

A DIRECT TEST OF THE BUFFER-STOCK MODEL OF SAVING

Tullio Jappelli

University of Naples Federico II

Luigi Pistaferri

Stanford University

Mario Padula

University "Ca' Foscari" of Venice

Abstract

Recent models with liquidity constraints and impatience emphasize that consumers use savings to buffer income fluctuations. When wealth is below an optimal target, consumers try to increase their buffer stock of wealth by saving more. When it is above target, they increase consumption. This important implication of the buffer stock model of saving has not been subject to direct empirical testing. We derive from the model an appropriate theoretical restriction and test it using data on working-age individuals drawn from the 2002 and 2004 Italian Surveys of Household Income and Wealth. One of the most appealing features of the survey is that it has data on the amount of wealth held for precautionary purposes, which we interpret as target wealth in a buffer stock model. The test results do not support buffer stock behavior, even among population groups that are more likely, a priori, to display such behavior. The saving behavior of young households is instead consistent with models in which impatience, relative to prudence, is not as high as in buffer stock models. (JEL: D91)

1. Introduction

Recent intertemporal consumption models emphasize the role of savings as a buffer stock against income fluctuations. Deaton (1991) and Carroll (1997) have solved sophisticated versions of such models. Although the specific details of the models differ, emphasizing liquidity constraints or the probability of low income realizations, they share similar predictions. In both models, consumers have a

The editor in charge of this paper was Orazio Attanasio.

Acknowledgments: We thank Chris Carroll for extensive discussions and suggestions and, for comments, the editor, two anonymous referees, Erich Battistin, Annamaria Lusardi, Enrico Rettore, Jon Skinner, Steve Pischke, and seminar participants at the NBER Summer Institute, the 7th Workshop of the RTN Project on Economics of Aging in Venice, the University of Bologna, the University of Brescia, the University of Napoli Federico II, the University of Padua, the University of Reading, and Stanford University. Work supported in part by the European Community's Human Potential Programme under contract HPRN-CT-2002-00235 [AGE: The Economics of Aging in Europe] and by the Italian Ministry of University and Research (MUR). Jappelli is a member of CSEF and CEPR; Padula is a member of CSEF; and Pistaferri is a member of SIEPR and CEPR.

E-mail: Jappelli: tullioj@tin.it; Padula: mpadula@unive.it; Pistaferri: pista@leland.stanford.edu

unique and stable ratio between cash on hand (wealth plus disposable income) and the permanent component of income, which we term the “target wealth to permanent income ratio.” As stated by Carroll, buffer-stock savers have a target wealth to permanent income ratio “such that if [actual] wealth is below the target, the precautionary saving motive will dominate impatience, and the consumer will save, while if [actual] wealth is above the target, impatience will dominate prudence, and the consumer will dissave” (p. 2).

This key implication of the buffer-stock saving model has not been subject to empirical scrutiny. Current evidence of buffer-stock behavior is based on two model’s implications: that consumption tracks income closely, and that precautionary saving represents an important reason for wealth accumulation. Several simulations of intertemporal consumption models predict consumption-income tracking in the early part of the life-cycle (Attanasio et al. 1999; Laibson, Repetto and Tobacman 1998; Gourinchas and Parker 2002; Cagetti 2003). Empirical evidence on the importance of precautionary saving is mostly based on reduced form regressions of net worth or financial assets on proxies for income risk. Some studies report that precautionary wealth represents a small portion of total wealth, e.g. Guiso, Jappelli, and Terlizzese (1992), and Hurst, Lusardi et al. (2005); others find a large impact of income risk e.g., Carroll and Samwick (1997), and Gruber and Yelowitz (1999). These studies differ in many respects, such as the definition of wealth, the measure of risk, and institutional features. But even findings of large effects of income risk on saving are not conclusive evidence of buffer stock behavior, because life-cycle models with income risk also provide an important role for precautionary saving (see Hubbard, Skinner, and Zeldes 1995). In short, the literature still lacks a convincing test of the buffer-stock model.

In this paper we use a survey question on precautionary wealth available in the 2002 and 2004 Bank of Italy Surveys on Household Income and Wealth (SHIW) to propose a direct test of buffer stock behavior. The question asks people how much savings they think they need for future emergencies, and is similar to a question contained in the 1995 and 1998 Surveys of Consumer Finances described in Kennickell and Lusardi (2004). We interpret this question as providing information on target wealth in a buffer-stock model and test the proposition that people with a ratio of actual wealth to permanent income below the target intend to save, whereas those with a ratio above target intend to dissave.

Although we focus on Carroll’s (1991) version of the buffer stock model, the test applies equally well to Deaton’s (1997) case. In Carroll, buffer stock behavior emerges from the tension between impatience, prudence, and the chance of zero earnings. Impatient individuals would like to anticipate consumption, but the chance of zero future earnings generates a demand for wealth. In Deaton, there is an explicit liquidity constraint, but the insights are similar, and buffer-stock behavior emerges again as the optimal policy.

The rest of the paper is organized as follows. Section 2 derives our test of the buffer stock model, and evaluates it with simulated data. The test can be used to check whether or not buffer stock behavior is rejected by the data, and relies on a specific survey question on precautionary wealth. Section 3 describes such a question and compares it with a similar question asked in the U.S. Survey of Consumer Finances. The test results, presented in Section 4, are inconsistent with the buffer stock model. In Section 5 we therefore present estimates of the age–wealth profile obtained with Italian repeated cross-sectional data to provide further evidence on the validity of the buffer-stock model. In particular, the model suggests that the ratio of wealth to permanent income of young consumers should be stationary, whereas models in which prudence dominates impatience suggest that the ratio grows even in young ages because consumers save for retirement. In addition this evidence runs contrary to the predictions of the buffer-stock model. Section 6 summarizes our findings.

2. Deriving Testable Implications of Buffer-Stock Behavior

We take as our point of departure Carroll's (1997) buffer-stock saving model to derive testable predictions and explain our empirical strategy. Consumers have finite horizons and choose consumption to maximize the following objective function:

$$E_0 \sum_{t=0}^T \beta^t u(C_t),$$

where β is the subjective discount factor, the instantaneous utility function is isoelastic, $u(C_t) = C_t^{1-\rho} / (1-\rho)$, and $\rho > 0$ is the coefficient of relative risk aversion. The dynamic budget constraint is

$$W_{t+1} = R[W_t - C_t + Y_t],$$

where $R = 1 + r$ is the constant interest rate factor, and W_t , Y_t , and C_t are, respectively, non-human wealth, labor income, and consumption at time t . Labor income shifts due to transitory and permanent shocks, assumed to be log-normally distributed, namely,

$$Y_{t+1} = P_{t+1} V_{t+1}, \quad (1)$$

$$P_{t+1} = GP_t N_{t+1}, \quad (2)$$

where G is the growth rate of income, P_{t+1} is permanent income, and V_{t+1} and N_{t+1} are i.i.d. shocks with mean equal to 1.¹ The model also assumes that in each

