely to be well identified empirically. In the empirical section, therefore, I will estimate the degree of risk sharing of permanent shocks only.

Summarizing, there is substantial variation of the model moments with respect to changes in the income process parameters, behavioral and risk-sharing parameters. It appears possible to identify the model parameters by matching the data moments to the same moments estimated within the model.

4 Estimation of the Model

In this section, I use a life cycle model of consumption to estimate the parameters of the income process and the behavioral parameters. I assume that model households are married couples who maximize expected utility from consumption over the life cycle. The only source of uncertainty in the model before retirement is uncertainty over the flows of income, arising from transitory and permanent income shocks. I assume that all households start working at age 26 and retire at age 66.

As in previous section, I assume that households have access to one instrument for saving and consumption smoothing—a riskless bond with the deterministic gross interest rate R . Cash-on-hand accumulation constraint and the income process are given in equations (2), and (4)–(6)

respectively. I assume that households are subject to liquidity constraints so that their total consumption is constrained to be below their total cash-on-hand in each period.

Cash-on-hand and consumption at age t can be expressed in terms of the ratios to the permanent component of income at age t , and the state space of the corresponding dynamic programming problem reduces to one variable, cash-on-hand relative to the permanent income, x_{it} . The details of the model solution are provided in Appendix B.

Matching Empirical Moments

In this section, I describe the method used to estimate the structural parameters of the model. The vector of structural parameters θ consists of the behavioral parameters— β , ρ ; the parameters of the income process— σ_{u^T} , σ_{u^P} and $\text{corr}(u_{it}^P u_{it}^T)$; and the partial risk sharing parameters, ω_P and ω_T . I estimate the model parameters by the method of simulated moments.

I recover the parameters by matching the empirical moments listed in Table 5. I match fifteen moments in total (enumerated in the table). Since the model does not provide a closed-form solution for these moments, I simulate the moments and estimate the parameters of the model by matching these simulated moments to the data moments. I estimate the model in two stages. In the first stage I estimate the exogenous parameters χ , I then fix them in the MSM optimization routine; in the second stage I estimate, within the MSM routine, the model parameters θ . χ consists of the life cycle profile of the (deterministic) gross growth rates of disposable income, $\{G_t\}_{t=26}^{65}$; the life cycle profile of the effective family size, $\{n_t\}_{t=26}^{66}$; the mean and standard deviation of the distribution of the permanent component of household disposable income at age 26;¹⁵ and the variance of measurement error in household total expenditures. I set the gross real interest rate on safe liquid assets to 1.03.

Given the estimates of the first stage parameters, the MSM estimates of the second stage parameters θ are such that the weighted distance between the vector of simulated moments and the vector of empirical moments is as close to zero as “possible.” $\hat{\theta}$ is the solution to the minimization of the following criterion function

$$\left[\log m^s(\theta; \hat{\chi}) - \log m^d \right]' W \left[\log m^s(\theta; \hat{\chi}) - \log m^d \right] = g'_{I_s} W g_{I_s},$$

where superscript d denotes data; s denotes simulation; I_d is the number of households in the data contributing towards estimation of the second-stage moments; I_s is the number of simulated households; m^d is a vector of the second-stage moments estimated from the data; $m^s(\theta; \hat{\chi})$ is

¹⁵I set those to the mean and variance of the distribution of household disposable income at age 26.

a vector of simulated moments; W is a positive definite weighting matrix; $\hat{\chi}$ is a vector of the preestimated first-stage moments; θ is a vector of the second-stage parameters.

Construction of Empirical Moments and Life Cycle Profiles

In this section, I describe estimation of the empirical moments I match. I first briefly describe the data sources used. I obtain consumption information from two data sources, the CEX and the PSID. The CEX contains detailed information on total expenditures and its components, and the demographics for representative cross sections of the US population. I use extracts from the 1980–2003 waves of the CEX available at the National Bureau of Economic Research (NBER) webpage. Unlike the CEX, the PSID provides panel data yet limits its coverage of consumer expenditures to food at home and away from home. Since I am interested in the link between changes of household disposable income and total household consumption, I impute the total consumption to the sample PSID households using information on household food consumption in the PSID and the CEX, and matched demographics from the CEX and the PSID. PSID data are taken from 1981–1997, 1999, 2001, and 2003 waves. I follow the methodology of Blundell, Pistaferri, and Preston (2005) to impute total consumption to the PSID households. The full details on sample selection of CEX and PSID households are provided in Appendix C. Briefly, from the PSID, I choose married couples headed by males of ages 26–70 born between 1912 and 1978, with no changes in family composition (no changes at all or changes in family members other than the head and wife). I drop income outliers, observations with missing or zero records on food at home and, for each household, keep the longest period with consecutive information on household disposable income and no missing demographics. From the CEX, I choose households who are complete income and expenditure reporters, with heads who belong to the same age groups and cohorts as in the PSID sample.

In the PSID, federal income taxes are calculated by staff until 1991. To have a consistent measure of federal income taxes for the data that extend beyond 1991, one needs to impute them to the PSID households. I use the TAXSIM tool at the NBER to calculate federal income taxes and social security withholdings for the head and wife and all other family members if present. I use information on imputed household disposable income for estimation of the moments listed in Table 5. I use CEX data to construct the profile of the life-cycle growth in household disposable income. In the CEX, federal income taxes and taxable household income are reported rather than imputed. Thus, the profile of the deterministic life-cycle growth in household disposable income can be more reliably estimated using CEX data.

I decompose household disposable log income into cohort, time, and age effects, controlling

for the effect of family size. As is well known, age, cohort, and time effects are not separately identified. I follow Deaton (1997) and restrict the time dummies to be orthogonal to a time trend and to add up to zero. The age effects from such regression, smoothed using a fifth-degree polynomial, are depicted in panel (a) of Figure 1.¹⁶ Household disposable income peaks at age 52. The profile of deterministic growth in household disposable income, G_t , is obtained by taking the difference between the adjacent points in the figure.

Attanasio, Banks, Meghir, and Weber (1999) showed the importance of controlling for changing family size over the life cycle in a model with income uncertainty and self-insurance. I calculate the effective family size for household i at age t as $n_{it} = (\text{no. adults}_{it} + 0.7 \times \text{no. children}_{it})^{0.7}$, following Scholz, Seshadri, and Khitatrakun (2006). I use PSID data to construct the life cycle profile of the effective family size.¹⁷ I run a regression that controls for household fixed effects, age effects, and time effects.¹⁸ As with household income data, I assume that the time dummies are orthogonal to a time trend and sum to zero. The age effects, smoothed using a fifth-degree polynomial, are depicted in panel (b) of Figure (1). The effective family size peaks at age 40.

In Table 3, I present the autocovariance function of idiosyncratic growth rate of household total (imputed) consumption.¹⁹ I utilize data from the 1981–1997 surveys of the PSID. I run cross-sectional regressions of the first difference in household log consumption on a quadratic polynomial in the head’s age and the change in family size. The residuals from those regressions are labeled idiosyncratic consumption growth. Only the first-order autocovariance of consumption growth is significant; higher-order autocovariances are small in magnitude and not significant. The variance of idiosyncratic consumption growth is large in magnitude, which can be partly explained by the variance of measurement and imputation error in consumption. In theory, household consumption is a martingale unless consumption is measured with error, in which case the first-order autocovariance of consumption growth will be negative. I assume that the variance of measurement and imputation error in consumption is equal to 0.08, the negative of the estimated value of the first-order autocovariance in idiosyncratic consumption growth.

In panel (c) of Figure 1, I plot the life cycle profile of the wealth-to-income ratio. I use household net worth exclusive of business wealth, and household disposable income to construct this measure. Relative to the household income and family size data, I do not have as much data for construction of the wealth profile. For the time span of my sample, the wealth data

¹⁶The omitted age group in the regression are the heads of age 26. Thus, the depicted age effects should be interpreted as the average log income at ages 27, . . . , 65 relative to the average log income at age 26.

¹⁷The profile is similar if CEX data are used instead.

¹⁸If household fixed effects are included, the cohort effects are not identified.

¹⁹It is similar to the autocovariance function of consumption growth in Blundell, Pistaferri, and Preston (2008).

are recorded only in six supplements, collected every five years from 1984 to 1999, and every other year afterwards. The time effects, restricted as above, are unlikely to be identified.²⁰ Because of a small number of wealth observations at different ages, especially at older ages, I run a regression of the wealth-to-income ratio on a limited set of age dummies, controlling for household fixed effects. I consider 8 age dummies. Each age dummy comprises households who fall into one of the 8 five-year age intervals: 26 – 30, . . . , 61 – 65. The results of this regression are provided in Table 5 and depicted in panel (c) of Figure 1. Households start with low wealth in the beginning of the life cycle and have wealth exceeding their income by more than 4.5 times when they approach retirement. In the model, I assume that households start with zero wealth at age 26. To eliminate the influence of this assumption on the results, I do not match the wealth moment calculated for ages 26–30.

