SUPERIOR INFORMATION, INCOME SHOCKS, AND THE PERMANENT INCOME HYPOTHESIS

Luigi Pistaferri*

Abstract—According to the permanent income hypothesis with quadratic preferences, households save for a rainy day the transitory component of income innovations and consume entirely the permanent one. The model also rules out precautionary saving. Typically, income shock components are not separately observable, and information on the conditional variance of income is hard to come by. We show how to combine income realizations with subjective expectations to identify separately the transitory and the permanent shock to income and to obtain a measure of idiosyncratic uncertainty, thus providing a powerful test of the theory in short panels. The empirical analysis is performed on a sample of Italian households drawn from the 1989–1991 Survey of Household Income and Wealth.

I. Introduction

ACCORDING to the textbook version of the permanent income hypothesis (PIH), household consumption responds on a one-to-one basis to permanent income shocks, but is nearly insensitive to transitory income shocks. Equivalently, households save the transitory component of the income innovation and consume entirely the permanent one. Yet, testing for the separate effect of income shocks on consumption or saving is difficult. The main problem is that, although the agent may be able to discriminate between a transitory and a permanent shock, the econometrician is not. As a result, econometric identification of separate income shock components is difficult in the extreme.[1]

In this paper, we show how to combine subjective income expectations with income realizations to identify separately the transitory and the permanent income shock. This has two main advantages. First, it allows the examination of the cross-sectional distribution of separate income shocks; second, it provides a powerful test of the theory in short panels.[2] In particular, we test whether households save for a rainy day using data available for a panel of Italian households drawn from the 1989–1991 Bank of Italy Survey of Household Income and Wealth (SHIW). The survey we use is representative of the Italian population, contains a measure of total nondurable consumption, and has good quality income data.

To assess the validity of the PIH, we regress savings on income shocks. If the theory is true, only transitory shocks should explain saving. However, households save for a rainy day only if they display quadratic preferences. If instead preferences are mispecified, precautionary saving can represent a likely source of failure of the theory. In practice, the estimates of the effect of income shocks on saving will be inconsistent if the omitted higher moments of the distribution of income shocks are correlated with the realization of the shocks. But this also suggests that one might test for the departure from the certainty equivalence assumption augmenting the saving equation with the conditional variance term: if the PIH with quadratic preference is true, the conditional variance should not explain saving.

The contribution of the paper is that once income shocks become separately identifiable the consistency of empirical estimates does not rely on a long time series of observations for each individual, a problem that plagues most of the empirical studies (Chamberlain, 1984). Unlike previous studies, the consistency of the saving equation estimates relies only on a large cross-sectional dimension, not a large time series; in other words, the availability of income expectations implies that Chamberlain's critique does not apply in this context. This allows us to estimate structural parameters (the marginal propensities to save out of income shocks) that are robust to short panel inconsistency and useful for policy analysis.

The rest of the paper is organized as follows. Section II presents a formal decomposition of the income innovation into a permanent and a transitory component and shows how subjective expectations of income can help to identify separately the two components. In section III, we present the data used in this study and examine the validity of the assumptions concerning the stochastic structure of the income process; in section IV, we examine the empirical distribution of the two shocks. Section V presents the empirical results and shows that the evidence supports an extended version of the PIH. In particular, savings do respond to transitory income shocks, but also to permanent income shocks and higher moments of the distribution of earnings. The positive impact of permanent shocks on savings can be explained by measurement error or be the result of the interaction between prudence and impatience in a buffer stock model of saving. The positive impact of the variance term on savings reflects precautionary behavior. Section VI concludes.

Received for publication October 22, 1998. Revision accepted for publication August 23, 2000.

* Stanford University, CEPR, and SIEPR.

This is a revised version of chapter 3 of my doctoral dissertation at University College London. I am grateful to the two anonymous referees, the editor (John Campbell), my UCL supervisors (Orazio Attanasio and Costas Meghir), Joe Altonji, Victoria Chick, Jonathan Parker, Danny Quah, seminar participants in Tilburg and Pompeu Fabra, and especially Tullio Jappelli for useful comments. All errors are mine. Financial support from the Italian Ministry of Universities and Scientific Research, the James Lanner Memorial Fund, and the Taube Faculty Research Fund at the Stanford Institute for Economic Policy Research is gratefully acknowledged.

[1] Attempts in the direction of estimating the separate effect of transitory and permanent income shocks on consumption include Hall and Mishkin (1982) on PSID data, and Flavin (1981) on aggregate U.S. data.

[2] Campbell (1987) shows that it is still feasible to test whether households save for a rainy day by replacing the information set available to the agent with that available to the econometrician. Although consistent under some regularity conditions (see the discussion by Deaton (1992) and Flavin (1993))), estimates based on the econometrician's information set are inefficient.

The Review of Economics and Statistics, August 2001, 83(3): 465–476

II. The Estimation Strategy

In this section, we show how to decompose income shocks into a transitory and a permanent component, and how to determine their separate effect on saving using the saving equation firstly derived by Campbell (1987). We also discuss identification issues and consider some extensions.

A. *The Income Process*

Suppose that current disposable income admits the following decomposition (as in Muth (1960) and Blundell and Preston (1998)):

$$y_{it} = \theta' Z_{it} + \pi'\Gamma_i + p_{it} + \epsilon_{it}, \tag{1}$$

where $\theta' Z_{it}$ is a deterministic time-varying component (which may include age and other characteristics that change over the life cycle);

$\pi'\Gamma_i$ is a deterministic time-invariant component (which may include education, gender, region of residence, and perhaps fixed individual effects);
p_{it} is the permanent component; and
ϵ_{it} is a transitory shock. For the sake of simplicity, assume that the latter is i.i.d. with constant variance σ^2_ϵ.

The permanent component of income follows a random walk process of the form

$$p_{it} = p_{it-1} + \zeta_{it} \tag{2}$$

where ζ_{it} is the permanent shock; this is assumed to be i.i.d. with constant variance σ^2_ζ. We also assume that the transitory and the permanent shock are orthogonal to each other at all leads and lags.[3] Combining equation (1) and (2), one obtains

$$\Delta y_{it} = \theta' \Delta Z_{it} + \zeta_{it} + \Delta\epsilon_{it}, \tag{3}$$

where Δ is the first-difference operator. Notice that first-differencing has eliminated all time-invariant effects. To control for predictable life cycle effects, in the empirical analysis we will assume that the deterministic component depends on a quadratic polynomial in age; that is, $\theta' Z_{it} = \theta_0 + \theta_1 age_{it} + \theta_2 age^2_{it}$. Equation (3) can then be rewritten as

$$\Delta y_{it} = \gamma_0 + \gamma_1 age_{it} + \zeta_{it} + \Delta\epsilon_{it} \tag{4}$$

with $\gamma_0 = (\theta_1 - \theta_2)$ and $\gamma_1 = 2\theta_2$.

B. *The Identification of Income Shocks*

In this section we will show that, given the income process (equation (1) and (2)), the transitory and the permanent shock to income can be identified if one observes, for at least two consecutive time periods, both the conditional expectation and the realization of the variable of interest (disposable family income, say). This is of course unthinkable in the presence of realization data only.