1. More precisely, P_{t+1} is the permanent component of income. We use the two terms interchangeably.

period there is a small chance $p > 0$ that transitory income is zero. The Bellman equation of the problem is

$$V_t(W_t, P_t) = \max_{C_t} u(C_t) + \beta E_t V_{t+1}(W_{t+1}, P_{t+1}) \tag{3}$$

$$\text{subject to: } P_{t+1} = GP_t N_{t+1},$$

$$W_{t+1} = R[W_t - C_t + Y_t].$$

To exploit the homogeneity of the instantaneous utility function, let's define cash-on-hand X_t as the sum of non-human wealth and income ($X_t = W_t + Y_t$), and write question (3) as

$$v_t(x_t) = \max_{c_t} u(c_t) + \beta E_t G^{1-\rho} N_{t+1}^{1-\rho} v_{t+1}(x_{t+1}) \tag{4}$$

$$\text{subject to: } x_{t+1} = R[x_t - c_t] \frac{1}{GN_{t+1}} + V_{t+1}, \tag{5}$$

where $c_t = C_t/P_t$, $v_t(x_t) = V_t(W_t, Y_t)/P_t^{1-\rho}$, and $x_t = (W_t + Y_t)/P_t$ is what we call, for brevity, "wealth to permanent income ratio."²

Carroll (2004) shows that for specific ranges of parameter values, the problem has a solution (i.e., the functional defined in equation (4) has a fixed point), optimal consumption is an increasing and concave function of cash on hand, and the marginal propensity to consume out of cash on hand is bounded from above and from below. Furthermore, there exists a unique and stable value of x (which we call x^* and term, again for brevity, "target wealth to permanent income ratio") such that, "if actual wealth is greater than the target, impatience will outweigh prudence, and wealth will fall, while if wealth is below the target, the precautionary saving motive will outweigh impatience and the consumer will try to build wealth back up toward to target" (Carroll 2001, p. 33).³ In our notation, if $(x_t - x^*) < 0$, then x_t grows in expectation. If instead $(x_t - x^*) > 0$, x_t falls (again, in expectation). Using cross-section data, we construct a test of the theory based on this crucial insight.

At any given point in time, households differ in their value of the wealth gap $(x_t - x^*)$. A first source of heterogeneity concerns preferences and the parameters of the income-generating process, which set different values of x^* for each individual. Income shocks are a second source of heterogeneity: Even if two identical consumers have the same preferences and the same income-generating

2. This is to avoid the rather cumbersome terminology of "ratio of cash-on-hand to the permanent component of income."

3. Carroll (2004) shows also that, at the target, expected consumption growth is less than expected permanent income growth and that expected consumption growth is declining in cash on hand.

process—and therefore the same x^* —they receive different income shocks and have therefore different x_t and wealth gaps.⁴

Thus in a cross-section, the model implies that

$$\text{Cov}(x_{ht} - x_h^*, E_{ht}(x_{ht+1} - x_{ht})) < 0, \quad (6)$$

where $\text{Cov}(\cdot, \cdot)$ is a population covariance and h is a household index. This notation makes explicit that $E_{ht}(x_{ht+1} - x_{ht})$ is the time- t expectation of household h 's next period change in the wealth-permanent income ratio, and the covariance is taken with respect to the cross-sectional distribution of the wealth gap and of expected asset accumulation.

To use equation (6) as a basis for an empirical test, one needs to observe x_h^* , x_{ht} and $E_{ht}(x_{ht+1})$. As we shall see, we have data on actual wealth and on a proxy of target wealth, but not on the expected value of the change in the wealth-permanent income ratio x_{ht} . To evaluate $E_{ht}(x_{t+1})$, let's take the expectation as of time t of equation (5) for household h , and recall that $E_{ht}(N_{t+1}) = 1$, $E_{ht}(V_{t+1}) = 1$, and $\text{Var}_{ht}(\ln N_{t+1}) = \sigma_N^2$:

$$\begin{aligned} E_{ht}(x_{ht+1}) &= R[x_{ht} - c_{ht}] \times E_{ht}\left(\frac{1}{GN_{t+1}}\right) + E_{ht}V_{ht+1} \\ &\approx \frac{R}{G}[x_{ht} - c_{ht}] \times e^{\sigma_N^2} + 1, \end{aligned} \quad (7)$$

where the second equality follows from a second-order Taylor expansion of $1/N_{ht+1}$ around the mean of N_{ht+1} .

Substituting equation (7) in equation (6) and defining $\gamma = e^{\sigma_N^2}$, we can restate equation (6) in terms of observable variables as

$$\theta = \frac{\text{Cov}(x_{ht} - x_h^*, c_{ht})}{\text{Cov}(x_{ht} - x_h^*, x_{ht})} > \left(1 - \frac{G}{R\gamma}\right). \quad (8)$$

The sign of θ (which from now on we term the ‘‘covariance ratio’’) is a priori ambiguous. In fact, $\sigma_N > 0$ implies $\gamma > 1$. If the growth rate of income is lower than the interest rate ($G < R$) the covariance ratio is a positive number. However, if $G > R$ the covariance ratio must exceed a negative number, and so it might itself be negative. In the buffer stock model both cases might arise. Indeed, Carroll (1997) shows that the parameters must satisfy the following inequality:

$$\frac{r - \delta}{\rho} + \frac{\rho}{2}\sigma_N^2 < g - \frac{1}{2}\sigma_N^2, \quad (9)$$

where $\ln R \approx r$, $\ln \beta \approx -\delta$, and $\ln G \approx g$.

4. These are not the only possible sources of heterogeneity. In Section 2.2 we use simulation analysis to explore the effect of heterogeneity in income risk, income growth, and interest rates.

2.1. Test Interpretation and Implementation

Our test strategy is as follows. First note that the sample analog of the left-hand side of the inequality (8) is

$$\hat{\theta} = \frac{\text{Cov}(x_h - x_h^*, c_h)}{\text{Cov}(x_h - x_h^*, x_h)} = \frac{\sum_{h=1}^H ((x_h - x_h^*) - \overline{(x_h - x_h^*)})(c_h - \bar{c}_h)}{\sum_{h=1}^H ((x_h - x_h^*) - \overline{(x_h - x_h^*)})(x_h - \bar{x}_h)},$$

where we have dropped the time subscripts, $\text{Cov}(\cdot, \cdot)$ is a sample covariance, and a bar over a variable denotes its cross-sectional mean. This is simply the instrumental variables (IV) estimate of a regression of c_h on x_h using the wealth gap $(x_h - x_h^*)$ as an instrument. The advantage of the regression framework is that it naturally delivers standard errors which allow us to conduct statistical inference on the value of θ .

There are two ways to test the implications of the buffer-stock model underlying equation (8) (and hence equation (6)). One way would be to choose values of G , R , and γ and test whether θ satisfies the bound restriction implied by equation (8). The problem with this strategy is that the bound restriction may be satisfied for implausible values of the parameters. Instead, our test strategy is based on the comparison between the empirical and theoretical values of θ . Because we don't know the true parameters, we simulate the buffer-stock model for a variety of plausible parameter configurations. In particular, we explore cases in which $G < R$, as well as cases in which individuals expect high earnings growth relative to the interest rate ($G > R$). Finally, we compare the estimated with the simulated covariance ratio.

2.2. The Simulated Covariance Ratio

To implement our test, we start by simulating the model for an economy populated by heterogeneous consumers. In the baseline scenario, we posit two sources of heterogeneity. Each individual has a different discount factor, uniformly distributed between 0.86 and 0.96. This guarantees that each consumer has different target wealth. Secondly, although in the baseline scenario the income process is the same, in each period consumers are hit by different realizations of the shocks.