The other moments used for matching measure the extent of smoothness of consumption with respect to income changes, and the variance of idiosyncratic income growth and the persistence of household disposable income. The consumption smoothness is measured by the coefficients $\hat{\beta}_j$, $j = 1, \dots, 4$, from the following panel regressions estimated by OLS:

$$\Delta_j \log c_{it} = \beta_0 + \beta_j \Delta_j \log y_{it} + \gamma' x_{it} + \epsilon_{it},$$

where $z_{it} \equiv Z_{it} - \bar{Z}_t$ for any variable z in the regression, $\Delta_j \log z_{it} \equiv \log z_{it} - \log z_{it-j}$, and x_{it} is a vector that comprises a quadratic polynomial in the head's age, and family size. I take out the time-specific averages from the variables since I do not have aggregate uncertainty in the model. The estimated value of β_1 suggests that about 12 percent of the shocks to current income are translated into consumption—row (8) of Table 5. This can be due to the presence of a large transitory component in income, measurement error in income, or different insurance mechanisms available to households for smoothing out fluctuations in disposable income. Households react to about 23 percent of income shocks cumulated over the four-year horizon, as indicated by the estimated value of β_4 —row (11) of Table 5.

If permanent and transitory shocks are uncorrelated and the transitory component is an iid process, the following respective moments will identify α_P and α_T in equation (8) (see, e.g., Kaufmann and Pistaferri 2009): $\frac{E(\epsilon_{it}^c \epsilon_{it}^y) + E(\epsilon_{it}^c \epsilon_{it+1}^y)}{E(\epsilon_{it}^y)^2 + 2E(\epsilon_{it}^y \epsilon_{it+1}^y)}$ and $\frac{E(\epsilon_{it}^c \epsilon_{it+1}^y)}{E(\epsilon_{it}^y \epsilon_{it+1}^y)}$, where ϵ_{it}^c (ϵ_{it}^y) is the first difference in residuals from a regression of log consumption (income) on dummy variables for

²⁰The life cycle profile of the wealth-to-income ratio is similar if I simply take out the time-specific means from the ratio and then run a regression controlling for age and household fixed effects.

the head’s year-of-birth, high school and college, race, family size, Census region, number of kids, employment status, as well as interactions of those dummies (but race) with year dummies. For matching, I use $(1-\hat{\alpha}_P)$ and $(1-\hat{\alpha}_T)$ as the data estimates of partial insurance against permanent and transitory shocks; I denote them as $\hat{\phi}_P$ and $\hat{\phi}_T$, respectively. I find that households smooth out about 40% (100%) of permanent (transitory) shocks to their disposable incomes—rows (12)–(13) of Table 5. Similar estimates, but for a different sample, have been found in Blundell, Pistaferri, and Preston (2008). Ideally, a successful model should match both the reaction of consumption to income shocks at different horizons (β -coefficients) and the partial insurance moments (ϕ_P and ϕ_T).

I match the AR(1) coefficient of the reduced form process rather than an MA(1) estimate, since an AR(1) process is less time consuming to estimate. This proves to be very important when repeated estimations are performed on simulated data. Specifically, I match the value of $\hat{\alpha}$ estimated from the following regression:

$$\Delta \log y_{it} = \alpha_0 + \alpha \Delta \log y_{it-1} + \gamma' x_{it} + \epsilon_{it},$$

where x_{it} is a vector that includes a quadratic polynomial in the head’s age and dummies for the head’s high school graduation and college completion. As in the previous regressions, I take out the time-specific means from each variable prior to running the regression.

The size of income risk over the life cycle is calculated as the variance of idiosyncratic income growth. For its estimation, I first run cross-sectional regressions of the difference in household log disposable income on a quadratic polynomial in the head’s age and dummies for high school and college using data from the 1981–1997 surveys, when household income was continuously recorded each year. I limit the regression sample to the households with heads of ages 26–65. The unconditional variance of the residuals from those regressions provides an estimate of the proportional risk to household disposable income over the life cycle. The estimated variance along with its standard error are shown in row (14) of Table 5.

Further details on the model solution, and calculation of standard errors are provided in Appendix B.

5 Results

Table 6 contains the main results. The values of the data moments reported in Table 5 are replicated for convenience in square brackets in the bottom of Table 6.

I first assume that transitory and permanent shocks are uncorrelated and target the wealth moments, income moments, and the reaction of consumption to current income shocks—column (1). The coefficient of relative risk aversion is estimated at about 7, and the time discount factor at about 0.81: households exhibit low patience and high degree of risk aversion. The estimated value of the time discount factor compares well with Cagetti (2003); the relative risk aversion parameter is comparable in magnitude with the estimates in Cagetti (2003) and Nielsen and Vissing-Jorgensen (2006). The size of permanent shocks to household disposable income and the size of transitory shocks are tightly estimated: one standard deviation in permanent shocks equals about 14 percent, while one standard deviation in transitory white-noise shocks equals 12 percent.²¹ The wealth moments, the variance of income risk and the persistence of risk are well fitted to the corresponding data moments.²² Consumption, however, appears to be excessively smooth: consumption is more than twice as sensitive to current income shocks and shocks cumulated over the four-year horizon as in the data; in the data, households appear to smooth away twice as much permanent shocks as in the model. As emphasized in the simulation section, the consumption smoothness moments can be better explained if households have access to some insurance mechanisms other than self-insurance via accumulation of wealth, or if they react to smoother income components, the dynamics of which agree with the reduced form income moments but cannot be detected by using the income data alone. I first examine if the second mechanism is enough to explain the consumption smoothness moments found in the data.

In column (2), I allow components of income to be contemporaneously correlated. I precisely estimate the negative correlation between permanent and transitory shocks of about -0.61 . The fit of the model improves, as reflected in a lower value of the goodness-of-fit statistics, and the model is able to explain the reaction of consumption to the shocks to current income, along with the wealth moments and income moments. The model, however, fails at predicting the reaction of household consumption to the income shocks cumulated over the four-year horizon as well as the moment relating to partial insurance against permanent shocks, $\hat{\phi}_P$. That is, consumption

²¹These values are similar to the estimates of the income process using data on household idiosyncratic income growth alone.

²²In columns (1) and (3) of Tables 6–7, to better match the income moments I placed relatively larger weights to the income moments in the weighting matrix while using equal weights for the wealth and consumption smoothness moments. The reason for this choice is that I wanted model households to face exactly the same amount of risk and its persistence as seen in the data, and to explore the performance of the model with respect to the consumption smoothness moments, while fitting the income and wealth moments as close as possible.

is still excessively smooth.

There are some plausible explanations for the correlation between structural income shocks found in the data. The sign of the covariance may indicate that unfavorable permanent shocks to disposable household income, such as the head's long-term unemployment, are partially offset by increases in the transitory income such as unemployment compensation from the government. It is also likely that this offsetting effect will manifest itself at the annual frequency, the frequency I use for modeling consumption in the life cycle model. Consider another explanation for this finding. Household income derives from multiple sources: the wage of wife and head, transfer income of various sorts, (labor part of) business and farm income, (labor part of) income from roomers and boarders, bonuses, overtime and tips. As an example, if a household experiences a negative shock to the head's wages, plausibly assumed to be in the list of permanent shocks, it may compensate the adverse effect by temporarily leasing available housing. In a recent paper, Belzil and Bognanno (2008) find that increases in the base pay (positive permanent shocks) for American executives are followed by bonus cuts (negative transitory shocks). They argue that this phenomenon (of the negative correlation between the shocks) may reflect a compensation-smoothing strategy on the part of firms' managers.²³

In column (3) of Table 6 I match all moments in Table 5, allowing for both household information about the (potentially correlated) income components and partial risk sharing against permanent income shocks. Previous results are based on fitting empirical data moments to the same moments from a life cycle model that features self-insurance. In the model, households have access only to one vehicle of consumption smoothing over the life cycle, the risk-free bond. In reality, households may rely on other insurance mechanisms, e.g., state contingent assets and transfers of different sorts, that is, their permanent and transitory idiosyncratic shocks may be partially insured. In a recent paper, Attanasio and Pavoni (2007) showed, for a partial risk sharing model with moral hazard and hidden asset accumulation, that the self-insurance Euler equation still holds if applied to the "after-risk-sharing" income.²⁴ To account

²³Jacobson, LaLonde, and Sullivan (1993), for a sample of high-tenure workers, find that job displacement results into an initial drop of about 50% of pre-displacement earnings; eventually earnings recover but they are still 25% below their pre-displacement levels in 6 years. If one is willing to make an inference that the permanent shock equals 25%, then, at the arrival of the displacement event, this negative permanent shock should be accompanied by a negative transitory shock. Thus, contrary to this paper's finding the correlation between the shocks should be positive. However, the displacement event is just one among the very many events that cause permanent variations in incomes. Moreover, it is quite infrequent, with the annual likelihood of occurrence of about 4%. See Krebs (2003) for a review of the literature on the effect of displacements on earnings. In this paper I find that, on average, permanent and transitory shocks comove in different directions, and this negative comovement helps reconcile the reaction of consumption to income observed in the data. The negative correlation is not inconsistent with an observation that some permanent and transitory shocks move in the same direction.