Let us assume that individual agents form rational expectations. We define $E(x_{it}|\Omega_{it-1})$, the subjective expectation of x_{it} given the individual's information set at time $t - 1$. It is worth pointing out that Ω_{it-1} is the set of information possessed at individual level; the econometrician's information set is generally less rich. Flavin (1993) describes the "discrepancy between the econometrician's and the agent's information sets" as "the omitted information issue (from the perspective of the econometrician) or the superior information issue (from the perspective of the agent)" (p. 652). Indeed, some information is available to the individual but not to the econometrician; alternatively, the econometrician holds information that can be irrelevant to individual choices (such as some kind of aggregate information). Equation (3) and the assumption of rational expectations imply that the transitory shock at time t can be exactly identified by

$$\epsilon_{it} = -E(\Delta y_{it+1}|\Omega_{it}) + (\gamma_0 + \gamma_1 age_{it+1}). \tag{5}$$

We defer discussion of identification and estimation of the parameters γ_0 and γ_1 and proceed as if they were known. Using equation (3) and (5), the permanent shock at time t is exactly identified by the expression

$$\begin{aligned}\zeta_{it} = {} & [\Delta y_{it} - E(\Delta y_{it}|\Omega_{it-1})] + E(\Delta y_{it+1}|\Omega_{it}) \\ & - [(\gamma_0 + \gamma_1) + \gamma_1 age_{it}],\end{aligned} \tag{6}$$

that is, the income innovation at time t adjusted by a factor that takes into account the arrival of new information concerning the change in income between t and $t + 1$ and the contribution of life cycle predictable effects.[4] It is worth noting that time-invariant income attributes do not affect the identification strategy, because income shocks are identified by income changes.

As we will show in the data section, the 1989 and 1991 SHIW data provide a unique opportunity to implement equation (5) and (6), as they elicit subjective expectations of future income and the corresponding realizations.

C. *The Effect of Transitory and Permanent Income Shocks on Saving*

As shown by Campbell (1987), under stringent assumptions concerning preferences and technology (in particular, quadratic preferences, intertemporal separability, infinite horizon, a rate of intertemporal discount set equal to the real

[3] The assumption of constant variance for both the transitory and the permanent income shock and that of independence can all be removed without altering the identification strategy.

[4] It turns out that, if transitory shocks are serially correlated, subjective expectations no longer identify income shocks. This is our main identifying assumption. We comment on the validity of such an assumption in subsection IIIB.

interest rate, and perfect credit markets), one obtains the following saving equation:

$$s_{it} = -\sum_{\tau=1}^{\infty} \frac{1}{(1+r)^{\tau}} E(\Delta y_{it+\tau}|\Omega_{it}). \tag{7}$$

Equation (7) implies that savings mirror the present discounted value of expected income declines. Using the income process (4) above, it is easy to prove that equation (7) simplifies to

$$s_{it} = \alpha_0 + \alpha_1 age_{it} + \frac{1}{1+r}\epsilon_{it}, \tag{8}$$

where $\alpha_0 = -\frac{1}{r}\left[\gamma_0 + \gamma_1\left(\frac{1+r}{r}\right)\right]$ and $\alpha_1 = -\frac{\gamma_1}{r}$. The implications one derives from equation (8) are clear: permanent shocks do not matter (because, under the conditions recalled above, the optimal rule is to consume them all), and only transitory income shocks explain saving. Note further that time-invariant income attributes do not affect savings.

If one relaxes the assumption of quadratic preferences, there is no longer a closed-form solution for consumption or savings. Moreover, the error term of the intractable saving equation will contain higher moments of the distribution of income shocks that are likely to be correlated with the realization of the shocks; if this is the case, estimates will be inconsistent. But this also suggests that one can test for the validity of the quadratic preferences assumption by augmenting the saving equation by higher moments of the distribution of income; under the null hypothesis of the permanent income hypothesis with quadratic preferences, higher moments should not explain saving. We will implement such a test by including the subjective variance of future income var $(y_{it+1}|\Omega_{it})$ in the saving equation.

If transitory and permanent income shocks are separately identifiable according to equation (5) and (6), one can thus implement the regression

$$\begin{aligned} s_{it} = {} & \psi' D_{it} - \beta_1 E(\Delta y_{it+1}|\Omega_{it}) \\ & + \beta_2[\Delta y_{it} - E(\Delta y_{it}|\Omega_{it-1}) + E(\Delta y_{it+1}|\Omega_{it})] \\ & + \beta_3 \operatorname{var}(y_{it+1}|\Omega_{it}) + e_{it}, \end{aligned} \tag{9}$$

where D_{it} contains a constant term, age and other demographics to control for preference heterogeneity, and the error term e_{it} reflects measurement error in saving. The implications of the permanent income hypothesis we will test in the empirical analysis are $\beta_1 \cong 1$, $\beta_2 = 0$, and $\beta_3 = 0$.

The identification strategy is robust to an income process of the form $y_{it} = \theta' Z_{it} + \pi'\Gamma_i + \mu_t + p_{it} + \epsilon_{it}$, where μ_t captures stochastic business cycle effects (that is, $E(\mu_{t+\tau}|\Omega_{it}) = 0$ for all $\tau > 0$). In this case, it is easy to show that, omitting for simplicity age effects, the saving equation is $s_{it} = \frac{1}{1+r}(\mu_t + \epsilon_{it})$. Equation (5) is used to identify the composite error term $(\mu_t + \epsilon_{it})$, and the permanent shock is still identified by equation (6). A regression of s_{it} on the composite error term $(\mu_t + \epsilon_{it})$ delivers an estimate of the marginal propensity to save out of a transitory shock to income, one of the structural parameters of interest.

To our knowledge, tests of the "saving for a rainy day" equation based on microeconomic data have been performed only by Deaton (1992) and Alessie and Lusardi (1997). In both cases, the authors had available short panels (two years in Deaton and three years in Alessie and Lusardi). Thus, their estimates are likely to suffer from the problem of inconsistency in short panels remarked by Chamberlain (1984), even if aggregate shocks are controlled for.

D. *Consistency*

In the absence of information on income shocks, the disturbance term of the saving equation (8) is a forecast error, the difference between realized and expected saving $s_{it} - E(s_{it}|\Omega_{it-1})$. According to the permanent income hypothesis with rational expectations, the conditional expectation of a forecast error must be zero. The empirical analog of this expectation is an average taken over long periods of time, not across a large number of households. In fact, as pointed out by Chamberlain (1984), there is no guarantee that the cross-sectional average of forecast errors (or the cross-sectional correlation between forecast errors and lagged instruments) will converge to zero as the dimension of the cross section gets large. The problem is sometimes handled by including time dummies in the Euler equation. This approach is restrictive, because it rules out that aggregate shocks are unevenly distributed in the population. For this reason, tests of the PIH performed on short panels are in fact implicitly assuming that the stochastic structure of the forecast error has a known form (so that the distance between the true forecast error and its empirical analog can be suitably adjusted). Rejection of the null need not be interpreted as the failure of the theory, but could also be attributed to misspecification of the stochastic structure of the forecast error.

