

Identifying Household Income Processes Using a Life Cycle Model of Consumption *

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 simulations of a life cycle model of consumption, I first show that households' information on (potentially correlated) individual income components can be revealed in the sensitivity of consumption growth to income growth. I further use a structural life cycle model of consumption to estimate the parameters of the income process by the method of simulated moments. I find significant negative contemporaneous correlation between permanent and transitory shocks, and reasonable, precise estimates for the time discount factor and the relative risk aversion parameter.

KEYWORDS: Unobserved components models, income processes, buffer stock model of savings, method of simulated moments, consumption dynamics, life cycle.

*I owe special thanks to Bent Sørensen for extensive comments, advice, and encouragement. I am also grateful to Chris Carroll, Eric French, Alex Michaelides, Maria Luengo-Prado, Chris Murray, and Nat Wilcox. For useful comments, I thank seminar participants at the Universities of Alberta and Connecticut, the EERC in Kiev, New Economic School, conference participants at the 2006 NBER Summer Institute, the 2006 SED Meetings in Vancouver, and the 2006 North American Econometric Society Summer Meetings in Minneapolis.

[†]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

Households face a variety of income shocks. Promotions, layoffs, long term and temporary unemployment, health shocks, and lump-sum bonuses are a few in the list of events that make disposable household income volatile. In a world of imperfect insurance markets, idiosyncratic labor income risk is important for household decisions over consumption and savings, portfolio choice, and even the choice of career. Economists refer to persistent, or long lasting, shocks as permanent and temporary, or short-lived, shocks as transitory. 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 Obviously, households may have better information about (potentially correlated) income components, and therefore about the stochastic processes that govern the dynamics of each component.² In this paper, I use a structural life 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. 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. The permanent income hypothesis (PIH), for example, predicts that households adjust consumption fully to the newly arrived permanent shocks, and change consumption only by the annuity value of the transitory shocks, a very small adjustment in economic terms. 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.³ Thus, Quah (1990) implicitly shows that the “correct” decompo-

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, identify and estimate household or group-specific volatility of permanent and transitory shocks.

²In general, any structural decomposition of non-stationary income processes is not unique. E.g., Quah (1992) shows how to decompose an integrated time series process into permanent and transitory components of different relative sizes.

³If income is non-stationary and income growth exhibits positive serial correlation—as supported by aggregate data—the 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, though, 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).

sition of income is the one that helps reconcile the joint dynamics of consumption and income with the PIH predictions. 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.

In this paper, I explore an idea similar to that in Quah (1990, 1992) in the context of the buffer stock model of savings. I first simulate a life cycle buffer stock models that only differ in terms of unobserved components (UC) decompositions of the same reduced form income process, and analyze the simulated economies at the household level. Specifically, I estimate different model statistics: the sensitivity of consumption growth to current and lagged income growth; the sensitivity of consumption growth to income growth over the five-year horizon; and the variance and persistence of the reduced form income process. 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 and lagged income, and a lower MPC out of shocks to income cumulated over the five-year horizon. 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.⁴

Having established these results, I suggest that the MPCs estimated from empirical micro data should help 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 processes⁵ and must show how well the or-

⁴See Section 3.3.1 for details.

⁵Zero covariance between permanent and transitory shocks is also typically assumed in the literature modelling income processes at the aggregate level (e.g., Clark (1987)). Reduced form dynamics for aggregate income processes is richer than reduced form dynamics of household income processes, and requires more complicated models for the transitory component. Morley et al. (2003) show that if quarterly US GDP follows ARIMA(2,1,2), permanent component is a random walk, and transitory component is an AR(2), the covariance between permanent and transitory shocks can be identified from the univariate dynamics of GDP. Abowd and Card (1989) and Meghir and Pistaferri (2004), using data from the Panel Study of Income Dynamics, find that the growth rate of the head's labor income can be described by a moving average process of order 1 or 2. If the true order is 1, the data allow identification of at most two parameters—the variance of shocks to the random walk component, and the variance of purely transitory i.i.d. shocks, the covariance between the shocks is not identified. In general, if the

thogonal decomposition of income done by econometricians describes the joint dynamics of household consumption and income. In other words, an estimate of the correlation between structural shocks should reveal the extent to what (income) information sets of econometricians may differ from the ones held by households.

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 consumption profile 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 between transitory and permanent income shocks of about -0.60 , and precise estimates of the time discount factor and the relative risk aversion parameter.⁶

In the literature, it has been shown 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 (e.g., Quah (1990), Pischke (1995), Ludvigson and Michaelides (2001)). In this paper, I find that, using a plausible structure of the model economy, it is possible to (parametrically) identify the unique information held by households from their consumption choices over the life cycle.

Correct identification of the components of the income process improves specification of the consumption function of a life cycle dynamic optimization problem and estimation of the behavioral parameters. For the no-correlation case, I find a substantially larger and less plausible point estimate for the relative risk aversion parameter, and a smaller point estimate for the time discount factor. I can reject the no-correlation model in favor of the model with correlated permanent and transitory shocks at any standard level of statistical significance. Another important result of this paper is that an adequate fit of the joint dynamics of consumption and

growth rate of idiosyncratic household income is described by $MA(q)$, income can be decomposed into a permanent random walk component and a transitory moving average process of order $q - 1$. The auto-covariance structure of income in first differences will feature $q + 1$ unique moments, necessary for identification of the variance of permanent shocks, the variance of transitory shocks, and $q - 1$ moving average parameters; the contemporaneous covariance between permanent and transitory shocks, if present, is not identified from the income data alone.

⁶Friedman (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). This negative co-movement between changes in the base pay and bonuses may represent a compensation smoothing strategy adopted by firms.

income requires the proper identification of the income process. If the covariance between the shocks is ignored, the no-correlation model fits well the relative log-consumption profile over the life cycle but fails to fit the empirical sensitivity of consumption growth to current and lagged income growth and sensitivity of consumption to income growth over the five-year horizon. Thus, the results of this paper can prove to be important for investigations of wealth accumulation and inequality, portfolio choice, career choice, and other fields where correct identification of the income process and behavioral parameters are of prime importance.

The main results of the paper are obtained using a life cycle model with a specific market structure: households can transfer resources inter-temporally using only one asset, the risk-free bond. In reality, however, households may have access to a wider array of assets that can be used to smooth consumption over time and across states of nature. Introducing partial risk sharing of permanent and transitory income shocks into the model, I still find the negative correlation between the shocks to income.

The rest of the paper is organized in the following way. I elaborate further on the main idea in Section 2. In Section 3, first, I estimate an ARIMA(0,1,1) income process using the PSID data; second, I decompose this income process into permanent and transitory components with different correlation between them; third, I simulate a life cycle buffer stock model and present results on the sensitivity of (simulated) consumption to informational assumptions. In Section 4, I lay out the procedure of estimating the income and behavioral parameters by the MSM. In Section 5, I discuss the main results, and relate them to the available literature. In Section 6, I explore robustness of the results to the introduction of partial insurance into the model; Section 7 concludes.

2 Information Sets of Econometricians and Households

In this section, I set up a model of household consumption over the life cycle, present the unobserved components (UC) income model used in the literature, emphasize that households and econometricians may use different UC models that imply different information sets, and discuss the potential importance of different UC models, and therefore (income) information sets for consumption dynamics.

Assume that households value consumption, supply labor inelastically, face income uncertainty over the life cycle, and are subject to liquidity constraints. Households start their life

cycle at period 0, and die at period T . Thus, a household's problem is:

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

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:

$$C_{it} \leq X_{it}, \quad \forall t \in [0, T]. \quad (3)$$

Cash-on-hand available to household i in period $t+1$, X_{it+1} , consists of labor income realized in period $t+1$, Y_{it+1} , and resources brought from previous period, accumulated at a possibly stochastic gross interest rate on a risk-free asset, R_{t+1} . β is the common pure time discount factor, Z_{it} is a vector of household i 's time preference shifters, C_{it+1} is household i 's consumption in period $t+1$, and E_{i0} denotes the expectation of household i based on the information available to it in the beginning of the life cycle.

If preferences are CRRA, and income is stochastic, the consumption problem cannot be solved analytically, and one must rely on computational methods to obtain the consumption function. Under certain regularity conditions on preferences, interest and the growth rate of income, Deaton (1991), in an infinite-horizon setting, has shown that this model generates a buffer stock behavior—a household targets a certain level of liquid wealth to buffer bad income shocks. If shocks to income are unfavorable, households smooth consumption by running down available assets, and gradually rebuild wealth to meet the desired target level. The model is called the buffer stock model; it was originally proposed by Deaton (1991) and later refined by Carroll (1992, 1997). The model proved to fit well consumption facts from micro data (e.g., co-movement of consumption and income, and low household wealth holdings over the life cycle).

In this section, I examine the sensitivity of consumption to informational assumptions on the income processes. A popular, intuitively appealing, and empirically justifiable income model is an unobserved components (UC) model, where household income, Y_{it+1} , consists of a random

walk permanent component, P_{it+1} , and a transitory component, ϵ_{it+1}^T .⁷

$$Y_{it+1} = P_{it+1}\epsilon_{it+1}^T, \quad (4)$$

$$P_{it+1} = G_{it+1}P_{it}\epsilon_{it+1}^P, \quad (5)$$

where ϵ_{it+1}^P is an innovation to the permanent component, and G_{it+1} is the gross growth rate of household i 's income at time $t + 1$.

Taking natural logs, the first difference of income is:⁸

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

where $\log Y_{it+1}$ is household i 's log-income at time $t + 1$; g_{it+1} is the log of its gross growth rate at time $t + 1$; 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_{it+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_{it+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 information about income and its components is generated exactly by this model, households can differentiate between permanent and transitory shocks,

⁷In the context of computational consumption models, this model was first used by Zeldes (1989*b*) and Carroll (1992), Carroll (1997).

⁸For some evidence that idiosyncratic household log-income is a difference stationary process see, e.g., Meghir and Pistaferri (2004) and Guiso et al. (2005). Another model of idiosyncratic household income advanced in the literature is the heterogenous growth-rate model (see, e.g., Baker (1997) and Guvenen (2007)) where idiosyncratic household log-income, y_{it} , is a person-specific function of experience. 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. In this paper, I do not intend to test one model against the other, rather, I am interested in the correct decomposition of the difference stationary model, which has been traditionally used in the consumption literature.

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 informational sets of econometricians and households? To fix ideas, consider a simple example. 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 income growth, and adjust consumption appropriately. Econometricians, in turn, do not differentiate between income news known to households, and 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. In this case the household's information set is finer than the econometrician's.⁹

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 constructs different UC representations of several reduced form models of the aggregate US income, and finds that there always exists an UC model consistent with the relative pattern of variances of consumption and income observed in the aggregate US data, and consistent with the PIH. The intuition behind this result is that 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

⁹Throughout 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. E.g., 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 of each component.

permanent component within an UC model without distortion of the properties of the reduced form process.

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

$$\Delta C_{it} = \frac{r}{1+r} \frac{\theta_q(\frac{1}{1+r})}{(1 - \frac{1}{1+r})} \epsilon_{it} = \theta_q(\frac{1}{1+r}) \epsilon_{it}, \quad (7)$$

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 UC 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:¹⁰

$$\Delta C_{it} = \frac{r}{1+r} \theta_{q_0}(\frac{1}{1+r}) \epsilon_{it}^T + \theta_{q_1}(\frac{1}{1+r}) \epsilon_{it}^P. \quad (8)$$

¹⁰Note that Quah (1990) considers linear difference stationary processes, while equation (6) features log-linear income processes. 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 (8), 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. I will be explicit when I switch to log-linear income processes outlined in equations (4)-(6) and commonly used in the literature on household income processes.

Take $q_1 = 0$ and $q_0 = 0$, so that the order of auto-covariance 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. It is obvious that 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 auto-covariance functions of reduced and structural form processes are identical. Since households have better information on the sequences of permanent and transitory shocks, 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 underlines the main theme of the paper and 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 the 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 this issue.

3 The Same Reduced Form But Different Components: Sensitivity of Consumption to Informational Assumptions

In this section, I use the PSID to estimate a reduced form ARIMA(0,1,1) income model. I further decompose it into permanent and transitory components with different correlations between them. Specifically, I construct nine decompositions of idiosyncratic household income that differ in the volatility of transitory shocks, and contemporaneous correlation between per-

manent and transitory shocks. 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 examine the effect of different UC decompositions on consumption dynamics in the buffer stock model. Specifically, for different decompositions of the reduced form process, holding other relevant parameters fixed, I simulate life cycle economies and estimate different model statistics: sensitivity of consumption growth to income growth at one and five-year horizons, relative log-consumption profile, and the variance and persistence of income growth.

I consider the orthogonal decomposition of income adopted in the literature, along with other potentially valid UC decompositions. Thus, different implications arising from different decompositions may be attributed to differences in information sets held by households and econometricians. More precisely, econometricians cannot identify correctly the joint distribution of permanent and transitory shocks if the shocks are correlated and household income in first differences is a moving average process. Households, to the contrary, make their consumption decisions based on the knowledge of the correctly specified joint distribution of permanent and transitory shocks, be they correlated or not.

Suppose that the reduced form process for log household income is an ARIMA(0,1, q) process:

$$\Delta \log Y_{it} = \theta_q(L)u_{it}, \tag{9}$$

where $\theta_q(L)$ is a lag polynomial of order q in L .

Further assume that the structural log-income process is the sum of a difference stationary permanent component, $\log Y_{it}^P$, and a transitory component, a stationary process in log-levels, $\log Y_{it}^T$. The reduced and structural forms of observed series should agree in time (and frequency) domain:

$$\Delta \log Y_{it} = \Delta \log Y_{it}^P + \Delta \log Y_{it}^T = A(L)u_{it}^P + (1 - L)B(L)u_{it}^T, \tag{10}$$

where $A(L)$ and $B(L)$ are the lag polynomials that describe dynamics of the first difference of the permanent component and the level of the stationary component, respectively; u_{it}^P and u_{it}^T

are permanent and transitory innovations, respectively.

The variance of the permanent component, $[A(L)]^2\sigma_{u^P}^2$, is equal to the spectral density of the series at frequency zero, and is determined by estimates of $\theta_q(L)$, and the variance of the innovation from the reduced form process of equation (9).

The auto-covariance function of the reduced form process has $q+1$ non-zero auto-covariances, which is sufficient to estimate q moving average parameters, along with the variance of the reduced form income shock. An estimable UC 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— $A(L)$ and $B(L)$. 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 structural shocks is not identifiable from the sole dynamics of household income.¹¹ Without estimation, though, for any known reduced form data generating process one may always construct infinitely many UC representations. In the next section, I estimate the reduced form income process using data on household income from the PSID. I then construct nine unobserved components models of income that imply different (income) information sets but have the same auto-covariance function as the reduced form.

3.1 Univariate Dynamics of Idiosyncratic Household Income

In this section, I present some results on the univariate dynamics of household income in the PSID data. It is important to know whether the income process in equations (4)-(6) is empirically justified.

The income measure I consider is the residuals from the cross-sectional regressions of household log-disposable income on the head's education, household state of residence, a second degree polynomial in the head's age, and the head's race.¹² In the literature, it is typically labeled idiosyncratic household income. For the cross-sectional regressions, I use information from the 1969–1997 annual family files of the PSID.¹³ Sample selection is described in the notes to Table 1. Table 1 presents the auto-covariance function for the growth in household idiosyncratic

¹¹For the issues of identification of structural form time series processes see, e.g., Harvey (1989).