²⁴There can be other market structures that allow for partial risk sharing. I consider the one proposed by Attanasio and Pavoni (2007) since it is easy to integrate it into the model in the "reduced-form" way.

for partial risk sharing, I model the after-risk-sharing household income as $Y_{it} = P_{it}\epsilon_{it}^T$, where $P_{it} = G_t P_{it-1} (\epsilon_{it}^P)^{\omega_P}$, and $1 - \omega_P$ is the fraction of permanent shocks that is smoothed out before households make their decisions on asset holdings. I do not consider partial risk sharing of transitory shocks since those are well insured by means of self-insurance.²⁵ The life-cycle path of the average wealth-to-income ratio in the data and model are plotted in Figure 2, panel (a). The profiles align quite well. In panel (b), I plot the profile of consumption over the life cycle implied by the model estimates in column (3) of Table 6. Consumption has a visible hump and peaks at around age 50. The hump and the timing of the peak in different measures of household consumption is documented in many studies—see, e.g., Fernández-Villaverde and Krueger (2007). The estimates of the time discount factor and the coefficient of relative risk aversion do not change appreciably.

Relative to column (2), I estimate a somewhat lower standard deviation of transitory shocks and correlation between the shocks, both significant at the 2% level. The model in column (3) fits well the profile of consumption adjustment to the shocks over different horizons (β -coefficients) but overestimates the moments relating to partial insurance of permanent and transitory shocks; I estimate that about 52 percent of permanent shocks are insured before self-insurance.

In Table 7, I repeat the same analysis assuming that measurement error in income explains 25 percent of idiosyncratic income growth.²⁶ This results in lower estimates of the size of transitory shocks, and less precisely estimated correlation between the shocks when partial risk sharing of permanent shocks is allowed for—the correlation is still significant at the 8% level. The size of permanent shocks is estimated at about the same value.²⁷ The estimates of the time discount factor, the coefficient of relative risk aversion, the correlation between the shocks and partial insurance of permanent shocks are very similar in magnitude.

It may be the case that a more superior fit of the wealth and β -moments is due to the equal weighting scheme and a relatively larger amount of those moments. In column (4) of Table 7, therefore, I fit two wealth moments (average wealth for ages 31–45 and ages 46–65), two β -moments (consumption reaction to income shocks cumulated over one and four-year horizons), two income process moments, and two partial insurance moments. As a result, the fit of the model to the moment relating to partial insurance of permanent shocks improves but the model

²⁵As shown in Table 4, partial risk sharing of transitory shocks is not likely to be empirically identified since the moments barely change when I allow for it in the model.

²⁶This is consistent with the literature on measurement error in *individual* earnings surveyed in Bound, Brown, and Mathiowetz (2001). I am not aware of research on the extent of measurement error in *household* disposable income.

²⁷This result is consistent with the theory: if log income contains a random-walk component, the volatility of permanent shocks should be identified for any covariance-stationary model of the transitory component, and any value of the contemporaneous correlation between permanent and transitory shocks—see, e.g., Cochrane (1988).

still overestimates this moment.

It is possible to provide direct evidence on the correlation between the shocks if both data on household expectations and realizations of income growth are available for several years. This information is not available in the U.S. but can be found in Italian data from the Survey of Household Income and Wealth (SHIW). Pistaferri (2001), using these data, identifies permanent and transitory shocks and tests the PIH by studying the reaction of household savings to the shocks. I use income and demographic data for individuals from the 1995, 1998, and 2000 waves of the SHIW. The first two waves contain individual records on one-year ahead expectations of non-financial income. I utilize the sample comprising individuals with non-missing information on income expectations in the 1995 and 1998 waves, and non-missing information on income realizations in all three waves. I find the correlation between transitory and permanent shocks of about -0.45 (-0.39) with a bootstrapped standard error of 0.11 (0.18) for the whole (heads-only) sample.²⁸ It appears that my findings using a matching exercise are in accord with the findings using survey data, with a qualification that they are based on data from different developed countries.

6 Conclusion

In an estimated life cycle model with self-insurance, I find that household consumption is excessively smooth. To reconcile the model with the data, I suggest that households have better information about income components than econometricians. In this case, the structure of the income process that econometricians can identify from the univariate dynamics of household income may differ from the true income structure.

When permanent and transitory shocks are contemporaneously negatively correlated, the unpredictable part of income growth, to which households react, is smoother compared with the case of zero and positive correlation between the shocks. Likewise, consumption becomes smoother and this is reflected in lower sensitivities of consumption to current income shocks, and shocks cumulated over longer horizons. This mechanism allows me to identify the correlation in the data and to better fit the consumption smoothness moments. Consumption, however, is still excessively smooth in the data compared to the model. The model is able to replicate the

²⁸For these data, I can directly identify the correlation between the sum of permanent shocks in years 1996–1998 and the transitory shock in 1998—details are available upon request. Assuming that the variance of permanent and transitory shocks, as well as the covariance between the shocks are time-invariant, I can further identify the correlation between the shocks, which is reported in the text. If the transitory component is an MA(1) process, an estimate of the reported correlation is biased towards zero.

reaction of consumption to current income growth, while the reaction of consumption to income growth over the four-year horizon is larger in the model than in the data.

The model is consistent with income, wealth, and consumption smoothness moments estimated from CEX and PSID data when I allow for household information about components of income and partial risk sharing against permanent income shocks. The precautionary savings motive found to be important for understanding household wealth accumulation in Cagett (2003) and Gourinchas and Parker (2002) is not enough to explain the smoothness of consumption that households are able to achieve in the data. More research is needed to understand insurance mechanisms available to U.S. households, besides those provided publicly in the form of miscellaneous public transfers and progressive taxation.