We set the other models' parameters following Carroll (2004): the growth factor $G = 1.03$, the interest rate factor $R = 1.04$, the coefficient of relative risk aversion $\rho = 2$, the standard deviation of permanent and transitory income shocks $\sigma_N = \sigma_V = 0.1$, and the probability of unemployment $p = 0.005$. Such parametrization satisfies equation (9) and guarantees a stationary target wealth to permanent income ratio x^* .

We assume that consumers start with zero wealth and simulate the model for 100 periods and 1,000 consumers. We then compute, for each consumer, the




FIGURE 1. Simulated target wealth to permanent income ratio. The figure plots the target wealth to permanent income ratio (x^*) as a function of the discount factor β . The target wealth ratios are obtained simulating the buffer stock model with the baseline parameters $G = 1.03$, $R = 1.04$, $\rho = 2$, $\sigma_N = 0.1$, $\sigma_V = 0.1$, and $p = 0.005$.

target wealth to permanent income ratio such that $E_{ht}(x_{ht+1}) = x_{ht}$. In Figure 1 we plot the cross-sectional distribution of target wealth against the intertemporal discount rate β for the 1,000 buffer stock consumers of our simulations. The figure highlights the positive association between the discount factor and the target wealth to permanent income ratio: x_h^* increases from about 1.2 for $\beta = 0.86$ to 1.45 for $\beta = 0.96$.

Based on the different values of x_h^* and different realizations of x_{ht} and c_{ht} , we compute in each of the 100 periods the cross-sectional covariance ratios, and summarize its distribution by the median value. For the baseline experiment we find a simulated $\theta = 0.62$.⁵

It is important to check that the simulated θ does not depend heavily on the specific parametrization of the model. Therefore we simulate the covariance ratio under a wide range of alternative, realistic parameter assumptions. Table 1 reports the simulated θ for different parameter values and sources of individual heterogeneity, separately for the case $G < R$ (Panel A) and $G > R$ (Panel B). Each parametrization satisfies condition (9) and is obtained from the baseline case by varying one parameter at the time.

5. The 25th and 75th percentiles of the simulated covariance ratio are 0.61 and 0.63, respectively.

TABLE 1. The simulated covariance ratio.

Panel A: $G < R$			
Growth factor	Covariance ratio θ	Interest rate factor	Covariance ratio θ
$G = 1.025$	0.605	$R = 1.035$	0.632
$G = 1.035$	0.641	$R = 1.045$	0.616
$G = [1.025, 1.035]$	0.629	$R = [1.035, 1.045]$	0.617
Probability of zero income		Relative risk aversion	
$p = 0.001$	0.704	$\rho = 1.5$	0.669
$p = 0.01$	0.572	$\rho = 4$	0.452
$p = [0.001, 0.01]$	0.594	$\rho = [1.5, 4]$	0.513
S.D. of permanent shocks		S.D. of transitory shocks	
$\sigma_N = 0.04$	0.663	$\sigma_V = 0.04$	0.691
$\sigma_N = 0.14$	0.524	$\sigma_V = 0.14$	0.586
$\sigma_N = [0.04, 0.14]$	0.561	$\sigma_V = [0.04, 0.14]$	0.596
Panel B: $G > R$			
Growth factor	Covariance ratio θ	Interest rate factor	Covariance ratio θ
$G = 1.035$	0.654	$R = 1.025$	0.669
$G = 1.045$	0.675	$R = 1.035$	0.660
$G = [1.035, 1.045]$	0.672	$R = [1.025, 1.035]$	0.661
Probability of zero income		Relative risk aversion	
$p = 0.001$	0.688	$\rho = 1.5$	0.716
$p = 0.01$	0.624	$\rho = 4$	0.485
$p = [0.001, 0.01]$	0.645	$\rho = [1.5, 4]$	0.567
S.D. of permanent shocks		S.D. of transitory shocks	
$\sigma_N = 0.04$	0.731	$\sigma_V = 0.04$	0.757
$\sigma_N = 0.14$	0.625	$\sigma_V = 0.14$	0.615
$\sigma_N = [0.04, 0.14]$	0.644	$\sigma_V = [0.04, 0.14]$	0.643

Notes: The table reports the median simulated covariance ratio under alternative parameterization of a buffer-stock economy populated by 1,000 individuals. In the baseline scenario of Panel A $G = 1.03$, $R = 1.04$, $\rho = 2$, $\sigma_N = 0.1$, $\sigma_V = 0.1$, and $p = 0.005$, β ranges from 0.86 to 0.96, and the median simulated covariance ratio is $\theta = 0.623$. In the baseline scenario of Panel B, $G = 1.04$, $R = 1.03$, $\rho = 2$, $\sigma_N = 0.1$, $\sigma_V = 0.1$, and $p = 0.005$, β ranges from 0.86 to 0.96, and the median simulated covariance ratio is $\theta = 0.664$.

Panel A shows that changing the growth factor to 1.025 or 1.035, and the interest factor to 1.035 or 1.045 does not affect the covariance ratio appreciably, while raising the coefficient of risk aversion to 4 reduces θ to 0.45. Changing the income process has a larger impact on the ratio. For instance, lowering the standard deviation of permanent or transitory income shocks to 0.04 raises the covariance ratio to 0.66 and 0.69, respectively. Finally, the simulated θ increases to 0.70 when the probability of unemployment is lowered to 0.1%. In all cases, the simulated θ ranges from 0.45 to 0.7.

In Panel A we also compute the simulated covariance ratio under different assumptions about the source of heterogeneity in the model. We consider cases in which the growth factors are uniformly distributed between 1.025 and 1.035 (obtaining $\theta = 0.63$), interest factors between 1.035 and 1.045 ($\theta = 0.62$), coefficients of relative risk aversion between 1.5 and 4 ($\theta = 0.51$), standard

deviations of permanent and transitory income shocks between 4% and 14% ($\theta = 0.56$ and $\theta = 0.60$, respectively), and probabilities of zero income between 0.1% and 1% ($\theta = 0.59$).

We then we repeat the simulations considering cases with $G > R$. In the baseline scenario we set $G = 1.04$ and $R = 1.03$; the other parameters are the same as in the baseline scenario with $G < R$ ($\rho = 2$, $\sigma_N = \sigma_V = 0.1$, $p = 0.005$). The simulated covariance ratio in this case is equal to 0.66. In Panel B we consider the same sources of heterogeneity as in Panel A. Each of the experiments assumes $G > R$, and satisfies condition (9). The simulated covariance ratios range from 0.48 to 0.76.

Finally, we compute θ choosing parameter values that fit the Italian economy. In the past two decade the productivity growth rates of Italian workers in the age group 20–50 has been 1.5%, and the real interest rate 2.5%; accordingly, we set $G = 1.015$ and $R = 1.025$. Jappelli and Pistaferri (2006) estimate the income process (equations (1)–(2)) with Italian panel data and find standard deviations of permanent and transitory income shocks of 0.16 and 0.28, respectively. For such parameterization, θ equals 0.46.

We conclude from these experiments that in realistically calibrated buffer stock models the covariance ratio is likely to be in the 0.5–0.6 range, and unlikely to fall below 0.4 or to exceed 0.8. Any empirical estimate of θ that is statistically significantly away from this range would therefore be hard to reconcile with buffer-stock behavior.

3. Data

To implement the empirical test of the buffer-stock model, we use the 2002 and 2004 Italian Surveys of Household Income and Wealth (SHIW). SHIW is a biannual representative sample of the Italian population conducted by the Bank of Italy.⁶ In each year, the sample includes about 8,000 households and 24,000 individuals. Details on the questionnaire, sample design, response rates, results, and comparison of survey data with macroeconomic data are given in Biancotti et al. (2004) and Faiella et al. (2006).⁷

For our purposes, the SHIW has several advantages. It has data on wealth, income, consumption, and detailed demographic characteristics of the household.

