

CSEF

Centre for Studies in
Economics and Finance

WORKING PAPER NO. 574

Permanent Income Shocks, Target Wealth, and the Wealth Gap

Tullio Jappelli and Luigi Pistaferri

August 2020



University of Naples Federico II



University of Salerno



Bocconi

Bocconi University, Milan

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS - UNIVERSITY OF NAPLES
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 - e-mail: csef@unina.it
ISSN: 2240-9696

WORKING PAPER NO. 574

Permanent Income Shocks, Target Wealth, and the Wealth Gap

Tullio Jappelli* and Luigi Pistaferri**

Abstract

We test the key implication of the buffer stock model, namely that any revision in permanent income leads to a proportionate revision in target wealth. We use panel data on the amount of wealth held for precautionary purposes available in the 2002-2016 SHIW. Using an instrumental variable approach to overcome measurement error issues and direct estimates of the permanent component of income, we find that households indeed revise approximately one-for-one their target wealth in response to permanent income shocks. We explore heterogeneity of the response across the cash-on-hand distribution, for positive and negative shocks, and for shocks of different size. We also find that the change in the ratio of cash-on-hand to permanent income is negatively correlated with the “wealth gap”, particularly for individuals whose wealth is substantially above target.

Keywords: Buffer Stock Model; Target Wealth; Wealth Gap; Permanent Income Shocks; Panel Data.

JEL Classification: D12, D14, E21

Acknowledgements: Jappelli acknowledges financial support from the Italian Ministry of University and Research (2017 PRIN Grant).

* Università di Napoli Federico II, CSEF and CEPR. E-mail: tullio.jappelli@unina.it

** Stanford University, SIEPR, NBER and CEPR. E-mail: pista@stanford.edu

Table of contents

1. *Introduction*

2. *The strength of precautionary saving*

3. *Target wealth and the wealth gap*

4. *The data*

5. *Regression results*

5.1. *Using consumption as a proxy for permanent income*

5.2. *Using the stochastic structure of income to isolate the permanent component of income*

5.3. *Convergence to the target*

6. *Conclusions*

References

Figures and Tables

1. Introduction

In the past three decades the buffer stock model has become one of the leading models for explaining saving behavior, and a vast body of evidence has confirmed many of its theoretical tenets. Several papers - using reduced form approaches, structural estimation, quasi-experimental evidence, and direct survey evidence - provide empirical tests of the model and show that precautionary assets represent a significant determinant of household wealth accumulation decisions. Other papers focus on the model's implication of a concave consumption function, leading to the prediction that the marginal propensity to consume of the poor is greater than the propensity of the rich. Still others have extended the model to consider liquid and illiquid assets. Carroll (2020) provides a rigorous treatment of the model. Carroll and Kimball (2007) and Jappelli and Pistaferri (2017) survey the empirical evidence on precautionary savings.

One of the key theoretical implications of the buffer stock model is that consumers' decisions balance two opposite forces. The first is prudence, which leads consumers to save for precautionary reasons; the second is impatience, which instead leads consumers to defer savings. As a result of these opposite forces, consumers choose to maintain a "target" level of wealth that is proportional to permanent income. This implies that, other things being equal, *any revision in permanent income leads to a proportionate revision in target wealth*. The intuition is simple: to insure a higher level of consumption, people need to scale up their precautionary assets, and *vice versa*. To our knowledge, this key implication of the model has not been directly tested in previous empirical literature.

To estimate the effect of revisions to permanent income on revisions of target wealth, one needs an operational measure of both. We rely on survey data on the amount of wealth held for precautionary purposes, which we interpret as target wealth in a buffer stock model. To measure permanent income, we consider two empirical strategies: (a) approximating permanent income with non-durable consumption, and (b) isolating the permanent component from the stochastic structure of the income process. Using Italian panel data from 2002 to 2016, both empirical strategies reveal that households indeed revise approximately one-for-one their target wealth in

response to permanent shocks. We also explore heterogeneity of the response across the cash-on-hand distribution, for positive and negative shocks, and for shocks of different size.

While this finding strongly supports the buffer stock model, suggesting that people understand what they should do in response to a shock, it still begs the question of whether people actually adjust their stock of wealth when it is off-target. This is a key issue, as in many household surveys (including the US Survey of Consumer Finances) people typically report that they are unprepared to meet even small financial emergencies. In the second part of the paper we thus check whether the ratio of cash-on-hand to permanent income is negatively correlated with the “wealth gap”, the deviation of cash-on-hand from target wealth (relative to permanent income). In our panel data we find a negative, sizable and statistically significant adjustment coefficient: on average, it takes about 3.5 years to close the gap between actual wealth and target wealth (absent additional shocks). We uncover substantial heterogeneity of responses, in particular for positive and negative wealth gaps: people appear much more likely to decumulate when they have too much wealth relative to target than to increase wealth when the gap is negative, possibly due to adjustment costs, consumption habits, or inertia.

The rest of the paper is organized as follows. In Section 2 we review the evidence on precautionary saving, and in Section 3 we discuss our empirical tests. Section 4 presents the data and our operational measure of target wealth. Section 5 focuses on the relation between target wealth and permanent income, and how the response varies by cash-on-hand, the sign and the size of permanent income shocks. We also provide several robustness checks, and study whether people actually adjust their wealth to target, by estimating how wealth accumulation responds to the wealth gap, its size and sign. Section 6 concludes.

2. The strength of precautionary saving

Tests of the validity of the buffer stock model rely on empirical studies of the strength of precautionary saving, which is what motivates consumers to accumulate assets in case of unexpected income shocks and other emergencies. A first research strategy is to estimate the degree

of prudence from the consumption Euler equation. This approach was first implemented by Dynan (1993), who approximated the conditional variance of consumption growth in the Euler equation with the realized variance, instrumented with socioeconomic variables. The approach was later refined by Bertola, Guiso, and Pistaferri (2005), who use the subjective variance of income as an instrument, Christelis, Georgarakos, and Jappelli (2019), who use subjective expectations on future consumption growth, and Fagereng, Guiso, and Pistaferri (2017), who use firm-related risk. These three studies all find a coefficient of relative prudence around 2, providing support for models in which consumers save for precautionary reasons.

A second strategy uses estimation-by-simulation methods to match empirical and theoretical moments of the consumption and wealth distributions under precautionary saving. Gourinchas and Parker (2002) estimate a rate of time preference of 4% and a coefficient of relative risk aversion of 2. The match between the theoretical life-cycle consumption profile generated by these estimates and the actual profile is quite good, supporting the precautionary saving model. Moreover, they find that young consumers accumulate little wealth (since for them impatience outweighs prudence), and that it is only in middle age that people start accumulating assets for retirement. Cagetti (2003), using a similar approach, confirms these findings.¹

A third strategy is to estimate a reduced form for wealth relying on various proxies for income risk. Skinner (1988) uses occupational dummies, while Guiso, Jappelli and Terlizzese (1992) rely on the subjective variance of future income changes. Carroll and Samwick (1997) find that wealth is positively associated with the variance of transitory and permanent income shocks estimated from the Panel Study of Income Dynamics. However, any risk indicator is subject to the problem of self-selection. Individuals in risky occupations may have sorted into those jobs because they are less risk-averse, in which case their wealth accumulation behavior might be no different from the average even if they have a precautionary saving motive. To address this selection issues, Fuchs-Fuchs-Schündeln and Schuendeln (2005) use the German unification shock as a quasi-experiment.

¹ A combination of the two strategies is Attanasio, Banks, Meghir and Weber (1999), who estimate preference parameters using the Euler equation but then use their estimates to simulate consumption profiles in alternative scenarios, showing that the data are best explained by a model with precautionary saving and preferences that depend on demographic variables.

Their results support the precautionary model and also demonstrate the importance of the selection bias induced by the correlation between risk preference and income risk.

The fourth approach, which is the one we follow in this paper, evaluates the importance of precautionary saving using survey questions that elicit the importance of the various motives for saving. The Survey of Consumer Finances in the US and the Survey of Household Income and Wealth (SHIW) in Italy ask individuals about their target wealth, namely how much they think that they should have in savings to face income risk and other emergencies. These data have been used by Kennickell and Lusardi (2004), Jappelli, Padula, and Pistaferri (2008), and more recently Fulford (2015). Kennickell and Lusardi (2004) report that the bulk of the distribution of target wealth in the US is between \$5,000 and \$10,000. Fulford (2015) shows that this measure of target wealth is much lower than predicted by standard modelling assumptions, and that perceived income uncertainty does not affect target wealth (while it should). Jappelli, Padula, and Pistaferri (2008) test the proposition that people with a below-target wealth expect to save, while those with above target expect to dissave, and reject this implication of the buffer stock model.