¹²My specification of the “predictable” component of labor income is quite flexible: it assumes, for example, time-varying returns to experience and education, differentiated by cohort.

¹³The PSID collected data biennially after 1997. Inclusion of data after 1997 would require a different modelling 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.

income. In this table, the auto-covariances and their respective standard errors are pooled over time. As can be seen from the table, the auto-covariance function is statistically significant up to order one, and the first order auto-covariance is negative. This is consistent with an integrated moving average process of order one. Thus, even if household income contains a unit root, it also contains a strong mean-reverting transitory component.

In Table 2, I test the null hypothesis that the auto-covariances of a given order are equal to zero in all time periods. Results of this test indicate that the transitory component may be a moving average process of order one. This is consistent with 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 assumption is not at odds with the data as Table 1 suggests.

In Table 3, I present estimates of a moving average process for idiosyncratic household income. Household idiosyncratic income is highly volatile, with a standard deviation of the reduced form shocks of about 28% per year, and contains a strong mean-reverting component.

3.2 Constructing Different UC Models

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

Assume that log income in differences, after the deterministic growth rate g_{it} has been removed, follows a stationary MA(1) process. A corresponding UC model may be 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. Following the above notation, the reduced and structural forms of the process for the first differences in income are:

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

where superscripts rf and sf denote reduced and structural form respectively.

I will use this process for simulating the life cycle buffer stock economy since it is easy enough to deal with computationally, and general enough to allow for decompositions of income into permanent and transitory components of different relative importance.¹⁴

Since the reduced form has only two pieces of information, the auto-covariances 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 information 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 3. 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 . The variance of the permanent component is determined by the spectral density of the reduced form series at frequency zero. Thus, for estimated income parameters, the estimate of the variance of innovations to the random walk permanent component is equal to $(1 + \hat{\theta})^2 \hat{\sigma}_u^2 = 0.02$. The variance of transitory innovations can be estimated by $-\hat{\gamma}(1) - cov(u_{it}^P, u_{it}^T)$, where $\hat{\gamma}(1)$ is the first order auto-covariance of the reduced form process and $cov(u_{it}^P, u_{it}^T)$ is the covariance between permanent and transitory innovations. Thus, for the covariance equal to -0.0328 (and the corresponding correlation between income shocks approximately equal to -1.0), the standard deviation of transitory innovations is 0.227 ; for the covariance equal to 0.00 , the standard deviation of transitory innovations is 0.136 .

The covariances between transitory and permanent shocks—and the corresponding correlations—determine the relative volatility and size of permanent and transitory shocks. Thus, the income model with the perfect negative correlation between permanent and transitory shocks has the most volatile transitory shocks, while the income model with the perfect positive correlation between the permanent and transitory shocks has the least volatile transitory shocks.

Correspondingly, I call the models built from these covariances as model (1)-model (9) in

¹⁴Ludvigson 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, down payments, and adjustment costs in the market for durable goods. Luengo-Prado and Sørensen (2008) use a generalization of this process to gauge the effects of different layers of uncertainty (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 et al. (2005) calibrate the parameters of this income process to investigate consumption and portfolio choice over the life cycle.

Appendix A1, with model (1) featuring the perfect negative correlation between permanent and transitory shocks and model (9)—the perfect positive correlation.

3.3 Results for Simulated Life Cycle Buffer Stock Economies

Solving the dynamic programming problem in Section 2 by iterating the Euler equation, I obtain a set of age-dependent consumption policy functions. I assume that the gross interest rate R_t is non-stochastic and that the joint probability density function of (potentially correlated) transitory and permanent shocks $f(u^P, u^T)$ is time invariant. In addition, shocks are assumed to be jointly log-normal, where the underlying joint normal distribution has a mean vector zero, and the variance-covariance matrix $\Sigma_{u^P u^T}$. To induce the correlation between independent normal draws, I use the Cholesky factorization of the variance-covariance matrix $\Sigma_{u^P u^T}$. Upon finding the converged policy functions $\{c_t(x_{it})\}_{t=24}^{65}$, I simulate the economy populated by 5,000 ex ante identical consumers, who are differentiated ex post due to different history of income draws.

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, 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 0.95, and the coefficient of relative risk aversion to 4.0. 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 detail in the previous subsection. The full details of the model solution are laid out in Appendix B.

I perform pooled panel regressions of the growth of (simulated) household consumption on the current and lagged growth of (simulated) household income. In addition, I examine the sensitivity of long differences in log-consumption to long differences in log-income; the relative log-consumption profile, defined by the ratio of the cross-sectional average of log-consumption at age 45 to age 30, and the cross-sectional average of log-consumption at age 65 to age 45; and the reduced form income parameters—persistence and the variance of income growth. The results for income models 3 (negative correlation between structural shocks equal -0.50), 5 (no correlation), and 7 (positive correlation equal 0.50) are presented in Table 4, Panel A.¹⁵

¹⁵The MPCs out of shocks to current and lagged income, and the shocks cumulated over the five-year horizon are larger for models with a higher correlation between the shocks. The ranking of models in terms of MPCs in ascending order is: model 1, model 2, ..., model 9. Thus, without losing valuable information, I chose to report only the results for income models (3), (5), and (7). The results of simulations for all income models are available upon request.

Motivation for the use of these statistics is fairly straightforward. The magnitude of the coefficient on the current income growth should depend on the smoothness of income innovations, while the coefficient on the lagged income growth should measure excess sensitivity of consumption growth to lagged income news. Long differences in log-income will be largely dominated by the contribution of permanent income shocks, which should be reflected into long differences in log-consumption. 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 (model 3). Similar results hold for the sensitivity of long consumption growth to long income growth, measured by the differences in current log-consumption (income) and log-consumption (income) observed five periods before. Due to borrowing restrictions, the models indicate that some fraction of the shocks to lagged income remains unsmoothed (4% for model 7 vs. 12% for model 3). The difference between sensitivities across income models is significantly different from zero at any conventional level of statistical significance.

3.3.1 Intuition Behind the Results

The basic intuition behind the results is the following. Absent borrowing restrictions, households react 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 $\Delta \log C_{it} = \alpha_P u_{it}^P + \alpha_T u_{it}^T$, where α_P and α_T are the (partial regression “insurance”) coefficients that depend 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, is equal to $u_{it}^P + u_{it}^T - u_{it-1}^T$. 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 $\Delta \log C_{it} = \beta_1 \Delta \log Y_{it} + error$, and β_k from an OLS regression $\Delta_k \log C_{it} = \beta_k \Delta_k \log Y_{it} + error$, where $\Delta_k \log x_{it} = \log x_{it} - \log x_{it-k}$. 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 correlation, a positive permanent shock is, on average, accompanied by a negative transitory

shock, smoothing out the sum of income innovations. Hence, consumption 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 five-year horizon (β_5). 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 reflected, consequently, in higher coefficients measuring the sensitivity of consumption to income growth at different horizons (β_1 and β_5). In statistical terms, $\hat{\beta}_1 = \frac{cov(\Delta \log C_{it}, \Delta \log Y_{it})}{var(\Delta \log Y_{it})} = \frac{\hat{\alpha}_P * \sigma_{u_P}^2 + \hat{\alpha}_T * \sigma_{u_T}^2 + (\hat{\alpha}_P + \hat{\alpha}_T) * cov(u_{it}^P, u_{it}^T)}{var(\Delta \log Y_{it})}$. The denominator is the same for all structural decompositions of the reduced form income model, the “smoothing” term is measured by $(\hat{\alpha}_P + \hat{\alpha}_T) * cov(u_{it}^P, u_{it}^T)$ in the numerator—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 $\hat{\beta}_k = \frac{k * \hat{\alpha}_P * \sigma_{u_P}^2 + \hat{\alpha}_T * \sigma_{u_T}^2 + (\hat{\alpha}_P + k * \hat{\alpha}_T) * cov(u_{it}^P, u_{it}^T)}{k * \sigma_{u_P}^2 + 2 * \sigma_{u_T}^2 + 2 * cov(u_{it}^P, u_{it}^T)}$. Again, the denominator is the same for different structural income processes while the numerator contains the “smoothing” term $(\hat{\alpha}_P + k * \hat{\alpha}_T) * cov(u_{it}^P, u_{it}^T)$, which is larger, in absolute value, for the processes with more negatively correlated permanent and transitory shocks. For the chosen parameters in the first three rows of Table 4, $\hat{\alpha}_P$ is estimated at about 0.60, and $\hat{\alpha}_T$ —at about 0.05.¹⁶ Using the above formulas, I can estimate $\hat{\beta}_1 = 0.10$ (0.23) and $\hat{\beta}_5 = 0.38$ (0.45) for the income models with the correlation between the shocks equal to -0.50 (0.0). Those are just slight underestimates of the values presented in rows (1) and (2) of Table 4.

4 The Life Cycle Model of Consumption, and Empirical Estimation of the Income Process

4.1 The Model

Simulations in the previous section show that different decompositions of the same reduced form income process lead to sizeable differences in the sensitivity of consumption growth to contemporaneous and lagged income growth, and the sensitivity of long consumption growth to

¹⁶These values of “insurance” parameters, α_P and α_T , indicate that a life cycle model with liquidity constraints and self-insurance against idiosyncratic shocks can go a long way towards explaining the extent of insurance of permanent and transitory shocks found recently by Blundell et al. (2008). For their whole sample, they find that about 36% (95%) of permanent (transitory) shocks are insured.

long income growth. Thus, the joint dynamics of consumption and income in real data may help identify parameters of the income process—the variances of permanent and transitory shocks, and the correlation between them.

In this section, I use a structural life cycle model of consumption to estimate parameters of the income process and the behavioral parameters. I assume that the model households are married couples that maximize expected utility from consumption over the life cycle. The only source of uncertainty in the model is uncertainty over income flows, arising from transitory and permanent income shocks.¹⁷ I assume that all households start working at age 24 and retire at age 66. Households maximize the expected utility from annual consumption flows:

$$E_{i,24} \left[\sum_{t=24}^T \beta^{t-24} v(Z_{it}) U(C_{it}) + \beta^{T+1} V_{iT+1}(X_{iT+1}) \right].$$

T —the last working period age—is set to 65; β is the time discount factor; V_{iT+1} is the value function at age 66, equal to the maximized expected utility at age 66 and onwards; C_{it} is household i 's consumption at age t ; Z_{it} denotes the variables that proxy household i 's taste shocks at time t ;¹⁸ X_{iT+1} is the cash-on-hand at age 66. The utility function is the time separable CRRA utility function.

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 (see equation (3)).

Cash-on-hand and consumption can be expressed in terms of the ratios to permanent income, and the state space reduces to one variable, cash-on-hand relative to the permanent income, x_{it} . As in Gourinchas and Parker (2002), I assume that the consumption function at retirement is linear in cash-on-hand, X_{iT+1} and illiquid wealth, H_{iT+1} : $C_{iT+1} = \kappa_x X_{iT+1} + \kappa_h H_{iT+1}$,

¹⁷Other poorly insured risks over the life cycle are health shocks. In this paper, I purposefully limit my analysis to 24–65 year olds, a subgroup of population for whom medical expenses and health shocks are relatively less important.

¹⁸In the literature, vector Z_{it} usually contains leisure time of a spouse, the number of adults, and the number of children over the life cycle. I follow Gourinchas and Parker (2002) and use family size for Z_{it} . I assume that family size affects household marginal utility exogenously and deterministically, and estimate the family-size adjustment factors from empirical data.

where κ_x (κ_h) is the marginal propensity to consume from liquid (illiquid) assets at retirement. Dividing both sides of the equation by the permanent income at age $T + 1$, it becomes $c_{iT+1} = \kappa_x x_{iT+1} + \kappa_h h_{iT+1}$, where c_{iT+1} is household i 's consumption relative to the permanent income at age $T + 1$ and h_{iT+1} is the level of household i 's illiquid assets relative to the permanent income at age $T + 1$. The age-dependent consumption functions $\{c_t(x_{it})_{t=24}^{65}\}$ are found recursively by iterating the Euler equation and utilizing the consumption function at retirement. The details of the model solution are provided in Appendix B.

4.2 Estimation by the Method of Simulated 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 retirement process parameters— κ_x , κ_h ; and the parameters of the income process— σ_{u^T} , σ_{u^P} and $cov(u_{it}^P, u_{it}^T)$. I estimate the model parameters by the method of simulated moments. Since my model is cast in terms of one state variable, cash-on-hand relative to permanent income, I reformulate the consumption rule at retirement as $c_{iT+1} = \kappa_x x_{iT+1} + \kappa_0$, where κ_0 is the product of the marginal propensity to consume from illiquid wealth and the average illiquid wealth (relative to the permanent income) at retirement. Since the model does not have enough identifying information for κ_0 and κ_x , and since it is hard to benchmark the marginal propensity to consume from illiquid assets using empirical data,¹⁹ I set κ_0 and κ_x to the values estimated in Gourinchas and Parker (2002), Table III, column 1: $\kappa_0 = 0.0015$, $\kappa_x = 0.0710$.²⁰ I estimate five parameters in total.

4.2.1 The Identification Scheme

In this section, I discuss identification of the model parameters.

In the absence of uncertainty or in the case of a perfect foresight, both the life cycle and dynastic consumption models predict that the shape of the consumption profile is determined by the behavioral parameters (the time discount factor and the relative risk aversion parameter) and the interest rate. Thus, the time discount factor and the relative risk aversion parameter may be identified from the life cycle consumption profile, or from the long-run features of consumption

¹⁹The PSID reports housing equity that may qualify for illiquid wealth, yet it does not provide separate records on other important components of illiquid wealth (e.g., pension wealth).

²⁰The same values of these parameters were used in Section 3.3.

data. When household incomes are uncertain and insurance markets are incomplete, the life cycle consumption profile is also affected by precautionary motives. Thus, the shape of the consumption profile may be also informative for the identification of the magnitude of income uncertainty over the life cycle. Gourinchas and Parker (2002) and Cagetti (2003), guided by these considerations and treating the volatility of permanent and transitory shocks as given, estimate the relative risk aversion parameter and the time discount factor for different occupations and education groups. Furthermore, the results from Table 4 Panel A suggest that the volatility of individual components of income and the correlation between them determines the joint short-run dynamics of income and consumption.