6. In the buffer-stock model, the marginal propensity to consume (MPC) is high because consumers are impatient. Carroll (2001) interprets the excess sensitivity of consumption found by Campbell and Mankiw (1991) and Jappelli and Pagano (1989) in time series data for several OECD countries, and in Italy in particular, as dependent on the prevalence of impatient households. He argues that in these countries there are “more households who are impatient and consequently inhabit the portion of the consumption function where the MPC is high, whether they are formally constrained or not” (Carroll, 2001, p. 38). Italy, therefore, provides a good testing ground for the buffer-stock model.

7. The SHIW started in 1977, but data on consumption have been collected only since 1984.

Net financial assets measure the liquid portion of wealth, and are the sum of transaction accounts, government bonds, CDs, corporate bonds, retirement accounts, life insurance, and stocks, less household debt (mortgage loans, consumer credit, and other personal loans). Total assets are the sum of net financial assets and real assets (real estate, unincorporated business holdings, valuables and art objects). The SHIW also includes a rotating panel component: in each year, about 45% of the households are also interviewed two years later. We will later use the panel section of the SHIW to recover individual-level variables available only in the 2000 survey and to assess the robustness of our results in the presence of fixed effects.

Most importantly for the present study, the 2002 and 2004 SHIW have a direct question on precautionary wealth, which we use to proxy target wealth in the buffer stock model:

People save in various ways (depositing money in a bank account, buying financial assets, property, or other assets) and for different reasons. A first reason is to prepare for a planned event, such as the purchase of a house, children's education, etc. Another reason is to protect against contingencies, such as uncertainty about future earnings or unexpected outlays (owing to health problems or other emergencies). About how much do you think you and your family need to have in savings to meet such unexpected events?

The question is patterned after a similar question in the Survey of Consumer Finances (SCF), described in Kennickell and Lusardi (2004).⁸

Table 2 reports sample means and quartiles of target wealth for various sample groups, pooling data for 2002 and 2004. The median value of target wealth is €25,000, and the mean is €55,137. Interestingly, these values are considerably higher than in the United States, where Kennickell and Lusardi (2004) report that the bulk of the distribution of target wealth is between \$US5,000 and \$US10,000. Target wealth is higher among high-school and college graduates, self-employed, households with multiple income recipients, and households living in the North.

The median ratio of target wealth to total wealth is 0.31, and 3.32 if wealth includes only financial assets. These numbers are higher than in Kennickell and Lusardi (2004), who report 0.08 and 0.2, respectively. This shows that in Italy precautionary wealth potentially accounts for a larger portion of wealth, possibly due to higher income risk or lower degree of development of financial and insurance markets. The Italian data also indicate that in 75% of the

8. The SCF question is: "About how much do you think you and your family need to have in savings for unanticipated emergencies and other unexpected things that may come up?" The question has been extensively tested in the SCF with focus groups. Nevertheless, the question may be criticized because some consumers might report what they wish to save rather than what they aim to save, and because some may act "as if" they behave according to a buffer-stock model, even though they have difficulty identifying what their target wealth-income ratio is. Because there are no objective measures of target wealth, we cannot check whether these criticisms are founded.

TABLE 2. Selected statistics for target wealth.

	Mean	First quartile	Median	Third quartile	Number of observations
20 ≤ age ≤ 30	47,661	9,429	23,572	47,143	617
30 < age ≤ 40	58,531	10,000	25,000	50,000	2,243
40 < age ≤ 50	57,022	10,000	28,285	52,000	3,051
Self-employed	69,350	15,000	47,143	94,286	1,078
Employee	51,967	10,000	25,000	50,000	4,833
Single earner	49,864	9,429	23,572	50,000	2,683
Multiple earners	60,302	10,000	28,286	60,000	3,228
North–Center	61,316	11,314	28,286	66,000	3,977
South	40,857	50,000	18,857	47,143	1,934
Entrepreneurs	69,399	14,143	40,000	94,286	1,209
Non-entrepreneurs	51,651	9,429	25,000	50,000	4,702
Low education	48,232	9,429	25,000	50,000	2,665
High education	61,447	10,000	28,286	60,000	3,246
Total sample	55,137	10,000	25,000	50,000	5,911

Notes: The table reports sample statistics of target wealth. The sample is obtained pooling the 2002 and 2004 SHIW. Sample statistics are estimated using population weights. Values are expressed in 2002 euros.

cases financial wealth is below target, and in 28% of cases total wealth is below target. Comparable figures for Kennickell and Lusardi are 48% and 17%, respectively.

In the empirical application we measure consumption as non durable expenditures.⁹ We define cash-on-hand as $Y + W_f + \lambda W_r$, where Y is household disposable income, W_f and W_r are, respectively, net financial assets and real assets, and $0 \leq \lambda \leq 1$ measures the portion of real assets that can be used in the current period to finance consumption.¹⁰ We focus on a sample where buffer-stock behavior is more likely to emerge, selecting households with heads aged between 20 and 50. The resulting sample consists of 5,911 observations (2,953 for 2002 and 2,958 for 2004).

Consumption, target wealth, cash-on-hand, and the wealth gap are all normalized by an estimate of the permanent component of income, that is, income during the working life purged from transitory components. We opt for a simple and straightforward definition, and estimate the permanent component of income by the fitted value of a regression of household non-financial income on age, education, dummies for occupation, region of residence, head gender, number of earners, and a year dummy.

Figure 2 plots the histogram of the target wealth to permanent income ratio. The median ratio represents slightly more than one year of income, and the bulk

9. Results are unchanged if one defines consumption as the sum of non-durable and durable expenditures, see Section 4.6.

10. Another reason to let λ vary is that the definition of cash-on-hand adds a flow (income) to a stock (wealth).




FIGURE 2. Target wealth to permanent income ratio. The figure plots the sample distribution of the target wealth to permanent income ratio in the pooled 2002–2004 SHIW.

of the distribution is between 2 months and 6 years. Figure 3 reports the density of the ratio splitting the sample by business ownership and number of income recipients. The figure shows that the distribution of the ratio is higher for business owners and single income households, which should indeed be the case if these households face more uncertainty relative to the other groups.

Table 3 reports median regressions for the logarithm of target wealth, and for the target wealth to permanent income ratio. In each case we control for the log of permanent income. Because regional factors are quite important in Italy, and might be correlated with background economic variables and preferences, we run the two specifications including a full set of 20 regional dummies. The coefficient estimates show that target wealth increases with age (slightly less than 1% for each year) and education (between 2% and 3% per year). Business owners have a target wealth that is about 20% higher than the reference group, confirming the graphical comparison in Figure 3. Single-income households have a target wealth that is about 10% higher, although the coefficient is not statistically different from zero when we control for regional effects.

4. Testing the Buffer-Stock Model

In this section we estimate the covariance ratio θ and compare it with the simulated values from Section 2.2. Recall that the simulated θ never falls below 0.4 or exceeds 0.8. In this section we also test if θ differs by economic and demographic characteristics of the household. We focus on households facing high income risk (such as business owners, single-income households, the self-employed) and




FIGURE 3. Target wealth to permanent income ratio, by business ownership and number of income recipients. The figure plots the kernel density of the target wealth to permanent income ratio using the pooled 2002–2004 SHIW by business ownership and number of income recipients.