In this paper we use the same set of questions to test another implication at the heart of the buffer stock model, that is whether people revise target wealth in response to permanent income shocks. We also point out that direct survey evidence is most useful with genuine panel data, while cross-sectional inference is plagued by bias arising from unobserved heterogeneity correlated with preference and individual income risk.

Our paper is part of a larger literature that exploits subjective expectations and direct survey questions to explore various features of consumption decisions. For instance, some surveys contain questions on how consumers have reacted to actual income changes or asking consumers to report how they would respond to hypothetical income changes. Shapiro and Slemrod (1995) pioneered this approach asking direct questions in the Michigan Survey of Consumers, focusing on spending in response to the various tax rebates and tax credit interventions taking place in the US in the past two decades (Shapiro and Slemrod, 2009; Sahm et al. 2015). Jappelli and Pistaferri (2014; 2020) use Italian survey data where consumers were asked to report the fraction of a positive income shock that they would consume. Christelis et al. (2019) analyze a survey of Dutch households in which respondents report how much their consumption would change in response to unexpected

shocks of different sign and size. Most of these papers highlight the importance of household heterogeneity for assessing the impact of macroeconomic policies and tax reforms on output and distributional issues (see Jappelli and Pistaferri, 2010, Heathcote, Storesletten and Violante, 2009, and Krueger, Mitman and Perri, 2016, for excellent surveys).

3. Target wealth and the wealth gap

There are two versions of the buffer stock model in the literature. Both emphasize the interaction between liquidity constraints and precautionary saving. One version of the model, developed by Deaton (1991), considers the possibility that a prudent and impatient consumer may face explicit credit constraints. The other, proposed by Carroll (1997), features the same type of consumer but allows for the possibility of income falling to zero and so generating a “natural” borrowing constraint. The two versions of the model deliver similar implications, and here we follow closely Carroll’s version. One important implication of the buffer stock model is that optimal consumption is an increasing and concave function of cash-on-hand, in contrast with the certainty equivalence version of the permanent income model, where the consumption function is linear. The second implication of the model is that there exists a unique and stable value of target wealth as a ratio to permanent income, such that, if actual wealth is greater than the target, impatience outweighs prudence, and wealth falls, while if wealth is below the target, the precautionary saving motive outweighs impatience, and the consumer accumulates wealth. This theoretical mechanism is precisely the focus of our paper.

Denote target wealth by w^* , permanent income by P , and let’s consider the ratio $x^* = w^*/P$ as the “unique and stable value of the ratio of target wealth to permanent income” in the parlance of the buffer stock model. There are two sources of heterogeneity in x^* . The first source is preferences, that is, more patient and more risk averse individuals tend to accumulate more wealth to protect against unforeseen events and have a higher x^* . The second source of heterogeneity concerns the parameters of the income generating process. In particular, higher income volatility and higher probability of unemployment require more precautionary saving and therefore a higher x^* . Carroll

(2020) proves these results in the general case, and also provides analytical solutions in some special cases.²

If we believe these sources of income and preference heterogeneity are stable over time for a given individual, they can be adequately captured by a fixed effect, and hence we can write:

$$\frac{w_{it}^*}{P_{it}} = \tilde{\theta}_i$$

And, in logs:

$$\ln w_{it}^* = \ln P_{it} + \theta_i, \tag{1}$$

where $\theta_i = \ln \tilde{\theta}_i$. Given the nature of our data, we focus on the steady-state implications of the model. The central prediction of equation (1) is that *any revision in permanent income leads to a proportionate revision in target wealth*, holding individual effects (preferences and the income generating process) constant. To exemplify, suppose that people receive a permanent and positive income shock, such as a job promotion. This leads to a revision in consumption, which will be permanently higher. To protect and insure this higher level of consumption from income shocks, the optimal plan requires now a higher level of target wealth. Symmetrically, if people receive a negative and permanent income shock, such as job demotion, consumption and the required target wealth will be permanently lower.

To make the model operational, assume we have an empirical counterpart of target wealth and permanent income, and let's consider the following regression:

² When the period utility function is isoelastic, it can be shown that $x^* = \frac{w^*}{P} = \frac{1}{[(\gamma-r)+\delta(1+\frac{\gamma}{\pi})(1-\frac{\gamma}{\pi}\omega)]}$, where γ is the growth rate of income, r is the interest rate, δ is the rate of time preference, π is the probability of falling into permanent unemployment, and ω is the coefficient of prudence. The expression shows that an increase in the growth rate of labor income, a reduction in the interest rate (equivalent to an increase in human wealth), and a reduction in the probability of unemployment reduces target wealth relative to permanent income. More impatient people have lower target wealth. On the other hand, an increase in prudence reduces the denominator and therefore is associated with higher target wealth.

$$\ln w_{it}^* = \alpha + \beta \ln P_{it} + \theta_i + v_{it} \quad (2)$$

where θ_i is a time-invariant individual fixed effect capturing unobserved heterogeneity correlated with target wealth (and potentially also with permanent income), and v_{it} an i.i.d. error term capturing classical measurement error in reported target wealth (or permanent income). The buffer stock model suggests that the ratio between target wealth and permanent income is a constant. This implies that, at the individual level, target wealth and permanent income move one-to-one, so that $\beta=1$. This test, which to the best of our knowledge has never been implemented, is a sufficient statistic for the validity of the steady-state solution of the buffer stock model.

Cross-sectional data are not adequate for testing (2), since the fixed effect is related to individual preferences and income risk, which are clearly correlated with permanent income. For instance, hard-working individuals with higher permanent income may also be more prudent and save more. Therefore, in the presence of unobserved heterogeneity potentially correlated with target wealth, the OLS estimate of β is biased and inconsistent. In particular, the estimated β coefficient is lower than the true value. From regression (2), the bias can be inferred by computing the probability limit of $\tilde{\beta}_{OLS}$:

$$\text{plim} \tilde{\beta}_{OLS} = \beta + \frac{\text{cov}(P_{it}, \theta_i)}{\text{var}(P_{it})}$$

The expression above shows that the bias generated by unobserved heterogeneity depends on the sign and magnitude of the covariance term $\text{cov}(P_{it}, \theta_i)$. Suppose that θ_i represents unobserved differences in rates of time preferences, implying that people with high θ_i have lower taste for saving, and that taste for saving and hard work are correlated. These individuals will report low target wealth and are also likely to have low permanent income. We hence expect $\text{cov}(P_{it}, \theta_i) > 0$. This implies that $\tilde{\beta}_{OLS}$ will be lower than the true β and the OLS estimate will underestimate the relation between permanent income and target wealth.

With panel data one can test whether individuals who have experienced a change in their permanent income also report a change in their target wealth, as opposed to testing whether (in a cross-section) high-permanent income individuals report higher target wealth, an association that may be related to risk-aversion or patience affecting both in the same direction. In particular, with panel data one can eliminate the bias by differencing the relationship (2), and hence estimate:

$$\Delta \ln w_{it}^* = \beta \Delta \ln P_{it} + \Delta v_{it} \quad (3)$$

As we discuss when we present our results, the validity of the test assumes that preferences or the income generating process (as indicated by the θ_i effects) don't change over time, and hence they can be conveniently differenced out. In our case, unobserved heterogeneity in target wealth should reflect preference traits (such as discount rate, risk tolerance, etc.) and parameters of the income process, which in the consumption literature are typically assumed to be constant over time (and parametrized as such in calibrated versions of the model). It is possible that changes in economic circumstances may shift such parameters. Our regressions control for age, family size, and other socio-economic characteristics in the attempt to minimize this possibility.

The test we propose in equation (3) requires empirical counterparts of target wealth and permanent income for several years. We use direct survey evidence on target wealth, relying on a question available between 2002 and 2016 in the SHIW, and described in Section 3. However, permanent income is unobserved. We discuss two different empirical strategies to tackle this issue. Under the assumption that people actually follow the buffer stock model, non-durable consumption is a good proxy for permanent income. A second approach is to use the stochastic structure of the income process to isolate the permanent component from transitory fluctuations.