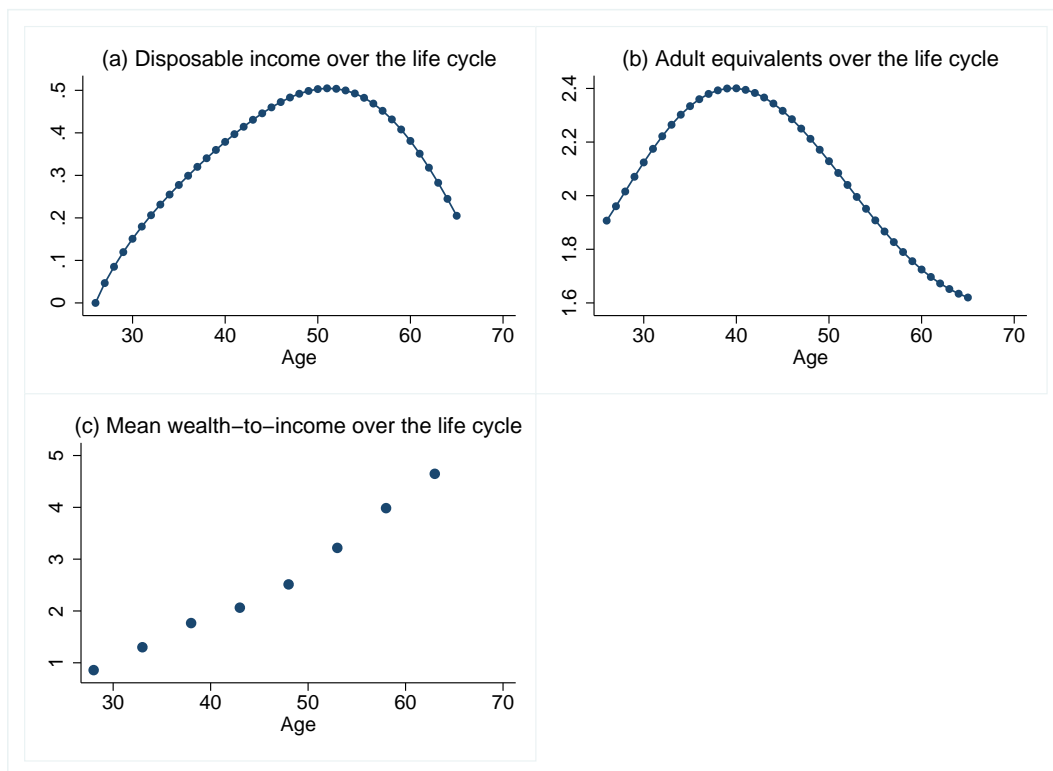
References

- ABOWD, J., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57, 411–445.
- ATTANASIO, O., J. BANKS, C. MEGHIR, AND G. WEBER (1999): "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economics Statistics*, 17, 22–35.
- ATTANASIO, O., AND N. PAVONI (2007): "Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption," NBER Working Paper # 12994.
- BAKER, M. (1997): "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings," *Journal of Labor Economics*, 15, 338–375.
- BELZIL, C., AND M. BOGNANNO (2008): "Promotions, Demotions, Halo Effects, and the Earnings Dynamics of American Executives," *Journal of Labor Economics*, 26, 287–310.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2005): "Imputing Consumption in the PSID Using Food Demand Estimates from the CEX," IFS Working Paper #04/27.
- (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.
- BOUND, J., C. BROWN, AND N. MATHIOWETZ (2001): "Measurement Error in Survey Data," in *Handbook of Econometrics*, Vol.5, chap. 59, pp. 3705–3843. North-Holland, New York.
- CAGETTI, M. (2003): "Wealth Accumulation Over the Life Cycle and Precautionary Savings," *Journal of Business and Economics Statistics*, 21, 339–353.
- CAMPBELL, J. Y., AND A. DEATON (1989): "Why is Consumption so Smooth?," *Review of Economic Studies*, 56, 357–374.
- CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–156.
- (2001): "Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income," NBER Working Paper # 8223.
- (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91, 312–320.
- CARROLL, C. D., AND A. A. SAMWICK (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40, 41–71.
- COCCO, J., F. GOMES, AND P. MAENHOUT (2005): "Consumption and Portfolio Choice Over the Life Cycle," *Review of Financial Studies*, 18, 491–533.
- COCHRANE, J. (1988): "How Big is the Random Walk in GNP?," *Journal of Political Economy*, 96, 893–920.
- DEATON, A. (1992): *Understanding Consumption*. Oxford University Press, New York, NY.

- (1997): *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. The Johns Hopkins University Press, Baltimore, Maryland.
- FERNÁNDEZ-VILLAVARDE, J., AND D. KRUEGER (2007): “Consumption Over the Life Cycle: Facts from Consumer Expenditure Survey Data,” *Review of Economics and Statistics*, 89, 552–565.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*. Princeton University Press, Princeton, NJ.
- (1963): “Windfalls, the “Horizon,” and Related Concepts in the Permanent-Income Hypothesis,” in *Measurement in Economics; Studies in Mathematical Economics and Econometrics in Memory of Yehuda Grunfeld*. Stanford University Press, Stanford, California.
- GOMES, F., AND A. MICHAELIDES (2005): “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *Journal of Finance*, 60, 869–904.
- GOURINCHAS, P.-O., AND J. PARKER (2002): “Consumption Over the Life Cycle,” *Econometrica*, 70, 47–89.
- GUIISO, L., L. PISTAFERRI, AND F. SCHIVARDI (2005): “Insurance within the Firm,” *Journal of Political Economy*, 113, 1054–1087.
- GUVENEN, F. (2009): “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, 12(1), 58–79.
- GUVENEN, F., AND A. SMITH (2010): “Inferring Labor Income Risk from Economic Choices: an Indirect Inference Approach,” NBER Working Paper # 16327.
- HALL, R. E. (1978): “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, 86, 971–987.
- HILL, M. (1992): *The Panel Study of Income Dynamics: A User’s Guide*. Sage Publications, Inc., Newbury Park, California.
- HRYSHKO, D. (2008): “RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles,” Mimeo, University of Alberta.
- HUBBARD, R. G., J. SKINNER, AND S. P. ZELDES (1994): “The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving,” *Carnegie-Rochester Conference Series on Public Policy*, 40, 59–125.
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): “Earnings Losses of Displaced Workers,” *American Economic Review*, 83, 685–709.
- KAPLAN, G., AND G. L. VIOLANTE (2010): “How Much Consumption Insurance Beyond Self-Insurance?,” *American Economic Journal: Macroeconomics*, 2, 53–87.
- KAUFMANN, K., AND L. PISTAFERRI (2009): “Disentangling Insurance and Information in Intertemporal Consumption Choices,” *American Economic Review: Papers and Proceedings*, 99, 387–392.

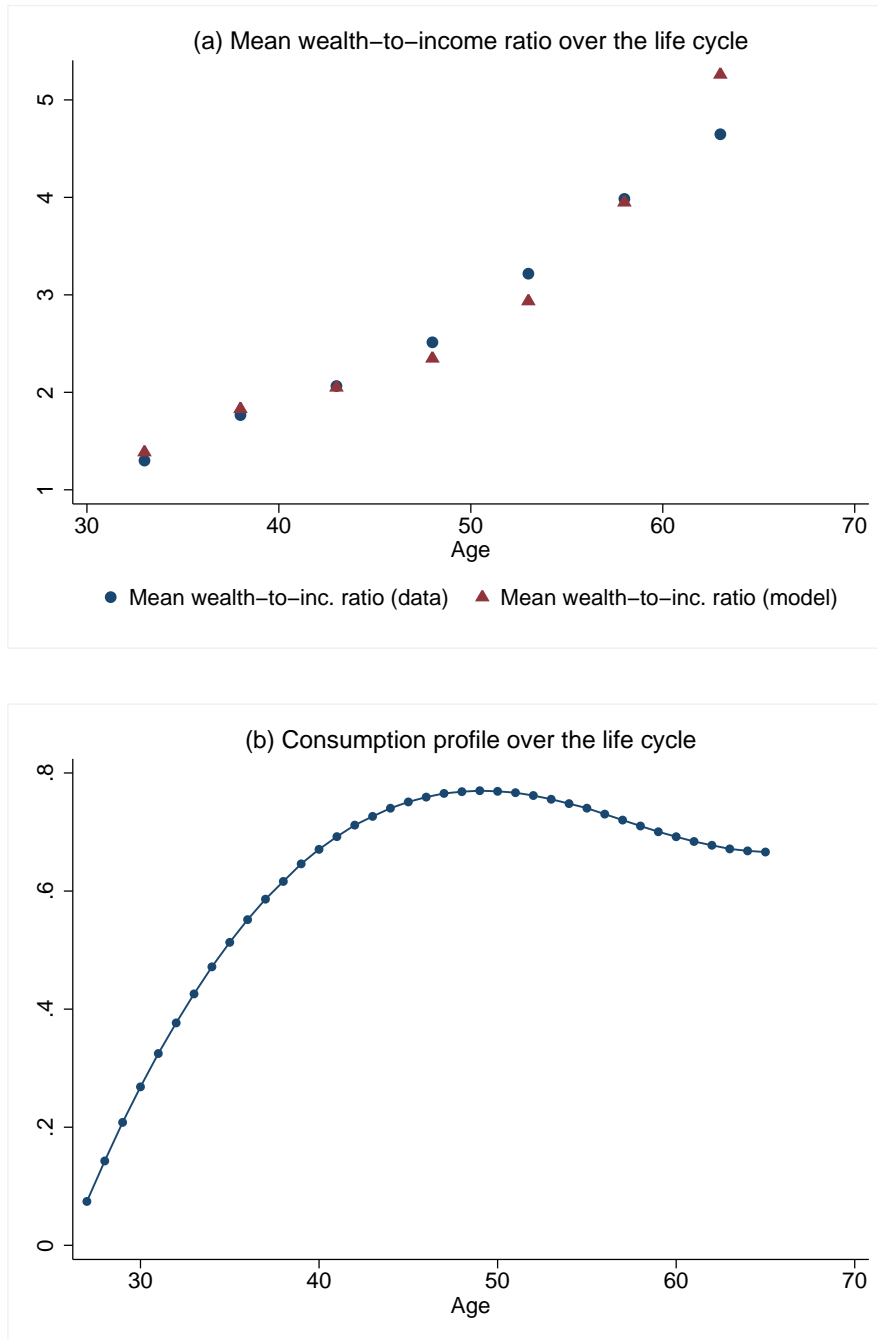
- KREBS, T. (2003): “Human Capital Risk and Economic Growth,” *Quarterly Journal of Economics*, 118(2), 709–744.
- LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2007): “Estimating Discount Functions with Consumption Choices Over the Lifecycle,” NBER Working Paper # 13314.
- LUDVIGSON, S., AND A. MICHAELIDES (2001): “Can Buffer Stock Saving Explain the Consumption Excesses?,” *American Economic Review*, 91, 631–647.
- LUENGO-PRADO, M. (2007): “Durables, Nondurables, Down Payments, and Consumption Excesses,” *Journal of Monetary Economics*, 53, 1509–1539.
- LUENGO-PRADO, M., AND B. SØRENSEN (2008): “What Can Explain Excess Smoothness and Sensitivity of State-Level Consumption?,” *Review of Economics and Statistics*, 91, 65–80.
- MACURDY, T. E. (1982): “The Use of Time-Series Processes to Model the Error Structure of Earnings in Longitudinal Data analysis,” *Journal of Econometrics*, 18, 83–114.
- MEGHIR, C., AND L. PISTAFERRI (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72, 1–32.
- MICHAELIDES, A. (2001): “Buffer Stock Saving and Habit Formation,” Mimeo, London School of Economics.
- NEWKEY, W. K. (1985): “Generalized Method of Moments Specification Testing,” *Journal of Econometrics*, 29, 229–256.
- NEWKEY, W. K., AND D. MCFADDEN (1994): “Large Sample Estimation and Hypotheses Testing,” in *Handbook of Econometrics, Vol.4*, chap. 36. Elsevier, Amsterdam.
- NIELSEN, H. S., AND A. VISSING-JØRGENSEN (2006): “The Impact of Labor Income Risk on Educational Choices: Estimates and Implied Risk Aversion,” Mimeo, Kellogg School of Management.
- PISCHKE, J.-S. (1995): “Individual Income, Incomplete Information and Aggregate Consumption,” *Econometrica*, 63, 805–840.
- PISTAFERRI, L. (2001): “Superior Information, Income Shocks, and the Permanent Income Hypothesis,” *Review of Economics and Statistics*, 83, 465–476.
- QUAH, D. (1990): “Permanent and Transitory Movements in Labor Income: An Explanation for “Excess Smoothness” in Consumption,” *Journal of Political Economy*, 98, 449–475.
- SCHOLZ, J. C., A. SESHADRI, AND S. KHITATRAKUN (2006): “Are Americans Saving “Optimally” for Retirement?,” *Journal of Political Economy*, 114(4), 607–643.
- SKINNER, J. (1987): “A Superior Measure of Consumption from the Panel Study of Income Dynamics,” *Economics Letters*, 23, 213–216.
- WANG, N. (2004): “Precautionary Saving and Partially Observed Income,” *Journal of Monetary Economics*, 51, 1645–1681.
- ZELDES, S. (1989): “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *Quarterly Journal of Economics*, 104, 275–298.