The results in Table 4 are supportive of these intuitions. In Table 4 Panel B, I compare model statistics for the benchmark model and its modifications. In the benchmark model, the time discount factor is 0.95, the relative risk aversion parameter is 4, the correlation between structural income shocks is -0.5 , and there are no measurement errors in log-income and log-consumption. In Panel B, I present the percentage changes in the benchmark model's predictions (first row of Panel A) after I change the behavioral parameters, or assume that measurement errors in log-income and log-consumption are not zero. In the first row, I assume a multiplicative measurement error in consumption such that the log-error is distributed as $iidN(0, \sigma_{u_{it}^C}^2)$. I set $\sigma_{u_{it}^C}$ to 0.24, an estimate in Blundell et al. (2005a) for their total sample. As can be seen, the measurement error in consumption does not alter significantly the predictions obtained from the benchmark model. In the next row, I add a multiplicative measurement error to income to the previous model. Following Meghir and Pistaferri (2004), I assume that the measurement error in income explains 25% of the total income growth. Since $\hat{\sigma}^2 \widetilde{\Delta \log Y_{it}} = 0.058$ for my income sample, I assume the log-error in income is distributed as $iidN(0, \sigma_{u_{it}^Y}^2)$, with $\sigma_{u_{it}^Y} = 0.085$.²¹ When I allow for the measurement error in income, I have to adjust the variance of transitory shocks downwards to fit the auto-covariance function of income growth. The standard deviation of transitory shocks, in this case, is 0.149 instead of 0.177. Since negative permanent shocks are now, on average, accompanied by positive transitory shocks of a smaller size, the unpredictable part of income growth becomes less smooth compared with the benchmark model and the sensitivity of consumption to income growth should be higher. The second row of Panel B

²¹If measurement error is i.i.d., it contributes twice its variance towards the variance of idiosyncratic income growth. For this choice of the variance of measurement error, the share of total variance of idiosyncratic income growth "explained" by measurement error is $(2 * 0.085^2) / 0.058 \approx 0.25$, or 25%.

supports the prediction: compared with the benchmark model's prediction, measurement error in income significantly increases the sensitivity of consumption growth to contemporaneous income growth. Thus, although measurement error in income is not separately identified from the variance of transitory innovations, it is important to take it into account while estimating the income process using a life cycle model. In the next row, I change the benchmark model's time discount factor to 0.90. The effect of the change is intuitive. Due to higher impatience, consumption becomes more sensitive to current and lagged income shocks, log-consumption profile becomes flatter in the beginning of the life cycle (since now, on average, consumers prefer larger consumption at early ages), and steeper in the end of the life cycle (since now consumers are more sensitive to unfavorable income shocks in the end of the life cycle), accentuating the well-known hump in the profile. The reverse is true if the benchmark model's time discount factor is changed to 1.0. In the fifth row, I set the coefficient of relative risk aversion to 1.0 instead of 4.0. Similar to the results for the high discount rate, consumers become significantly more sensitive to current and lagged income shocks. Also, the cumulative consumption growth over the five-year horizon becomes significantly more sensitive to the cumulative income growth over the same horizon. The reverse is true when the relative risk aversion parameter is set to 7.0. Reassuringly, the income process parameters, measured by an OLS auto-regressive coefficient in an AR(1) regression of income growth and the variance of income growth, are essentially the same across different specifications.

Summing up, the structure of the income process, the time discount factor and the relative risk aversion parameter can be identified using the long run features of consumption data (contained in the consumption life cycle profile); the short run features of consumption and income data (determined by the sensitivity of consumption growth to current and lagged income growth, and the sensitivity of long consumption growth to long income growth); and the features of income data (determined by the variance and persistence of income growth).

4.2.2 Matching the Moments

In the matching exercise, I look for the volatility of structural income shocks, the contemporaneous correlation between them, and behavioral parameters that lead to the closest match to the model moments in Table 4 estimated from empirical data. In several specifications, I also match the entire log-consumption profile relative to the cross-sectional average of log-consumption at

age 24, the beginning of the life cycle. I construct the consumption profile from the CEX and the PSID data.

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 gross growth rates of disposable income, $\{G_t\}_{t=25}^{65}$; the life cycle profile of utility shifters, $\{v(Z_t)\}_{t=24}^{66}$; the gross real interest rate on safe liquid assets, R ; the mean and the standard deviation of the distribution of wealth-to-permanent income ratio at age 24, $\{(\overline{W/Y^p})_{24}, \sigma_{(W/Y^p)_{24}}\}$; the mean and the standard deviation of the distribution of permanent income at age 24, $\{\overline{Y}_{24}^p, \sigma_{Y_{24}^p}\}$; the parameters of the consumption function at retirement, κ_0 and κ_x . I set κ_0 to 0.0015, κ_x to 0.0710, and the gross real interest rate on safe liquid assets to 1.0344.²²

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 criterion function

$$\left[\frac{1}{I_d} \sum_{i=1}^{I_d} h_i^d - h^s(\theta; \hat{\chi}) \right]' W \left[\frac{1}{I_d} \sum_{i=1}^{I_d} h_i^d - h^s(\theta; \hat{\chi}) \right] = g'_{I_s} W g_{I_s}, \quad (11)$$

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 moment; I_s is the number of simulated households; h_i^d is an estimated second stage moment for the data household i ; h^s is a simulated population moment; W is a positive definite weighting matrix; $\hat{\chi}$ is a 91×1 vector of the first stage parameters; θ is a vector of the second stage parameters.

Given $\hat{\chi}$ and an initial vector of the second stage parameters, θ_0 , the algorithm proceeds in following steps: 1) numerically derive the age-dependent policy functions $\{c_t(x_{it}; \theta_0, \hat{\chi})\}_{t=24}^{65}$; 2) given these policy functions, simulate the model economy populated by 5,000 households; 3) for these simulated data, calculate the model moments and the MSM criterion function; 4)

²²These estimates are from Gourinchas and Parker (2002). The real interest rate is the average real return on Moody’s AAA municipal bonds for monthly data spanning 1980–1993. When correcting standard errors of the model parameters for uncertainty in the first stage estimates, I use the standard errors of κ_0 , κ_x , and R provided in Gourinchas and Parker (2002).

if convergence is not achieved, update θ using some optimization method; repeat steps (1)-(3) until convergence is achieved.

For the weighting matrix W , I choose the diagonal matrix with diagonal elements equal to the diagonal elements of the inverse of $\widehat{\Sigma}_g$, a consistent sample estimate of the variance-covariance matrix of the second-stage moments. Further details on the asymptotic distribution of the model parameters and goodness-of-fit statistics are provided in Appendix C.

4.3 Construction of Empirical Moments

In this section, I describe estimation of the empirical moments I match. I first briefly describe the data sources used.

4.3.1 Data

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 the 1980–1998 waves of the CEX. The PSID features panel data and is considered to be the best source of income data for the US population. Unlike the CEX, the PSID limits its coverage of consumer expenditures to food at home and away from home. Since I am interested in the link between changes of disposable household income and total household consumption, I impute the total consumption to the PSID households using information on household food consumption in the PSID and the CEX, and matched demographics from the CEX and the PSID, and exploit superior (to the CEX) income data from the PSID. The PSID data are taken from 1980–1997, 1999, and 2001 waves. I follow the methodology of Blundell et al. (2005*b*) to impute the total consumption to the PSID households. The full details on sample selection of CEX and PSID households are provided in Appendix D. The imputation procedure is described in Appendix E.

4.3.2 Estimation of Life Cycle Profiles

At each point in time total consumption and disposable income are affected by cohort effects, time (business cycle) effects, life cycle (age) effects, and idiosyncratic effects. Since the age, cohort, and time effects cannot be estimated simultaneously, I assume that the time effects can be captured by regional unemployment, as in Gourinchas and Parker (2002) and French (2005).

Consumption (and, similarly, income) can be decomposed as follows:

$$\log C_{it} = \alpha'_c c_i + \alpha'_f f_{it} + \alpha_u ue_{it} + \alpha'_a Age_{it} + \eta_{it}, \quad (12)$$

where $\log C_{it}$ is the total real household log-consumption, c_i is the set of cohort dummies (with an omitted category comprising households with heads born between 1939–1944), f_{it} is the set of family size dummies (with an omitted category of households with 7 or more members),²³ Age_{it} is the full set of age dummies (created for households aged between 24–70), and η_{it} is the idiosyncratic effect that consists of the time-varying and time-invariant random effects, α_c is a vector of cohort coefficients, α_f is a vector of coefficients on family size dummies, α_u is the coefficient on regional unemployment, ue_{it} , and α_a is a vector of age coefficients.

I eliminate cohort, family size, and aggregate effects from consumption predicted by equation (12) using the following transformation:

$$\widetilde{\log C_{it}} = \widehat{\log C_{it}} - \widehat{\alpha}'_c c_i + \widehat{\alpha}'_f (f_a - f_{it}) + \widehat{\alpha}_u (\overline{ue} - ue_{it}) + \widehat{\eta}_{it}, \quad (13)$$

where $\widehat{\log C_{it}}$ is the predicted consumption from equation (12), $\widetilde{\log C_{it}}$ is the transformed consumption, $f_a = \frac{1}{I_a} \sum_{i=1}^{I_a} f_i$ is the average family size for a group of I_a people of a certain age, \overline{ue} is the average unemployment rate over all sample years and households, and $\widehat{\eta}_{it}$ is a household-specific residual at a given age.²⁴

The unsmoothed life cycle consumption profile is a plot of the age-average of $\widetilde{\log C_{it}}$ against age. The smoothed profile is constructed from a fractional polynomial regression of $\widetilde{\log C_{it}}$ on a fifth degree polynomial in the head's age and year of birth. Figure 1 plots raw, smoothed and unsmoothed life cycle profiles of consumption. Similarly, I construct the smoothed life-cycle profiles of utility shifters (determined by changes in family size over the life cycle), and household life cycle gross growth rates in disposable income.

In Figure 1, smoothed profile is net of cohort and time effects, and shows the average total

²³The full set of family size dummies is created for households with family size of 2, 3, 4, 5, 6, and 7 or more members.

²⁴I estimate equation (12) by pooled OLS. Thus, the residuals contain household effects. Alternatively, I could have used the fixed effect panel regression to construct the profiles. It should not matter in practice which method is used. If $\eta_{it} = \zeta_i + \nu_{it}$, household specific effects, ζ_i 's, will be averaged out to zero while constructing the average log-consumption profile. Similarly, they will vanish when I run regressions using log-consumption in first differences as the dependent variable.

consumption at a given age, with family size set to the average family size at this age. Consistent with Gourinchas and Parker (2002) and Fernández-Villaverde and Krueger (2007), consumption profiles exhibit a visible hump in the mid-forties. When the profile is constructed such that the family size is set to its average over the life cycle (3.34 household members), the ratio of the average total consumption at the peak to consumption at age 24 is about 1.3—a finding similar to that in Fernández-Villaverde and Krueger (2007). Unsmoothed profiles are quite noisy due to, among other things, sampling uncertainty and imputation error. Instead of relying on consumption levels and fitting the average log-consumption profile over the life cycle as in Gourinchas and Parker (2002), I fit the relative log-consumption profile in a matching exercise. The relative profile, presented in Figure 2, measures the ratio of the average log-consumption at a given age to the average log-consumption at age 24.²⁵

4.3.3 Persistence of Income, and the Joint Dynamics of Consumption and Income

The other moments, I am interested in, are the persistence in the growth rate of income, the variance of income growth; the coefficients from an OLS regression of consumption growth on contemporaneous and lagged income growth, and cumulative consumption growth over the five-year horizon on cumulative growth of disposable income over the same horizon. To reconcile empirical data with simulated data, I use $\widetilde{\log C_{it}}$ and $\widetilde{\log Y_{it}}$, log-household consumption and log-household disposable income from empirical data, “net” of cohort, aggregate, and family size effects. First, I regress the first difference in $\widetilde{\log C_{it}}$ on the current and lagged growth in $\widetilde{\log Y_{it}}$, changes in family size over the life cycle, and a quadratic polynomial in the head’s age. The results of this regression are presented in Table 5. The estimated relationship between current consumption and income growth is statistically significant: for every 10% increase in income, consumption increases by 0.6%. In accordance with the PIH, if log-income is a random walk and therefore shocks to log-income are permanent, consumption should have changed by 10% instead of the observed 0.6%. The magnitude of the contemporaneous association between consumption and income growth may be taken as another piece of evidence that household log-income contains, apart from a random walk permanent component, a strong mean reverting transitory component. The response of current consumption growth to lagged income growth is small and imprecisely estimated. This result is in line with the rational expectations tradition of

²⁵The ratio of the average log-consumption at the peak to the average log-consumption at age 24 is slightly above 1.05. It is different from 1.3 estimated for the average consumption in Figure 1 due to log-transformation.

consumption theory, which predicts that consumption is a martingale, and therefore consumption changes should be orthogonal to any past information, inclusive of the past income changes (e.g., Hall (1978) and Hall and Mishkin (1982)).

Table 6 contains results for an OLS regression of cumulative consumption growth over the five-year horizon on cumulative income growth over the same horizon. The relationship is much stronger compared to the contemporaneous relationship in Table 5: a 10% increase in income growth over the five-year horizon is associated with a 2.5% increase in consumption growth over the same horizon.

To benchmark the persistence of idiosyncratic income, I estimate an AR(1) process for income growth by OLS. The results are presented in Table 7. The OLS estimate of an AR(1) coefficient is equal to -0.347 with a standard error of 0.01. 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.²⁶ This proves to be very important when repeated estimations are performed on simulated data.

5 Discussion of Results

5.1 Results

Table 8 contains the main results. They are based on minimization of the MSM criterion function, using information from the total sample of households.

It may be hard to identify the relative risk aversion parameter and the time discount factor jointly: rows 3 and 5 of Table 4 Panel B, for particular simulations, show that low risk aversion and a low time discount factor may fit data equally well.²⁷ In column (1) of Table 8 I set the relative risk aversion parameter to 4 and assume that there are no measurement errors in consumption and income. For matching, I utilize only the sensitivity of consumption growth to income growth at one and five-year horizons, and the variance and persistence of income growth. The model is over-identified with one overidentifying restriction. The time discount

²⁶It is well-known that any invertible MA(1) process has an autoregressive representation of the infinite order (see, e.g., Hamilton (1994)). Galbraith and Zinde-Walsh (1994) show that low order auto-regressive approximations of an MA(1) process with a moving average parameter of 0.5 and less in absolute value—as low as order one and up to the third order—perform the best in terms of minimizing the mean squared error. Thus, using the estimated AR(1) process as an approximation to (the infinite order representation of) MA(1) process may be sufficient to benchmark the persistence of household idiosyncratic income growth.

²⁷For a discussion of the difficulties involved in estimation of the relative risk aversion parameter in macro data see, e.g., a monograph by Hall (2005).

factor is estimated to be close to 1 but not significant, the variance of transitory and permanent innovations are both significant, the covariance between the innovations is negative but imprecisely estimated. Apparently, the moments used do not provide enough information to precisely identify the time discount factor, and the covariance between structural shocks. Standard deviations of permanent and transitory innovations are estimated at about 12% and 22%, respectively. The over-identification statistic rejects the model. In column (2), to capture the life cycle hump, I add to the previous set of moments the relative average log-consumption at age 45 to age 30, and at age 65 to age 45. Results show that these additional moments bring crucial information for the identification of the time discount factor—it becomes lower in magnitude and statistically significant. To fit the hump, consumers should be impatient enough and face substantial uncertainty in their income—both the volatility of transitory and permanent shocks are estimated to be larger compared with the results in column (1). I estimate the correlation between the structural shocks at about -0.63 . The model fit slightly improves compared with the model in column (1) but still remains unsatisfactory.