Both empirical strategies present pros and cons. Using consumption as a proxy for permanent income is transparent and in agreement with theoretical predictions since Friedman's permanent income hypothesis (1957). However, a potential problem with using consumption is measurement error, which leads to an attenuation bias in β . This is likely to be an important issue in our context, since in the SHIW only a few consumption questions are available. Moreover, it is well known that first-differencing exacerbates the attenuation bias induced by measurement error if the right-hand

side variable (consumption) is correlated over time (Griliches and Hausman, 1986). Labor income provides a valid instrument. Empirically, it is strongly correlated with consumption. And in the buffer stock model, once we control for permanent income (as proxied by consumption), target wealth should not be affected by labor income, since permanent income is a “sufficient statistics” for the household’s current and future resources. We use labor income (net of taxes, and including transfers) because capital income may be potentially correlated with target wealth.

As for the alternative empirical strategy, the main advantage is that it does not rely on a proxy for permanent income. On the other hand, since it is based on the assumptions about the income process being correct, it imposes more structure on the data. Furthermore, the approach is more demanding in terms of data, and therefore one can less easily explore heterogeneity of responses of target wealth with respect to permanent income shocks of different size and sign. In practice, as we shall see, both strategies deliver similar findings.

One possible criticism of evidence based on self-reported target wealth is that some consumers might report what they wish to save rather than what they aim to save. Furthermore, some may provide answers “as if” they behave according to a buffer stock model, even though they have difficulty identifying what their target wealth-income ratio is, or will actually follow a different model. To shed light on this issue, one can check whether people who report below-target wealth attempt to close the wealth gap in the following periods. We thus estimate the following regression for the change in the actual wealth to permanent income ratio x :

$$x_{it} - x_{it-1} = \alpha + \delta(x_{it-1} - x_{it-1}^*) + e_{it} \quad (4)$$

A negative δ coefficient signals convergence towards target wealth. Since the regression differences out individual effects, this convergence concept is reminiscent of the convergence criteria of growth regressions: each individual has a different steady-state target wealth, and each converges to its own steady state. Furthermore, the magnitude of δ measures the speed of convergence. For instance, $\delta=0.5$ means that half of the wealth gap at time $t-1$ is filled between period $t-1$ and period t .

Parker and Souleles (2017) perform a similar check between reported and actual MPCs. In particular, they compare the “revealed preference” approach (in which inference is based on actual data) with the “reported response” approach, which consists of asking people to report their choices. They find that households reporting that they “mostly spent” their economic stimulus payments in 2008, had indeed spent twice as much as those reporting that they used their payments “mostly to save or pay down debt”. As we will see in Section 5, we find strong convergence, particularly when the gap is positive and large.

4. The data

To implement our empirical test of the buffer stock model, we use six waves of the Italian Surveys of Household Income and Wealth (SHIW). SHIW is a biannual representative sample of the Italian population conducted by the Bank of Italy. In each year, the sample includes about 8,000 households. The surveys provide detailed information on demographic variables, income, consumption, wealth (broken down into real assets and various components of financial assets and debt). The survey has also a rotating panel component: each year close to 50% of the sample is composed of households interviewed in the previous wave, while 50% represents new interviews.

In 2002, 2004, 2010, 2012, 2014, 2016 SHIW has a direct question on target wealth for precautionary reasons,³ which we take as a proxy of 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), and has been used also by Jappelli, Padula and

³ The question was not asked in 2006 and 2008 due to revisions in the special modules of the questionnaire.

Pistaferri (2007) to study some of the implications of the buffer stock model.⁴ As mentioned, the main advantage of the present paper is that we can exploit the panel structure of the data to test the key implication of the model.

The wording of the question is the same in the six surveys, but the way the question is introduced is abridged in 2014 and 2016 (see the Appendix for a description of the question in the various years). Since the 2014 and 2016 formulation might reduce attention to the question, we explore the stability of the results limiting the analysis to the period 2002-12.

Table 1 reports sample statistics for the whole 2002-2016 sample (46,569 households) and for the panel sample (25,738 households). To make sure that the question on target wealth (our empirical proxy for w^* in Section 2) is answered by the same person, our panel sample selects households with a stable demographic structure (the same household head with the same level of education, and no change in marital status or region of residence across the waves). Given the rotating structure of the panel, 20% of the households are interviewed five or six times, 44% three of four times, and 36% only in two consecutive waves.

The sample mean of target wealth (deflated using the consumer price index) is 35,667 euro, slightly higher than the corresponding values in the panel sample (34,932 euro). The median is actually the same in the two samples (20,030 euro). There are essentially no differences between the two samples in terms of gender, family size and marital status. Years of education, disposable income, non-durable consumption and financial wealth are 4 to 8 percentage points higher in the panel sample, as should be expected given panel attrition and our requirements to focus only on households with a stable demographic structure. The median target wealth to consumption ratio (our proxy for x in Section 2) is about 1, showing that precautionary saving is potentially quite important for Italian households, and higher than in the U.S. Indeed, using a similar proxy for target wealth, Kennickell and Lusardi (2004) and Fulford (2015) report that the bulk of the distribution of target wealth is between 1 and 2 months' worth of income.

Another interesting statistic is the wealth gap, $(x - x^*)$, where x and x^* are, respectively, the actual and target wealth to permanent income ratio (as proxied by consumption). In keeping with the buffer stock model, actual wealth is defined as cash-on-hand. The latter includes resources

⁴ Fulford (2015) uses the SCF question to look at the link between perceived income uncertainty and target wealth.

that are liquid and available to face emergencies, that is financial assets (transaction accounts, CDs, bonds, mutual funds, investment accounts, and stocks) plus monthly household disposable income. Most households are below target (the median wealth gap is -30% in the total sample and -26% in the panel sample), but approximately one third of respondents report a positive wealth gap. Later in the paper we will check if those who are off target actually attempt to close the gap by increasing or reducing cash-on-hand over time.

Figure 1 plots the histogram of the cross-sectional distribution of $\ln w^*$ from 2002 to 2016 using all sample observations, showing that in each year there is considerable heterogeneity of target wealth. The distributions are concentrated between values in the 10,000-50,000 euro range, corresponding to the first and third quartiles of the distributions, but 5% of households report values of less than 1,500 euro, and 10% over 100,000.

Figure 2 compares the distributions of $\ln w^*$ and of the wealth gap in the total sample and in the panel sample. Given the presence of outliers for target wealth and cash-on-hand, we winsorize the wealth gap at the top and bottom 1% of the distribution. The distributions of $\ln w^*$ and of the wealth gap are remarkably similar in the total and panel samples. As shown in Figure 3, except for an increase in 2002-04, households tend to report decreasing values of target wealth, particularly after the 2008 financial crisis. This applies not only to the mean and median, but also to other percentiles of target wealth.

One possible explanation for this reduction in planned precautionary savings is the stagnant productivity growth of the Italian economy of the last two decades, and weak, persistent expectations about the state of the economy. National accounts indicate that in 2000 per capita GDP was 29,157 euro, and after a moderate increase until 2007, due to the financial crisis and the sovereign debt crisis, in 2019 it fell to 28,583 euro, an average decline of approximately -0.1% per year.

Figure 4 provides a first check of the hypothesis that the ratio of target wealth to consumption (our proxy for permanent income) is constant. In the top panels we plot the log of target wealth for each consumption percentile. The two variables are strongly and positively correlated, although one can notice that the relation is weaker for low and high consumption percentiles. The bottom panel presents the same relation in first differences, plotting the growth

rate of target wealth against the percentiles of the growth rate of consumption. The graph supports the hypothesis that households increase/decrease their target wealth when they see changes in their permanent income (as proxied by consumption).

Figure 5 explores two dimensions of heterogeneity of the relation between growth of target wealth and consumption growth, showing that the positive association exists both at high and low levels of financial wealth. The two lower figures show that households who experience a reduction in consumption adjust downward target wealth, while households with positive consumption growth also tend to report an increase in target wealth. Overall, Figures 4 and 5 provide important qualitative support for the key implication of the buffer stock model that revisions in permanent income are associated with changes in target wealth.

What remains to be seen is whether households actually adjust cash-on-hand when they are off target. Figure 6 plots the change in cash-on-hand (the adjustment) against the lagged ratio of the wealth gap (both expressed as a ratio to consumption). The graph clearly depicts two situations. When their asset gap is positive, consumers actually reduce cash-on-hand in the next period. Instead, when the asset gap is negative, the relation is flat, and consumers hardly increase their cash-on-hand. In our regression analysis we will try to distinguish between these two regions.