FIGURE 1: DISPOSABLE INCOME, WEALTH-TO-INCOME RATIO AND ADULT EQUIVALENTS OVER THE LIFE CYCLE



Notes: Panel (a) depicts the life cycle profile of household log disposable income relative to household log disposable income at age 26. The profile is estimated using CEX data. Panel (b) depicts the average number of adult equivalents over the life cycle calculated using PSID data. Panel (c) depicts the mean of wealth-to-income ratio over the life cycle calculated using PSID data.

FIGURE 2: WEALTH-TO-INCOME RATIO AND CONSUMPTION OVER THE LIFE CYCLE



Notes: Panel (a) depicts the mean of wealth-to-income ratio in the data (circles) and in the model of column (3) in Table 6 (diamonds). Panel (b) depicts the life cycle profile of household log consumption relative to household log consumption at age 26 implied by the estimates of the model in column (3) of Table 6.

TABLE 1: TEST OF THE NULL HYPOTHESIS OF ZERO AUTOCOVARANCE IN ALL TIME PERIODS.
HOUSEHOLD DISPOSABLE INCOME DATA FROM THE PSID.

	(1)	(2)	(3)	(4)
Order	Avg. autocov.	Test statistic	DF	p-value
0	0.046	1065.17	16	0.00
1	-0.013	272.17	15	0.00
2	-0.002	38.01	14	0.00
3	-0.001	9.95	13	0.70
4	0.00007	11.61	12	0.48

Notes: I run cross-sectional regressions of the first difference in household disposable income on a quadratic polynomial in the head's age and education dummies. The residuals from those regressions are labeled the idiosyncratic growth in household disposable income. I use data from the 1981–1997 surveys of the PSID. Column (1) reports the average autocovariance in income growth rates of a given order across all individuals and survey years. The test statistic tabulated in column (2) is distributed as χ^2 with degrees of freedom equal to the number of (zero) restrictions (the number of unique autocovariances of a given order in the estimated variance-covariance matrix).

TABLE 2: ESTIMATES OF THE REDUCED FORM INCOME PROCESS. HOUSEHOLD DISPOSABLE INCOME
DATA FROM THE PSID.

Moving average parameter, $\hat{\theta}$	-0.30 (0.016)
Variance of the innovation, $\hat{\sigma}_u^2$	0.041 (0.001)
Degrees of freedom	134
χ^2	274.73
p-value of the model	0.00

Notes: The estimated process is $\Delta \log y_{it} = (1 + \theta L)u_{it}$, where $\Delta \log y_{it}$ denotes the change in idiosyncratic household log disposable income, θ is a moving average parameter, and $u_{it} \sim iid(0, \sigma_u^2)$. The model is estimated by the equally weighted minimum distance method. Standard errors in parentheses. The results are based on PSID data from the 1981–1997 surveys.

TABLE 3: TEST OF THE NULL HYPOTHESIS OF ZERO AUTOCOVARANCE IN ALL TIME PERIODS.
TOTAL IMPUTED CONSUMPTION DATA FROM THE PSID.

	(1)	(2)	(3)	(4)
Order	Avg. autocov.	Test statistic	DF	p-value
0	0.193	1514.22	13	0.00
1	-0.083	588.45	11	0.00
2	-0.0038	7.41	9	0.59
3	-0.00002	6.08	7	0.53
4	-0.0002	2.43	6	0.88

Notes: I run cross-sectional regressions of the first difference in total (imputed) household consumption on a quadratic polynomial in the head's age and the change in family size. The residuals from those regressions are labeled the idiosyncratic growth in household total consumption. I use data from the 1981–1997 surveys of the PSID. Consumption data for the 1988 and 1989 surveys are not available since the PSID did not collect food consumption data in those years.

TABLE 4: ESTIMATED MOMENTS FROM A SIMULATED MODEL

	$\overline{W/Y}_{33}$	$\overline{W/Y}_{63}$	$\hat{\beta}_1$	$\hat{\beta}_4$	$\hat{\alpha}$	$\text{var}(\Delta \log y_{it})$
<u>Income parameters</u>						
(1) $\rho = 6, \beta = 0.85$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$	0.92 (0.34)	4.29 (0.84)	0.212 (0.019)	0.565 (0.029)	-0.274 (0.002)	0.045 (0.0002)
(2) $\rho = 6, \beta = 0.85$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = 0.0$	1.07 (0.31)	4.60 (0.76)	0.398 (0.019)	0.652 (0.028)	-0.274 (0.002)	0.045 (0.0002)
(3) $\rho = 6, \beta = 0.85$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = 0.5$	1.19 (0.36)	4.85 (0.88)	0.491 (0.028)	0.694 (0.036)	-0.274 (0.002)	0.045 (0.0002)
<u>Risk aversion</u>						
(4) $\rho = 2, \beta = 0.85$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$	0.05 (0.008)	0.32 (0.05)	0.605 (0.031)	0.843 (0.017)	-0.274 (0.002)	0.045 (0.0001)
(5) $\rho = 10, \beta = 0.85$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$	2.28 (0.32)	8.88 (0.99)	0.172 (0.008)	0.470 (0.021)	-0.274 (0.002)	0.045 (0.0002)
<u>Time discount factor</u>						
(6) $\rho = 6, \beta = 0.75$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$	0.25 (0.15)	1.74 (0.40)	0.239 (0.049)	0.363 (0.032)	-0.274 (0.002)	0.045 (0.0002)
(7) $\rho = 6, \beta = 0.95$ $\text{corr}(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$	2.01 (0.41)	9.96 (1.33)	0.10 (0.006)	0.242 (0.012)	-0.274 (0.002)	0.045 (0.0002)
<u>Partial risk sharing</u>						
(8) $\rho = 6, \beta = 0.85$ $\omega_P = \omega_T = 0.5$	0.93 (0.33)	4.54 (0.84)	0.103 (0.008)	0.281 (0.013)	-0.274 (0.002)	0.045 (0.0002)
(9) $\rho = 6, \beta = 0.85$ $\omega_P = 0.5, \omega_T = 1$	0.97 (0.41)	4.71 (1.07)	0.124 (0.020)	0.280 (0.021)	-0.274 (0.002)	0.045 (0.0002)
(10) $\rho = 6, \beta = 0.85$ $\omega_P = 1, \omega_T = 0.5$	0.85 (0.36)	4.08 (0.88)	0.196 (0.014)	0.572 (0.028)	-0.274 (0.002)	0.045 (0.0002)

Notes: $\hat{\beta}_k$ is the slope coefficient from an OLS regression $\Delta_k \log c_{it} = \beta_0 + \beta_k \Delta_k \log y_{it} + \epsilon_{it}$; where $\Delta_k \log c_{it} \equiv \log c_{it} - \log c_{it-k}$, $\Delta_k \log y_{it} \equiv \log y_{it} - \log y_{it-k}$ and $k = 1, \dots, 4$. Income persistence is the slope coefficient from an OLS regression $\Delta \log y_{it} = \alpha_0 + \alpha \Delta \log y_{it-1} + \epsilon_{it}$. Each simulated economy is populated by 5,000 ex ante identical consumers observed at ages 26–65. $\overline{W/Y}_{33} = \frac{1}{5} \sum_{a=31}^{35} \frac{1}{5,000} \sum_{i=1}^{5,000} (W/Y)_{ia}$, $\overline{W/Y}_{63} = \frac{1}{5} \sum_{a=61}^{65} \frac{1}{5,000} \sum_{i=1}^{5,000} (W/Y)_{ia}$ where $(W/Y)_{ia}$ is individual i 's wealth-to-income ratio at age a . The results are the averages of the corresponding statistics over 100 model repetitions. Standard errors, calculated as the standard deviations of the corresponding statistics over 100 repetitions, in parentheses. In all models, the standard deviation of permanent shocks is 0.1426; the standard deviation of transitory shocks is 0.1525 when permanent and transitory shocks are negatively correlated, 0.0812 when the shocks are positively correlated, 0.1112 when the shocks are uncorrelated. The parameters of the reduced-form income process are taken from Table 2. When partial risk sharing is allowed, idiosyncratic “after-risk-sharing” labor income is modeled as $Y_{it} = P_{it}(\epsilon_{it}^T)^{\omega_T}$, where $P_{it} = G_t P_{it-1}(\epsilon_{it}^P)^{\omega_P}$, and $1 - \omega_T$ ($1 - \omega_P$) is the fraction of transitory (permanent) shocks that is insured.