Unlike previous studies, the consistency of our saving equation estimates does not rely on a large time-series dimension, but on a large cross-sectional dimension. This is essentially because we do observe the innovation in savings; for example, we can condition on it. Indeed, under the null hypothesis of the PIH, the residual term of equation (9) is assumed to reflect only (additive) measurement error in saving and perhaps measurement error in the independent variables. Hence, the consistency of our estimates rests only on the weak assumption that measurement errors in saving are not correlated across individuals in the cross section. These are, of course, weaker conditions than the ones usually required in tests of the Euler equation or of the permanent income hypothesis. Indeed, the availability of income expectations implies that Chamberlain's critique (1984) does not apply in our context.[5]

III. The Data

A. The SHIW

We estimate the saving equation using the 1989–1991 panel section of the Bank of Italy Survey of Household Income and Wealth (SHIW). This data set contains measures of consumption, income, and demographic characteristics of households. The SHIW surveys a representative sample of the Italian resident population. From 1987 through 1995, the survey was conducted every other year and covered approximately 8,000 households, defined as groups of individuals related by blood, marriage, or adoption, and sharing the same dwelling. Starting in 1989, each SHIW has reinterviewed some households from the previous surveys. The panel component has increased over time: 15% of the sample was reinterviewed in 1989, 27% in 1991, 43% in 1993, and 45% in 1995. The response rate (ratio of responses to contacted households net of ineligible units) was 64% in 1987, 38% in 1989, 33% in 1991, 58% in 1993, and 57% in 1995. For the panel section of the SHIW, the response rate has been steadily increasing over time: 23% in 1989, 53% in 1991, 64% in 1993, and 75% in 1995. Details on sampling, response rates, processing of results, and comparison of survey data with macroeconomic data are provided by Brandolini and Cannari (1994). The 1989–1991 panel section of the SHIW includes 2,187 households. After excluding those with missing reports on subjective expectations in 1989, 1991, or both, and other variables used in the empirical analysis, we are left with a sample of 1,125 households. Finally, to avoid having our estimates be contaminated by influential outliers, we trim the sample at the bottom and top percentile of the distribution of saving; this implies a loss of 23 households and leaves us with a final sample of 1,102 households.[6]

For the purpose of this paper, the most important feature of the SHIW is that it collected subjective information on future income in both 1989 and 1991. The 1989 and 1991 SHIW have been used by Guiso, Jappelli, and Terlizzese (1992) and Jappelli and Pistaferri (2000) to test for precautionary saving and excess sensitivity of consumption to predicted income growth, respectively.[7]

Several surveys contain subjective income expectations, but vary considerably as to the way expectations are elicited. In the case of the SHIW, in 1989 and 1991 each labor income and pension recipient interviewed was asked to attribute probability weights, summing to 100, to given intervals of inflation and nominal income increases one year ahead. Appendix A details the wording of the survey questions and how these are used to construct the subjective mean and variance of future earnings.

A problem with the SHIW data is that they are not available for consecutive years, but only at two-year intervals; moreover, subjective expectations stretch over a single calendar year. The interviews take place between March and September, although income, consumption, and wealth data refer to the previous calendar year.[8] We thus need to assume that people do not update their information set between the end of 1989 (1991) and the date of the interview, or that their updating does not affect subjective expectations of income. This can be a strong assumption if people receive important news about the evolution of their future income between the end of 1989 (1991) and the date of the interview. This assumption has been made implicitly in all the papers quoted in this section.

More precisely, the SHIW data provide information on age, income realizations (y_{it} and y_{it-2}), and the subjective expectations of income changes ($E(\Delta y_{it-1}|\Omega_{it-2})$ and $E(\Delta y_{it+1}|\Omega_{it})$), with $t = 1991$). It can be shown that, modifying equation (5) and (6) to incorporate the SHIW data features, the following expressions

$$-E(\Delta y_{it+1}|\Omega_{it}) + \hat{\gamma}_0 + \hat{\gamma}_1 age_{it+1} = \epsilon_{it}$$

$$-E(\Delta y_{it-1}|\Omega_{it-2}) + \hat{\gamma}_0 + \hat{\gamma}_1 age_{it-1} = \epsilon_{it-2}$$

identify the transitory shock at times t and $t - 2$, respectively, while the expression

$$\begin{aligned}(y_{it} - y_{it-2}) &- E(\Delta y_{it-1}|\Omega_{it-2}) + E(\Delta y_{it+1}|\Omega_{it}) \\ &- [2\hat{\gamma}_0 + \hat{\gamma}_1 + 2\hat{\gamma}_1 age_{it}] \\ &= (\zeta_{it} + \zeta_{it-1})\end{aligned} \tag{10}$$

identifies the sum of the permanent shocks at time t and $t - 1$. The estimates $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are obtained regressing the

[5] Pistaferri (1998) tests the relevance of Chamberlain's critique noting that, under the null of the permanent income hypothesis, the innovation in savings is observed. This allows the calculation of the empirical covariance between the latter and lagged instruments.

[6] A potential problem with the SHIW panel is that of endogenous attrition. To address this issue, we divide the 1989 SHIW households into three groups, those that are both in the 1989–1991 panel and in our sample (1,102 households), those that are in the 1989–1991 panel but not in our sample (1,085 households), and those that are not in the 1989–1991 panel (6,079 households). If attrition were endogenous, one should find saving rates to be substantially different for the three groups. In fact, there is little difference in saving behavior between groups: the median saving rates for the three groups are 24.65%, 22.36%, and 24.71%, respectively. All differences are not statistically significant.

[7] Jappelli and Pistaferri (2000) test whether consumption is excessively sensitive to anticipated changes in income, whereas this paper tests whether consumption and savings react to unanticipated changes. The two tests are obviously related. On the other hand, this paper's contribution to the literature is that we obtain estimates of truly structural parameters (the marginal propensities to save or consume out of income shocks) that are robust to short panel inconsistency. Structural parameters are of course of paramount importance for policy analysis.

[8] The reason for surveying in May is that previous experience has shown that people report income more accurately when filing the income tax forms, which must be returned by May 31.

subjective expectation on a constant term and age, respectively. Because under the null of the permanent income hypothesis savings depend only upon transitory innovations, that is all we need to implement the estimation of equation (9). The strategy we use to test for the null hypothesis of no effect of permanent shocks on savings is described below.

Although each labor income recipient is asked to answer the survey question, we rely only on the information provided by the head of the household or, if the latter are lacking, on those provided by the spouse. The reason is that, in most cases, information on income recipients other than the head or spouse is lacking. Thus, we regress saving on the head's earnings shocks, rather than on the shocks referring to disposable family income. For single-earner households, the two measures of income will roughly coincide.

Saving is defined as the difference between family disposable income (including asset income) and nondurable consumption. The age variable refers to the head of the household (or the spouse, if this is missing). Head's (or spouse's) earnings are net of taxes and social security contributions.

B. *The Stochastic Structure of the Income Process*

In subsection IIA, a specific structure is placed on the income process; the identification of the transitory and the permanent shocks is thus heavily dependent on whether such structure is correct. A possible argument is that the simple partition we have considered may fail to identify shocks of different nature if, say, the transitory shock is itself serially correlated.

To assess whether the data on disposable family income are indeed consistent with these assumptions, we consider the 1987–1995 SHIW rotating panel (15,065 observations). Following equation (1), we first remove predictable effects by regressing income on a quadratic polynomial in age, family size, and dummies for education, region of residence, and self-employment. We save the residuals (let's call them u_{it}) and carefully examine their covariance properties. Autocovariances and autocorrelation coefficients are estimated using equally weighted minimum distance methods, as suggested by Altonji and Segal (1996).[9]

If the change in income is well described by equation (3), the unanticipated component is an $MA(1)$ process. The presence of i.i.d. measurement error in earnings adds a further error component, but does not change the $MA(1)$ structure of the process. Thus, in general, $\Delta u_{it} = v_{it} - \theta v_{it-1}$, where the MA coefficient θ is a function of the variance of the transitory shock, the variance of the permanent shock, and the variance of the measurement error. If this is the case, one should find Δu_{it} to be serially correlated at the first order, but not at the second or higher orders. In other words, the autocovariance at order zero should be $E[(\Delta u_{it})^2] = (1 + \theta^2)\sigma_v^2$, the autocovariance at the first order $E(\Delta u_{it}\Delta u_{it-1}) = -\theta\sigma_v^2$, and the autocovariances $E(\Delta u_{it}\Delta u_{it-j}) = 0$ for all $j \geq 2$. The corresponding autocorrelation coefficients (the ratio between the autocovariance at a given order and that at order zero) should be 1, $-\frac{\theta}{(1+\theta^2)}$, and 0, respectively.