5. Regression results

This section contains our empirical findings related to how target wealth responds to changes in permanent income. We present results using two strategies for capturing permanent income and discuss various extensions and robustness checks. Finally, we test the hypothesis that people actually adjust their wealth when they are off target.

5.1. Using consumption as a proxy for permanent income

Our first empirical strategy consists of proxying permanent income in (3) with consumption. The discussion in Section 2 suggests three possible regression methods to estimate the relation

between target wealth and our proxy of permanent income: OLS (exploiting the cross-sectional variability of the two variables), first differences (removing the fixed component of the wealth-permanent income ratio), and an IV first-difference estimator (using labor income as instrument for consumption). After presenting results using these three estimators (and discuss their potential biases), we perform various robustness checks and, most importantly, explore possible heterogeneity of β in response to shocks of different directions and size.

We start by pooling all available data (from 2002 to 2016) and run an OLS regression of the log of target wealth against log consumption and a set of demographic variables (age and its square, family size, marital status, education, and region of residence) using the full sample of 46,569 observations. The regression is reported in column 1 of Table 2, and includes also a full set of year dummies that control for aggregate shocks and differences across surveys. Standard errors in this and other regressions are clustered at the household level.

The estimated coefficient in the pooled OLS regression is $\tilde{\beta}_{OLS} = 0.69$. The coefficient is quite precisely estimated, and shows that individuals with relatively high consumption also report relatively high target wealth, in line with the descriptive evidence of Figure 3. However, as shown in the last row of Table 1, the restriction $\beta=1$ implied by the buffer stock model is strongly rejected at the 1% confidence level. The other coefficients indicate that the relation between age and target wealth is concave, reaching a maximum at age 55. As is usual in these types of analyses, without further restrictions one cannot disentangle a pure age effect from cohort effects. Target wealth increases by 1% for each year of education of the household head. Married couples report 8% more target wealth relative to singles, and males 3.4% more than females. Residents in the South report lower target wealth relative to the control group (the North), while for residents in the Centre target wealth is 11% higher.

Column (2) restricts the sample to the 25,738 panel households, but still performs pooled OLS estimation. Results are quite similar to the full sample regression. The estimated β is 0.70, almost identical to the whole sample estimate, and the other coefficients are similarly hardly affected. We take this as an indication that sample selection of panel observations is not a major issue in this context.

Column (3) reports the results of estimating the first differences specification. In this case the estimated $\tilde{\beta}_{FD} = 0.42$, still positive and significant, but substantially lower than the OLS estimates. Note that in this specification education and region of residence are time-invariant and are therefore dropped from the estimation. As we point out in Section 2, measurement error could be an important issue. In fact, if measurement error is classical and consumption is positively correlated over time (which is almost a certainty), panel data exacerbate the standard attenuation bias (Griliches and Hausman, 1986). Thus, measurement error in consumption is a potential explanation for the difference between the OLS estimates of columns 1 and 2, and the estimation in first difference of column 3.

Table 3 contains the key takeaway of the paper. In column (1) we report an IV regression for the specification in first difference, using labor income as instrument.⁵ The estimated coefficient is $\tilde{\beta}_{FD,IV} = 0.98$, and a formal test of the restriction $\beta=1$ is not rejected at standard statistical levels (p-value of 89%), strongly supporting the buffer stock model.

This main result could hide some asymmetries with respect to permanent shocks of different sign or size. The regression of column (2) distinguishes between positive and negative consumption growth, instrumenting them with positive and negative income growth.⁶ Two results are worth mentioning. First, the β coefficients are both close to one (0.85 for positive consumption growth and 1.09 for negative consumption growth). Second, one cannot reject the hypothesis that the coefficients are equal (p-value 62%). The results suggest that individuals revise target wealth upward when they receive a positive permanent income shock, and downwards when they receive a negative shock, exactly the main implication of the buffer stock model.

The regression of column (3) distinguishes between shocks of different size. In our sample, one third of households experiences large shocks (more than 15% of annual income growth in absolute value), and the remaining two thirds more moderate shocks. We interact consumption growth with a dummy for these relatively large shocks, and use large and small income growth as instruments. The point estimates of the β coefficients suggest that large shocks induce larger

⁵ This is the sum of earnings from dependent employment, self-employment, and transfers. Excluding transfers has no effect on the results.

⁶ Recent consumption literature finds that the consumption response to income shocks varies with the sign of the shock, the size of the shock, and individual resources (Jappelli and Pistaferri, 2014).

adjustments of target wealth (a coefficient of 1.18). In contrast, for individuals who experience relatively small shocks the adjustment of target wealth is only 0.23 and insignificant. These results may reflect salience or the presence of adjustment costs (smaller shocks are easier to miss than larger ones or induce only modest utility losses if ignored). However, perhaps because of the imprecision in this estimate, formally one cannot reject the hypothesis that the two coefficients are equal at the 5% level (although the hypothesis would be rejected at the 10% level).

Table 4 explores another source of possible heterogeneous response of target wealth to permanent income shocks. In particular, cash-poor consumers, who face potential liquidity constraints, might be more reluctant to increase their saving and target wealth in response to positive income shocks than households who do not face such constraints. In the first regression of Table 4 we interact beginning-of-period financial wealth quintiles with consumption growth. We find that the sensitivity of target wealth with respect to permanent income shocks is close to 1 (or statistically indistinguishable from 1) for each of the five financial wealth quintiles. In the other two regressions of Table 4 we split the sample among “poor” and “rich” households (columns 2 and 3, respectively), using as a threshold whether financial wealth exceeds two-months’ income (as in Zeldes, 1989). For both groups, we find again β coefficients that are quite close to 1 (0.94 for the poor and 1.02 for the rich).

Table 5 presents three robustness results. First, the estimated β is still close to 1 when we restrict the sample to individuals interviewed in three or five consecutive waves (columns 1 and 2). While the number of observations is clearly reduced, the coefficient equals 1.12 even in the most restricted sample, using only 4,258 observations. Second, in column (3) we restrict the sample to 2002-2012, dropping years in which the question on target wealth was asked in slightly different form. Results are again unchanged. The last regression in column (4) introduces lagged consumption growth as an additional regressor, to allow for some sluggishness in response to shocks. The relevant test is now that the sum of the coefficients of consumption growth and lagged consumption growth sum to 1, and again this restriction is not rejected (p-value 81%).

5.2. Using the stochastic structure of income to isolate the permanent component of income

In this section we use the stochastic structure of the income process to obtain a different measure of permanent income, thus avoiding using consumption as a proxy or the use of possibly invalid instruments. Suppose that the log income process can be written as:

$$\ln y_{it} = X'_{it}\gamma^y + \ln P_{it} + \varepsilon_{it}$$

where X is a vector of observables, P_{it} the permanent component of income, and ε_{it} an i.i.d. transitory shock. As typically done in the buffer-stock literature (Carroll, 2020), assume that the permanent component follows a random walk process:

$$\ln P_{it} = \ln P_{it-1} + \zeta_{it}$$

where ζ_{it} is an i.i.d. innovation. Taking first differences gives:

$$\Delta \ln y_{it} = \Delta X'_{it}\gamma^y + \zeta_{it} + \Delta \varepsilon_{it} \quad (5)$$

The buffer-stock model outlined in Section 3 predicts that the log of target wealth and the permanent component of income move one-to-one, i.e., that in the regression (2) (augmented to control for observable characteristics):

$$\ln w_{it}^* = X'_{it}\gamma^{w^*} + \beta \ln P_{it} + \theta_i + v_{it}$$

the coefficient $\beta = 1$. First differencing allows us to get rid of the fixed effect θ_i and yields:

$$\Delta \ln w_{it}^* = \Delta X'_{it}\gamma^{w^*} + \beta \zeta_{it} + \Delta v_{it} \quad (6)$$

Longitudinal data on target wealth and income can be used to estimate β by GMM using moment restrictions imposed by the model on the residuals of income growth (g_{it}^y) and target wealth growth ($g_{it}^{w^*}$). The latter are defined, from estimation of (5) and (6), as:

$$\begin{aligned} g_{it}^y &= \Delta \ln y_{it} - \Delta X_{it}' \gamma^y = \zeta_{it} + \Delta \varepsilon_{it} \\ g_{it}^{w^*} &= \Delta \ln w_{it}^* - \Delta X_{it}' \gamma^{w^*} = \beta \zeta_{it} + \Delta v_{it} \end{aligned}$$