TABLE 5: MOMENTS USED FOR FITTING THE MODEL

	stat.	s.e.	no. obs.
Average wealth-to-income			
(0) Ages 26–30	0.86	0.074	925
(1) Ages 31–35	1.30	0.063	1,178
(2) Ages 36–40	1.77	0.065	1,134
(3) Ages 41–45	2.06	0.079	1,049
(4) Ages 46–50	2.51	0.102	826
(5) Ages 51–55	3.22	0.152	662
(6) Ages 56–60	3.98	0.233	413
(7) Ages 61–65	4.65	0.320	223
Consumption smoothness moments			
$\Delta \log c_{it} = \beta_0 + \beta_1 \Delta \log y_{it} + \gamma' x_{it} + \epsilon_{it}$			
(8) $\hat{\beta}_1$	0.119	0.027	1,905
$\Delta_2 \log c_{it} = \beta_0 + \beta_2 \Delta_2 \log y_{it} + \gamma' x_{it} + \epsilon_{it}$			
(9) $\hat{\beta}_2$	0.176	0.029	1,905
$\Delta_3 \log c_{it} = \beta_0 + \beta_3 \Delta_3 \log y_{it} + \gamma' x_{it} + \epsilon_{it}$			
(10) $\hat{\beta}_3$	0.198	0.031	1,905
$\Delta_4 \log c_{it} = \beta_0 + \beta_4 \Delta_4 \log y_{it} + \gamma' x_{it} + \epsilon_{it}$			
(11) $\hat{\beta}_4$	0.227	0.03	1,357
(12) $\hat{\phi}_P: 1 - [E(\epsilon_{it}^c \epsilon_{it}^y) + E(\epsilon_{it}^c \epsilon_{it+1}^y)] / [E(\epsilon_{it}^y)^2 + 2E(\epsilon_{it}^y \epsilon_{it+1}^y)]$	0.397	0.073	1,714
(13) $\hat{\phi}_T: 1 - E(\epsilon_{it}^c \epsilon_{it+1}^y) / E(\epsilon_{it}^y \epsilon_{it+1}^y)$	1.003	0.082	1,714
Income moments			
(14) $\text{var}(\Delta \log y_{it})$	0.048	0.002	1,971
$\Delta \log y_{it} = \alpha_0 + \alpha \Delta \log y_{it-1} + \gamma' x_{it} + \epsilon_{it}$			
(15) $\hat{\alpha}$	-0.298	0.014	1,768

Notes: Wealth data are from the 1984, 1989, 1994, 1999, 2001 and 2003 wealth supplements of the PSID. The mean wealth-to-income ratios are the coefficients from a panel regression of the wealth-to-income ratio on the set of dummies for the age categories listed in the table. The regression controls for household fixed effects. Total consumption is imputed to the sample PSID households. Data for household disposable income and consumption are from the 1981–1997 surveys of the PSID. In the second panel, x_{it} includes a quadratic polynomial in the head's age, and family size. I subtract the time-specific mean from each variable in the regressions (e.g., $\Delta \log y_{it} = \Delta \log Y_{it} - \Delta \overline{\log Y}_t$, where $\Delta \log Y_{it}$ is the difference in household log disposable income before the time-specific average, $\Delta \overline{\log Y}_t$, is taken out). In the last-panel regression, x_{it} includes a quadratic polynomial in the head's age, and dummies for high school and college. $\text{var}(\Delta \log y_{it})$ is calculated as the unconditional variance of residuals from the cross-sectional regressions of the difference in household log disposable income on a quadratic polynomial in the head's age, and dummies for high school and college. Standard errors are calculated by bootstrap.

TABLE 6: ESTIMATION RESULTS

	(1)	(2)	(3)
Time Discount Factor, $\hat{\beta}$	0.806 (0.061)	0.835 (0.059)	0.811 (0.045)
CRRA, $\hat{\rho}$	6.761 (1.301)	6.426 (1.660)	6.63 (1.044)
$\hat{\sigma}_P$	0.137 (0.006)	0.139 (0.012)	0.139 (0.006)
$\hat{\sigma}_T$	0.120 (0.005)	0.168 (0.012)	0.151 (0.014)
$\widehat{\text{corr}}(u_{it}^P, u_{it}^T)$	0 —	-0.607 (0.069)	-0.407 (0.162)
Risk sharing param., $\hat{\omega}_P$	1 —	1 —	0.476 (0.075)
χ^2 [d.f.]	23.05 [6]	8.12 [5]	37.20 [9]
Moments used for matching	(1)–(8), (14), (15)	(1)–(8), (14), (15)	(1)–(15)
Implied moments			
$\hat{\beta}_1^s$ [0.119]	0.334	0.119	0.118
$\hat{\beta}_4^s$ [0.227]	0.532	0.474	0.247
$\hat{\alpha}^s$ [-0.298]	-0.302	-0.298	-0.299
$\text{var}(\Delta \log y_{it}^s)$ [0.048]	0.048	0.048	0.048
$\hat{\phi}_P^s$ [0.397]	0.198	0.214	0.626
$\hat{\phi}_T^s$ [1.003]	0.974	1.70	1.14

Notes: Standard errors in parentheses. Idiosyncratic “after-risk-sharing” labor income is modeled as $Y_{it} = P_{it}\epsilon_{it}^T$, where $P_{it} = G_t P_{it-1} (\epsilon_{it}^P)^{\omega_P}$, and $(1 - \omega_P)$ is the fraction of permanent shocks that is insured before households make their savings decisions. In the panel labeled “Moments,” in square brackets I report the data estimates of the moments used for fitting the model—see Table 5. The goodness-of-fit statistic is based on Newey (1985).

TABLE 7: ESTIMATION RESULTS. MEASUREMENT ERROR IN INCOME

	(1)	(2)	(3)	(4)
Time Discount Factor, $\hat{\beta}$	0.776 (0.07)	0.835 (0.059)	0.812 (0.058)	0.832 (0.056)
CRRA, $\hat{\rho}$	7.462 (1.467)	6.489 (1.682)	6.615 (1.295)	6.187 (1.324)
$\hat{\sigma}_P$	0.137 (0.006)	0.138 (0.014)	0.138 (0.006)	0.138 (0.006)
$\hat{\sigma}_T$	0.092 (0.006)	0.146 (0.014)	0.120 (0.018)	0.124 (0.015)
$\widehat{\text{corr}}(u_{it}^P, u_{it}^T)$	0 —	-0.644 (0.079)	-0.367 (0.208)	-0.407 (0.179)
Risk sharing param., $\hat{\omega}_P$	1 —	1 —	0.473 (0.079)	0.527 (0.09)
χ^2 [d.f.]	21.82 [6]	8.03 [5]	28.34 [9]	23.48 [2]
Moments used for matching	(1)–(8) (14), (15)	(1)–(8) (14), (15)	(1)–(15)	avg. (1)–(3), avg. (4)–(7) (8), (11)–(15)
Implied moments				
$\hat{\beta}_1^s$ [0.119]	0.313	0.119	0.118	0.122
$\hat{\beta}_4^s$ [0.227]	0.517	0.473	0.250	0.276
$\hat{\alpha}^s$ [-0.298]	-0.302	-0.298	-0.299	-0.299
$\text{var}(\Delta \log y_{it}^s)$ [0.048]	0.048	0.048	0.048	0.048
$\hat{\phi}_P^s$ [0.397]	0.220	0.403	0.610	0.572
$\hat{\phi}_T^s$ [1.003]	1.01	1.67	1.13	1.17

Notes: See notes to Table 6. I assume that measurement error in household disposable income contributes 25% towards the observed variance of the growth in household disposable income.