TABLE 3. Regressions for target wealth.

	Log target wealth		Target wealth to permanent income ratio	
Age	0.007 (0.003)**	0.010 (0.003)***	0.008 (0.004)**	0.009 (0.003)***
Years of education	0.023 (0.007)***	0.028 (0.006)***	0.024 (0.007)***	0.030 (0.007)***
Self-employed	0.029 (0.081)	0.047 (0.080)	0.017 (0.090)	0.031 (0.082)
Business owners	0.216 (0.078)***	0.210 (0.076)***	0.229 (0.086)***	0.260 (0.078)***
Family size	-0.006 (0.018)	0.050 (0.020)**	-0.006 (0.020)	0.050 (0.020)**
Single income	0.142 (0.058)**	0.071 (0.058)	0.171 (0.064)***	0.081 (0.059)
Log permanent income	0.645 (0.071)***	0.324 (0.078)***	-0.367 (0.078)***	-0.725 (0.079)***
Regional dummies	No	Yes	No	Yes

Notes: The table reports median regressions for the log of target wealth and the ratio of target wealth to permanent income using the pooled 2002–2004 SHIW. The second and fourth regressions include a set of 20 regional dummies. The permanent component of income is estimated by the fitted value of a regression of household non-financial income on age, education, dummies for occupation, region of residence, head gender, and number of earners in the household and year dummies. Standard errors are reported in parentheses.

TABLE 4. Estimated covariance ratio: baseline regression and group estimates.

	$Y + \lambda W_r + W_f$				Number of observations
	$\lambda = 1$	$\lambda = 0.75$	$\lambda = 0.5$	$\lambda = 0.25$	
Total sample	0.012 (0.000)	0.015 (0.001)	0.017 (0.001)	0.015 (0.001)	5,911
20 ≤ Age ≤ 30	0.020 (0.002)	0.023 (0.003)	0.024 (0.005)	0.005 (0.010)	617
30 < Age ≤ 40	0.011 (0.001)	0.013 (0.001)	0.014 (0.001)	0.011 (0.003)	2,243
40 < Age ≤ 50	0.012 (0.001)	0.015 (0.001)	0.018 (0.001)	0.017 (0.002)	3,051
Self-employed	0.008 (0.001)	0.010 (0.001)	0.013 (0.001)	0.015 (0.002)	1,078
Employees	0.017 (0.001)	0.020 (0.001)	0.022 (0.001)	0.017 (0.002)	4,833
Single earners	0.013 (0.001)	0.015 (0.001)	0.017 (0.001)	0.014 (0.002)	2,683
Multiple earners	0.012 (0.001)	0.015 (0.001)	0.020 (0.001)	0.023 (0.002)	3,228
North–Center	0.013 (0.000)	0.016 (0.001)	0.020 (0.001)	0.023 (0.002)	3,977
South	0.011 (0.001)	0.012 (0.001)	0.013 (0.002)	0.008 (0.002)	1,934
Entrepreneurs	0.009 (0.001)	0.011 (0.001)	0.014 (0.001)	0.015 (0.002)	1,209
Non-entrepreneurs	0.016 (0.001)	0.019 (0.001)	0.021 (0.001)	0.016 (0.002)	4,702
Low education	0.009 (0.001)	0.011 (0.001)	0.013 (0.002)	0.009 (0.003)	2,665
High education	0.014 (0.001)	0.017 (0.001)	0.020 (0.001)	0.018 (0.001)	3,246

Notes: The sample is based on the pooled 2002 and 2004 SHIW. W_r and W_f are, respectively, real and financial wealth. Y is disposable income, and λ is the share of real assets that can be used in the current period to finance consumption. Standard errors are reported in parentheses.

impatient consumers (using a direct survey question on the rate of time preference). Finally, we check robustness of our results to measurement error, different definitions of income and consumption, and unobserved heterogeneity.

4.1. Baseline Estimates

The first row of Table 4 displays baseline estimates for the whole sample on the pooled 2002–2004 sample. They are obtained regressing consumption on cash on hand, using the wealth gap as an instrument. All variables are divided by our estimate of the permanent component of income. In the first column, we set $\lambda = 1$ and cash on hand is just $Y + W_f + W_r$, on the assumption that households can use

all assets to buffer income shocks. The point estimate of θ is 0.012, much lower than the values consistent with buffer stock behavior. Because the estimate has a small standard error, we formally reject the hypothesis that θ equals any one of the simulated values of Table 2. Setting $\lambda = \{0.75, 0.50, 0.25\}$, we find that θ ranges from 0.015 to 0.017, much below the range of admissible values of the simulated ratio.

4.2. Group Estimates

Even if our baseline results do not support it, the buffer-stock model might still characterize the behavior of some population groups that face high income volatility or are more impatient. We are particularly interested in detecting buffer-stock behavior for groups that, a priori or based on previous evidence, are more likely to exhibit such behavior. The self-employed clearly face greater income risk than employees. If the incomes of households with multiple earners are not perfectly correlated, single income households face more risk than households where both spouses work. The young might face more income uncertainty, or be more impatient than the middle-aged because do not yet perceive the need to accumulate for old age. In Italian regions with better functioning credit and insurance markets (the North and the Center), employment shocks and other risks are more likely to be insured. And in the case of education, we have hard evidence with the same dataset that income risk differs by level of education.

To check if buffer stock behavior characterizes some population groups, in Table 4 we present estimates of θ splitting the sample by age (less than 30, between 30 and 40, between 40 and 50), number of income earners, employment status, business ownership, and region of residence. In the first column $\hat{\theta}$ never exceeds 0.02, confirming the full sample estimates for each of the group considered. The $\hat{\theta}$ are precisely estimated, and in each case we reject the hypothesis that they are consistent with the simulated covariance ratio. The estimates for different definitions of cash on hand do not change the pattern of results.

Recent work on the extent of precautionary motive for saving has focused on business owners. Business owners and entrepreneurs face higher income risk, but their wealth holdings are also higher than average. Hurst, Lusardi, Kennickell, and Torralba (2005) provide evidence that tests of precautionary saving are considerably affected by the treatment of entrepreneurs. In the total sample, they find a strong, positive relation between wealth and permanent income shocks, as in Carroll and Samwick (1997). But the result is almost entirely due to business owners: when these are excluded from the sample, there is hardly evidence for precautionary saving. Table 4 reports $\hat{\theta}$ distinguishing by entrepreneurship, defined as positive business wealth. The results are again at variance with the buffer-stock model for both groups. We also find a lower value of $\hat{\theta}$ in the group

with lower education, but the point estimates are again far away from the simulated values. Different definitions of cash on hand do not change appreciably the pattern of results for each of the sample splits considered.

4.3. Impatience

The rate of time preference is a critical parameter of models of intertemporal choice, but microeconomic data seldom allow to pin down particular features of this and other preference parameters. The 2000 SHIW attempts at providing data on time preference through a lottery question. Frederick, Loewenstein, and O'Donoghue (2002) survey theoretical and empirical research on time preferences, and classify the various methods by elicitation methodology (choice, matching, rating or pricing), type of instrument used to elicit preferences (field versus experiment), and time frame (less than one day to many years). They report that a widely used way to elicit the rate of time preference is through survey questions asking the respondent to report how much he or she is willing the pay to receive a lottery winnings today instead of later in time. The 2000 SHIW has precisely such a question: "Suppose that you win €5,000, payable for certain in a year's time. What is the *maximum amount* that you are willing to pay to have the €5,000 immediately?"