We find that the estimated autocovariances and autocorrelations are not inconsistent with the income process in equation (3), that is, that the stochastic component of the income process is well described by the sum of a random walk permanent component and a serially uncorrelated transitory shock. Recall that, because of the sample design of the SHIW data, we can only construct the covariance matrix for two-years-apart income residuals, $u_{it} - u_{it-2} = v_{it} + (1-\theta)v_{it-1} - \theta v_{it-2}$. To check the consistency of the estimated income process with the model in equation (3), note that the model implies the following restrictions on the covariance matrix of the first difference of the income residuals in level:

$$E[(u_{it} - u_{it-2})^2] = 2(1 - \theta + \theta^2)\sigma_v^2$$

$$E[(u_{it} - u_{it-2})(u_{it-2} - u_{it-4})] = -\theta\sigma_v^2$$

$$E[(u_{it} - u_{it-2})(u_{it-j} - u_{it-j-2})] = 0 \text{ for all } j \geq 4$$

The corresponding autocorrelation coefficients are 1, $-\frac{\theta}{2(1-\theta+\theta^2)}$, and 0, respectively. The restrictions above are not rejected by the data. We find that the first-order autocorrelation coefficient (an estimate of $-\frac{\theta}{2(1-\theta+\theta^2)}$) is negative (−0.26) and statistically significant (a t-statistic of −3.2). In contrast, higher-order autocorrelation coefficients are not statistically different from zero: the second-order autocorrelation is −0.09 (with an insignificant t-statistic of −1.2), and the third-order autocorrelation −0.10 (with an insignificant t-statistic of −1).[10]

IV. The Empirical Distribution of the Income Shocks

Table 1 displays the cross-sectional distribution of income shocks for the sample that includes heads or spouses (1,102 households).[11] Income shocks are divided by current earnings. Because we have available only the sum of permanent shocks in 1990 and 1991, the figures in the first column are the ratio of average permanent shock between

[9] Covariances can be estimated by equally weighted minimum distance or optimal minimum distance. As shown by Altonji and Segal (1996), the latter can produce inconsistent estimates in finite samples, so we adopt the former.

[10] If the process is estimated in logs rather than levels, the size and the statistical significance of autocorrelation coefficients are very similar.

[11] For 95% of our sample, we use information directly pertaining to the head of the household.

TABLE 1.—THE DISTRIBUTION OF INCOME SHOCKS

	Average Permanent Shock	Transitory Shock in 1991	Transitory Shock in 1989
	Earnings in 1991	Earnings in 1991	Earnings in 1989
Age in 1991			
≤35	−0.0193	0.0133	−0.0117
	(0.0153)	(0.0046)	(0.0043)
35–55	−0.0345	0.0080	−0.0087
	(0.0108)	(0.0034)	(0.0023)
>55	0.0060	−0.0052	−0.0182
	(0.0098)	(0.0044)	(0.0035)
Education			
Compulsory	−0.0031	−0.0002	−0.0118
	(0.0061)	(0.0030)	(0.0025)
High school	−0.0393	0.0103	−0.0112
	(0.0112)	(0.0047)	(0.0028)
University	−0.0334	0.0093	−0.0232
	(0.0189)	(0.0077)	(0.0062)
Region			
North	−0.0258	0.0104	−0.0099
	(0.0101)	(0.0043)	(0.0029)
South	−0.0243	−0.0038	−0.0114
	(0.0113)	(0.0029)	(0.0025)
Occupation			
Employed	−0.0177	−0.0004	−0.0077
	(0.0107)	(0.0036)	(0.0021)
Self-employed	−0.0665	0.0101	−0.0059
	(0.0235)	(0.0076)	(0.0057)
Sector			
Private	−0.0190	0.0005	−0.0063
	(0.0092)	(0.0030)	(0.0029)
Public	−0.0209	0.0122	−0.0089
	(0.0098)	(0.0042)	(0.0030)
Family income			
1^{st} quartile	−0.0325	−0.0137	−0.0171
	(0.0162)	(0.0052)	(0.0036)
2^{nd} quartile	−0.0182	0.0049	−0.0088
	(0.0131)	(0.0035)	(0.0041)
3^{rd} quartile	−0.0100	0.0103	−0.0132
	(0.0112)	(0.0054)	(0.0032)
4^{th} quartile	−0.0077	0.0148	−0.0115
	(0.0130)	(0.0047)	(0.0036)
Whole sample	−0.0171	0.0039	−0.0127
	(0.0068)	(0.0024)	(0.0018)

Standard deviations are reported in parenthesis.

1990 and 1991 and earnings in 1991. The other two columns focus on the relative transitory shocks in 1991 and 1989, respectively.

In 1991, average earnings feature a negative innovation of approximately 1.4% in real terms; the decomposition into transitory and permanent shocks, however, reveals that, although the permanent component plays a negative role (−1.7% on average), the transitory shock is positive (+0.4% on average).

With the exception of the elderly, permanent shocks are on average negative for all population groups; however, the effect is stronger for the self-employed, the middle aged, those with more than compulsory education, and the poor (as measured by family income quartiles). As for the transitory shocks in 1991, these are higher for the young, the self-employed, the public employees, and those living in the north. Note that, unlike other groups, the poor face a negative transitory innovation in 1991. Finally, the transitory earnings shock in 1989 is negative (−1.3% vis-à-vis + 0.4% in 1991), and it is particularly strong for those approaching retirement and those with a college degree.

An average permanent shock of −1.7% is not negligible. There is evidence (Miniaci & Weber, 1999; Bertola & Ichino, 1996) showing that in the early 1990s Italian households perceived a negative permanent change in their lifetime income. This was due to various reasons: radical political change, pay freezing in the public sector that spread to the private sector through income policy experiments, increasing taxation aimed at meeting the Maastricht Treaty criteria, and pension and labor market reforms.

In particular, in 1991 the wage indexation clause (scala mobile) was abolished, and the laws regulating the hiring process were dramatically renewed with the aim of relaxing labor market regulations. It has been argued that the former had the effect of increasing earnings inequality after decades of compression in the earnings distribution, whereas the latter had the effect of increasing earnings uncertainty because of greater job instability (Bertola & Ichino, 1996).

The income policy experiments were introduced as a transitory measure aimed at freezing pay rise after years of unnecessary adjustments; ex post, some of these measures seem to have permanently reduced wages purchasing power.

In our context, pension reforms can be important to an extent that depends on how much the prospective income power of those who are currently working is affected. Due to the unprecedented imbalance between contributors and beneficiaries in the Italian pay-as-you-go social security system, both the Amato and the Dini reforms (the two main reforms implemented in the early 1990s, named after the prime ministers who signed them) went in the direction of cutting future benefits and increasing contributions.