In column (3), I use the same set of moments but estimate the relative risk aversion parameter. The estimate of the relative risk aversion is about 6, significant at the 5% level. The rest of the parameters are precisely estimated, with an estimate of the correlation between structural income shocks of about -0.5 . In column (4), I allow for the measurement and imputation error in consumption, assuming that the log-error is distributed as $iidN(0, \sigma_{u_{it}^C}^2)$, where $\sigma_{u_{it}^C} = 0.24$ (see Section 4.2.1). The point estimates of the parameters are essentially the same as in column (3). Even though the measurement error is large, it is averaged out while constructing the average log-consumption profile and does not affect appreciably the other moments. Since it does not affect the identifying information (see also row 1 of Table 4 Panel B), it also does not affect the estimated parameters. In column (5), I allow for measurement error in income only. I assume that the measurement error in income explains 25% of the variance of income growth, and that the log-error is distributed as $iidN(0, \sigma_{u_{it}^Y}^2)$. The estimated variance of income growth for the sample of households with income and consumption information is 0.082 (see Table 7). Thus, I set $\sigma_{u_{it}^Y}$ to 0.101. Allowance for the measurement error in income leads, quite predictably, to a significant reduction of an estimated volatility of transitory shocks. The rest of the parameters do not change much compared with the results in column (4). In column (6), I allow for measurement errors both in consumption and income. The results are essentially the same as

in column (5).

In column (7), I utilize a larger set of moments which includes the entire relative log-consumption profile over the life cycle, measured by the ratio of the average log-consumption at ages 25 and higher to the average log-consumption at age 24. The relative risk aversion parameter is now significant at the 1% level, and the model cannot be rejected using the J -test. The time discount factor is estimated at about 0.90; the relative risk aversion—at about 7.0; the volatility of permanent and transitory shocks—at 12% and 17%, respectively; and the correlation between the shocks—at negative 0.65. If the model is true, imposing zero correlation between permanent and transitory shocks will bias estimates of the structural parameters. I explore this issue in column (8) of Table 8. Relative to column (7), I obtain much lower point estimates for the time discount factor, the volatility of permanent and transitory shocks, and a much larger estimate for the relative risk aversion parameter. The difference between the estimates of permanent and transitory shocks in column (7) and (8) is significant at the 10% and 5% levels, respectively. Apparently, zero correlation between structural shocks allows for a lower volatility of transitory income shocks (see Appendix A1, e.g., models 3 and 5). Since the bulk of the shocks to log-income is smoothed, the effect of smoothing in the model, reflected in low coefficients on income growth in regressions of consumption growth on income growth at one and five-year horizons, is achieved by a much larger point estimate of the relative risk aversion, rather than by the presence of a negative correlation between the structural shocks. A larger estimate of the relative risk aversion, in turn, “flattens out” the hump in the second part of the life cycle. Apparently, to fit the life cycle profile of consumption, the matching procedure now requires lower patience on the part of households (see Table 4 Panel B, rows 3 and 6). These estimates of behavioral parameters, coupled with lower estimates of idiosyncratic permanent and transitory income volatility, fit well the hump-shaped relative log-consumption profile. I find the estimates in column (7), where structural shocks are allowed to be contemporaneously correlated, more plausible. I will later test the models of column (7) and (8) formally.

The results in column (7) may be affected by an insignificant and low estimate of the sensitivity of current consumption growth to lagged income growth (see Table 5). In column (9) of Table 8 I show that this is not the case: the point estimates for the model in column (9), where I do not utilize the information on the sensitivity of current consumption growth to lagged income growth, are practically the same as in column (7).

In column (10) of Table 8, I fix the standard deviation of permanent shocks at 0.1219, the estimate from an unrestricted estimation in column (7) of Table 8: 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 for the contemporaneous correlation between the structural shocks (see, e.g., Cochrane (1988)). The results are the following. The volatility of transitory shocks is estimated at 14.0%, about 3% lower than the value obtained in the unrestricted estimation. The result corresponds to the decomposition exercise performed in Appendix A1: for the same volatility of permanent shocks, income models with negative covariance between structural shocks feature transitory shocks with a larger volatility. Simulations in the first three rows of Table 4 Panel A show that models with the same behavioral parameters, the same reduced form income parameters but different correlations between the shocks should result in a similar fit of the log-consumption profile but different joint dynamics of consumption and income. In particular, models with more negatively correlated innovations, and therefore with a smoother unpredictable part of income growth, result in lower estimates of consumption sensitivity to income growth. Since the estimates of the time discount factor and the relative risk aversion parameter in column (10) are virtually the same as those in column (7), one may expect that the model in column (10) fails at predicting the joint dynamics of income and consumption. Indeed, this turns out to be true: the model in column (10) substantially over-predicts sensitivity of consumption growth to current and lagged income growth, and consumption to income growth cumulated over the five-year horizon, while having a similar fit of the relative log-consumption profile as the model in column (7).²⁸ The relative log-consumption profiles from the data and the model in column (7) of Table 8 are depicted in Figure 7.²⁹

Comparison of the models in columns (7) and (10) brings an important result: to adequately fit the joint dynamics of income and consumption, one needs to properly identify the income process—the volatilities of permanent and transitory shocks, and the covariance between them.

In Table 9, I re-estimate the models of columns (7) and (8) of Table 8, using an optimal weighting matrix in the MSM estimation. The optimal weighting matrix is constructed at the parameter estimates of the unrestricted model of column (7). Using the optimal weighting

²⁸In the data, β_1 is estimated at 0.06, β_2 —at 0.017, and β_5 —at 0.247. In the model in column (7) of Table 8, $\hat{\beta}_1 = 0.041$, $\hat{\beta}_2 = 0.018$, and $\hat{\beta}_5 = 0.297$. In the model in column (10) of Table 8, $\hat{\beta}_1 = 0.145$, $\hat{\beta}_2 = 0.063$, and $\hat{\beta}_5 = 0.346$.

²⁹The models in columns (8) and (10) produce the profiles that are hardly distinguishable from the profile of the model in column (7) of Table 8.

matrix, for both models, I obtain slightly larger estimates for the time discount factor, and the variance of transitory and permanent income shocks, and a lower estimate for the relative risk aversion parameter. The correlation between structural income shocks is now estimated at -0.60 and is significant at the 1% level. The distance statistic indicates that one can reject the no-correlation model in favor of the model with correlated permanent and transitory shocks at any conventional level of significance.

5.2 Discussion and Comparison to the Literature

How do my results compare to results in the literature? The time discount factor is estimated to be lower than in Gourinchas and Parker (2002) but compares well with the recent estimates of exponential discounting models in Laibson et al. (2007). The estimated relative risk aversion parameter is about 6, and is comparable with the estimates in Cagetti (2003) and Nielsen and Vissing-Jorgensen (2006).³⁰

The parameters of the income process are precisely estimated and indicate substantial idiosyncratic income uncertainty over the life cycle. Consistent with the literature on the univariate income dynamics at the household level, the volatility of permanent shocks is smaller than the volatility of transitory shocks.³¹ I find that the contemporaneous covariance between transitory and permanent shocks is negative, with the correlation of about -0.60 . The null hypothesis of zero correlation between transitory and permanent shocks can be easily rejected.

³⁰Using the results of Section 4.2.1, high risk aversion is needed to match low sensitivity of consumption to income growth (Panel B, row 7), while a lower time discount factor makes the relative log-consumption profile steeper in the end of the life-cycle (Panel B, row 3). Perhaps, Gourinchas and Parker (2002) find a lower estimate for the relative risk aversion parameter since, first, they do not consider the sensitivity of consumption to income shocks in their estimations, and, second, their log-consumption profile is relatively flat over the life cycle.

³¹Storesletten et al. (2004) estimate an AR(1) process for idiosyncratic log-income and find that it is highly persistent, with an auto-correlation coefficient of about 0.95. Their estimates (p. 705, Table 2) can be used to infer that the (unconditional) standard deviation of persistent (“permanent”) shocks ranges from 0.13 to 0.21 while the standard deviation of the white noise transitory shock ranges from 0.24 to 0.56. Carroll and Samwick (1997), for their full sample, estimate the standard deviation of the permanent component to be equal to 0.15, and the standard deviation of the transitory component to be equal to 0.21. Estimates of Gourinchas and Parker (2002) for their total sample are: 0.15—for the standard deviation of permanent shocks and 0.21—for the standard deviation of transitory white noise shocks. The estimates of Blundell et al. (2008) can be used to infer that the (unconditional) standard deviation of permanent shocks is equal to 0.15 while the (unconditional) standard deviation of transitory shocks (to the moving average transitory component) is equal to 0.20. The estimates of Meghir and Pistaferri (2004) for their pooled sample are: 0.18—for the (unconditional) standard deviation of permanent shocks, and 0.21—for the upper bound of the (unconditional) standard deviation of transitory shocks (to the moving average transitory component). The estimates of Cocco et al. (2005) are: 0.10—for the standard deviation of permanent shocks and 0.27—for the standard deviation of white noise transitory shocks. These papers assume that the covariance between transitory and permanent shocks is equal to zero. They use data from the PSID with different sample selection criteria and time span. My estimates of the volatility of permanent and transitory shocks fall within the interval of estimates in the just cited literature.

There are many plausible explanations for the correlation between structural income shocks found in the data. In economic terms, the sign of the covariance may indicate that unfavorable permanent shocks to disposable household income, like the head's long-term unemployment, are partially offset by increases in the transitory income, like unemployment compensations from the government. It is also likely that this offsetting effect will manifest itself at the annual frequency, the frequency I use for modelling the consumption behavior in the life cycle buffer stock 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.³²

From a forecasting perspective, the presence of correlation between structural income shocks helps improve the forecast of one type of shock when a shock of another type arrives.³³ This forecasting ability may be revealed in household choices of consumption over the life cycle.

6 Robustness

In this section, I first briefly discuss the results from a life cycle buffer stock model when other measures of income are used in a matching exercise; I then present the results of a model with partial risk sharing against permanent and transitory shocks.

³²Jacobson et al. (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 is equal 25%, then, at the arrival of the displacement event, this negative permanent shock should be accompanied by the 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.) This paper finds that, on average, permanent and transitory shocks co-move in different directions, and this negative co-movement 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.

³³E.g., if a household observes the permanent shock u_{it}^P , it may forecast the transitory shock using a forecasting regression $u_{it}^T = \rho \frac{\sigma_{u^T}}{\sigma_{u^P}} u_{it}^P + \epsilon_{it}$, where ρ is the correlation between the shocks, and ϵ_{it} is an error term orthogonal to u_{it}^P .

6.1 Other Measures of Income

The main analysis was based on a measure of household real disposable income defined by the sum of the head's and wife's labor income net of taxes and their combined transfer income. Table 10 contains the results for other measures of income: head's and wife's combined labor income—column (2); head's income only—column (3). Compared with the income measure used previously, the measure of column (2) excludes the value of taxes and transfers; the measure of column (3), in addition, excludes the labor income of wife. For these different income measures, I fit the same relative log-consumption profile; the sensitivity of total consumption to income shocks, however, differs across the measures.³⁴ I use an identity weighting matrix in estimations.³⁵ In column (1), I repeat the estimation of Table 8 column (7), now with the identity weighting matrix. The results are very similar. Estimations for income measures not affected by smoothing implicitly present in taxes and transfers (column (2)), and the labor income of wife (column (3)) result in larger estimates of the variances of permanent and transitory shocks; the result of significant negative correlation between the shocks is unaffected by the income measure used.

6.2 Partial Risk Sharing

The main results of the paper 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, i.e., their permanent and transitory idiosyncratic shocks may be partially insured. In a recent paper, Attanasio and Pavoni (2007), for a partial risk sharing model with moral hazard and hidden asset accumulation, showed that the self-insurance Euler equation still holds if one considers the “after-risk-sharing” income. To account for partial risk sharing, I model the after-risk-sharing household income as $Y_{it} = P_{it}(\epsilon_{it}^T)^{\omega_T}$, where $P_{it} = P_{it-1}(\epsilon_{it}^P)^{\omega_P}$, ω_T (ω_P) is the fraction of transitory (permanent) shocks that is insured through different channels other than self-insurance.

³⁴For the measure of column (2) [column (3)], $\hat{\beta}_1 = 0.092$ [0.046]; $\hat{\beta}_2 = 0.032$ [0.017]; $\hat{\beta}_5 = 0.232$ [0.188]; $var(\Delta \log \widehat{Y}_{it}) = 0.102$ [0.174]; $\hat{\alpha} = -0.298$ [-0.298]. For both estimations, I assume that measurement error explains 25% of variation of the growth rate in the respective income measure.

³⁵This choice of the weighting matrix is not dependent on sampling uncertainty of the estimates of the moments used for matching in different estimations.

In Panel C of Table 4, I first show that if $\omega_T = \omega_P = 0.5$, the coefficients that measure the sensitivity of consumption growth to income growth at one and five-year horizons are almost halved relative to the benchmark model where only self-insurance is allowed—row 1. The coefficients are further reduced if 90% of permanent income shocks are insured ($\omega_P = 0.1$) and there is no insurance of transitory shocks beyond self-insurance ($\omega_T = 1.0$)—row 2. If, to the contrary, 90% of transitory shocks are insured before self-insurance ($\omega_T = 0.1$) and there is no partial risk sharing of permanent shocks ($\omega_P = 1.0$), the sensitivity of consumption growth to income growth over the five-year horizon does not change significantly, while the contemporaneous sensitivity of consumption growth to income growth is reduced relative to the benchmark—row 3. These results are expected for this particular choice of partial risk sharing parameters since cumulative income growth over the five-year horizon is largely affected by permanent income shocks, while current income growth is largely affected by the transitory innovation which, on average, has a larger size than the permanent one.

In Table 11, I estimate a life-cycle model with self-insurance and partial risk sharing for different income measures. When net household income is used in estimation, the variance of permanent shocks is substantially above the estimate in Table 10 for the same income measure; the rest of the parameters, while different in terms of point estimates, are within two standard errors from their counterparts in Table 10—column (1). Importantly, the result of the negative correlation between permanent and transitory shocks is robust to the introduction of partial risk sharing into the model. I find that about 37% of permanent shocks are insured through different channels other than self-insurance; I do not find additional insurance of transitory shocks beyond that provided by self-insurance. Self-insurance provides additional smoothing of permanent income shocks. Blundell et al. (2008), using different techniques, for their whole sample find that about 36% of permanent shocks are insured, and about 95% of transitory shocks are insured. My results indicate more insurance of permanent shocks and almost the same insurance of transitory shocks, the latter done through personal savings only.³⁶ Consistent with the results in Blundell et al. (2008), I find more insurance of permanent income shocks when I use income measures that exclude insurance mechanisms provided by taxes and transfers, and

³⁶Note, however, that the estimated time discount factor, the relative risk aversion (and the partial insurance parameters) have larger standard errors compared with the estimates in Table 10 column (1). Perhaps, additional information is needed to identify these parameters more accurately and precisely in a model that allows for partial insurance. I can replicate the point estimate for partial insurance of permanent income shocks in Blundell et al. (2008) if I set ω_P two standard errors above its point estimate (less partial insurance), and the relative risk aversion parameter one standard error below its point estimate.

the labor income of wife—columns (2) and (3), respectively.

7 Conclusions

I analyze the plausible situation when 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, more complicated, income structure.

I argue that households' unique information about income components should be reflected in their consumption behavior. I simulate a life cycle buffer stock model and find that different decompositions of the same reduced form income process imply different model moments such as those related to the sensitivity of consumption growth to current and lagged income growth; the sensitivity of cumulative consumption growth to cumulative income growth over the five-year horizon; and the relative life cycle log-consumption profile. I use these moments estimated from empirical data to identify the parameters of the income process, including the correlation between structural income shocks. The latter is not identified from the univariate dynamics of household income.