The question is asked only to half of the sample (household heads born in odd-numbered years), and about 15% don't answer it. The 2000 data can be merged with 2002 and 2004 data using the panel component of SHIW (1,749 households are interviewed in 2000, 2002, and 2004). After merging the data, and considering that we select individuals less than 50 years old in 2002, we are left with 797 valid observations with data on both target wealth and time preference. On average, to cash the lottery one year in advance, respondents are willing to pay €150, implying a quite standard rate of time preference of 3%. Several studies use questions similar to this, as documented in Frederick, Loewenstein, and O'Donoghue (2002), who also reviews pros and cons of various methods for eliciting time preference.¹¹

We therefore split the sample according to whether the rate of time preference is above or below 3%. Table 5 reports $\hat{\theta}$ in the two sub-samples. It is important to keep in mind that in this case we have a limited number of observations. For $\lambda < 0.75$, the $\hat{\theta}$ for the high impatience group is higher for all measures of cash-on-hand. However, buffer stock behavior is rejected in both groups, as $\hat{\theta}$ ranges between 1.1% and 2.5%.

11. Frederick, Loewenstein, and O'Donoghue (2002) emphasize that measurement of time preference can be affected by confounding factors, such as uncertainty, intertemporal arbitrage, and consumption smoothing.

TABLE 5. Estimated covariance ratio: sample splits by rate of time preference.

	$Y + \lambda W_r + W_f$				Number of observations
	$\lambda = 1$	$\lambda = 0.75$	$\lambda = 0.5$	$\lambda = 0.25$	
High impatience	0.011 (0.001)	0.014 (0.002)	0.020 (0.003)	0.025 (0.006)	188
Low impatience	0.012 (0.001)	0.014 (0.002)	0.017 (0.003)	0.011 (0.006)	

Notes: The sample splits are based on a question in the 2000 SHIW asking the respondent to report how much he or she is willing to pay to receive a lottery winnings today instead of later in time. "High" and "Low impatience" refer to values of the reported rate of time preference greater or lower than 3%. Observations for 2000 are then merged with data from 2002–2004, and estimation is performed on the resulting pooled data. W_r and W_f are, respectively, real and financial wealth, Y is disposable income, and λ is the share of real assets that can be used in the current period to finance consumption. Standard errors are reported in parentheses.

4.4. Panel Estimates

In Section 2.1 we stress that the covariance ratio θ can be obtained by the cross-sectional regression $c_h = \eta + \theta x_h + u_h$. Any valid estimation of such an equation requires instruments that are correlated with x but uncorrelated with c , except through their effect on x . Indeed, our framework posits that the wealth gap ($x_h - x_h^*$) provides such an instrument. In the buffer-stock model, the only reasons why the wealth gap might differ across consumers are the history of household-specific shocks, and the parameters determining target wealth (risk aversion, rate of time preference relative to the interest rate, growth rate of income, and standard deviation of permanent income shocks).

In practice, however, one cannot rule out that there is some systematic relationship between the error term in consumption and what people report about their target wealth. As an example, suppose that some households are patient and have larger target wealth than others, and that it takes them longer to reach their target wealth than impatient consumers. Then during the transition period following an income shock one would observe that households with large negative values of the wealth gap (the patient ones) consume less. More generally, any omitted variable might lead to biased estimates of θ .

To address the potential source of bias arising from unobserved heterogeneity, we can rely on the panel section of the SHIW. A total of 4,408 households were interviewed in 2002 and 2004, providing data on all the relevant variables in two time periods. Excluding households where the head is older than 50 years, results in a two-year panel of 1,087 households for a total of 2,174 observations.

The first row of Table 6 reports fixed effect IV estimates of c_h using $(x_h - x_h^*)$ as instrument for x_h . In the total sample $\hat{\theta} = 0.007$. The parameter is precisely estimated and even lower than the cross-sectional estimate. The other cells of the table report $\hat{\theta}$ for various definitions of wealth. Overall, the panel results suggest

TABLE 6. Estimated covariance ratio: sensitivity checks.

	$Y + \lambda W_r + W_f$				Number of observations
	$\lambda = 1$	$\lambda = 0.75$	$\lambda = 0.5$	$\lambda = 0.25$	
IV fixed effect	0.007 (0.001)	0.008 (0.002)	0.008 (0.002)	0.006 (0.003)	2,174
Permanent income (2000–2004 average)	−0.001 (0.001)	−0.003 (0.001)	−0.004 (0.001)	−0.006 (0.002)	1,601
Consumption includes durables	0.014 (0.001)	0.016 (0.001)	0.019 (0.001)	0.017 (0.002)	5,911

Notes: The table reports the estimated covariance ratio obtained, in turn, as fixed effect IV estimate in the 2002–2004 panel; using the 2000–2004 panel to estimate permanent income; adopting the definition of consumption including durables. W_r and W_f are, respectively, real and financial wealth, Y is disposable income, and λ is the share of real assets that can be used in the current period to finance consumption. Standard errors are reported in parentheses.

that, if any bias exists, it cannot explain rejection of the model obtained in the absence of unobserved heterogeneity correlated with the instrument.

Panel data also offer the opportunity to test model predictions that do not depend on the target wealth variable directly. Under the buffer stock model, low wealth households hit by negative income shocks should save and high wealth households dissave. Accordingly, we compute median change in cash-on-hand for poor and rich households. In our sample median saving is positive in both groups (€4,000 and €6,482, respectively).¹² Although it is hard to draw inference with a two-year panel, this represents further evidence against the buffer-stock model.

4.5. Further Sensitivity Checks

Our measure of permanent income, which we use to normalize cash-on-hand, consumption, and target wealth, is obtained through cross-sectional regressions, and may not be purged from transitory components. In the second row of Table 6 we report the estimated covariance ratio using an alternative measure, obtained averaging household disposable income (net of financial income) over time. Averaging should remove income components that are purely transitory and mean-reverting. This measure can only be computed for the panel section of the SHIW. Accordingly, we report estimates using the panel section of the last three surveys (2000, 2002, and 2004). The drawback is that the number of observations is considerably reduced and that the evolution of income in the short time period covered by the panel is a function of shocks that arrived during that period. In practice, the estimated covariance ratio appears to be very similar.

12. Poor households have less than €10,000 in cash-on-hand, and rich more than €60,000.

All our tests have been conducted defining consumption as non-durable expenditure. SHIW has also data on expenditures on durable goods, and therefore we can use total expenditure as an alternative measure of consumption. As shown in the last row of Table 6, the results barely change.

4.6. Measurement Error

In our baseline estimates of Table 4 we estimate a covariance ratio of 0.012, to be contrasted with a simulated value of 0.62. Can measurement error account for such a large difference between the theoretical benchmark and empirical estimates? To explore the robustness of our findings in the presence of measurement error, suppose that consumption, cash on hand and target wealth are all measured with error:

$$\tilde{c}_h = c_h + \varepsilon_h^c, \quad \tilde{x}_h = x_h + \varepsilon_h^x, \quad \tilde{x}_h^* = x_h^* + \varepsilon_h^{x^*},$$

where tilded variables are observed, untilded are true, unobserved values, and ε_h^k is a measurement error in variable k having mean zero. Under the assumptions that the errors are uncorrelated with each other and with true consumption, cash-on-hand and target wealth, one can show that the probability limit of our IV estimator of θ is

$$\theta(\nu, \xi) = \text{plim } \hat{\theta} = \theta \left(\frac{1 - \nu - \xi}{1 - \xi} \right) < \theta,$$

where ν is the percent variation in measured cash on hand explained by measurement error, and $\xi = \sigma_{\tilde{x}\tilde{x}^*} / \sigma_x^2$. The expression shows that in the presence of measurement error our estimator can be indeed downward-biased. Thus, one could reject the buffer-stock model even when the model is true, at least in principle.