We use the following moment restrictions to identify the parameters of interest:

$$\begin{aligned} E\left((g_{it}^y)^2\right) &= \sigma_\zeta^2 + 2\sigma_\varepsilon^2 \\ E(g_{it}^y g_{it-1}^y) &= -\sigma_\varepsilon^2 \\ E(g_{it}^y g_{it}^{w^*}) &= \beta \sigma_\zeta^2 \\ E\left((g_{it}^{w^*})^2\right) &= \beta^2 \sigma_\zeta^2 + 2\sigma_v^2 \end{aligned}$$

The model is exactly identified, since we use four moment restrictions to estimate four parameters (the variance of transitory and permanent income shocks, the variance of target wealth and the β coefficient). In estimation we account for the fact that the data (instead of being annual as in the model above) are spaced two years apart (and for 2010, six years apart). Moreover, since estimation involves higher moments that are more likely to be influenced by extreme values, we also estimate the model winsorizing the top and bottom 0.5% and 1% of the distributions of g^y and g^{w^*} .

The results are reported in Table 6. Even using a completely different measure of revisions in permanent income (as opposed to consumption growth) delivers similar results to those found above with consumption as a proxy of permanent income. The estimates of β across the three specifications range between 0.76 and 0.87 and we never reject the null hypothesis that $\beta=1$ at the 1% level.⁷ Hence, our results confirm that reported target wealth adjusts approximately one-for-one in response to permanent income innovations.

⁷ Note that with this alternative empirical strategy we cannot replicate the heterogeneity analysis of Section 5.1, since we do not observe permanent income shocks directly.

As for the other parameters of the model, the variance of permanent income shocks and that of transitory income shocks are precisely estimated and broadly comparable to previous estimates obtained using the SHIW (Jappelli and Pistaferri, 2006) and with evidence from the US (Carroll and Samwick, 1997; Gourinchas and Parker, 2002). Finally, the estimate of unobserved heterogeneity in target wealth is fairly large, reflecting sources of heterogeneity not captured by the set of rather parsimonious demographic variables we include in the regression, measurement error, as well as the fact that target wealth - as any measure of wealth - has a much larger cross-sectional variance than flow variables such as income and consumption.

5.3. Convergence to the target

In Table 7 we regress the change in the wealth to permanent income ratio across two waves on the lagged wealth (over consumption) gap, as in equation (4). The regression includes the same controls as in the previous tables (changes in age, age squared, marital status, and family size). In column (1) the estimated convergence coefficient is -0.59, and statistically different from zero at the 1% level. This estimate implies that, on average, it takes about 3.4 years to close the gap between actual wealth and target wealth (absent additional shocks).⁸

Column (2) distinguishes between the response to positive and negative wealth gaps. The coefficient for positive gaps is 1.11 and precisely estimated, while the coefficient for negative gaps is much smaller (0.08). This implies that individuals who report to be above their preferred target have reduced quite substantially their wealth in the following period (in fact, on average they close the gap within one year), while those who report to be below target have increased wealth only slightly. The third column explores a different source of heterogeneity, distinguishing between relatively large and small wealth gaps (more than twice the level of consumption, or less, respectively). The coefficient for large wealth gaps is similar to the full sample estimate (0.60), while the coefficient for small gaps is 0.30.

⁸ Since the data are biannual, the adjustment is $(0.59/2)$ on an annual basis. Hence, absent other shocks it would take approximately $1/(0.59/2)=3.4$ years to entirely close the gap.

Overall, Table 7 supports the hypothesis that people actually adjust their wealth towards their precautionary target, particularly in case they are substantially above target. The adjustment is modest when people are below target and when they are relatively close to the target to start with. This asymmetry may be driven by consumption commitments (such as housing, educational expenses, health insurance, etc.) that make consumption cuts harder to implement than consumption increases (Chetty and Szeidl, 2007).⁹ A variant of the consumption commitments idea is the consumption thresholds model proposed by by Miranda-Pinto et al. (2020). When adverse shocks hit a household, the household chooses to accumulate debt rather than let consumption fall below a threshold level (for example, rather than move out of a house or slash food consumption, households accumulate debt when faced with a large, unanticipated expenditure).

6. Conclusions

One of the most important implications of the buffer stock model of saving is that consumers choose to maintain a target level of wealth that is proportional to permanent income. This implies that, other things being equal, any revision in permanent income leads to a proportionate revision in target wealth. While many papers provide direct or indirect evidence on the importance of precautionary saving, this key implication of the model has thus far not been directly tested. In this paper, we propose such test. Doing so requires operational measures of target wealth and permanent income in panel data.

We rely on survey data available in the Italian 2002-2016 SHIW on the amount of wealth held for precautionary purposes, which we interpret as target wealth in a buffer stock model. To measure permanent income, we consider two empirical strategies: (a) approximating permanent income with non-durable consumption, and (b) isolating the permanent component from the stochastic structure of the income process. Both empirical strategies reveal that households indeed revise approximately one-for-one their target wealth in response to permanent shocks. This result appears robust to various checks on sample definitions and regression specifications.

⁹ Indeed, Chetty and Szeidl (2007) define a commitment as a consumption category that “is difficult to cut”.

We also explore heterogeneity of the response across the cash-on-hand distribution, for positive and negative income shocks, and for shocks of different size. We find no significant differences for positive and negative permanent income shocks, namely individuals revise target wealth upward approximately one-for-one when they receive a positive shock, and downwards when they receive a negative shock. However, people are more likely to revise their target when shocks are relatively large, which may reflect salience or the presence of adjustment costs, meaning that relatively small shocks are easier to miss than larger ones, or induce only modest utility losses if ignored.

The survey question that we use requires household to provide a “normative” answer (“About how much do you think you and your family *need to have* in savings to meet [...] unexpected events?” – italics added). In the last part of the paper we leverage our longitudinal data to check whether people *actually* adjust their wealth when they are off target, and find a negative, sizable and statistically significant adjustment coefficient, revealing that households align their saving decisions in the direction of the buffer stock model in response to shocks to their permanent income. On average, it takes about 3.4 years to close the gap between actual wealth and target wealth.

This response, however, masks important heterogeneity. Households that have accumulated excess wealth relative to the target decumulate assets at a fast pace and close the gap quickly, while those that have insufficient wealth relative to the target make only modest adjustments towards the target. This asymmetry is not present in the traditional buffer stock model, but could be made consistent with models in which frictions prevent households from adjusting their saving upwards at the same rate as they adjust it downward. For example, consumption commitments of the type studied by Chetty and Szeidl (2007) would make consumption cuts more costly to implement than consumption increases. Exploring the link between buffer stock behaviour and consumption commitments is therefore an important avenue for future research.

References

- Attanasio, Orazio P., James Banks, Costas Meghir, Guglielmo Weber (1999), “Humps and Bumps in Lifetime Consumption,” *Journal of Business Economics and Statistics* 17, 22-35.
- Bertola, Giuseppe, Luigi Guiso, Luigi Pistaferri (2005), “Uncertainty and Consumer Durables Adjustment,” *Review of Economic Studies* 72, 973-1007.
- Cagetti, Marco (2003), “Wealth Accumulation over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics* 21, 339-53.
- Carroll, Christopher D. (1997), “Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics* 112, 1–55.
- Carroll, Christopher D. (2020), “Theoretical Foundations of Buffer Stock Saving,” forthcoming, *Quantitative Economics*.
- Carroll, Christopher D., Miles S. Kimball (2007), “Precautionary Saving and Precautionary Wealth,” *Palgrave Dictionary of Economics and Finance*.
- Carroll, Christopher D., Andrew A. Samwick (1997), “The Nature of Precautionary Wealth,” *Journal of Monetary Economics* 40, 41-71.
- Chetty, Raj, Adam Szeidl (2007), “Consumption Commitments and Risk Preferences,” *The Quarterly Journal of Economics* 122, 831-77.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, Maarten van Rooij (2019), “Asymmetric Consumption Effects of Transitory Income Shocks,” *The Economic Journal* 129, 2322-41.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, and Maarten van Rooij (2020), “Consumption Uncertainty and Precautionary Saving,” *The Review of Economics and Statistics* 102, 148-61.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij L. (2019), “Asymmetric Consumption Effects of Transitory Income Shocks,” *The Economic Journal* 129, 2322-41.
- Carroll, Cristopher D. and Andrew Samwick (1997), “The Nature of Precautionary Saving,” *Journal of Monetary Economics* 40, 41-71.
- Deaton, Angus, (1991), “Saving and Liquidity Constraints,” *Econometrica* 59, 1221-48.