Finally, labor market and pension reforms were accompanied by an increase in taxation. The self-employed are likely to have suffered more from the introduction of new fiscal measures. Although a privileged category because of the possibility of evading taxes more easily than the employed, the self-employed were hit by the introduction of a minimum tax, which based tax payments on the presumption of a minimum annual income. The radical change in political attitudes towards tax noncompliance and the introduction of stricter measures for tax enforcement might have contributed to strengthen the perception of a decline in the permanent income for this group.

A final remark is that one observes only a snapshot of the distribution of earnings shocks in 1989 and 1991; a thorough analysis of how people form and change their expectations in the face of idiosyncratic and aggregate events would require a longer period of observations, which would ease the task of disentangling life cycle from business cycle related shocks. Unfortunately, subjective expectations are rarely asked in survey data, and, in the case of the SHIW, they were asked in the format used in this paper only in 1989 and 1991.

TABLE 2.—THE "SAVING FOR A RAINY DAY" EQUATION

	(1)	(2)	(3)
Transitory shock$_t$	0.6384	1.2165	1.2201
	(0.2590)	(0.2979)	(0.2977)
Permanent shock$_t$		0.1662	0.1613
		(0.0743)	(0.0740)
$\frac{\text{Conditional variance}_t}{1000}$			0.9866
			(0.3676)
Age	211.30	219.41	212.58
	(114.35)	(148.23)	(147.95)
Age^2	−2.49	−2.40	−2.29
	(1.07)	(1.38)	(1.37)
Family size	2253.68	2435.85	2441.49
	(300.49)	(361.72)	(358.13)
Children aged 0–5	−3469.48	−3517.38	−3430.82
	(646.32)	(956.08)	(954.64)
Children aged 6–13	−3023.84	−2748.13	−2711.16
	(469.08)	(647.16)	(643.42)
Children aged 14–17	−3521.75	−3014.45	−2978.25
	(639.92)	(890.71)	(887.80)
# of observations	2,204	1,102	1,102
R^2	0.0560	0.0986	0.1039

The sample includes 1,044 heads (94.74% of the sample) and 58 (5.26%) spouses. Saving is defined as the difference between family disposable income (including asset income) and nondurable consumption. Standard errors robust to heteroskedasticity of unknown form are reported in parenthesis. All regressions include a constant.

V. Empirical Results

A. The Saving Equation

In this section, we present the results of estimating the saving equation (9). We check the robustness of our main findings and discuss the issue of measurement error in subsections B and C, respectively.

Saving is defined as the difference between family disposable income (including income from assets) and nondurable consumption. This definition implies that durables are a form of capital that produces a service flow. In this case, the service flow should be included in asset income and added to consumption, thus canceling out of the definition of saving we adopt here. Descriptive statistics show that average saving is about 10.9 million lire in 1989 (with a standard deviation of 12.2 million lire) and 11.8 million lire in 1991 (with a standard deviation of 11.6 million lire). In both years, the distribution of saving is positively skewed (sample medians are about 7.2 and 8.6 million lire, respectively, in 1989 and 1991), with few households reporting negative savings (4.72% in 1989 and 6.72% in 1991). Now we turn to the discussion of our regression analysis.

Table 2 presents the results of estimating the saving equation (9) for the sample of heads and spouses (1,102 households). Three basic regressions are estimated: (i) the one strictly implied by the PIH with only the transitory income shock as an additional explanatory variable, and then including (ii) the permanent income shock, and (iii) the conditional variance of income. The latter can be derived from subjective expectation data. (see appendix A.) Note that the OLS regressions for specification (i) can be estimated for both 1989 and 1991 as it does not involve lagged variables; thus, in this case, the sample size is twice as large as the one for specifications (ii) and (iii). OLS estimates for the three models above are presented in columns 1 through 3 of table 2. Standard errors are robust to the presence of heteroskedasticity of unknown form. We control for preference heterogeneity by assuming that the bliss point of household utility is a function of age, age squared, family size, and the number of children in three age bands (0–5, 6–13, 14–17).

The results reported in column 1 support the permanent income hypothesis with rational expectations: savings react strongly to transitory income shocks (a point estimate of 0.64). The null hypothesis that saving reacts on a one-for-one basis to transitory earnings shocks has a p-value of 14%. The null hypothesis that the coefficient on the transitory shock equals $(1 + r)^{-1}$ can be tested by considering a grid of possible values for the real interest rate ranging from 0% to 10%;[12] the null hypothesis is never rejected. As for the effect of demographics, family size increases saving, and the presence of children in any age band reduces it, the effect being stronger for infants and adolescents.

As discussed in subsection IIIA, the biennial nature of the SHIW data prevents the identification of the permanent shock ζ_{it}, while the sum $(\zeta_{it} + \zeta_{it-1})$ is identified by equation (10). In column 2 of table 2, we thus add to the main specification the sum of the permanent income shocks in periods t and $t - 1$. The coefficient attached to the latter is an estimate of the marginal propensity to save out of a permanent shock. The results show that the null hypothesis that the latter is zero can be rejected: permanent income shocks are significant predictors of household savings. Taken at face value, these results suggest that households save not only the transitory income shocks, but also a

[12] The average real interest rates in 1991 were 0.58% (deposits), 5.58% (Treasury bonds), and 4.32% (other assets, including shares). Interest rates in 1989 were very similar to those for 1991.

TABLE 3.—THE "SAVING FOR A RAINY DAY" EQUATION: SENSITIVITY ANALYSIS

	Excluding Those Aged over 65 (1)	Excluding the Self-employed (2)	Excluding Multiple-earner Households (3)	Including Durables (4)
Transitory shock$_t$	1.1245	0.9474	0.9899	1.0575
	(0.3205)	(0.3642)	(0.3576)	(0.2630)
Permanent shock$_t$	0.1643	0.1941	0.1682	0.1739
	(0.0799)	(0.0709)	(0.0757)	(0.0722)
$\frac{\text{Conditional variance}_t}{1000}$	0.9948	0.7938	0.7659	0.7945
	(0.4174)	(0.3517)	(0.3867)	(0.4753)
Age	308.96	263.50	168.77	182.03
	(309.98)	(151.09)	(140.91)	(148.13)
Age^2	−3.41	−2.83	−1.91	−1.49
	(3.41)	(1.38)	(1.33)	(1.37)
Family size	2169.64	2825.75	−226.82	1963.58
	(396.71)	(369.88)	(356.97)	(354.66)
Children aged 0–5	−3251.97	−3367.53	−965.47	−2534.68
	(982.45)	(985.01)	(1004.65)	(1025.01)
Children aged 6–13	−2408.48	−2775.69	−146.93	−2162.19
	(689.32)	(712.85)	(609.04)	(644.59)
Children aged 14–17	−2780.64	−3829.58	525.40	−3205.93
	(923.34)	(963.15)	(909.67)	(920.02)
# of observations	911	903	494	1,102
R^2	0.0879	0.1130	0.0875	0.0786

In columns 1 through 3, saving is defined as the difference between family disposable income (including asset income) and nondurable consumption; in column 4, as the difference between family disposable income and total consumption. Standard errors robust to heteroskedasticity of unknown form are reported in brackets. All regressions include a constant.

sizeable portion of the permanent shocks.[13] Therefore, the certainty equivalence model seems to fail in the sense of predicting a lower standard deviation of saving than we see in the data.[14]

In column 3, we test for precautionary saving including the conditional variance of head's earnings[15] alongside the transitory shock and the permanent shock. The version of the permanent income hypothesis we have tested so far might fail because preferences are not quadratic. If individual preferences admit a positive third derivative (for instance, if consumers are prudent in the sense clarified by Kimball (1992)), then the estimates of the "saving for a rainy day" equation are inconsistent because of the omission of higher moments of the distribution of income shocks that are likely to be correlated with the realization of the shocks. The test we conduct is simple. Under the null of the permanent income hypothesis with quadratic preferences, higher moments of the distribution of earnings should not matter. The hypothesis is firmly rejected: the conditional variance of earnings has the expected sign (more uncertainty increases current saving) and it is statistically significant, thus suggesting that the assumption of quadratic preferences is inappropriate. This conclusion is supported by previous empirical evidence available for Italy (Guiso et al., 1992; Jappelli & Pistaferri, 2000).