To establish how large should measurement error be in order to reconcile our results with the buffer-stock model, we plot the probability limit of $\hat{\theta}$ as a function of ν and compare it with $\hat{\theta}$. As long as the probability limit of $\hat{\theta}$ is larger than estimated $\hat{\theta}$, measurement error cannot account for the model rejection. To compute the probability limit of $\hat{\theta}$ we set θ to its baseline simulated value (0.62), and note that ξ can be obtained as the slope coefficient in the OLS regression of \tilde{x}_h^* on \tilde{x}_h , which we find to be 0.0349.

In Figure 4 we plot the estimated covariance ratio (the horizontal line $\hat{\theta} = 0.012$) and the probability limit of $\hat{\theta}$ against measurement error ν . The figure shows that only for very large values of ν the probability limit of $\hat{\theta}$ falls below $\hat{\theta}$: that is, measurement error leads to false rejection of the buffer-stock model only if $\nu > 0.95$. In other words, only extremely large measurement error in wealth can possibly make $\hat{\theta} = 0.012$ consistent with buffer-stock behavior. Because the reliability index of income and wealth in the SHIW exceeds 80%, it is extremely




FIGURE 4. The effect of measurement error in the estimated covariance ratio. The dashed line is the estimated covariance ratio ($\hat{\theta} = 0.012$) from the total sample regression in Table 4 with cash on hand defined as $Y + W_r + W_f$. The continuous line is the probability limit of $\hat{\theta}$ as function of v , the percent variation in measured cash on hand explained by measurement error. The probability limit is computed setting $\theta = 0.623$ and $\xi = 0.0349$.

unlikely that measurement error invalidates our test, regardless of the wealth definitions.¹³

5. The Wealth–Income Ratio of Young Households

The version of the buffer-stock model that we analyze is one with impatient consumers, uncertainty about future earnings, and no borrowing constraints. If such consumers are sufficiently prudent and expect their earnings to grow over time, they will never borrow and keep their consumption within their current incomes, thus inducing “tracking” between consumption and income. In other versions of the buffer-stock model, impatient consumers would like to borrow but are prevented to do so because of credit market imperfections, as in Deaton (1991). The implications for the behavior of consumption and wealth are similar, however, and “consumption is smoothed, not over the whole life-cycle, but over

13. Biancotti, D’Alessio, and Neri (2004) give extensive account of the quality of the main variables in SHIW. Exploiting the panel section of the survey, they compute the reliability index for a broad range of variables. The index is the fraction of total variability of the measured characteristic accounted by its true variability.

much shorter periods of a few years at a time" (Deaton 2005, p. 103). In the literature, this is often referred to as "high-frequency" smoothing of income, as opposed to the "low-frequency" or "life-cycle frequency" smoothing that was postulated by Modigliani and Brumberg (see Browning and Crossley 2001).

Tracking of income and consumption and buffer stock behavior stand in sharp contrast with one of the most important implications of the life-cycle hypothesis, according to which young people save for post-retirement expenditures, and accumulate wealth up to retirement. In the certainty version of the model, the wealth to permanent income ratio increases during the working span, target wealth to permanent income ratio is reached at retirement age, and the consumption and income profiles are completely detached. If income is expected to increase over the working life, consumers borrow early in life, and start accumulating wealth only when debt is repaid, which might be even after several years of work, depending on preferences and the growth rate of individual incomes (Hubbard and Judd 1986).

In a more sophisticated version of the life-cycle model with income risk and life uncertainty, Hubbard, Skinner, and Zeldes (1995) show that sufficiently patient consumers save even earlier in life. In these life-cycle models with income risk, uncertainty generates a demand for precautionary saving during the working span. But, as noted by Modigliani (1986), cash on hand can serve the double purpose of providing resources for retirement and a buffer against unexpected emergencies.

Gourinchas and Parker (2002) results fall in between these two polar cases. They estimate that the behavior of young consumers exhibits buffer-stock behavior, at least in the United States. These consumers would like to borrow but cannot, or are too prudent to borrow. One way or another, their consumption tracks income closely and the wealth-income ratio is approximately constant. Once consumers reach middle-age, however, they follow the standard life-cycle model and the wealth-income ratio increases until retirement. Similar tracking of income and consumption arises in models with hyperbolic discounting, see Laibson, Repetto, and Tobacman (1998).¹⁴ The age profile of the wealth-income ratio of working age consumers provides therefore a useful avenue to distinguish different classes of models of intertemporal choice.

In the previous section we established that Italian wealth data are at variance with the buffer stock model. Even though we select a sample where buffer stock behavior is most likely to arise (individuals aged 20 to 50, or individuals with relatively high rates of time preference), we do not find evidence that deviations of wealth from target are offset by changes in consumption. What then explains the

14. The composition of wealth, however, differs between models with exponential and hyperbolic discounting, because hyperbolic consumers hold a smaller share of assets in liquid form, see Angeletos et al. (2001).




FIGURE 5. The age profile of the median wealth–income ratio. The two graphs on the left report the median wealth–income ratio of selected cohorts using the 1989–2002 SHIW. We use two definitions of wealth: the broad definition is $Y + W_r + W_f$, and the narrow definition is $Y + 0.25W_r + W_f$. The two graphs on the right report the estimated age profiles, obtained by regressions of the wealth–income ratio on age dummies, cohort dummies, and restricted year dummies summing to zero and orthogonal to a time trend.

saving decisions of young households? To provide further evidence of the validity of the buffer stock model, we estimate the age profile of the wealth–income ratio with seven SHIW waves, running from 1989 to 2002 (almost 60,000 households). To account for the fact that some of the wealth is illiquid and cannot be used for precautionary purposes, we use two definitions of $x = (W_f + \lambda W_r + Y)/P$, with $\lambda = 0.25$ and $\lambda = 1$.

We sort the data by the year of birth of the head of the household. The first cohort includes all households whose head was born in 1939 (50 years old in 1989, the year of the first survey). The second includes those born in 1940, and so on up to the last cohort, which includes those born in 1980 (22 years old in 2002, the last year). As with other survey data, the distribution of x is skewed. We therefore report only results for the median ratio; results for the 25th and 75th percentiles exhibit similar patterns.

The left graphs in Figure 5 offer important insights into the process of wealth accumulation of young Italian households. To make the graphs more readable, we plot x for selected cohorts. The numbers in the graph refer to the year of birth,

from 50 (individuals born in 1950) to 65 (individuals born in 1965). Except for the youngest and the oldest generations, each cohort is observed at seven different points in times, one for each cross-section. As said, the cross-sections run from 1989 to 2002. Thus, each generation is observed for 13 years with each line being broken (for instance, cohort 60 is sampled 7 times from age 29 in 1989 to age 42 in 2002). The x ratio is potentially affected by age, cohort, and time effects.