I estimate a structural life cycle buffer stock model utilizing household consumption and income data from the CEX and the PSID. I find a significant negative contemporaneous correlation between permanent and transitory shocks of about -0.60 . The estimates of the time discount factor and the relative risk aversion parameter are precise and plausible. I find that the relative risk aversion parameter is about 6.0 and the time discount factor is 0.92. The result of the negative correlation between the shocks is robust to the introduction of partial insurance of permanent and transitory shocks into the model 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 growth to current income growth, and long consumption growth to long income growth. This mechanism allows me to identify the correlation in the data. Correct identification of the components of the income process, in turn, improves specification of the consumption function and estimation of the behavioral parameters. I can reject the no-correlation model in favor of the model with correlated shocks.

The negative correlation between the shocks found in this paper may be the result of deliberate household actions, or some income-smoothing strategies adopted by governments (through taxes and transfers) and firms. New research on income components and income dynamics may shed light on this phenomenon.

References

- Abowd, J. and Card, D. (1989), ‘On the Covariance Structure of Earnings and Hours Changes’, *Econometrica* **57**, 411–445.
- Attanasio, O. and Pavoni, N. (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 Bognanno, M. (2008), ‘Promotions, Demotions, Halo Effects, and the Earnings Dynamics of American Executives’, *Journal of Labor Economics* **26**, 287–310.
- Blundell, R., Pistaferri, L. and Preston, I. (2005*a*), ‘Consumption Inequality and Partial Insurance’. IFS Working Paper #04/28.
- Blundell, R., Pistaferri, L. and Preston, I. (2005*b*), ‘Imputing Consumption in the PSID Using Food Demand Estimates from the CEX’. IFS Working Paper #04/27.
- Blundell, R., Pistaferri, L. and Preston, I. (2008), ‘Consumption Inequality and Partial Insurance’, *American Economic Review*. Forthcoming.
- 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 Deaton, A. (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.
- Carroll, C. D. (1997), ‘Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis’, *Quarterly Journal of Economics* **112**, 1–55.
- Carroll, C. D. (2006), ‘The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems’, *Economics Letters* **91**, 312–320.
- Carroll, C. D. and Samwick, A. A. (1997), ‘The Nature of Precautionary Wealth’, *Journal of Monetary Economics* **40**, 41–71.
- Clark, P. (1987), ‘The Cyclical Component of US Economic Activity’, *Quarterly Journal of Economics* **102**, 797–814.

- Cocco, J., Gomes, F. and Maenhout, P. (2005), ‘Consumption and Portfolio Choice Over the Life Cycle’, *Review of Financial Studies* **18**, 491–533.
- Cochrane, J. H. (1988), ‘How Big Is the Random Walk in GNP?’, *Journal of Political Economy* **96**, 893–920.
- Deaton, A. (1991), ‘Saving and Liquidity Constraints’, *Econometrica* **59**, 1121–1248.
- Deaton, A. (1992), *Understanding Consumption*, Oxford University Press, New York, NY.
- Fernández-Villaverde, J. and Krueger, D. (2007), ‘Consumption Over the Life Cycle: Facts from Consumer Expenditure Survey Data’, *Review of Economics and Statistics* **89**, 552–565.
- French, E. (2005), ‘The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behavior’, *Review of Economic Studies* **72**, 395–427.
- Friedman, M. (1957), *A Theory of the Consumption Function*, Princeton University Press, Princeton, NJ.
- Friedman, M. (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.
- Galbraith, J. and Zinde-Walsh, V. (1994), ‘A Simple Noniterative Estimator for Moving Average Models’, *Biometrika* **1**, 143–155.
- Gomes, F. and Michaelides, A. (2005), ‘Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence’, *Journal of Finance* **60**, 869–904.
- Gourinchas, P.-O. and Parker, J. (2002), ‘Consumption Over the Life Cycle’, *Econometrica* **70**, 47–89.
- Guiso, L., Pistaferri, L. and Schivardi, F. (2005), ‘Insurance within the Firm’, *Journal of Political Economy* **113**, 1054–1087.
- Guvenen, F. (2007), ‘An Empirical Investigation of Labor Income Processes’. NBER Working Paper # 13394.
- Hall, A. R. (2005), *Generalized Method of Moments*, Oxford University Press, New York, NY.
- Hall, R. E. (1978), ‘Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence’, *Journal of Political Economy* **86**, 971–987.
- Hall, R. and Mishkin, F. (1982), ‘The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data of Households’, *Econometrica* **50**, 461–481.
- Hamilton, J. (1994), *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.
- Harvey, A. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.
- 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.
- Jacobson, L. S., LaLonde, R. J. and Sullivan, D. G. (1993), ‘Earnings Losses of Displaced Workers’, *American Economic Review* **83**, 685–709.
- Krebs, T. (2003), ‘Human Capital Risk and Economic Growth’, *Quarterly Journal of Economics* **118**(2), 709–744.
- Laibson, D., Repetto, A. and Tobacman, J. (2007), ‘Estimating Discount Functions with Consumption Choices Over the Lifecycle’. NBER Working Paper # 13314.
- Ludvigson, S. and Michaelides, A. (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 Sørensen, B. (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 Pistaferri, L. (2004), ‘Income Variance Dynamics and Heterogeneity’, *Econometrica* **72**, 1–32.
- Michaelides, A. (2001), ‘Buffer Stock Saving and Habit Formation’. Mimeo, London School of Economics.
- Morley, J., Nelson, C. and Zivot, E. (2003), ‘Why are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?’, *Review of Economics and Statistics* **85**, 235–243.
- Newey, W. K. and McFadden, D. (1994), Large Sample Estimation and Hypotheses Testing, *in* ‘Handbook of Econometrics, Vol.4’, Elsevier, Amsterdam, chapter 36.
- Nielsen, H. S. and Vissing-Jorgensen, A. (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.
- Quah, D. (1990), ‘Permanent and Transitory Movements in Labor Income: An Explanation for “Excess Smoothness” in Consumption’, *Journal of Political Economy* **98**, 449–475.
- Quah, D. (1992), ‘The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds’, *Econometrica* **60**, 107–118.
- Storesletten, K., Telmer, C. and Yaron, A. (2004), ‘Cyclical Dynamics in Idiosyncratic Labor Market Risk’, *Journal of Political Economy* **112**, 213–216.

- Wang, N. (2004), 'Precautionary Saving and Partially Observed Income', *Journal of Monetary Economics* **51**, 1645–1681.
- Zeldes, S. (1989*a*), 'Consumption and Liquidity Constraints: An Empirical Investigation', *Journal of Political Economy* **142**, 305–346.
- Zeldes, S. (1989*b*), 'Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence', *Quarterly Journal of Economics* **104**, 275–298.
- Ziliak, J. (1998), 'Does the Choice of Consumption Matter? An Application to the Permanent Income Hypothesis', *Journal of Monetary Economics* **41**, 201–216.

FIGURE 1: HOUSEHOLD CONSUMPTION PROFILE OVER THE LIFE CYCLE. TOTAL PSID SAMPLE.




FIGURE 2: HOUSEHOLD RELATIVE LOG-CONSUMPTION PROFILE OVER THE LIFE CYCLE. AVERAGE LOG-CONSUMPTION AT AGES 25-64 RELATIVE TO THE AVERAGE LOG-CONSUMPTION AT AGE 24.

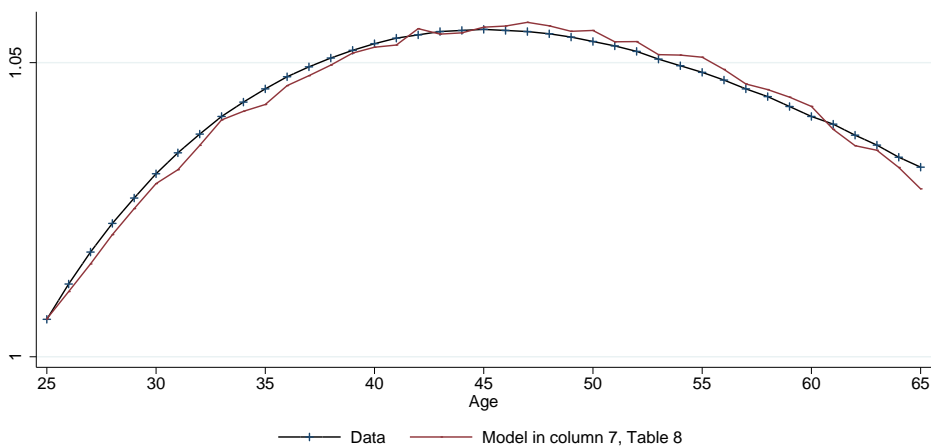


TABLE 1: AUTO-COVARIANCES OF THE FIRST DIFFERENCES IN LOG-IDIOSYNCRATIC INCOME

Auto-Covariance of Order	Value	St. Error
0	0.0864	0.0062
1	-0.0242	0.0040
2	-0.0044	0.0033
3	-0.0026	0.0032
4	-0.0008	0.0031

Notes: Income data are the residuals from cross-sectional regressions of the sum of head’s and wife’s real disposable income on the head’s education, household state of residence, a second degree polynomial in the head’s age, and the head’s race. The data are drawn from 1969–1997 annual family files of the PSID. I select households with married heads aged between 24 and 65. For each year, cross sectional regressions are performed for seven age groups: heads of age 24–29, heads of age 30–35, . . . , heads of age 60–65. I drop top-coded observations, and observations for households with reported zero head’s labor income. In addition, I drop observations with an absolute percentage change in the income residual greater than or equal to 300% or with real income below 1,000 of 1982–1984 dollars. A household contributes an observation on income difference if it has stable family composition between years t and $t - 1$. A household is present in the data if at least one income difference is non-missing. Auto-covariances represent the averages of unique elements of the variance-covariance matrix of log income changes corresponding to the theoretical auto-covariances of a given order.

TABLE 2: TEST OF THE NULL HYPOTHESIS OF ZERO AUTO-COVARIANCE IN ALL TIME PERIODS

Order	Statistics	DF	p-value
1	750	27	0.00
2	68.91	26	$9.55 * 10^{-6}$
3	30.46	25	0.21
4	31.02	24	0.15

Notes: The test statistic is distributed as χ^2 with degrees of freedom equal to the number of (zero) restrictions (the number of unique auto-covariances of a given order in the estimated variance-covariance matrix). See Appendix A of Abowd and Card (1989) for details of its construction.

TABLE 3: ESTIMATES OF THE REDUCED FORM INCOME PROCESS

Parameter	Value
$\hat{\theta}$	-0.3624 (0.0076)
$\hat{\sigma}_u^2$	0.0514 (0.0009)
DF	404
χ^2	807.29 (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-income, θ is a moving average parameter, and $u_{it} \sim iid(0, \sigma_u^2)$. The model is estimated by the method of optimal minimum distance. Standard errors in parentheses; p-value for the quality of the fit in parentheses, under the value of the χ^2 statistics.

TABLE 4: SENSITIVITY OF CONSUMPTION TO INCOME, RELATIVE LIFE CYCLE LOG-CONSUMPTION PROFILE, AND THE REDUCED FORM INCOME PROCESS PARAMETERS FOR THE DATA FROM A SIMULATED MODEL

Model [‡]	$\overline{\log C}_{45}/\overline{\log C}_{30}$	$\overline{\log C}_{65}/\overline{\log C}_{45}$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\hat{\beta}_5$	$\hat{\alpha}$	var. inc. growth
<u>PANEL A</u>							
Benchmark $corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$	1.05 (0.016)	1.005 (0.014)	0.121 (0.004)	0.044 (0.001)	0.421 (0.012)	-0.317 (0.002)	0.059 (0.0002)
Benchmark $corr(\epsilon_{it}^P, \epsilon_{it}^T) = 0.0$	1.05 (0.017)	1.004 (0.015)	0.268 (0.013)	0.089 (0.003)	0.486 (0.019)	-0.314 (0.002)	0.058 (0.0002)
Benchmark $corr(\epsilon_{it}^P, \epsilon_{it}^T) = 0.5$	1.05 (0.011)	1.003 (0.009)	0.353 (0.011)	0.116 (0.003)	0.522 (0.013)	-0.314 (0.002)	0.058 (0.0002)
<u>PANEL B</u>							
$\sigma_{it}^C = 0.24$	-0.24%	-0.23%	0.84%	0.51%	0.41%	-0.16%	0.3%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
$\sigma_{it}^C = 0.24, \sigma_{it}^Y = 0.085$	-0.20%	-0.20%	13.98%**	11.03%	2.02%	-0.48%	-0.12%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
$\beta = 0.90$	-2.92%	-4.63%**	17.90%	17.44%*	10.83%	-0.28%	0.09%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
$\beta = 1.0$	3.0%	5.12%**	-6.78%	-9.81%***	-3.57%	-0.20%	-0.15%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
$\rho = 1.0$	-5.23%***	-6.86%***	58.88%*	43.09%*	26.47%**	-0.17%	0.10%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
$\rho = 7.0$	2.78%	2.24%	-63.69%***	-9.81%***	-3.38%	-0.04%	-0.13%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
<u>PANEL C</u>							
Benchmark, $\omega_P = \omega_T = 0.5$	-4.0%*	-2.28%	-44.85%***	-46.66%***	-45.04%***	-0.07%	-0.15%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
Benchmark, $\omega_P = 0.1, \omega_T = 1.0$	-5.62%***	-3.09%*	-60.50%***	-61.76%***	-86.59%***	-0.11%	-0.11%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							
Benchmark, $\omega_P = 1.0, \omega_T = 0.1$	0.23%	0.15%	-15.67%***	-13.80%***	1.27%	-0.09%	-0.07%
$corr(\epsilon_{it}^P, \epsilon_{it}^T) = -0.5$							

Notes: [‡] $\hat{\beta}_1$, and $\hat{\beta}_2$ are the slope coefficients from an OLS regression $\Delta \log C_{it} = const + \beta_1 * \Delta \log Y_{it} + \beta_2 * \Delta \log Y_{it-1} + error$. $\hat{\beta}_5$ is the slope coefficient from an OLS regression $\Delta_5 \log C_{it} = const + \beta_5 * \Delta_5 \log Y_{it} + error$, where $\Delta_k \log x_t = \log x_t - \log x_{t-k}$. Income persistence is the slope coefficient from an OLS AR(1) regression $\Delta \log Y_{it} = const + \alpha * \Delta \log Y_{it-1} + error$. Each simulated economy is populated by 5,000 ex ante identical consumers observed during 42 periods (ages 24–65). $\overline{\log C}_a = (1/5000) \sum_{i=1}^{5000} \log C_{ia}$, where a denotes age. The results in rows 1-3 of Panel A are the averages of the corresponding statistics over 50 repetitions. Standard errors, calculated as the standard deviations of the corresponding statistics over 50 repetitions, in parentheses. In panel B, the cell entries are the percentage changes of the model statistics relative to those in the first row of Panel A. In the benchmark model of the first row of Panel A, the time discount factor is 0.95, the coefficient of relative risk aversion is 4.0, the standard deviation of permanent shocks is 0.1445, the standard deviation of transitory shocks is 0.1773, the covariance between the shocks is -0.0128, there are no measurement errors in income and consumption, the reduced form income model's parameters are as in Table 3. When measurement error in income is allowed, $\sigma_{uT} = 0.1487$, and $cov(u_{it}^P, u_{it}^T) = -0.0107$. In Panel C, idiosyncratic “after-risk-sharing” labor income is $Y_{it} = P_{it}(\epsilon_{it}^T)^{\omega_T}$, where $P_{it} = P_{it-1}(\epsilon_{it}^P)^{\omega_P}$, ω_T (ω_P) is the fraction of transitory (permanent) shocks that is insured. *** significant at the 1% level, ** significant at the 5% level, * significant at the 10% level.