- Dynan, Karen E. (1993), “How Prudent Are Consumers?,” *Journal of Political Economy* 101, 1104-13.
- Fagereng, Andreas, Luigi Guiso, Luigi Pistaferri (2017), “Firm-Related Risk and Precautionary Saving Response”, *American Economic Review Papers and Proceedings* 107, 393-97.
- Friedman, Milton (1957), *A Theory of the Consumption Function*. Princeton University Press: Princeton.
- Fulford, Scott L. (2015), “The Surprisingly Low Importance of Income Uncertainty for Precautionary Saving,” *European Economic Review* 79, 151–71.
- Fuchs-Schündeln, Nicola, Matthias Schündeln (2005), “Precautionary Savings and Self-selection: Evidence from the German Reunification Experiment,” *Quarterly Journal of Economics* 120, 1085-1120.
- Griliches, Zvi, Jerry Hausman (1986), “Errors in Variables in Panel Data”, *Journal of Econometrics* 31, 93-118.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002), “Consumption over the Life Cycle,” *Econometrica* 70, 47-89.
- Guiso, Luigi, Tullio Jappelli and Daniele Terlizzese (1992), “Earnings Uncertainty and Precautionary Saving,” *Journal of Monetary Economics* 30, 307-37.
- Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante (2009), “Quantitative Macroeconomics with Heterogeneous Households”, *Annual Review of Economics*, 1, 319-354.
- Hurst, Erik, Annamaria Lusardi, Arthur Kennickell, Francisco Torralba (2010), “The Importance of Business Owners in Assessing the Size of Precautionary Savings,” *Review of Economics and Statistics* 92), 61–69.
- Kaplan Greg, Giovanni L. Violante (2014), “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica* 82, 1199–1239.
- Krueger, Dirk, Kurt Mitman, and Fabrizio Perri (2017), “Macroeconomics and Household Heterogeneity”, *Handbook of Macroeconomics*, Vol. 2A, John B. Taylor, Harald Uhlig (Eds), North Holland.
- Jappelli, Tullio, Mario Padula, Luigi Pistaferri (2008), “A Direct Test of the Buffer-Stock Model of Saving,” *Journal of the European Economic Association* 6,1186–1210.
- Jappelli, Tullio, Luigi Pistaferri (2006), “Intertemporal Choice and Consumption Mobility,” *Journal of the European Economic Association* 4, 75–115.

- Jappelli, Tullio, Luigi Pistaferri (2010), “The Consumption Response to Income Changes,” *Annual Review of Economics* 2, 479-506.
- Jappelli, Tullio, Luigi Pistaferri (2014), “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics* 6, 107-36.
- Jappelli, Tullio, Luigi Pistaferri (2017), *The Economics of Consumption*. Oxford University Press: Oxford, UK.
- Kennickell, Arthur, Annamaria Lusardi (2004), “Disentangling the Importance of the Precautionary Saving Motive,” NBER WP 10888.
- Lusardi, Annamaria (1998), “On the Importance of the Precautionary Saving Motive,” *American Economic Review Papers and Proceedings* 88, 449–53.
- Miranda-Pinto, Jorge, Daniel Murphy, Kieran James Walsh, Eric R. Young (2020), “A Model of Expenditure Shocks,” Federal Reserve Bank of Cleveland, Working Paper 2020/04.
- Parker, Jonathan A. (2017), “Why Don't Households Smooth Consumption? Evidence from a \$25 Million Experiment,” *American Economic Journal: Macroeconomics* 9, 153-83.
- Parker, Jonathan, Nicholas Souleles (2017), “Reported Preference vs. Revealed Preference: Evidence from the Propensity to Spend Tax Rebates,” *American Economic Review: Insights*, forthcoming.
- Sahm, Claudia R., Matthew D. Shapiro, Joel Slemrod (2015), “Balance-Sheet Households and Fiscal Stimulus: Lessons from the Payroll Tax Cut and Its Expiration,” NBER Working Paper No. 21220.
- Shapiro, Matthew D., Joel Slemrod (1995), “Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding,” *The American Economic Review* 85, 274-83.
- Shapiro, Matthew D., Joel Slemrod (2009), “Did the 2008 Tax Rebates Stimulate Spending?” *American Economic Review Papers and Proceedings* 99, 374-79.
- Skinner, Jonathan (1988), “Risky Income, Life-Cycle Consumption and Precautionary Saving,” *Journal of Monetary Economics* 22, 237-55.
- Zeldes, Stephen P. (1989), “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy* 97, 305-46.

Appendix – Target Wealth Questions in the 2002-2016 SHIW

2002, 2004, 2010

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?

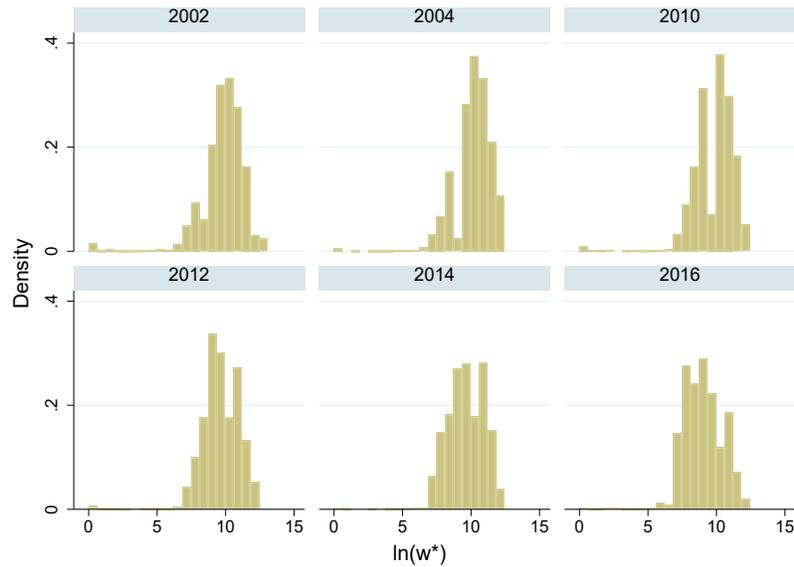
2012

People save in various ways (depositing money in a bank, buying financial assets, property, or other assets) and for different reasons. One 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?

2014 and 2016

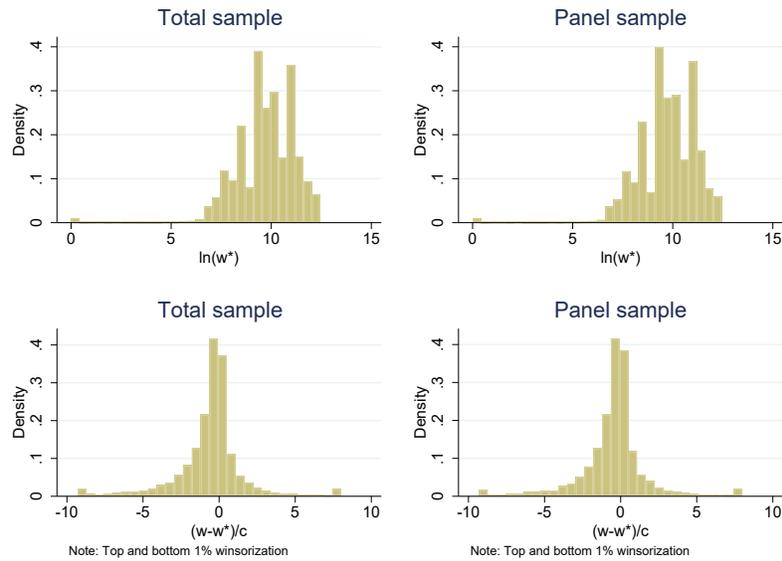
About how much do you think you and your family need to have in savings to meet unexpected events, such as health problems or other emergencies?