B. *Sensitivity Analysis*

We performed several tests to check the robustness of the results. Here we briefly comment on sample selection, the definition of saving, the relevant time horizon for saving choices, and the evidence for consumption growth.

We experiment by excluding the elderly (those aged more than 65, the standard retirement age for males), the self-employed, and multiple-earner households. The reason to exclude the elderly is that the decomposition of income shocks between a transitory and a permanent shock is possibly no longer valid for the retired or those approaching retirement. The reason for excluding the self-employed is that, as reported by Brandolini and Cannari (1994), they tend to understate or misreport their current earnings; moreover, for this group is more difficult to separate labor income from asset income. Finally, given that we rely only on the subjective expectations of the head of the household to identify both the shocks to income and the conditional variance, it is worth assessing whether the results change once multiple-earner households are removed from the sample.

[13] Our results are fairly close to Paxson (1992). Using a sample of Thai households, she finds that savings respond both to transitory and permanent shocks but not to regional rainfall variability (which she interprets as a measure of income uncertainty). Also, in her case, the response to transitory shocks is much higher than that to permanent shocks.

[14] Cambell (1987) notices that, at least in principle, the PIH restricts the average level of saving along with its variability. In particular, the PIH implies that the average level of saving should be $-r^{-1}$ times the average income change. This restriction is derived under the assumption that $T \to \infty$, and is therefore hard to test with household panel data, where one exploits for consistency a large cross-sectional dimension with a fixed time dimension.

[15] This is defined as var $(y_{it+1} - y_{it}|\Omega_{it}) =$ var $(y_{it+1}|\Omega_{it})$. The variance term is obtained from the conditional variance of the rate of growth of future earnings by noting that

$$\text{var}\left[\left(\frac{y_{it+1}-y_{it}}{y_{it}}\right)\bigg|\,\Omega_{it}\right] = y_{it}^{-2}\ \text{var}\ (y_{it+1} - y_{it}|\Omega_{it}).$$

Note that we cannot distinguish between the variance of the transitory shock and the variance of the permanent shock.

The results obtained from excluding these groups are presented in the first three columns of table 3. It is worth noting that the magnitude of the various effects is not much affected by such exclusions; on the other hand, the precision of the estimates (in particular the effect of demographics) is slightly attenuated if either population group is excluded.[16]

As noted in subsection VA our definition of saving excludes durable consumption. If the PIH applies to total consumption rather than nondurables, our measure of saving equals the true measure plus an error, the fraction of nondurables over total consumption times total consumption. Because the error is not orthogonal to the true measure of saving, there would be a bias in the OLS estimates. In theory, the bias depends on two factors: the fraction of nondurables over total consumption (the smaller the fraction the larger the OLS bias), and the correlation between total consumption and the true measure of saving. If the two are negatively correlated, there is a standard attenuation bias; if they are positively correlated, there is an upward bias. Empirically, however, the OLS regressions for the two definitions of savings are very similar, as shown in column 4 of table 3.

The textbook version of the PIH tested so far is based on the assumption of an infinite horizon; this can be justified by the presence of an operative bequest motive. If instead the horizon is finite (for instance, if individuals live T periods with certainty) the saving equation (8), omitting for simplicity age and other demographic effects, can be written as

$$s_{it} = \left\{1 - \frac{r}{1+r}\left[1 - \frac{1}{(1+r)^{T-t+1}}\right]^{-1}\right\}\epsilon_{it},$$

where the term in brackets is an annuitization factor. In this case, the effect of the transitory shock is declining with age. (In the last period of life, all shocks are alike and thus consumed.) Because the optimal rule for the permanent shock is invariant to the horizon length, this is still entirely consumed, not saved or dissaved. To test this variant of the PIH, we interact the transitory shock with the annuitization factor $\left\{1 - \frac{r}{1+r}\left[1 - \frac{1}{(1+r)^{T-t+1}}\right]^{-1}\right\}$ (assuming $r = 0.02$, $T = 100$, and t as the age of the respondent). The coefficient on such interaction is 1.26 with a standard error of 0.31; the coefficients on the other variables (including the permanent shock and the conditional variance) are very similar to those presented in table 2.[17] Thus, the results do not seem to depend on the assumption about the horizon length.

[16] Unobserved heterogeneity in savings may be a further cause of concern, in particular if it is correlated with shocks to income. A standard procedure to eliminate fixed unobserved heterogeneity is to first-difference the data. However, given the peculiar nature of the SHIW data, we would be able to test the effect of only the transitory shock on saving, because we do not have two observations on the permanent shock. As a result, the test implied by equation (9) could not be implemented.

[17] Experimenting with different values for T or r produces similar results.

TABLE 4.—THE CONSUMPTION CHANGE EQUATION

Transitory shock$_t$	0.0874
	(0.0812)
Permanent shock$_t$	0.5746
	(0.1089)
$\frac{\text{Conditional variance}_t}{1000}$	0.1421
	(0.3188)
Age	23.02
	(18.40)
ΔFamily size	3603.68
	(699.75)
ΔChildren aged 0–5	−466.46
	(969.35)
ΔChildren aged 6–13	−157.74
	(814.94)
ΔChildren aged 14–17	−227.77
	(771.13)
# of observations	1,102
R^2	0.0828

The sample includes 1,044 heads (94.74% of the sample) and 58 (5.26%) spouses. The dependent variable is the change in nondurable consumption. Standard errors robust to heteroskedasticity of unknown form are reported in parenthesis. The regression includes a constant.

As a final check of internal consistency, we look at nondurable consumption data. In equation (9), the coefficients β_1 and β_2 have in fact a structural interpretation: they are the propensity to save out of a transitory shock and the propensity to save out of a permanent shock, respectively. The theory predicts $\beta_1 = (1 + r)^{-1}$ and $\beta_2 = 0$. The same parameters can be estimated in a consumption change equation (that is, the Euler equation for consumption), which predicts, omitting for simplicity age and other demographic effects

$$\begin{aligned}\Delta c_{it} &= \frac{r}{1+r}\epsilon_{it} + \zeta_{it} \\ &= (1-\beta_1)\epsilon_{it} + (1-\beta_2)\zeta_{it}.\end{aligned} \tag{11}$$

The results of estimating equation (11) are reported in table 4, where we also control for changes in demographics and the subjective variance. Given the biennial nature of the SHIW data, our estimating equation is (omitting again the demographics and the conditional variance term)

$$\begin{aligned}c_{it} - c_{it-2} = {}& (1-\beta_1)(\epsilon_{it} + \epsilon_{it-1}) \\ &+ (1-\beta_2)(\zeta_{it} + \zeta_{it-1}) + \xi_{it}\end{aligned}$$

where ξ_{it} is an error term.[18] We observe ϵ_{it} and $(\zeta_{it} + \zeta_{it-1})$, but do not observe ϵ_{it-1}, which is therefore subsumed in the error term. The results are in line with those in table 2. The transitory shock does not impact consumption because it only impacts savings from the analysis above. The coefficient is small in magnitude (0.09) and statistically insignificant (a standard error of 0.08), implying that one would not

[18] As noted in subsection VA, our definition of nondurable consumption should include the service flow of durables. To avoid biases arising from the omission of the latter (on which we lack information), we need to assume that the service flow of durables is constant over time, or alternatively that its change is orthogonal to the left-hand side of equation (11).

reject the theoretically consistent null hypothesis that $(1 - \beta_1) = 0$. As for the permanent shock, the estimated effect is statistically significantly below 1. The variance has the expected sign, but is less precisely estimated than in the saving equation. Other demographics, with the exception of family size, are also less precisely measured.[19]

C. Measurement Error

If the variables we use to measure the transitory and the permanent shock are subject to measurement error, OLS estimates will be invalid. A way to tackle this issue would be to find suitable instruments for the income shocks, a strategy that we find difficult to implement given the available data. Absent a good empirical strategy to confront this problem, here we limit ourselves to a discussion of the bias we are likely to face and the related theoretical implications.