TABLE 5: SENSITIVITY OF CONSUMPTION GROWTH TO CURRENT AND LAGGED INCOME GROWTH

Dependent Variable: $\Delta \widetilde{\log C_{it}}$	
Variable	Coefficient
$\Delta \widetilde{\log Y_{it}}$	0.06 (0.02)
$\Delta \widetilde{\log Y_{it-1}}$	0.01 (0.02)
Δf_a	-0.016 (0.098)
Age_{it}	-0.003 (0.005)
$Age_{it}^2/100$	0.002 (0.005)
Constant	0.137 (0.123)
Number of households	2,359
Number of observations	10,923
Adj. R^2	0.001

Notes: Income data are from the 1980–1999 surveys of the PSID. Income is measured as the combined labor income of head and wife, net of federal income taxes, social security, and Medicare taxes, plus the combined transfer income of head and wife. Total consumption is imputed to the PSID households. Consumption and income data are “cleaned” of cohort, time, and idiosyncratic family size effects. Family size is the sample average of family size across households of a certain age. Standard errors in parentheses.

TABLE 6: SENSITIVITY OF CUMULATIVE CONSUMPTION TO CUMULATIVE INCOME GROWTH OVER THE 5-YEAR HORIZON

Dependent Variable: $\widetilde{\Delta_5 \log C_{it}}$	
Variable	Coefficient
$\widetilde{\Delta_5 \log Y_{it}}$	0.247 (0.03)
$\Delta_5 f_a$	0.193 (0.083)
Age_{it}	0.037 (0.02)
$Age_{it}^2/100$	-0.039 (0.019)
Constant	-0.649 (0.492)
Number of households	1,307
Number of observations	3,314
Adj. R ²	0.032

Notes: See notes to Table 5. $\Delta_5 \log x_{it}$ is defined by $\log x_{it} - \log x_{it-5}$. Standard errors in parentheses.

TABLE 7: PARAMETERS OF THE REDUCED FORM INCOME PROCESS

Dependent Variable: $\widetilde{\Delta \log Y_{it}}$	
Variable	Coefficient
$\widetilde{\Delta \log Y_{it-1}}$	-0.347 (0.01)
Constant	-0.012 (0.003)
Number of households	2,359
Number of observations	10,923
Adj. R ²	0.108
$var(\widetilde{\Delta \log Y_{it}})$	0.082 (0.249)
Number of households	2,453

Notes: See notes to Table 5. Standard errors in parentheses.

TABLE 8: ESTIMATION RESULTS: DIAGONAL WEIGHTING MATRIX

Moments*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	$\hat{\beta}_1, \hat{\beta}_2, \hat{\alpha}, \hat{\beta}_5,$	as in (1), and $\overline{\log C_{45}}/\overline{\log C_{30}},$	as in (2)	as in (2)	as in (2)	as in (2)	as in (1), and	as in (7)	as in (7), without	as in (7)
	$var(\Delta \log Y_{it})$	$\overline{\log C_{65}}/\overline{\log C_{45}}$					life cycle profile		$\hat{\beta}_2$	
CRRRA estimated?	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cons. measur. error	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
Inc. measur. error	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Time Discount Factor, $\hat{\beta}$	0.9997	0.9364	0.9128	0.9128	0.9143	0.9153	0.8939	0.8337	0.8938	0.8980
	(1.5281)	(0.0305)	(0.0592)	(0.0591)	(0.0594)	(0.0585)	(0.0698)	(0.1345)	(0.0699)	(0.0559)
CRRRA, $\hat{\rho}$	4.00	4.00	6.3905	6.3867	6.3877	6.3765	7.2111	12.7004	7.2165	6.4666
	—	—	(2.6629)	(2.6578)	(2.704)	(2.6586)	(2.7064)	(6.1133)	(2.7053)	(1.1408)
$\hat{\sigma}_P$	0.1236	0.1455	0.1205	0.1205	0.1181	0.1194	0.1219	0.0856	0.1218	0.1219
	(0.0165)	(0.0157)	(0.0156)	(0.0156)	(0.0156)	(0.0156)	(0.0151)	(0.0146)	(0.0149)	—
$\hat{\sigma}_T$	0.2152	0.2322	0.2188	0.2188	0.1589	0.1605	0.1662	0.1329	0.1661	0.1399
	(0.0769)	(0.0155)	(0.0098)	(0.0099)	(0.0127)	(0.0127)	(0.0116)	(0.0116)	(0.0116)	(0.0136)
$\widehat{cov}(u_{it}^P, u_{it}^T)$	-0.0132	-0.0214	-0.0143	-0.0143	-0.0115	-0.0122	-0.0131	0	-0.0131	0
	(0.0339)	(0.0066)	(0.0041)	(0.0042)	(0.0038)	(0.0039)	(0.0038)	—	(0.0038)	—
Implied $\widehat{corr}(u_{it}^P, u_{it}^T)$	-0.496	-0.633	-0.541	-0.541	-0.612	-0.634	-0.649	0	-0.647	0
	(1.0726)	(0.1105)	(0.1053)	(0.1079)	(0.1316)	(0.1308)	(0.1253)	—	(0.1277)	—
χ^2	9.18	12.61	11.62	10.76	11.54	10.84	18.11	35.61	18.05	32.24
Degrees of freedom	1	3	2	2	2	2	41	42	40	43
P-value	0.002	0.006	0.003	0.005	0.003	0.004	0.99	0.75	0.99	0.89

Notes:

* $\hat{\beta}_1$, and $\hat{\beta}_2$ are the slope coefficients from an OLS regression $\Delta \log C_{it} = const + \beta_1 * \Delta \log Y_{it-1} + \beta_2 * \Delta \log Y_{it-1} + error$. $\hat{\beta}_5$ is the slope coefficient from an OLS regression $\Delta_5 \log C_{it} = const + \beta_5 * \Delta_5 \log Y_{it} + error$, where $\Delta_k \log x_t = \log x_t - \log x_{t-k}$. Income persistence is the slope coefficient from an OLS AR(1) regression $\Delta \log Y_{it} = const + \alpha * \Delta \log Y_{it-1} + error$. Models are estimated by the method of simulated moments. Standard errors corrected for uncertainty in the first stage estimation in parentheses. Standard error for an implied correlation is calculated by the delta method. Weighting matrix used for estimation is the diagonal matrix, where diagonal elements are equal to the diagonal from the inverse of the sample variance-covariance matrix of the second stage moments.

TABLE 9: ESTIMATION RESULTS: OPTIMAL WEIGHTING MATRIX

Moments	(1)	(2)
	life cycle profile, $var(\widetilde{\Delta \log Y_{it}})$ $\hat{\beta}_1, \hat{\beta}_2, \hat{\alpha}, \hat{\beta}_5$	as in (1)
CRRA estimated?	Yes	Yes
Cons. measur. error	Yes	Yes
Inc. measur. error	Yes	Yes
Time Discount Factor, $\hat{\beta}$	0.9208 (0.0368)	0.8956 (0.0589)
CRRA, $\hat{\rho}$	5.9729 (1.9566)	8.5734 (3.1797)
$\hat{\sigma}_P$	0.1285 (0.0194)	0.1021 (0.0171)
$\hat{\sigma}_T$	0.1717 (0.0201)	0.1482 (0.0163)
$\widehat{cov}(u_{it}^P, u_{it}^T)$	-0.0131 (0.0069)	0 —
Implied $\widehat{corr}(u_{it}^P, u_{it}^T)$	-0.5937 (0.2170)	0 —
χ^2	15.10	26.45
Degrees of freedom	41	42
P-value	0.99	0.97
Distance statistic: (2) vs. (1)	11.35 (0.003)	

Notes: See notes to Table 8. Optimal weighting matrix is constructed using the parameter estimates in column (7) of Table 8. p-value for the distance statistic in parentheses, next to its value.

TABLE 10: ESTIMATION RESULTS: DIFFERENT INCOME MEASURES; IDENTITY WEIGHTING MATRIX

Inc. measure/moments	(1) Net head's+wife's income/ life cycle profile, $\widehat{var}(\Delta \log Y_{it})$ $\hat{\beta}_1, \hat{\beta}_2, \hat{\alpha}, \hat{\beta}_5$	(2) Head's+wife's labor income/ as in (1)	(3) Head's labor income/ as in (1)
CRRA estimated?	Yes	Yes	Yes
Cons. measur. error	Yes	Yes	Yes
Inc. measur. error	Yes	Yes	Yes
Time Discount Factor, $\hat{\beta}$	0.912 (0.054)	0.928 (0.041)	0.931 (0.033)
CRRA, $\hat{\rho}$	6.924 (2.424)	5.279 (1.752)	3.989 (1.303)
$\hat{\sigma}_P$	0.117 (0.016)	0.131 (0.018)	0.161 (0.022)
$\hat{\sigma}_T$	0.159 (0.013)	0.164 (0.014)	0.238 (0.014)
$\widehat{cov}(u_{it}^P, u_{it}^T)$	-0.012 (0.004)	-0.013 (0.005)	-0.028 (0.008)
Implied $\widehat{corr}(u_{it}^P, u_{it}^T)$	-0.659 (0.132)	-0.591 (0.144)	-0.732 (0.124)

Notes: See notes to Table 8.

TABLE 11: ESTIMATION RESULTS: DIFFERENT INCOME MEASURES; PARTIAL RISK SHARING;
IDENTITY WEIGHTING MATRIX

Inc. measure/moments	(1) Net head's+wife's income/ life cycle profile, $\widehat{var}(\Delta \log \widehat{Y}_{it})$ $\hat{\beta}_1, \hat{\beta}_2, \hat{\alpha}, \hat{\beta}_5$	(2) Head's+wife's labor income/ as in (1)	(3) Head's labor income/ as in (1)
CRRA estimated?	Yes	Yes	Yes
Cons. measur. error	Yes	Yes	Yes
Inc. measur. error	Yes	Yes	Yes
Time Discount Factor, $\hat{\beta}$	0.869 (0.102)	0.861 (0.112)	0.895 (0.071)
CRRA, $\hat{\rho}$	10.389 (5.497)	10.555 (5.858)	7.647 (4.143)
$\hat{\sigma}_P$	0.160 (0.013)	0.202 (0.012)	0.267 (0.017)
$\hat{\sigma}_T$	0.165 (0.017)	0.146 (0.031)	0.243 (0.029)
$\widehat{cov}(u_{it}^P, u_{it}^T)$	-0.018 (0.006)	-0.016 (0.010)	-0.051 (0.016)
Implied $\widehat{corr}(u_{it}^P, u_{it}^T)$	-0.691 (0.158)	-0.554 (0.228)	-0.792 (0.134)
$\hat{\omega}_P$	0.630 (0.117)	0.492 (0.099)	0.451 (0.084)
$\hat{\omega}_T$	1.0 (1.109)	1.0 (1.671)	1.0 (0.382)

Notes: See notes to Table 8. ω_T (ω_P) is the fraction of transitory (permanent) shocks that is insured.

Appendix A: Matching the Moments of the Reduced Form and UC Models

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

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

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

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

subject to the constraint that the spectrum of the reduced form series and the spectrum of 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, $cov(u_{it}^P, u_{it}^T)$ and σ_{u^T} . Since the two equations have three unknowns, there is no unique solution for $cov(u_{it}^P, u_{it}^T)$ 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:

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

Appendix A1: Structural Income Processes

The reduced form process for the growth in idiosyncratic income: $\Delta \log Y_{it} = (1 + \theta L)u_{it}$; $\hat{\theta} = -0.3624$, $\hat{\sigma}_u^2 = 0.0514$. The structural income processes are of the form: $\Delta \log Y_{k,it} = u_{k,it}^P + \Delta u_{k,it}^T$, $k = 1, \dots, 9$.

- Model (1): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.2267$, $cov(u_{it}^P, u_{it}^T) = -0.0328$, $\rho_{u^T u^P} = -1.00$.
- Model (2): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.2010$, $cov(u_{it}^P, u_{it}^T) = -0.0218$, $\rho_{u^T u^P} = -0.75$.
- Model (3): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.1773$, $cov(u_{it}^P, u_{it}^T) = -0.0128$, $\rho_{u^T u^P} = -0.50$.
- Model (4): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.1557$, $cov(u_{it}^P, u_{it}^T) = -0.0056$, $\rho_{u^T u^P} = -0.25$.
- Model (5): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.1364$, $cov(u_{it}^P, u_{it}^T) = 0.00$, $\rho_{u^T u^P} = 0.00$.
- Model (6): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.1196$, $cov(u_{it}^P, u_{it}^T) = 0.0043$, $\rho_{u^T u^P} = 0.25$.
- Model (7): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.1051$, $cov(u_{it}^P, u_{it}^T) = 0.0076$, $\rho_{u^T u^P} = 0.50$.
- Model (8): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.0926$, $cov(u_{it}^P, u_{it}^T) = 0.01$, $\rho_{u^T u^P} = 0.75$.
- Model (9): $\sigma_{u^P} = 0.1445$, $\sigma_{u^T} = 0.0822$, $cov(u_{it}^P, u_{it}^T) = 0.0119$, $\rho_{u^T u^P} = 1.00$.

For each period t , permanent and transitory income shocks are the respective entries of the 2×1 vector, $exp(chol(\Sigma_{u^T u^P}) \times \epsilon_t)$, where $chol(\Sigma_{u^T u^P})$ is the Cholesky factor of the variance-covariance matrix of structural log-income shocks, and ϵ_t is the 2×1 vector of independent random normal draws.

Appendix B: Details of Solution of the Life Cycle Consumption Model, and Numerical Procedures

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

(i) Model Solution.