To estimate the age profile of x , we use 203 age/year/cohort cells and proceed as Deaton and Paxson (1994), regressing x on age dummies, cohort dummies, and restricted year dummies, summing to zero and orthogonal to a time trend.¹⁵ Given the structure of our sample, the regressors include 28 age dummies (from age 22 to age 50), 41 cohort dummies (from 1939 to 1980), a set of restricted time dummies, and a constant term. The estimated age dummies can be interpreted as an individual age-wealth profile, purged from cohort effects. They are plotted on the right-hand side of Figure 5.

Using the broad definition of cash-on-hand ($\lambda = 1$), between age 20 and 50 there is a three-fold increase in x (from 2 to 6); using a narrow definition ($\lambda = 0$), a two-fold increase (from 1.5 to 3). Overall, the graphs suggest that models in which consumption and income of young households track each other closely are not an adequate description of the behavior of Italian households. Rather, consumers start saving early in life, and accumulate assets at the rate of around 10% of their income, or €3,000 per year.¹⁶

6. Conclusions

Intertemporal models with liquidity constraints, income risk, and impatience emphasize that consumers use savings to buffer income fluctuations. These models deliver a stationary distribution of the ratio of target wealth to permanent income. When actual wealth, relative to income, is below the optimal target, consumers try to increase their saving. When wealth is above target, they increase consumption. This important implication of the buffer-stock model has not been subject to direct empirical testing.

We derive from the model an appropriate theoretical restriction and test it using data drawn from the 2002 and 2004 SHIW. One of the most appealing features of these surveys is that people report the amount of wealth held for

15. An alternative identification assumption is to express the ratio as a function of age dummies and unrestricted time dummies (eliminating cohorts effects). This alternative decomposition delivers qualitative similar results, for example, an increasing x during the early part of the life-cycle. Both normalizations rule out time-age or time-cohort interaction terms.

16. Because wealth accumulation does not depend only on age and cohort, we also experiment adding to the basic regression household size and composition, a dummy for retirement, education, gender, and region of residence. The qualitative results of increasing wealth-income ratio in Figure 4 is unchanged.

precautionary purposes, which we interpret as target wealth in the buffer-stock model. The test results do not support buffer-stock behavior, even among population groups that are more likely, a priori, to display such behavior (the young and the self-employed). Unobserved heterogeneity and measurement error in target wealth or consumption are unlikely to explain the model's failure.

Our test rejects the buffer-stock model, but cannot necessarily be interpreted as suggesting that alternative consumption theories are valid. In the final part of the paper we therefore use estimates of the age-wealth profile obtained with Italian repeated cross-sectional data to provide further evidence on the validity of the buffer-stock model. Indeed, the model predicts that for the young consumption tracks income closely and that the wealth-income ratio is approximately constant. In fact, we find that the wealth-income ratio of young Italian households increases substantially with age.

Overall, the saving behavior of young households in our data is hard to reconcile with models predicting a close parallel between consumption and income, such as hyperbolic discounting and preference reversal models. The evidence is instead consistent with models in which impatience, relative to prudence, is not as high as in buffer stock models, and with models where life-cycle considerations (such as saving for home purchase or for retirement) are of paramount importance to understand the behaviour of the young.

References

- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001). "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives*, 15, 47–68.
- Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber (1999). "Humps and Bumps in Lifetime Consumption." *Journal of Business and Economic Statistics*, 17, 22–35.
- Biancotti, Claudia, Giovanni D'Alessio, Ivan Faiella, and Andrea Neri (2004). "I bilanci delle famiglie italiane nell'anno 2002." *Supplementi al Bollettino Statistico* 14, no. 12. Bank of Italy.
- Biancotti, Claudia, Giovanni D'Alessio, and Andrea Neri (2004). "Errori di misura nell'indagine sui bilanci delle famiglie italiane." *Temi di discussione*, no. 520. Bank of Italy.
- Browning, Martin, and Thomas F. Crossley (2001). "The Life-Cycle Model of Consumption and Saving." *Journal of Economic Perspectives*, 15(3), 3–22.
- Cagetti, Marco (2003). "Wealth Accumulation over the Life Cycle and Precautionary Savings." *Journal of Business and Economic Statistics*, 21, 339–353.
- Campbell, John Y., and Gregory N. Mankiw (1991). "The Response of Consumption to Income: A Cross-Country Investigation." *European Economic Review*, 35, 723–767.
- Carroll, Christopher D. (1997). "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, 112, 1–56.
- Carroll, Christopher D. (2001). "A Theory of the Consumption Function, with and without Liquidity Constraints." *Journal of Economic Perspectives*, 15(3), 23–45.
- Carroll, Christopher D. (2004). "Theoretical Foundations of Buffer Stock Saving." Working Paper no. 517, John Hopkins University.

- Carroll, Christopher D., and Andrew A. Samwick (1997). "The Nature of Precautionary Wealth." *Journal of Monetary Economics*, 40, 41–71.
- Deaton, Angus S. (1991). "Saving and Liquidity Constraints." *Econometrica*, 59, 1221–1248.
- Deaton, Angus S. (2005). "Franco Modigliani and the Life Cycle Theory of Consumption." *Banca Nazionale del Lavoro Quarterly Review*, 58(233–234), 91–108.
- Deaton, Angus, and Christina Paxson (1994). "Saving, Aging and Growth in Taiwan." In *Studies in the Economics of Aging*, edited by David Wise. University of Chicago Press.
- Faiella, Ivan, Romina Gambacorta, Stefano Iezzi, and Andrea Neri (2006). "I bilanci delle famiglie italiane nell'anno 2004." *Supplementi al Bollettino Statistico* 16, no. 7. Bank of Italy.
- Frederick, Shane, George F. Loewenstein, and Ted O'Donoghue (2002). "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40, 351–401.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker (2002). "Consumption over the Life Cycle." *Econometrica*, 70, 47–89.
- Gruber, Jonathan, and Aaron Yelowitz (1999). "Public Health Insurance and Private Savings." *Journal of Political Economy*, 107, 1249–1274.
- Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (1992). "Earnings Uncertainty and Precautionary Saving." *Journal Monetary Economics*, 36, 307–337.
- Hubbard, Glenn R., and Kenneth L. Judd (1986). "Liquidity Constraints, Fiscal Policy, and Consumption." *Brookings Papers on Economic Activity*, 1, 1–50.
- Hubbard, Glenn R., Jonathan Skinner, and Stephen P. Zeldes (1995). "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103, 360–399.
- Hurst, Erik, Annamaria Lusardi, Arthur Kennickell, and Francisco Torralba (2005). "Precautionary Savings and the Importance of Business Owners." NBER Working Paper no. 11731.
- Jappelli, Tullio, and Marco Pagano (1989). "Consumption and Capital Market Imperfections: an International Comparison." *American Economic Review*, 79, 1088–1105.
- Jappelli, Tullio, and Luigi Pistaferri (2006). "Intertemporal Choice and Consumption Mobility." *Journal of the European Economic Association*, 4, 75–115.
- Kennickell, Arthur, and Annamaria Lusardi (2004). "Disentangling the Importance of the Precautionary Saving Motive." NBER Working Paper no. 10888.
- Laibson, David I., Andrea Repetto, and Jeremy Tobacman (1998). "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity*, 1, 91–196.
- Modigliani, Franco (1986). "Life-Cycle, Individual Thrift, and the Wealth of Nations." *American Economic Review*, 76, 297–313.