In what follows, we neglect measurement error in subjective expectations. If these were reported with error, one should find β_1 to be biased towards zero. However, from the estimation results presented in column 1 of table 2, one can infer that such attenuation bias is not serious.

We consider two sources of measurement error bias. First, income may be measured with error. In this case, the OLS estimates will be biased because the dependent variable (saving, defined as disposable income minus consumption) and the permanent shock (defined as in equation (6), and therefor depending on the head's earnings) may have measurement errors that are not independent. The consequence is that the estimate of β_2 in equation (9) will be upward biased (a nonstandard effect) because of the positive correlation between the measurement error in disposable family income and the measurement error in the head's earnings. In other words, one may find β_2 higher than predicted by the PIH. This problem is not solved by estimating the parameters β_1 and β_2 from the Euler equation (11). In fact, even though the measurement error in consumption may be independent on the measurement error in the head's earnings, the permanent shock variable may still be measured with error (because earnings are). This means that the estimate of $(1 - \beta_2)$ in equation (11) will be biased towards zero, a standard attenuation effect.

The second source of bias is due to the fact that, with biennial data, the right-hand side of equation (10) measures the permanent shock ζ_{it} with an error, ζ_{it-1}. Under the null of the PIH, this measurement error does not matter because $\beta_2 = 0$. Under the alternative, however, a downward bias arises.

To understand how the two sources of measurement error bias impact our estimates, recall that $s_t = Y_t - c_t$, with disposable income Y_t being the sum of labor income y_t and asset income k_t. Assume also that labor income is measured with error, $y_t^* = y_t + \eta_t$, so that $s_t^* = s_t + \eta_t$ (the subscript for the individual is omitted for simplicity). Our saving regression is

$$\begin{aligned} s_t^* &= X\delta + \beta_2\zeta_t + \eta_t \\ &= X\delta + \beta_2(\zeta_t + \zeta_{t-1}) + [\eta_t - \beta_2\zeta_{t-1}], \end{aligned} \tag{12}$$

where X contains the constant term, the demographics, the transitory shock and the conditional variance term, and the second equality takes into account the fact that we approximate ζ_t with $(\zeta_t + \zeta_{t-1})$. Rewrite equation (12) as

$$\begin{aligned} s_t^* &= X\delta + \beta_2(y_t - y_{t-2} - E_{t-2} + E_t) + [\eta_t - \beta_2\zeta_{t-1}] \\ &= X\delta + \beta_2(y_t^* - y_{t-2}^* - E_{t-2} + E_t) \\ &\quad + [(1 - \beta_2)\eta_t + \beta_2\eta_{t-2} - \beta_2\zeta_{t-1}] \end{aligned}$$

where $(y_t - y_{t-2} - E_{t-2} + E_t) = (\zeta_t + \zeta_{t-1})$ from equation (10) and E_{t-2} and E_t are the subjective expectations of future earnings. The expression in square bracket is the error term of the OLS regression.

Assume that X is measured without error and that it is orthogonal to the permanent shock variable. Although this is admittedly unrealistic, examining this very simple case will give us at least a feeling of the problems involved. Finally, assume that ζ and η are serially and mutually uncorrelated i.i.d. processes with mean zero and variances σ_ζ^2 and σ_η^2, respectively.

The probability limit of the OLS estimator of β_2 is

$$p \lim \hat{\beta}_2 = \beta_2\left[\frac{\sigma_\zeta^2}{2(\sigma_\zeta^2 + \sigma_\eta^2)}\right] + \frac{\sigma_\eta^2}{2(\sigma_\zeta^2 + \sigma_\eta^2)}, \tag{13}$$

where we use the fact that $(y_t^* - y_{t-2}^* - E_{t-2} + E_t) = (\zeta_t + \zeta_{t-1} + \eta_t - \eta_{t-2})$.

If the PIH is true ($\beta_2 = 0$) this expression converges to $\frac{\sigma_\eta^2}{2(\sigma_\zeta^2 + \sigma_\eta^2)} > 0$. Thus, under the null of the PIH, $\hat{\beta}_2$ is upward biased. Empirically, we find that $\hat{\beta}_2$ is about 0.2, a result that is consistent with the PIH ($\beta_2 = 0$) and the presence of the two sources of measurement error bias considered here. For realistic values of $\frac{\sigma_\eta^2}{\sigma_\zeta^2 + \sigma_\eta^2}$ (around 0.4, say), the PIH holds approximately true ($\beta_2 \simeq 0$). It is easy to check that higher values of $\frac{\sigma_\eta^2}{\sigma_\zeta^2 + \sigma_\eta^2}$ imply $\beta_2 < 0$, a result inconsistent with the PIH but also with alternative models of consumption choice. Thus, one may justify an estimated marginal propensity to save out of a permanent shock strictly greater than zero just by appealing to measurement error. However, although the presence of measurement error can explain this finding, it cannot explain why the conditional variance of earnings is a statistically significant determinant of household saving.

Both findings can instead be made consistent with models of consumption that are alternative to the PIH irrespective

[19] Note that the demographics that enter the saving equation in level format must enter the consumption change equation in first-difference format.

of measurement error. Deaton (1991) presents simulation results for prudent households facing constraints on net wealth. He shows that, while a considerable amount of smoothing takes place, there are downward spikes in consumption when the constraint binds. In this model, saving reacts positively to both permanent shocks and income risk, a prediction that is consistent with our empirical results. Carroll (2000) refines Deaton's insight along the following lines. He shows by means of simulations that a marginal propensity to save out of a permanent shock strictly greater than zero and a positive reaction of saving to income uncertainty are consistent with a buffer-stock model in which prudent and impatient individuals face both transitory and permanent shocks. The reason is that, in this model, consumers attempt to maintain a target wealth-to-permanent income ratio. A positive permanent shock to the level of permanent income reduces this ratio, thus inducing the consumer to increase saving to reachieve the target. A negative permanent shock works exactly in the opposite direction. Additionally, prudence prompts precautionary saving. In Carroll's simulations, in particular, the marginal propensity to save out of a permanent shock ranges between 0 and 0.15.

VI. Conclusions

This paper has presented tests of the permanent income hypothesis with quadratic preferences. We have shown that the availability of subjective income expectations and income realizations allows the exact identification of transitory and permanent income shocks if data are available for at least two consecutive time periods. Subjective income expectations are then used to test the hypothesis that households save only in response to transitory shocks, a prediction that derives from the assumption that the stochastic component of income is the sum of a random-walk permanent component and a serially uncorrelated transitory shock. Given that the saving innovation is explicitly controlled for, Chamberlain's critique does not apply in our context; estimates are thus consistent even in a short panel.