I assume that the consumption function at age 66 is linear in total wealth: $c_{T+1} = \kappa_x x_{T+1} + \kappa_0$. The Euler equation that links consumption at age 65 and 66 is: $v(Z_T)U'(c_T) = v(Z_{T+1})\beta RE_T U'(c_{T+1})$, and, therefore, consumption at age 65 is $c_T = (\frac{v(Z_T)}{v(Z_{T+1})})^{\frac{1}{\rho}} (\beta R)^{-\frac{1}{\rho}} (\kappa_x x_{T+1} + \kappa_0)$, or $c_T = (\frac{v(Z_T)}{v(Z_{T+1})})^{\frac{1}{\rho}} (\beta R)^{-\frac{1}{\rho}} (\kappa_x R(x_T - c_T) + \kappa_x + \kappa_0)$. The last equality follows from the assumption that income at age 66 is non-stochastic, and is equal to the permanent income at age 65. For each value of $s_T = (x_T - c_T)$, the last equation uniquely determines c_T . Cash-on-hand x_T is obtained by summing s_T and c_T . The consumption function at age 65, the last working period, $c_T(x_T)$ is obtained by linear interpolation. For ages 24-64, the consumption function is determined by iterating the Euler equation, $c_t = \min\langle x_t, \gamma^* \left[E_t(c_{t+1} (\frac{R(x_t - c_t)}{G_{t+1} u_{t+1}^P} + u_{t+1}^T))^{-\rho} (G_{t+1} u_{t+1}^P)^{-\rho} \right]^{-1/\rho} \rangle$,

where $\gamma^* = (\frac{v(Z_t)}{v(Z_{t+1})})^{\frac{1}{\rho}} (\beta R)^{-\frac{1}{\rho}}$.³⁷ I create a grid of 120 equally spaced points between 0 and 40 in $s_t = x_t - c_t$, and a 250×2 matrix of pre-seeded independent random normal deviates ϵ_1 , and ϵ_0 . Expectation is taken over potentially correlated shocks u^P and u^T , permanent and transitory shocks respectively. Correlated log-normal random draws are obtained by taking the elements of the vector $\exp[\text{chol}(\Sigma_{u^T u^P})\epsilon]$, where $\text{chol}(\Sigma_{u^T u^P})$ is the Cholesky factor of the variance-covariance matrix of log-permanent and log-transitory shocks, and ϵ is the vector of independent random standard normal deviates. This way I obtain a 250×2 matrix of correlated random log-normal deviates. $E_t(\cdot)$ is calculated as $\frac{1}{250} \sum_{j=1}^{250} \left[c_{t+1} (\frac{R s_t}{G_{t+1} u_j^P} + u_j^T) \right]^{-\rho} [G_{t+1} u_j^P]^{-\rho}$ 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. x_t is 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 x_t and c_t . If x_t falls outside an upper bound of the grid, the corresponding consumption is calculated as the value of consumption from the life cycle model without uncertainty. Using the budget constraint and the Euler equation, it is straightforward to show that this value is defined

as $C_t = \frac{1 - (\beta R)^{\frac{1}{\rho}} R^{-1}}{1 - [(\beta R)^{\frac{1}{\rho}} R^{-1}]^{D-t+1}} \left(A_t + \sum_{i=0}^D R^{-i} Y_{t+i} \right)$, where D is the age of death (assumed to be 90), and A_t is the level of liquid assets at age t . Dividing the equation by P_t , and assuming that income grows deterministically at the gross growth rate G_{t+1} from t to D (set to be 1.0 for all time periods), the value of consumption per permanent income is calculated as $c_t = \frac{1 - (\beta R)^{\frac{1}{\rho}} R^{-1}}{1 - [(\beta R)^{\frac{1}{\rho}} R^{-1}]^{D-t+1}} \left(a_t + \frac{(1 - \frac{G_{t+1}}{R})^{D-t+1}}{1 - \frac{G_{t+1}}{R}} \right)$, where a_t is the value of liquid assets relative to the permanent income at age t .

(ii) Initialization of the Model Economy.

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 $\{c_t(x_{it})\}_{t=24}^{65}$. In the beginning period of the life cycle, at age 24, each household obtains a draw from the estimated empirical distribution of the wealth-to-permanent income ratio at age 24, and a transitory shock. The sum of these draws defines initial cash-on-hand. I use the estimates of mean and variance of the distribution of initial liquid wealth-to-permanent income ratio to generate a 5000×1 vector of pre-seeded log-normal draws from this distribution. To initialize the permanent income, for each household I take a draw from the estimated distribution of permanent income at age 24. I describe estimation of parameters of both initial distributions in Appendix D. Once the economy is generated, I calculate the (second stage) model moments used in optimization. During optimization search, created sequences of random draws are kept the same.

(iii) Optimization Procedure.

I use the constrained optimization module (CO) in GAUSS to find the model parameters $\hat{\theta}$. I impose the following bounds on the parameters: $\beta \in [0, 1]$, $\rho \in [0, \infty)$, $\sigma_{u^T} \in [0, \infty)$, $\sigma_{u^P} \in [0, \infty)$, and

³⁷The solution method is an application of the endogenous gridpoints method in Carroll (2006).

$cov(u_{it}^P, u_{it}^T) \in (-\infty, \infty)$. To obtain a “proper” variance-covariance matrix of income shocks within an optimization routine, I constrain it to be positive definite. Specifically, I require that the eigenvalues of the variance-covariance matrix are positive.

Appendix C: MSM Standard Errors and an Over-identification Test

In this appendix, I discuss estimation of standard errors of the MSM parameters, and construction of an over-identification statistic.

For each household in the data, the difference between the data moments and simulated moments is defined by the vector:

$$h_i^d - h^s = \begin{pmatrix} \log \widetilde{C}_{i,25} / \log C_{24} - \log C_{25}^s / \log C_{24}^s(\theta, \hat{\chi}) \\ \log \widetilde{C}_{i,26} / \log C_{24} - \log C_{26}^s / \log C_{24}^s(\theta, \hat{\chi}) \\ \vdots \\ \log \widetilde{C}_{i,65} / \log C_{24} - \log C_{65}^s / \log C_{24}^s(\theta, \hat{\chi}) \\ \widehat{\beta}_{1,i} - \beta_1^s(\theta, \hat{\chi}) \\ \widehat{\beta}_{2,i} - \beta_2^s(\theta, \hat{\chi}) \\ \widehat{\alpha}_i - \alpha^s(\theta, \hat{\chi}) \\ \widehat{\beta}_{5,i} - \beta_5^s(\theta, \hat{\chi}) \\ \frac{1}{T} \sum_{t=1}^T (\Delta \log \widetilde{Y}_{it})^2 - \frac{1}{I_s * T} \sum_{t=1}^T \sum_{i=1}^{I_s} (\Delta \log Y_{it}^s)^2(\theta, \hat{\chi}) \end{pmatrix}.$$

The superscript d denotes data; s denotes simulation; I_d is the number of households in the data contributing towards estimation of the moment; I_s is the number of simulated households; and $\hat{\chi}$ is a vector of the pre-estimated first stage moments.

The cross-sectional population mean of the vector, $g_{I_s} = \frac{1}{I_d} \sum_{i=1}^{I_d} h_i^d - h^s$, is equal to a vector of zeros. The MSM estimate $\hat{\theta}$ is the minimizer of the criterion function:

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

Note, for example, that persistence of the reduced form income process is the second element of the vector

$$\hat{\alpha} = \frac{1}{I_d} \sum_{i=1}^{I_d} \hat{\alpha}_i = \left(\sum_{i=1}^{I_d} \sum_{t=1}^T x'_{it} x_{it} \right)^{-1} \left(\sum_{i=1}^{I_d} \sum_{t=1}^T x'_{it} y_{it} \right),$$

where $x_{it} \equiv (1 \quad \Delta \log \widetilde{Y}_{it-1})'$, $y_{it} \equiv \Delta \log \widetilde{Y}_{it}$, and

$$\hat{\alpha}_i = I_d \left(\sum_{i=1}^{I_d} \sum_{t=1}^T x'_{it} x_{it} \right)^{-1} \left(\sum_{t=1}^T x'_{it} y_{it} \right),$$

such that it is ensured that $\frac{1}{I_d} \sum_{i=1}^{I_d} \hat{\alpha}_i$ returns $\hat{\alpha}$.

Following Newey and McFadden (1994), Laibson et al. (2007), and Gourinchas and Parker (2002), it can be shown that

$$\sqrt{I_d}(\hat{\theta} - \theta_0) = - \underbrace{[g_{I_s\theta}(\hat{\theta}; \hat{\chi})' W g_{I_s\theta}(\bar{\theta}; \hat{\chi})]^{-1} g_{I_s\theta}(\hat{\theta}; \hat{\chi})' W}_{\Pi} \sqrt{I_d} g_{I_s}(\theta_0; \hat{\chi}),$$

where θ_0 is a vector of unknown (“true”) structural parameters; $\bar{\theta}$ is the mean value; and $g_{I_s\theta}(\cdot; \cdot)$ is a vector of the first partial derivatives of $g_{I_s}(\theta; \hat{\chi})$ with respect to θ .

Expand $g_{I_s}(\theta_0; \hat{\chi})$ around χ_0 , a vector of unknown (“true”) first stage moments, to obtain:

$$\sqrt{I_d}(\hat{\theta} - \theta_0) = -\Pi \sqrt{I_d} [g_{I_s}(\theta_0; \chi_0) + g_{I_s\chi}(\theta_0; \chi_0)(\hat{\chi} - \chi_0)],$$

where $g_{I_s\chi}(\theta; \hat{\chi})$ is a vector of the first partial derivatives of $g_{I_s}(\cdot; \cdot)$ with respect to χ .

By definition,

$$\begin{aligned} \sqrt{I_d} g_{I_s}(\theta_0; \chi_0) &\equiv \sqrt{I_d} \left[\frac{1}{I_d} \sum_{i=1}^{I_d} h_i^d - h^s(\theta_0; \chi_0) \right] \\ &= \underbrace{\sqrt{I_d} \left[\frac{1}{I_d} \sum_{i=1}^{I_d} h_i^d - h(\theta_0; \chi_0) \right]}_A + \\ &\quad + \underbrace{\sqrt{I_d} [h(\theta_0; \chi_0) - h^s(\theta_0; \chi_0)]}_B, \end{aligned}$$

where A is the (scaled) difference between the sample and “true” moments, and B is the (scaled) difference between the “true” and simulated moments.

Note that

$$A \equiv \sqrt{I_d} g(\theta_0; \chi_0) \longrightarrow N(0, \Sigma_g).$$

For the moments defining the relative log-consumption profile, the typical diagonal element of Σ_g is:

$$E \left(\sqrt{I_d} \left[\frac{1}{I_d} \sum_{i=1}^{I_d} \left(\frac{\log \widetilde{C}_{i,65}}{\log C_{45}} - \frac{\log \overline{C}_{65}}{\log C_{45}} \right) \right] \right) \left(\sqrt{I_d} \left[\frac{1}{I_d} \sum_{i=1}^{I_d} \left(\frac{\log \widetilde{C}_{i,65}}{\log C_{45}} - \frac{\log \overline{C}_{65}}{\log C_{45}} \right) \right] \right)'$$

Define $\frac{\log \widetilde{C}_{i,65}}{\log C_{45}} - \frac{\log \overline{C}_{65}}{\log C_{45}} \equiv \widetilde{\log C}_{i,65/45}$. Then, assuming that household deviations $\widetilde{\log C}_{i,65/45}$ are independent, an empirical analog of the above equation can be calculated by:

$$\begin{aligned} &\frac{1}{I_d} E \left(\sum_{i=1}^{I_d} \widetilde{\log C}_{i,65/45} \right) \left(\sum_{i=1}^{I_d} \widetilde{\log C}'_{i,65/45} \right) \\ &= \frac{1}{I_d} \left[E(\widetilde{\log C}_{1,65/45} \widetilde{\log C}'_{1,65/45}) + E(\widetilde{\log C}_{2,65/45} \widetilde{\log C}'_{2,65/45}) + \dots + E(\widetilde{\log C}_{I_d,65/45} \widetilde{\log C}'_{I_d,65/45}) \right] \\ &= \text{var}(\widetilde{\log C}_{i,65/45}). \end{aligned}$$

Similarly, for the OLS moments, the typical diagonal element of Σ_g is:

$$E \left(\sqrt{I_d} \left[\frac{1}{I_d} \sum_{i=1}^{I_d} (\hat{\alpha}_i - E(\alpha_i)) \right] \right) \left(\sqrt{I_d} \left[\frac{1}{I_d} \sum_{i=1}^{I_d} (\hat{\alpha}_i - E(\alpha_i)) \right] \right)'.$$

Define $\hat{\alpha}_i - E(\alpha_i) \equiv \tilde{\alpha}_i$. Then, assuming that household OLS errors are independent, an empirical analog of the above equation can be calculated as:

$$\begin{aligned} E \left(\sqrt{I_d} \frac{1}{I_d} \sum_{i=1}^{I_d} \tilde{\alpha}_i \right) \left(\sqrt{I_d} \frac{1}{I_d} \sum_{i=1}^{I_d} \tilde{\alpha}_i \right)' &= \frac{1}{I_d} E \left[\left(\sum_{i=1}^{I_d} \tilde{\alpha}_i \right) \left(\sum_{i=1}^{I_d} \tilde{\alpha}_i \right)' \right] = \frac{1}{I_d} [E(\tilde{\alpha}_1 \tilde{\alpha}_1') + \dots + E(\tilde{\alpha}_{I_d} \tilde{\alpha}_{I_d}')] \\ &= \text{var}(\hat{\alpha}_i). \end{aligned}$$

Also,

$$B \equiv \sqrt{I_d} \frac{\sqrt{I_s}}{\sqrt{I_s}} [h(\theta_0; \chi_0) - h^s(\theta_0; \chi_0)] = \frac{\sqrt{I_d}}{\sqrt{I_s}} \sqrt{I_s} [h(\theta_0; \chi_0) - h^s(\theta_0; \chi_0)] \longrightarrow N(0, \Sigma_g \frac{I_d}{I_s}).$$

Thus,

$$\sqrt{I_d} g_{I_s}(\theta_0; \chi_0) \longrightarrow N(0, \Sigma_g) + N(0, \Sigma_g \frac{I_d}{I_s}) \longrightarrow N(0, \Sigma_g [1 + \frac{I_d}{I_s}]).$$

Note also that

$$\sqrt{I_d} g_{I_s \chi}(\theta_0; \chi_0) (\hat{\chi} - \chi_0) = \frac{\sqrt{I_d}}{\sqrt{I_1}} G_\chi \sqrt{I_1} (\hat{\chi} - \chi_0) \equiv C \longrightarrow N(0, \frac{I_d}{I_1} G_\chi \Sigma_\chi G_\chi'),$$

where I_1 is the scalar rate of convergence of the first stage moments and $G_\chi \equiv g_{I_s \chi}(\theta_0; \chi_0)$, consistently estimated by $g_{I_s \chi}(\hat{\theta}; \hat{\chi})$.

For the first stage moments defining deviations of an individual income growth from the average income growth profile, the typical diagonal element of Σ_χ is:

$$E \left[\left(\sqrt{I_1} \frac{1}{I_1} \sum_{i=1}^{I_1} (\log \widetilde{Y}_{i,25} - \log \widetilde{Y}_{i,24} - (\overline{\log Y}_{25} - \overline{\log Y}_{24})) \right) \left(\sqrt{I_1} \frac{1}{I_1} \sum_{i=1}^{I_1} (\log \widetilde{Y}_{i,25} - \log \widetilde{Y}_{i,24} - (\overline{\log Y}_{25} - \overline{\log Y}_{24})) \right)' \right].$$

Define $\log \widetilde{Y}_{i,25} - \log \widetilde{Y}_{i,24} - (\overline{\log Y}_{25} - \overline{\log Y}_{24}) \equiv \widetilde{\log Y}_{i,25/24}$. Then, assuming that household deviations around the average income growth path are independent, the above equation becomes:

$$\begin{aligned} E \left[\left(\sqrt{I_1} \frac{1}{I_1} \sum_{i=1}^{I_1} \widetilde{\log Y}_{i,25/24} \right) \left(\sqrt{I_1} \frac{1}{I_1} \sum_{i=1}^{I_1} \widetilde{\log Y}_{i,25/24} \right)' \right] \\ = \frac{1}{I_1} \left[E \left(\widetilde{\log Y}_{1,25/24} \widetilde{\log Y}'_{1,25/24} \right) + \dots + E \left(\widetilde{\log Y}_{I_1,25/24} \widetilde{\log Y}'_{I_1,25/24} \right) \right] \\ = \text{var}(\widetilde{\log Y}_{i,25/24}). \end{aligned}$$

Summarizing,

$$\sqrt{I_d}(\hat{\theta} - \theta_0) = -\Pi(A + B + C) \longrightarrow -\Pi \left[N \left(0, \Sigma_g \left(1 + \frac{I_d}{I_s} \right) \right) + N \left(0, \frac{I_d}{I_1} G_\chi \Sigma_\chi G_\chi' \right) \right] \longrightarrow N(0, V_\theta),$$

where

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

and $G_\theta = g_{I_s \theta}(\cdot; \cdot)$.