Figure 1: The distribution of target wealth



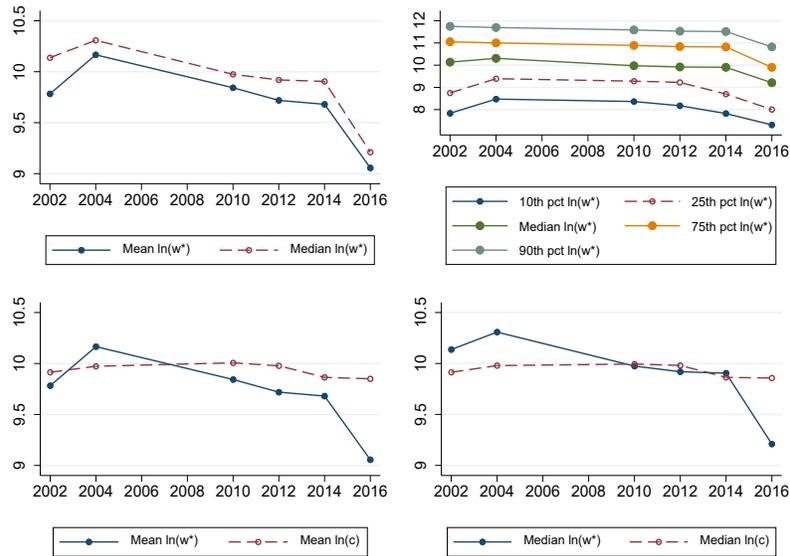
Note. The figure plots the cross-sectional distribution of the log of target wealth from 2002 to 2016.

Figure 2: Target wealth and the wealth gap



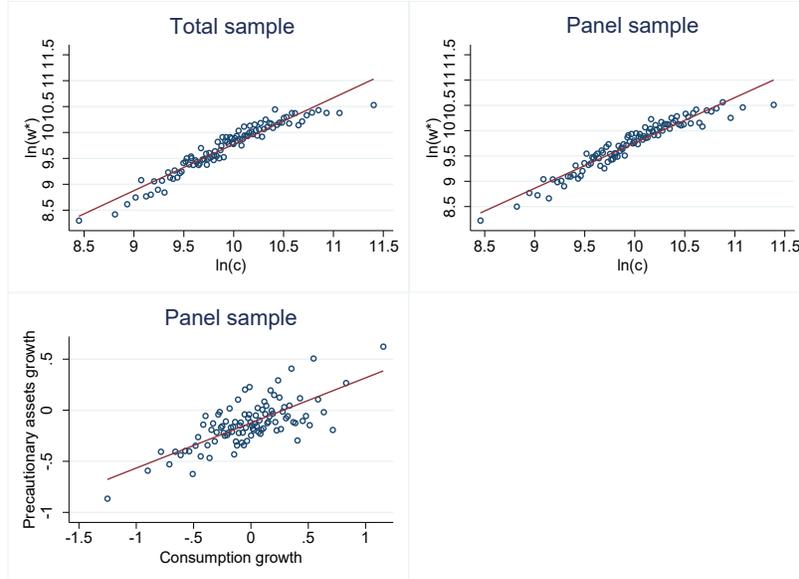
Note. The top two panels plot the cross-sectional distribution of the log of target wealth (w^*) in the total sample (2002-2016) and in the panel sample. The bottom two panels plot the cross-sectional distributions of the wealth gap relative to consumption, $(w-w^*)/c$, in the total sample (2002-2016) and in the panel sample.

Figure 3: Mean, medians and percentiles of target wealth and consumption



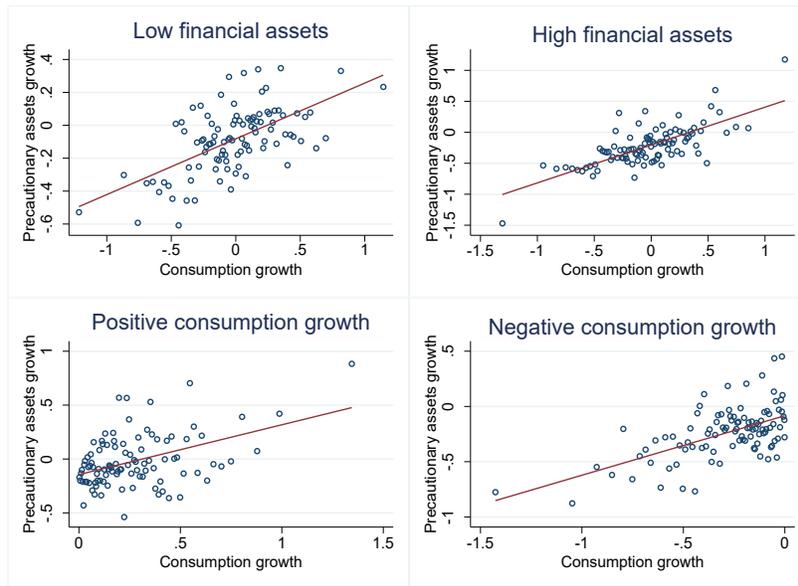
Note. The top two panels plot the average, median and percentiles of log target wealth from 2002 to 2016 in the total sample. The bottom two panels plot the means and medians of log target wealth and log consumption from 2002 to 2016 in the total sample.

Figure 4: Mean, medians and percentiles of target wealth and consumption



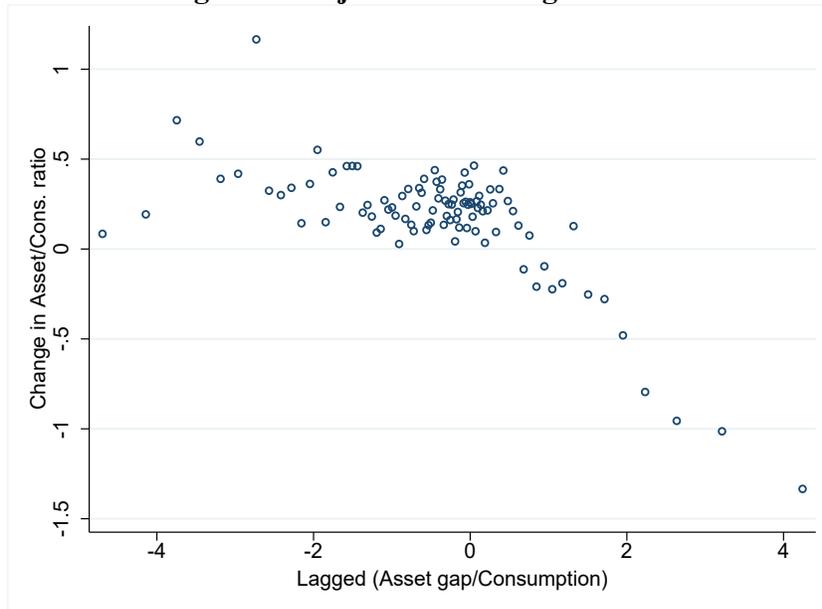
Note. The top two panels plot log target wealth against the percentiles of log consumption in the total sample and in the panel sample. The bottom panels plots the growth of target wealth against the percentiles of consumption growth in the panel sample.

Figure 5: Target wealth by financial assets and consumption growth



Note. The top two panels plot log target wealth against the consumption growth separately for observations with financial assets less or more than two-months' income. The bottom two panels plot the growth of target wealth separately for observations with positive and negative consumption growth.

Figure 6: Adjustment to target wealth



Note. The figure plots the change in cash-on-hand against the wealth gap. Both variables as expressed as a ratio to consumption.