The empirical analysis shows that, although the reaction of savings to transitory income shocks agrees with the predictions of the theory (a coefficient of one), the arrival of a permanent shock prompts saving as well (although the effect is tinier), contrary to the theory that predicts no effect. In addition, we have shown that the assumption of quadratic preferences is inappropriate: higher moments of the distribution of earnings should not matter, but they do. This finding is supported by previous evidence, is in agreement with the existence of a precautionary motive for saving, and is consistent with the theoretical lack of plausibility of the assumption of increasing risk aversion implied by quadratic preferences.

The evidence that savings react to permanent shock can have at least two explanations. One possibility is that the PIH is a good description of the data, but measurement errors bias the estimated coefficient away from zero. There is, however, also a strictly theoretical explanation in which the marginal propensity to save out of a permanent shock is greater than zero. Buffer stock behavior of the type described by Deaton (1991) and Carroll (2000), in which consumers are both impatient and prudent, implies that savings responds not only to transitory shocks but also to permanent shocks, although to a much lower extent. Given the evidence in favor of a precautionary motive for saving, the results obtained in this paper are probably best read as the combination of measurement error problems and theoretical amendments to the standard certainty equivalence framework.

REFERENCES

Alessie, Rob, and Annamaria Lusardi, "Saving and Income Smoothing: Evidence from Panel Data," *European Economic Review* 41(7) (1997), 1251–1279.

Altonji, Joseph G., and Lewis M. Segal, "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business and Economic Statistics* 14(3) (1996), 353–366.

Bertola, Giuseppe, and Andrea Ichino, "Crossing the River: A Comparative Perspective of Italian Employment Dynamics," *Economic Policy* 21 (1996), 361–420.

Blundell, Richard W., and Ian Preston, "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics* 113(2) (1998), 603–640.

Brandolini, Andrea, and Luigi Cannari, "Methodological Appendix" (pp. 282–299), in Albert Ando, Luigi Guiso, and Ignazio Visco (Eds.), *Saving and the Accumulation of Wealth* (Cambridge: Cambridge University Press, 1994).

Campbell, John Y., "Does Saving Anticipates Declining Labor Income? An alternative test of the permanent income hypothesis," *Econometrica* 55(6) (1987), 1249–1273.

Carroll, Cristopher D., "Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income," Johns Hopkins University *mimeograph (2000).

Chamberlain, Gary, "Panel Data," (pp. 1247–1318). in Zvi Griliches and Michael D. Intriligator (Eds.), *Handbook of Econometrics,* vol. 2, (Amsterdam: North-Holland, 1984).

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica* 59(5) (1991), 1221–1248.

———, "Saving and Income Smoothing in the Côte d'Ivoire," *Journal of African Economics* 1(1) (1992), 1–24.

Flavin, Marjorie, "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy* 89(5) (1981), 974–1009.

———, "The Excess Smoothness of Consumption: Identification and Interpretation," *Review of Economic Studies* 60(3) (1993), 651–666.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, "Earnings Uncertainty and Precautionary Saving," *Journal of Monetary Economics* 30(2) (1992), 307–337.

Hall, Robert E., and Frederick S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data of Households," *Econometrica* 50(2) (1982), 261–281.

Jappelli, Tullio, and Luigi Pistaferri, "Using Subjective Income Expectations to Test for the Excess Sensitivity of Consumption to Predicted Income Growth," *European Economic Review* 44(2) (2000), 337–358.

Kimball, Miles, "Precautionary Saving in the Small and in the Large," *Econometrica* 58(1) (1992), 53–73.

Miniaci, Raffaele, and Guglielmo Weber, "The Italian Recession of 1993: Aggregate Implications of Microeconomic Evidence," this REVIEW 81(2) (1999), 237–249.

Muth, J. F., "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association* 55(290) (1960), 299–306.

Paxson, Christina H., "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," *American Economic Review* 1 (1992), 15–33.

Pistaferri, Luigi, "Superior Information, Income Shocks and the Permanent Income Hypothesis," UCL Discussion Papers in *Economics,* 5 (1998), 1–31.

APPENDIX

Survey Questions on Income Prospects

In both 1989 and 1991, each labor income and pension recipient interviewed SHIW was asked the following two questions.

Inflation expectations: "On this table [a table is shown to the respondent] we have indicated some classes of inflation. We are interested in knowing your opinion about inflation twelve months from now. Suppose that you have 100 points to be distributed between these intervals. Are there intervals you definitely exclude? Assign zero points to these intervals. How many points do you assign to each of the remaining intervals?" For this and the following question, the intervals on the table shown to the person interviewed are: less than 0, 0–3, 3–5, 5–6, 6–7, 7–8, 8–10, 10–13, 13–15, 15–20, 20–25, >25%. If the response is "less than 0," the person is asked: "How much less than zero? How many points would you assign to this class?"

Income expectations: "We are also interested in knowing your opinion about your labor earnings or pensions twelve months from now. Suppose that you have 100 points to be distributed between these intervals. [A table is shown again.] Are there intervals you definitely exclude? Assign zero points to these intervals. How many points do you assign to each of the remaining intervals?".

To construct subjective expectations and variances of the variable of interest (either the rate of growth of nominal earnings or the rate of inflation), we set the upper bound of the distribution—the open interval—at 35%. Let x_{it} be the variable of interest. The subjective expectation of x_{it} at time $t - 1$ is then given by

$$E(x_{it}|\Omega_{it-1}) = \sum_{k=1}^{K} [\Pr(x_{k-1} \leq x_{it} \leq x_k)]2^{-1}(x_k + x_{k+1})$$

and the subjective variance by

$$\mathrm{Var}(x_{it}|\Omega_{it-1}) = \sum_{k=1}^{K} [\Pr(x_{k-1} \leq x_{it} \leq x_k)] \times [2^{-1}(x_k + x_{k+1}) - E(x_{it}|\Omega_{it-1})]^2,$$

where x_K and x_0 are, respectively, the upper and the lower bound of the distribution. Note that the intervals are not of the same size. More precisely, $x_0 = 0$ for those assigning zero probability to a negative earnings growth event; otherwise, it is a value chosen by the respondent; $x_1 = 0.03$; $x_2 = 0.05$; $x_3 = 0.06$; $x_4 = 0.07$; $x_5 = 0.08$; $x_6 = 0.1$; $x_7 = 0.13$; $x_8 = 0.15$; $x_9 = 0.2$; $x_{10} = 0.25$; $x_K = x_{11} = 0.35$. Because we do not attempt to estimate the variance within each interval, the conditional variance Var $(x_{it}|\Omega_{it-1})$ is equal to zero for those reporting point expectations.

People report one-year-ahead expectations referring to the rate of growth of their earnings; to obtain the one-year-ahead expectations of changes in earnings that would identify the transitory earnings shocks, we simply solve for the expected change in earnings. Given the assumptions on the timing of the expectations, the computation of the latter is simple. First, we define $E(Gy_{it,t+1}|\Omega_{it})$, the expected growth rate of income between year t and year $t + 1$ (where $t = 1989$ or 1991). This rate is defined in real terms. Then, we derive $E(\Delta y_{it+1}|\Omega_{it}) = y_{it}E(Gy_{it,t+1}|\Omega_{it})$.