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

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

For most estimations, I use a diagonal weighting matrix, with the diagonal elements equal to the inverses of the diagonal elements of Σ_g , consistently estimated from the data. When the optimal weighting matrix, $W_{opt} = [\Sigma_g(1 + \frac{I_d}{I_s}) + \frac{I_d}{I_1} G_\chi \Sigma_\chi G_\chi']^{-1}$, is used standard errors are calculated using $V_\theta = [G_\theta' W_{opt} G_\theta]^{-1}$.

The over-identifying restrictions test statistic is calculated as

$$I_d g_{I_s}' W_{opt} g_{I_s},$$

and is distributed as χ^2 with $\dim(g_{I_s}) - \dim(\theta)$ degrees of freedom. The number of households contributing towards estimation of the data moments is different. When calculating the over-identification statistic and standard errors I use “conservative” values of I_d . For the over-identification statistic, I set I_d to the maximum of $I_d^1, I_d^2, \dots, I_d^J$, where I_d^j is the number of households contributing towards estimation of the j -th moment, and there are J second stage moments in total used in the MSM. For calculation of standard errors, I set I_d to the average of I_d^1, \dots, I_d^J .

Appendix D: Data Used and Sample Selection

(i) *Data from the Consumer Expenditure Survey.*

I use the CEX data on total consumer expenditures and food consumption, available at the NBER website. The data set spans the period 1980–1998. 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.

Households may enter the survey in the same year but in different quarters. If a household enters the survey, say, in the first quarter of 1981, the earliest consumption information it brings will reflect consumption in the first quarter of 1981.

Total consumption is defined as household total expenditures less expenditures on education, medical expenses, mortgage interest payments, taxes on housing, mortgage principal and mortgage lump-sum payments. It includes food at home and food away from home, clothing, expenses on personal care items, rent, or imputed rent for homeowners, housing operation expenses, personal business expenses (life insurance and business services), transportation expenses (inclusive of the purchases of vehicles), recreation and charity expenses.

I assume that household data belong to year t if a household starts survey in quarter 1 or 2 of year t , or in quarter 3 or 4 of year $t - 1$. This way a household may have as few as six (if it enters the survey in quarter 3 of year $t - 1$), and as many as twelve months of information on spending (if it enters the

survey in quarter 1 of year t) as of year t . I deflate annual food consumption by the BLS food CPI with the base 100 in 1982–1984. I deflate total consumption by the BLS CPI for all items with the same base period. I use the CPI indices for the last month of the first quarter of the respective year.

In the CEX, the head of a household is the person who owns or rents the unit of a household residence. In the PSID, the head of a household is male, unless he is permanently disabled (Hill (1992)). These definitions are not directly comparable and it is not clear how to select the CEX households that match the PSID criterion of a male-headed household. Thus, I keep the CEX households “headed” either by a married female or male.

I keep households reporting consumption in all four quarters of the year, whose heads are not part-time or full-time in school, and who are classified as full income reporters. Although I do not use income information from the CEX, the latter restriction is done for comparability with Gourinchas and Parker (2002), and Cagetti (2003). I keep households whose heads are between 24-70 years old, whose real total consumption is greater than real food consumption, and whose real total consumption and real income are greater than 1,000 of 1982 dollars. I drop households with family size below 2, whose head does not have education, race, age, or state of residence records. Finally, I drop households located above (below) the top (bottom) percentile of real food and consumption distributions. My final CEX sample consists of 21,216 households with complete demographic, food and total consumption data.

(ii) Data from the Panel Study of Income Dynamics.

The Panel Study of Income Dynamics (PSID) started in 1968, interviewing 4,802 households. Sixty per cent of the households interviewed in 1968 belong to the “core” representative sample, the other portion is known as the low-income SEO (Survey of Economic Opportunity) sample. The PSID followed these original households and households initiated by their offsprings over time, collecting a panel data set on income, wealth, demographic information, food consumption, and housing. I use the PSID data for 1980–1998, the same time span over which I have data on consumption and demographics in the CEX. To have a more representative sample, I drop SEO households and their offsprings from my sample. The PSID has 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 use the superior (to the CEX) income data from the PSID, and impute the *total* consumption to the PSID households, using demand for food equation estimated on the CEX data. 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 $t-1$ (annual food consumption is obtained by multiplying weekly food consumption by 52). In this paper, I adopt the same strategy. Over the time span considered, the PSID did not collect food consumption data in 1988 and 1989.³⁸ 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 1982–1994 CPIs taken from the BLS. I drop top-coded observations on food consumption. As with the CEX data, I use the CPI indices for the third month of the first quarter of a given year. Total real food consumption is obtained by summing real food consumption at home and away from home.

From the PSID, I choose households headed by married males, with heads aged between 24 and 70, with at least one observation on food consumption during 1980–1998. I drop households whose heads report more than a two year absolute difference in age in adjacent years. If a household passes this criterion, I use the data on age from adjacent years to impute the head’s age if a record on age is missing in a given year. I fill in missing state-of-residence data using records on the state of residence of a household in adjacent years (I use non-missing state-of-residence data as far as 4 years backward and forward from a missing record). I group households into three education categories: households with heads who did not complete high school (below 12 years of education), households with heads who completed high school (12 years), and households with heads whose education exceeds high school (above 12 years). I “allow” households to switch education categories over time if heads move to an upper education category. For example, I keep households whose heads have 12 years of minimum education attained, but attain additional education in subsequent years. I drop households whose heads report years of education in year t greater than years of education in any year after t and if, during the sample span, they switch education categories (that is, if they report minimum education attained below 12

³⁸The PSID does not have data on food consumption in the 1973 survey either.

years and maximum education attained equal to or above 12 years; or if they report minimum education attained equal to 12 years, and maximum attained above 12 years). If the head of a household has at least one non-missing race record in the sample considered, and has missing race records in any other year, I impute race records to this head. I retain observations for the years when households did not experience significant changes in family composition. Specifically, I keep observations if households did not have any changes in family composition, or had changes in members other than the head or wife (e.g., a child leaving a household). I drop observations for households with family size below 2.

Income data used in the paper are combined taxable income of head and wife. I subtract federal income, social security, and Medicare taxes from household income. The published PSID data on income and federal taxes refer to previous calendar year. I use the PSID estimate of federal income taxes for years 1981–1991. Starting in 1992, the PSID discontinued calculation of federal income taxes. I impute federal income taxes to each household for survey years 1992–1997 and 1999 using records on combined taxable income of head and wife and published tables of the federal tax schedules. I assume that head and wife file tax form jointly, and use the appropriate schedule for each year. I calculate Medicare taxes and social security taxes for head and wife separately, taking their wage income as the base for these taxes. My estimate of disposable after-tax household income is $\max\{0, (\text{Labor Income of Head and Wife} - \text{Federal Income Taxes of Head and Wife} - \text{Social Security Taxes of Head and Wife} - \text{Medicare Taxes of Head and Wife} + \text{Transfer Income of Head and Wife})\}$. I deflate disposable after-tax income by the BLS 1982–1984 CPI for all items, the same deflator I use for total consumption.

After imputing total consumption to the PSID households, my other sample selection criteria are as follows. I drop observations if disposable income or total imputed consumption is below 1,000 of real 1982–1984 dollars; if real disposable income, or real imputed total consumption is below real food consumption; if real imputed total consumption, or real food consumption is above (below) the top (bottom) percentile of the respective CEX distribution; if an absolute percentage change of real disposable income, or real food consumption, or real total imputed consumption is greater than 300%. Finally, since I am interested in the link between income and consumption dynamics, I keep observations with non-missing records on both total imputed consumption and disposable income. The final PSID sample used to construct the life cycle profiles of income and consumption consists of 23,503 income and consumption observations, along with complete information on the state of household residence, the head’s race, education, family size and the head’s age.

(iii) Distributions of Initial Income and Initial Wealth-to-Permanent Income: Estimates from the PSID.

To initialize the model economy, I need an estimate of the initial income distribution, and the distribution of wealth-to-permanent income ratio.

I assume that the wealth-to-permanent income ratio and the initial income are log-normally distributed. I take liquid wealth of 24 and 25 year old heads from the PSID wealth supplements available in 1984, 1989, 1994, 1999, and 2001. “Liquid” wealth is equal to the sum of the monetary value of checking and saving accounts, net value of real estate other than main housing, net value of vehicles, net value of shares, stocks and bonds, minus the credit debt. For each household with a record on liquid wealth and with the head aged between 24–25 years old, I estimate the permanent income as the average income during a 4-year span. I use the 1983–1986 data on real disposable income if a household with the wealth record is observed in 1984; 1988–1991 income data—if the wealth record is in 1989; 1993–1996 income data—if the wealth record is in 1994; income data in 1996, 1998, and 2000—if the wealth record is in 1999, or 2001. For estimation of the distribution of the wealth-to-permanent income ratio, I keep only households with a positive ratio. I have a cross-section of 163 observations on log wealth-to-permanent income ratio. Since households with records on wealth and income are observed at different points of business cycle, I eliminate the time (business cycle) effect with the following transformation: $\widehat{w}_{i,t} = w_{i,t} - (w_{.,t} - w_{.,.})$, where i stands for household i , t is any year from the set $\{1984, 1989, 1994, 1999, 2001\}$, $w_{.,t}$ is the average of wealth-to-permanent income ratio across observations in year t , $w_{.,.}$ is the average across all households and all years. I use a vector of $\widehat{w}_{i,t}$ ’s to estimate the mean and variance of the distribution of wealth-to-permanent income ratio. The mean is -1.22 , and the standard deviation is 1.16 .

For each simulated household, I set its value of the initial permanent income equal to a draw from the cross-sectional distribution of $\log \widehat{Y}_{i,24}$. I assume that $\log \widehat{Y}_{i,24}$ is distributed normally. The mean across

340 observations is estimated at 5.35, and the standard deviation—at 0.41.

Appendix E: Imputation of Total Consumption to the PSID Households

In this appendix, I describe the procedure used to impute total consumption to the PSID households. There are several imputation methods adopted in the literature. Absent data on total consumption in the PSID, imputation is usually done in order to exploit the panel structure of the PSID, and its superior (to the CEX) data on income. Skinner (1987), using the 1972–1973 and 1983 waves of the CEX, showed that total CEX 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. Several researchers, inspired by this finding, used coefficients from the CEX regression of total consumption on consumption items, also available in the PSID, and household data from the PSID, to impute total consumption to the PSID households. Another way of imputation is to use data on liquid assets and income in the PSID to construct savings and total consumption (Ziliak (1998), and Zeldes (1989a)). Blundell et al. (2005b) pioneered a structural approach—inverting the food demand equation estimated on the CEX data. They relate log-food consumption to log-total expenditures, household demographics, price variables, time dummies, cohort dummies, and 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 the total consumption to the PSID households. I use consumption and demographic data of merged CEX cross-sections for years 1980–1998. The detailed sample selection procedure is described in Appendix D. I assume that business cycle effects are captured by regional unemployment. The regions considered are the U.S. Census regions: Northeast, North Central, South, and West. The results of the food demand equation are presented in Table 12. The elasticity of food consumption with respect to total expenditures is positive, tightly estimated, and is falling over time. This result implies that the share of food consumption in total expenditures falls over time: given improvements in standards of living and increases in total expenditures over time, it reflects the well-known fact that food is a necessity. The family size, age, and black coefficients have expected signs and are statistically significant. The regression explains about 50% of the total variation in the log-food consumption.

TABLE 12: (APPENDIX E) FOOD EQUATION ESTIMATED ON CEX DATA

Dependent Variable: Log-Food Consumption. CEX Sample: 1980–1998

Variable	Coefficient	St. error
Log-total consumption	0.58***	(0.01)
Log-food CPI	0.08	(0.14)
Regional unemployment	-0.01**	(0.00)
Northeast	0.04***	(0.01)
North Central	-0.03***	(0.01)
South	-0.02***	(0.01)
Total consumption×1981	-0.00	(0.00)
Total consumption×1982	-0.03***	(0.01)
Total consumption×1983	-0.03***	(0.01)
Total consumption×1984	-0.04***	(0.01)
Total consumption×1985	-0.04***	(0.01)
Total consumption×1986	-0.04***	(0.01)
Total consumption×1987	-0.05***	(0.01)
Total consumption×1988	-0.11***	(0.03)
Total consumption×1989	-0.04***	(0.01)
Total consumption×1990	-0.04***	(0.01)
Total consumption×1991	-0.04***	(0.01)
Total consumption×1992	-0.05***	(0.01)
Total consumption×1993	-0.04***	(0.01)
Total consumption×1994	-0.05***	(0.01)
Total consumption×1995	-0.05***	(0.01)
Total consumption×1996	-0.06***	(0.02)
Total consumption×1997	-0.06***	(0.02)
Total consumption×1998	-0.06***	(0.02)
Born between 1915–1920	-0.00	(0.02)
Born between 1921–1926	-0.01	(0.03)
Born between 1927–1932	-0.01	(0.03)
Born between 1933–1938	-0.01	(0.04)
Born between 1939–1944	0.01	(0.04)
Born between 1945–1950	0.01	(0.05)
Born between 1951–1956	0.01	(0.05)
Born between 1957–1962	0.01	(0.06)
Born between 1963–1968	0.04	(0.07)
Born between 1969–1975	0.05	(0.07)
Family size	0.08***	(0.00)
Finished high school (HS)	0.11*	(0.06)
Did not finish HS	0.19***	(0.07)
Did not finish HS×Tot. cons.	-0.04***	(0.01)
Finished HS×Tot. cons.	-0.02**	(0.01)
Age	0.03***	(0.00)
Age ² /100	-0.03***	(0.00)
Black	-0.13***	(0.01)
Constant	-0.70	(0.67)
Number of observations		21,216
Adj. R ²		0.504

Notes: The sample consists of merged CEX cross-sections for years 1980–1998. Data are accessed from the NBER website. Food consumption is the sum of real food consumption at home and away from home, exclusive of food at work. Total consumption is the sum of real expenditures on all items, exclusive of education, medical expenses, mortgage interest, other interest, taxes on housing, mortgage principal and lump-sum mortgage payments. The regional unemployment variable is created by the interaction of a survey year and unemployment in the region of household residence at the time of a survey. The sample consists of households with married heads aged between 24 and 70. Heads enrolled part-time or full-time in school are dropped. The sample households are present in all four quarters of a survey and are complete income reporters. Family size is the average family size during a survey year (across four quarters). The head belongs to a high school category if he/she reports 12 years of education; “Did not finish high school” is an indicator variable equal to 1 if the head’s education is below 12 years; omitted category consists of heads with education above 12 years. Omitted cohort comprises households with heads born between 1910–1914; omitted region is West. Price variable is the log of the BLS CPI index for total food consumption with the base in 1982–1984. *** significant at the 1% level, ** significant at the 5% level, and * significant at the 10% level.