Appendix A. Matching the Moments of the Reduced Form and Structural Income Models

As in the text, suppose that the first difference of log income is described by the reduced form MA(1) process. Thus, the unique moments of the reduced form (rf) are the first and zero-order autocovariances:

$$\begin{aligned}\gamma^{rf}(0) &= (1 + \theta^2)\sigma_u^2, \\ \gamma^{rf}(1) &= \theta\sigma_u^2.\end{aligned}$$

For a structural model (sf) with correlated permanent and transitory components described in the text, I need to match

$$\begin{aligned}\gamma^{sf}(0) &= \sigma_{u^P}^2 + 2\sigma_{u^T}^2 + 2\text{cov}(u_{it}^P, u_{it}^T) = \gamma^{rf}(0), \\ \gamma^{sf}(1) &= -\sigma_{u^T}^2 - \text{cov}(u_{it}^P, u_{it}^T) = \gamma^{rf}(1),\end{aligned}$$

subject to the constraint that the spectrum of the reduced form series and the permanent component at frequency zero are equal:

$$\sigma_{u^P}^2 = \gamma^{rf}(0) + 2\gamma^{rf}(1) = (1 + \theta)^2\sigma_u^2.$$

The above equation is the identifying condition for the variance of the permanent component. The two preceding equations determine the other two unknowns, $\text{cov}(u_{it}^P, u_{it}^T)$ and σ_{u^T} . Since the two equations have three unknowns, there is no unique solution for $\text{cov}(u_{it}^P, u_{it}^T)$ and σ_{u^T} . Thus, I choose the grid of covariances between the shocks such that they return the correlation that is less than or equal to one in absolute value. This procedure uniquely determines the variance of the transitory innovation from the following condition:

$$\sigma_{u^T}^2 = -\gamma^{rf}(1) - \text{cov}(u_{it}^P, u_{it}^T).$$

The estimated parameters of the reduced form process for the growth in idiosyncratic income are in Table 2. The structural income processes are of the form: $\Delta \log y_{it} = u_{it}^P + \Delta u_{it}^T$, with potentially correlated permanent and transitory shocks.

Appendix B. Numerical Procedures

In this appendix, I lay out the details of model solution, and describe the choices made in numerical analysis.

Model Solution

At age 90, the last period of the problem, $C_T = X_T$, where C_T and X_T are household consumption and cash-on-hand at age T , respectively. At each age of the life cycle, the Euler equation linking consumption in adjacent periods is $n_t u' \left(\frac{C_t}{n_t} \right) = \beta R n_{t+1} E_t u' \left(\frac{C_{t+1} [R(X_t - C_t) + Y_{t+1}]}{n_{t+1}} \right)$, where expectation is taken with respect to future mortality risk at ages 65, ..., 89 and labor income risk at ages 26, ..., 64. Given that $C_{t+1}(X_{t+1})$ is known, the Euler equation provides the unique solution for C_t given X_t . I express the problem in terms of consumption and cash-on-hand relative to the permanent component of income (c_t and x_t , respectively).

I solve the problem using the endogenous gridpoints method in Carroll (2006). I create a grid of 120 points in $s_t = x_t - c_t$, and a 250×2 matrix of pre-seeded independent random normal deviates ϵ_1 ,

and ϵ_0 . The deviates are used to construct sequences of correlated shocks u^P and u^T , permanent and transitory shocks respectively. For ages 26–65, when income uncertainty is relevant, $E_t(\cdot)$ is calculated as $\frac{1}{250} \sum_{j=1}^{250} \left[c_{t+1} \left(\frac{R s_t}{G_{t+1} u_j^P} + u_j^T \right) \right]^{-\rho} [G_{t+1} u_j^P]^{-\rho}$ for each value of s_t from the grid, and each pair of (u_j^T, u_j^P) from the joint distribution of random log-normal shocks. Each value of s_t provides the choice of c_t , x_t is then calculated as the sum of s_t and the corresponding choice of consumption, c_t . The consumption function at age t is linearly interpolated, using the points from vectors x_t and c_t . I use extrapolation if x_t falls outside an upper bound of the grid.

I create a matrix of 5000×2 pre-seeded random draws for permanent and transitory shocks, and simulate the economy populated by 5,000 households, each with unique income history, using the age-dependent consumption functions $\{c_t(x_{it})\}_{t=26}^{65}$. In the beginning period of the life cycle, at age 26, each household obtains a draw of the transitory shock. To initialize the permanent income, for each household I take a draw from the estimated distribution of household disposable log income at age 26. Once the economy is generated, I calculate the model moments used in optimization. During optimization, initially drawn sequences of random shocks are kept the same.

MSM Standard Errors and the Goodness-of-fit Test

Denote the vector of moments used for matching as m^d , and simulated moments as $m^s(\theta; \hat{\chi})$, where superscripts d and s denote data and simulation, respectively, and $\hat{\chi}$ is a vector of preestimated first-stage moments. Let $g_{I_s} = \log m^d - \log m^s(\theta; \hat{\chi})$ be the distance between the data and simulated moments.

The MSM estimate $\hat{\theta}$ is the minimizer of the criterion function:

$$g_{I_s}(\theta, \hat{\chi})' W g_{I_s}(\theta, \hat{\chi}),$$

where W is a weighting matrix.

Following Newey and McFadden (1994), Laibson, Repetto, and Tobacman (2007), and Gourinchas and Parker (2002), it can be shown that

$$\sqrt{I_d}(\hat{\theta} - \theta_0) \rightarrow N(0, V_\theta),$$

where

$$V_\theta = [G'_\theta W G_\theta]^{-1} G'_\theta W \left[\Sigma_g \left(1 + \frac{I_d}{I_s} \right) + \frac{I_d}{I_1} G_\chi \Sigma_\chi G'_\chi \right] W G_\theta [G'_\theta W G_\theta]^{-1'}$$

$G_\theta = g_{I_s \theta}(\theta; \hat{\chi})$ is a vector of the first partial derivatives of $g_{I_s}(\theta; \hat{\chi})$ with respect to θ ; G_χ is a vector of the first derivatives of the (logged) moments with respect to the first-stage parameters χ ; I_d is the number of households contributing towards estimation of the matching moments; I_s is the number of simulated households used for estimation of the model moments; Σ_g is the variance-covariance matrix of the (logged) matching moments estimated from the data; and Σ_χ is the variance-covariance matrix of the first-stage moments.

The standard errors of $\hat{\theta}$ are calculated as the square roots of the diagonal elements of

$$\text{var}(\hat{\theta}) = \frac{V_\theta}{I_d} = [G'_\theta W G_\theta]^{-1} G'_\theta W \left[\Sigma_g \left(\frac{1}{I_d} + \frac{1}{I_s} \right) + \frac{1}{I_1} G_\chi \Sigma_\chi G'_\chi \right] W G_\theta [G'_\theta W G_\theta]^{-1'}$$

The overidentifying restrictions test statistic, distributed as χ^2 with degrees of freedom equal to $[\dim(g_{I_s}) - \dim(\theta)]$, is calculated as

$$g'_{I_s} R^- g_{I_s},$$

where R^- is a generalized inverse (calculated as the Moore-Penrose inverse) of the matrix $R = P \frac{V_\theta}{I_d} P'$, $P = I - G_\theta(G'_\theta W G_\theta)G'_\theta W$ (see Newey 1985).

Appendix C. Data Used and Sample Selection

The Consumer Expenditure Survey

I use CEX data on total consumer expenditures and food consumption, available at the NBER website. The data set spans the period 1980–2003. The CEX is designed by the Bureau of Labor Statistics to construct the CPI at different levels of aggregation. The survey publishes at most four quarters of information on individual consumption, along with demographic information. The NBER extracts lump quarterly records into one annual record.

Total consumption is defined as household total expenditures on nondurables (food at home, food at work, food away from home, clothes, personal care items, utility payments, transportation including gasoline and insurance, recreation services, gambling and charity), household supplies and equipment, rents, medical services, vehicles and auto parts, books and publications, education, interest payments, housing property taxes, contributions to private pensions and self-employed retirement. The measure is defined similarly to Blundell, Pistaferri, and Preston (2008).

Households may enter the survey in the same year but in different quarters and months of a quarter. If household consumption record relates both to years t and $t+1$, I assume that annual consumption refers to year t if that year contains at least six months of consumption records, and to year $t+1$ otherwise. In the CEX, the head of a household is the person who owns or rents the unit of household residence. In the PSID, the head of a household is male, unless he is permanently disabled (see Hill 1992). To make the definitions of heads comparable, I assume that heads are males in the CEX families with couples.