Table 1: Descriptive Statistics

	Whole sample			Panel sample		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
Target wealth	35,667.18	20,030.15	44,009.04	34,932.81	20,030.15	42,746.41
Non-durable consumption	23,443.91	20,469.18	13,335.86	24,022.64	21,157.76	13,477.51
Disposable income	30,201.36	24,960.22	22,597.57	31,340.37	26,024.92	22,276.74
Log target wealth	9.71	9.91	1.5	9.72	9.91	1.47
Log consumption	9.93	9.93	0.51	9.96	9.96	0.51
Log disp. income	10.1	10.13	0.7	10.14	10.17	0.69
Target wealth / consumption	1.65	0.93	2.13	1.57	0.9	2.01
Financial assets	27,538.9	6,980.99	92,653.54	29,711.66	7,511.3	98,124.08
Wealth gap	-0.51	-0.30	4.07	-0.38	-0.26	4.20
Years of education	9.2	8	4.58	9.47	8	4.57
Age	59.14	60	15.85	60.01	60	15.09
Male	0.56	1	0.5	0.56	1	0.5
Family size	2.45	2	1.27	2.48	2	1.27
Married	0.6	1	0.49	0.62	1	0.49
Observations		46,569			25,738	

Table 2: OLS and First-Difference Estimates

	Total sample OLS (1)	Panel sample OLS (2)	Panel sample FD (3)
Log(consumption)	0.692 (0.019)***	0.702 (0.026)***	0.423 (0.043)***
Age	0.019 (0.003)***	0.018 (0.005)***	-0.041 (0.043)
Age sq./100	-0.017 (0.003)***	-0.016 (0.004)***	-0.027 (0.013)**
Male	0.036 (0.015)**	0.045 (0.022)**	
Family size	-0.033 (0.008)***	-0.039 (0.012)***	-0.006 (0.026)
Married	0.066 (0.019)***	0.066 (0.027)**	0.028 (0.075)
Education	0.010 (0.002)***	0.013 (0.003)***	
Centre	0.108 (0.018)***	0.160 (0.026)***	
South	-0.357 (0.017)***	-0.320 (0.025)***	
<i>N</i>	46,569	25,738	16,883
P-value test $\beta=1$	0.000	0.000	0.000

Notes. In columns (1) and (2) the dependent variable is the log of target wealth. In column (3) all variables are in first differences. All regressions also include year dummies. ***, **, and * indicate statistical significance at the 1%, 5% and 10% confidence level, respectively. Standard errors are clustered at household level.

Table 3: First Difference Instrumental Variables Estimates

	Baseline specification (1)	Pos. and neg. cons. growth (2)	Large and small cons. growth (3)
Log(Cons.)	0.979 (0.147)***		
Log(Cons.)* 1 {Pos. cons. growth}		0.849 (0.316)***	
Log(Cons.)* 1 {Neg. cons. growth}		1.093 (0.262)***	
Log(Cons.)* 1 {Large cons. growth}			1.181 (0.214)***
Log(Cons.)* 1 {Small cons. growth}			0.232 (0.420)
Age	-0.064 (0.045)	-0.061 (0.045)	-0.051 (0.044)
Age sq./100	-0.022 (0.013)*	-0.023 (0.013)*	-0.026 (0.013)**
Family size	-0.078 (0.032)**	-0.080 (0.032)**	-0.069 (0.032)**
Married	0.006 (0.077)	0.002 (0.078)	0.007 (0.077)
<i>N</i>	16,554	16,554	16,554
P-value test $\beta=1$	0.889		
P-value test $\beta_p = \beta_n$		0.623	
P-value test $\beta_l = \beta_s$			0.087

Note. All variables are in first differences. The instrument for the change in log consumption is the change in log disposable income (net of financial income). Large (small) consumption growth refers to observations for which the growth rate of annual income exceeds (does not exceed) 15% in absolute value. All regressions also include year dummies. ***, **, and * indicate statistical significance at the 1%, 5% and 10% confidence level, respectively. Standard errors are clustered at household level.

Table 4: First Difference Instrumental Variables Estimates, by Financial Wealth

	Total sample (1)	Low financial wealth (2)	High financial wealth (3)
Log(Cons.)*1st quintile FW	0.955 (0.249)***		
Log(Cons.)*2nd quintile FW	0.788 (0.239)***		
Log(Cons.)*3rd quintile FW	1.186 (0.395)***		
Log(Cons.)*4th quintile FW	1.046 (0.241)***		
Log(Cons.)*5th quintile FW	1.014 (0.289)***		
Log(Cons.)		0.942 (0.221)***	1.021 (0.190)***
Age	-0.067 (0.045)	-0.187 (0.100)*	-0.048 (0.049)
Age sq./100	-0.021 (0.013)	-0.059 (0.024)**	-0.010 (0.017)
Family size	-0.081 (0.032)**	-0.082 (0.045)*	-0.078 (0.043)*
Married	0.007 (0.077)	0.035 (0.130)	-0.018 (0.094)
<i>N</i>	16,554	5,579	10,975
P-value test $\beta=1$	0.902	0.795	0.913

Note. All variables are in first differences. The instrument for the change in log consumption is the change in log disposable income (net of financial income). Low (high) financial assets are observations with financial assets less (more) than two-months' income. All regressions also include year dummies. ***, **, and * indicate statistical significance at the 1%, 5% and 10% confidence level, respectively. Standard errors are clustered at household level. In column (1) the test $\beta=1$ is a test that all interactions of log consumption with financial wealth are equal to 1.

Table 5: Robustness Checks

	Panel, FD IV 3 waves+ (1)	Panel, FD IV 5 waves+ (2)	Panel, FD IV 2002-2012 (3)	Panel, FD IV log(c_{t-1}) (4)
Log(Cons.)	0.957 (0.163)***	1.115 (0.270)***	1.020 (0.201)***	0.812 (0.174)***
Age	-0.072 (0.054)	-0.038 (0.142)	-0.108 (0.056)*	-0.046 (0.055)
Age sq./100	-0.012 (0.013)	0.032 (0.023)	-0.011 (0.016)	-0.019 (0.020)
Family size	-0.082 (0.035)**	-0.152 (0.061)**	-0.072 (0.044)*	-0.082 (0.043)*
Married	0.019 (0.087)	0.014 (0.134)	-0.036 (0.103)	-0.026 (0.101)
Log(Cons.), lagged				0.126 (0.167)
N	12,056	4,258	9,207	7,853
P-value test $\beta=1$	0.794	0.669	0.919	
P-value test $\beta+\beta_{-1}=1$				0.806

Note. All variables are in first differences. The instrument for the change in log consumption is the change in log disposable income (net of financial income). Column (1) includes only households interviewed in at least 3 waves, column (2) only households interviewed in at least 5 waves. Column (3) includes only households interviewed in 2002-12. Column (4) includes lagged consumption growth among the regressors. All regressions also include year dummies. ***, **, and * indicate statistical significance at the 1%, 5% and 10% confidence level, respectively. Standard errors are clustered at household level. In column (4) the test $\beta+\beta_{-1}=1$ is a test that the sum of the coefficients on log consumption and log consumption lagged is equal to 1.

Table 6: GMM Results

	Untrimmed (1)	0.5% Trimmed (2)	1% Trimmed (3)
β	0.7606 (0.2146)***	0.8654 (0.2058)***	0.8411 (0.1882)***
σ_{ζ}^2	0.0151 (0.0028)***	0.0137 (0.0018)***	0.0130 (0.0014)***
σ_{ε}^2	0.0657 (0.0054)***	0.0541 (0.0033)***	0.0467 (0.0025)***
σ_v^2	1.1763 (0.0520)***	1.0010 (0.0286)***	0.9034 (0.0189)***
N	7,853	7,853	7,853
P-value test $\beta=1$	0.265	0.513	0.399

Notes: In this table we report estimates of the parameters determining the relationship between the growth of target wealth target and the stochastic income process. Column (1) uses the whole sample; column (2) winsorizes observations in the top and bottom 0.5%; column (3) winsorizes observations in the top and bottom 1%. ***, **, and * indicate statistical significance at the 1%, 5% and 10% confidence level, respectively. Standard errors are clustered at household level.

Table 7: Adjustment to Target Wealth

	Gap adjustment	Gap adjustment, below/above target	Gap adjustment, large/small
Lagged(Wealth gap/Consumption)	-0.587 (0.104)***		
Age	-0.162 (0.109)	0.062 (0.102)	-0.167 (0.110)
Age sq./100	0.083 (0.034)**	0.081 (0.031)***	0.076 (0.033)**
Family size	-0.021 (0.061)	-0.011 (0.055)	-0.012 (0.062)
Married	0.240 (0.195)	0.199 (0.175)	0.242 (0.195)
Lagged(Wealth gap/Consumption) > 0		-1.113 (0.189)***	
Lagged(Wealth gap/Consumption) ≤ 0		0.080 (0.022)***	
Large lagged(Wealth gap/Consumption)			-0.602 (0.108)***
Small lagged(Wealth gap/Consumption)			-0.298 (0.033)***
<i>N</i>	16,883	16,883	16,883

Note. Demographic variables are in first differences. All regressions also include year dummies. ***, **, and * indicate statistical significance at the 1%, 5% and 10% confidence level, respectively. Standard errors are clustered at household level.