My sample selection steps are the following. First, I keep only the households that report expenditures in all four quarters of the year, and who are classified as full income reporters. I further keep married couples, with heads who are not self-employed, of age above 25 but below 71; and with family size equal or greater than 2 family members. I drop households whose heads attend college, part-time or full-time; or do not have education, race, age, or state of residence records. I also drop households that report zero expenditures on food at home. I deflate food at home, food away from home, total food as well as total household expenditures by the Bureau of Labor Statistics (BLS) food at home, food away, total food and all items CPI with the base 100 in 1982–1984, respectively. I use non-seasonally adjusted, U.S. city average CPI indices for March of the respective year. I further drop households whose expenditures on food at home and away from home exceed household disposable income. Household disposable income is defined as the difference between household taxable income and federal income taxes. Household taxable income, in turn, is the sum of wages and salaries, rents, dividends and interest, business and farm income, pensions, social security benefits, supplemental security income, unemployment compensation, and workers' compensation for all family members. To eliminate the influence of outliers, I drop observations below the 1st and above the 99th percentiles of the annual total food distributions.

I utilize information on 23,133 households whose heads are of ages 26–65 to estimate the relationship between food consumption and total consumption. The relationship is later used for imputation of the total consumption to the sample PSID households. I use income information for households whose heads are of ages 26–70 to estimate the deterministic profile of life-cycle growth in household disposable income.

The Panel Study of Income Dynamics

I use PSID data from the 1981–2003 waves, the same time span during which I have data on consumption and demographics in the CEX. To allow for a more representative sample, I drop households that are part of the Survey of Economic Opportunity (SEO) subsample. The PSID consistently collected only two items of consumption over time: food consumption at home and food consumption away from home (excluding food at work). Since I am interested in *total* household consumption and income dynamics over the life cycle, I impute *total* consumption to the PSID households. Most of the studies that use food consumption from the PSID assume that food consumption recorded in survey year t reflects the typical weekly food consumption flow in year $t-1$. In this paper, I adopt the same strategy. Over the time span

considered, the PSID did not collect food consumption data in 1988 and 1989. Correspondingly, my final sample of analysis lacks food and total consumption data for 1987 and 1988. Food away from home and food at home are deflated by their respective CPIs taken from the BLS.

The sample selection details are as follows. I first make the age series consistent throughout the survey years. An individual's age in adjacent years can be the same or differ by more than one year since households may be interviewed in different months of a year. I take the first record on age when an individual appears as the head of household and impute age in other years using that record. I then drop households whose heads are younger than 26 or older than 70, and all female-headed households. I then choose continuously married heads with the longest spell of an uninterrupted headship. I further drop households whose heads had a spell of self-employment, and keep those who never experienced significant changes in family composition, that is those who had no changes at all or changes in members other than the head or wife. I drop households with inconsistent race records (e.g., households whose heads report being white in one survey year and black in some other year), and households with heads who become permanently disabled or continue schooling after age 26. I set head's education to the maximum years of schooling reported by the head during sample years. Heads are then assigned to three educational categories: high school dropouts (with years of schooling below 12), high school graduates (with years of schooling above 11 but less than 16 years), and college graduates (with years of schooling equal to or exceeding 16 years). I further drop observations with missing records on head's or household disposable income; with zero records on head's income when heads are of ages 26 to 65, at their pre-retirement stage of the life cycle. I also drop income outliers and single-headed households. An income outlier is defined as an observation on household disposable income at time t when the growth rate in household disposable income between periods $t - 1$ and t is above 500 percent or below -80 percent. I then select, for each household, the longest spell of an uninterrupted headship with all relevant information. I keep those who were born between 1912 and 1978, and whose total food expenditures do not exceed household disposable income. I also drop observations with household food expenditures above the 99th or below the 1st percentiles of the annual food distributions. Similar to my selection rules for the CEX sample, I drop observations with expenditures on food at home equal to zero. In the PSID, income recorded in year t refers to income earned in year $t - 1$. Conformably with income observations, I assume that demographic information and household food expenditures recorded in year t refer to a previous year.

Before 1994, the PSID recorded annual household expenditures on food at home or away from home. In 1994, the PSID started recording food at home and food away from home at different frequencies—daily, weekly, biweekly, monthly, or annual frequency. For those years, I use household food records at the monthly or weekly frequency. I lose a small number of observations on household food expenditures reported at other frequencies. Most of them, when converted to annual amounts, were clear outliers.

The PSID estimated household federal income taxes for 1980–1991. Starting in 1992, the PSID discontinued calculation of federal income taxes. Since my data extend well beyond 1992, I use the TAXSIM tool at the NBER to calculate a consistent series of household federal income taxes and social security withholdings. I assume that family members other than the head and wife are filing their tax returns separately. Household disposable income is then calculated as the sum of the head's and wife's labor income, their combined transfer income, transfer income of all other family members, taxable income of other family members less federal income taxes and social security withholdings for the head and wife and all other family members. I also add the total family social security income for 1994–2003 to the measure of disposable income since the records on head's and wife's transfer income and transfer income of other family members exclude social security income in those years.

For the households in my final sample, I use wealth information from the 1984, 1989, 1994, 1999, 2001 and 2003 wealth supplements. The wealth measure utilized is the total household net worth exclusive of household business wealth.

Imputation of Consumption

In this appendix, I describe the procedure used to impute total consumption to the PSID households. Absent data on total consumption in the PSID, imputation is usually done in order to exploit the panel structure of the PSID. Skinner (1987), using the 1972–1973 and 1983 waves of the CEX, showed that total household consumption tightly relates to several consumption items, also available in the PSID

(food at home and away from home, number of vehicles owned, and housing rent). Moreover, he showed that this relationship is stable over time. Blundell, Pistaferri, and Preston (2005) pioneered a structural approach to imputation—inverting the food demand equation estimated on CEX data. They relate food consumption to nondurable expenditures, household demographics, price indices, time dummies, cohort dummies, and nondurable expenditures interacted with time dummies and the head’s education category. I run a similar regression, and use the coefficients from this regression to impute total consumption to the PSID households. In the OLS setting, the estimated elasticities may be biased due to measurement error in total expenditures, and endogeneity of food and total consumption. I therefore follow Blundell, Pistaferri, and Preston (2008), and instrument log total expenditures (and its interactions with year and education dummies) with the head’s and wife’s education-year-cohort specific averages of log hourly wages (and their interactions with year and education dummies). The results of an IV regression of food consumption on total expenditures are presented in Table C-1. The estimated elasticity of food consumption with respect to total expenditures is high in the 1980s, and drops steadily to about 0.60 in 2002.

TABLE C-1: IV REGRESSION OF FOOD EXPENDITURES ON TOTAL EXPENDITURES. CEX DATA:
1980–2002

Log total cons.	0.595*** (10.15)	Log total cons. x HS	0.025 (0.86)
Log total cons. x 1980	0.111*** (5.75)	Log total cons. x coll.	0.114*** (3.56)
Log total cons. x 1981	0.093*** (5.51)	Log food CPI	0.320 (1.28)
Log total cons. x 1982	0.081*** (5.19)	Log fuel-util. CPI	0.099 (1.51)
Log total cons. x 1983	0.076*** (5.08)	White	0.053*** (5.18)
Log total cons. x 1984	0.069*** (5.00)	Fam. size	0.073*** (37.05)
Log total cons. x 1985	0.067*** (5.06)	HS	−0.302 (−1.02)
Log total cons. x 1986	0.059*** (4.83)	Coll.	−1.288*** (−3.94)
Log total cons. x 1987	0.055*** (4.92)	Age	0.028*** (6.71)
Log total cons. x 1988	0.057*** (5.61)	Age sq./100	−0.024*** (−5.74)
Log total cons. x 1989	0.052*** (5.93)	Born 1924–1932	0.018 (1.12)
Log total cons. x 1990	0.046*** (6.21)	Born 1933–1941	0.025 (1.23)
Log total cons. x 1991	0.039*** (5.80)	Born 1942–1950	0.027 (1.06)
Log total cons. x 1992	0.036*** (5.71)	Born 1951–1959	0.033 (1.03)
Log total cons. x 1993	0.033*** (5.64)	Born 1960–1968	0.053 (1.41)
Log total cons. x 1994	0.031*** (5.83)	Born 1969–1978	0.082* (1.83)
Log total cons. x 1995	0.026*** (5.74)	Northeast	0.049*** (4.96)
Log total cons. x 1996	0.019*** (5.03)	Midwest	−0.008 (−0.84)
Log total cons. x 1997	0.020*** (6.41)	South	0.002 (0.25)
Log total cons. x 1998	0.015*** (5.56)	Adj. R sq.	0.408
Log total cons. x 1999	0.012*** (5.37)	F	216.4
Log total cons. x 2000	0.010*** (5.99)	N	23,133
Log total cons. x 2001	0.004*** (3.37)		

Notes: t-statistics in parentheses. Instruments for log total consumption (and its interaction with year and education dummies) are the averages of log head's and wife's wages specific to cohort, education, and year (and their interactions with year and education dummies). *** significant at the 1% level, ** significant at the 5% level, * significant at the 10% level.